

A Comparative Study of Bayesian Portfolio Optimization: Evidence from U.S. Large-Cap AI-Related Stocks

Mironov Mikhailm, Master's Student

Shanghai University, SILC Business School, China

[Doi:10.19044/esj.2026.v22n7p31](https://doi.org/10.19044/esj.2026.v22n7p31)

Submitted: 19 January 2026

Accepted: 03 March 2026

Published: 31 March 2026

Copyright 2026 Author(s)

Under Creative Commons CC-BY 4.0

OPEN ACCESS

Cite As:

Mikhailm, M. (2026). *A Comparative Study of Bayesian Portfolio Optimization: Evidence from U.S. Large-Cap AI-Related Stocks*. European Scientific Journal, ESJ, 22 (7), 31.

<https://doi.org/10.19044/esj.2026.v22n7p31>

Abstract

This paper conducts a comparative analysis of portfolio optimization methods, focusing on Bayesian approaches, applied to U.S. AI-related stocks (2020–2025). While the classical Markowitz model relies on fixed estimates of return and risk, the Bayesian framework incorporates parameter uncertainty through priors on expected returns, soft portfolio concentration constraints, and weights parameterized via Dirichlet or softmax transformations. Posterior inference is conducted using Hamiltonian Monte Carlo with the No-U-Turn Sampler (NUTS), allowing more adaptive and probabilistically informed decision-making. Portfolio performance is evaluated using risk-adjusted returns, measured by the Sharpe ratio, and supplemented with conditional volatility and beta dynamics via the conditional CAPM and DCC-GARCH framework. In the reported experiment, the Markowitz model achieved the highest Sharpe ratio (0.049 in-sample; 0.089 out-of-sample), but this result is limited by the narrow stock universe, daily frequency, zero risk-free rate assumption, and the exclusion of transaction costs. Advanced Bayesian models showed improved risk-adjusted performance relative to early Bayesian specifications, reaching up to 0.038 in-sample and 0.084 out-of-sample Sharpe ratios, while simultaneously reducing conditional volatility and market beta.

Keywords: Bayesian portfolio optimization; Markowitz Mean-Variance Optimization; AI Stocks; Conditional beta; DCC GARCH; Parameter uncertainty; Posterior inference; Regularization; Portfolio concentration

Introduction

1.1 Background

Artificial intelligence (AI) has rapidly transformed industries and financial markets, enabling new tools for portfolio management and asset pricing. In particular, AI-related companies have exhibited high growth and volatility, posing challenges for traditional portfolio optimization methods. Classical approaches, such as Markowitz Mean-Variance Optimization (MVO), often suffer from estimation errors, unstable allocations, and concentrated portfolios when applied to small or volatile stock universes.

Bayesian portfolio optimization offers a potential solution by incorporating parameter uncertainty, regularizing extreme weights, and allowing soft constraints on portfolio concentration. These features can improve out-of-sample performance, reduce portfolio volatility, and mitigate excessive exposure to single assets.

This study compares classical MVO with a series of Bayesian portfolio optimization models applied to a portfolio of 10 large-cap U.S. AI-related stocks using daily returns. In addition to evaluating risk-adjusted performance via the Sharpe ratio, the research examines conditional volatility and market risk exposure using the conditional CAPM and DCC-GARCH framework. By explicitly modeling parameter uncertainty and penalizing concentrated allocations, the Bayesian framework provides insights into more robust portfolio construction under uncertainty.

1.2 Objective

The primary objective of this study is to conduct a comparative analysis of Bayesian portfolio optimization models applied to U.S. AI companies. In addition to evaluating risk-adjusted performance, the study examines portfolio volatility and market risk exposure using the conditional CAPM framework, complemented by DCC-GARCH to capture time-varying dynamics.

1.3 Significance

This research provides insights for institutional and retail investors by demonstrating how Bayesian portfolio optimization can improve allocation decisions under parameter uncertainty. The inclusion of conditional CAPM and DCC-GARCH allows for a detailed assessment of time-varying market risk and volatility, making the findings relevant for stakeholders interested in both performance and dynamic risk management in AI-focused equity portfolios.

II Literature Review

2.1 Literature Review

To begin with, the Modern Portfolio Theory (MPT) was initially introduced by Harry Markowitz in his article "Portfolio Selection", which was published in the *Journal of Finance* in 1952. The theory advocates for a portfolio that is diversified by incorporating assets that are poorly correlated, meaning that they behave differently in different market conditions. Markowitz's most significant contribution lies in his ability to translate the concepts of "risk" and "profitability" into mathematical terms and the development of the Markowitz Mean-Variance Optimization (MVO) model, which is going to be tested in the current research as well.

The foundation of the Bayesian framework can be traced back to the 18th century with the work of Thomas Bayes, whose theorem was later formalized and extended by Pierre-Simon Laplace (Bayes, 1763; Laplace, 1812). In the context of modern statistics and financial modeling, Bayesian methods have gained prominence due to their ability to incorporate prior beliefs and update them with new information. Zellner and Chetty (1965) were among the first to apply Bayesian techniques to econometric models, demonstrating their flexibility and robustness in estimating uncertain parameters. The Bayesian framework treats unknown parameters, such as expected returns and covariances in finance, as random variables with probability distributions. This approach enables analysts to formally incorporate uncertainty and derive posterior distributions that reflect both prior beliefs and observed data, offering a dynamic alternative to classical estimation methods. Olivier Ledoit and Michael Wolf (2003) propose a shrinkage estimator of the covariance matrix that combines the sample covariance matrix with a structured target to reduce estimation error. They show that this approach produces more stable covariance estimates and improves out-of-sample portfolio performance compared to traditional mean-variance optimization based on the sample covariance matrix.

Recent advancements in machine learning have significantly expanded the toolkit available for financial data analysis, particularly through the integration of Bayesian methods. Fischer Black and Robert Litterman (1992) introduced a model, which combines market equilibrium returns with investor views within a Bayesian framework to generate more stable and intuitive portfolio allocations. Lorenzo Garlappi, Raman Uppal, and Tan Wang (2007) develop a portfolio selection framework that explicitly incorporates both parameter uncertainty and model uncertainty using a multi-prior (ambiguity-averse) approach. They show that when investors account for uncertainty about expected returns and model specification, optimal portfolios become more conservative and better diversified compared to standard mean-variance solutions. Bade, Frahm, and Jaekel (2008) applied the portfolio optimization

models under the Bayesian framework and compared this methodology with traditional portfolio optimization models, resulting in better performance and finding that prior investors' information has a crucial role in the model's outcomes. Pfarrhofer (2024) further advances this approach by using multivariate Bayesian machine learning models for scenario analysis in macro-financial environments, emphasizing the importance of accounting for nonlinearities and asymmetries in the relationships between economic and financial indicators. In a related line of research, Gonzalez et al. (2019) explore the application of Gaussian processes and Bayesian optimization in financial contexts, such as interest rate modeling and trend-based investment strategies, highlighting the ability of Bayesian methods to improve forecasting accuracy and decision-making under uncertainty. Hoffman and Gelman (2014) introduce the No-U-Turn Sampler (NUTS), an adaptive extension of Hamiltonian Monte Carlo that automatically tunes path lengths to improve sampling efficiency and convergence without manual parameter calibration. The method significantly reduces tuning complexity while maintaining accurate posterior inference in high-dimensional Bayesian models

Subsequently, the capital asset pricing model (CAPM) was developed by Sharpe (1964) and Lintner (1965), and their papers significantly contributed to financial analysis by providing a simple yet powerful framework to understand the relationship between risk and expected return. The model introduced the concept of systematic risk, captured by the beta coefficient, and established that the expected return of an asset is determined by its sensitivity to the overall market return, as well as the risk-free rate. Black (1972) presents a model of capital market equilibrium with restricted borrowing, extending the traditional Capital Asset Pricing Model (CAPM) by incorporating constraints on investor borrowing and its impact on asset prices. Fabozzi and Francis (1978) introduce one of the major changes to the standard capital asset pricing model (CAPM) is the replacement of the constant beta by a time-varying beta, which can be named as the emergence of the conditional CAPM.

Nevertheless, there are various methods to estimate the time-varying beta. Engle (2000) proposes the Dynamic Conditional Correlation model, a straightforward enhancement of multivariate GARCH models, which enables the modeling of time-varying correlations between multiple financial time series, which could aid in the estimation of beta. Jain (2011) proposes the Heterogeneous Autoregressive Beta model to capture time-varying betas, providing a more adaptable approach to modeling dynamic risk exposures by incorporating diverse responses across various assets. Corradi, Distaso, and Fernandes (2013) investigate the connection between conditional alphas and realized betas, emphasizing the significance of time-varying risk measures in explaining asset returns and performance beyond conventional asset pricing

models. Engle (2016) develops the Dynamic Conditional Beta model, which expands upon traditional asset pricing models by incorporating time-varying betas to capture dynamic risk exposures in financial markets. Zhang and Choudhry (2016) compare GARCH models and the Kalman Filter in forecasting the daily time-varying beta of European banks during the crisis period, highlighting the advantages and limitations of both approaches in capturing dynamic risk exposures. In conclusion, Aloy et al. (2020) conducted a comparative analysis of various techniques for modeling time-varying conditional betas, focusing on their application to Real Estate Investment Trusts (REITs), demonstrating the advantages and limitations of different approaches in capturing dynamic risk exposures.

III Empirical analysis

3.1 Methodology

While this study aims to analyze portfolio optimization approaches, it is essential to start with the traditional method, which is known as Mean-Variance Optimization (MVO) as proposed by Harry Markowitz in 1952. MVO is implemented as a benchmark model using sample estimates of expected returns and the covariance matrix computed from daily stock returns. The vector of expected returns is defined as the arithmetic sample mean of asset returns, while portfolio risk is measured using the empirical covariance matrix. The optimization problem maximizes the Sharpe ratio under a fully invested, long-only constraint.

$$\max \frac{w^T \mu}{\sqrt{w^T \Sigma w}} \quad (1)$$

Subject to:

$$\sum_{i=1}^N w_i = 1, \quad 0 \leq w_i \leq 1,$$

Where μ is the vector of sample mean returns and Σ is the sample covariance matrix. The risk-free rate is set to zero in the baseline specification, making the objective equivalent to maximizing the excess return per unit of volatility. The justification on the risk-free rate is given further, and for the robustness check the study implies Sharpe ratio with risk-free adjustment as well.

This formulation ensures that the resulting portfolio balances both expected return and volatility in a way that maximizes efficiency. In this study, the adjusted Markowitz model served as a benchmark to compare against Bayesian strategies.

Bayesian portfolio optimization in this study is implemented within a decision-theoretic framework rather than through a fully specified generative likelihood model for asset returns. Instead of modeling the return-generating process directly, the Bayesian approach is applied to portfolio weights by combining prior beliefs with a utility-based objective function.

In classical mean–variance optimization, portfolio weights are determined using point estimates of expected returns μ and covariance matrix Σ . In contrast, the Bayesian framework treats portfolio weights as random variables and introduces prior distributions to regularize allocation decisions and incorporate parameter uncertainty. Bayes’ rule is formally defined as:

$$P(\theta|X) \propto P(\theta)P(\theta) \tag{2}$$

Where θ denotes model parameters and X represents observed data. However, in the present study, the likelihood term does not correspond to a fully generative distribution of returns. Instead, a utility-based specification is constructed from the Sharpe ratio objective. Specifically, the portfolio utility function is defined as:

$$U(w) = \frac{w^T \mu}{\sqrt{w^T \Sigma w}} \tag{3}$$

The posterior distribution over portfolio weights is then defined as:

$$P(w|X) \propto \exp \{U(w)\} \times P(w) \tag{4}$$

Where $P(w)$ denotes the prior distribution on weights. Depending on the model specification, weights are parameterized using Dirichlet or softmax transformations to ensure they sum to one and remain within admissible bounds. Posterior inference is conducted using Hamiltonian Monte Carlo with the No-U-Turn Sampler (NUTS), allowing efficient exploration of the weight space under the specified utility structure. This formulation corresponds to Bayesian decision modeling, in which the posterior reflects optimal allocation under uncertainty and regularization, rather than standard Bayesian inference on a parametric likelihood for asset returns.

Furthermore, to examine the conditional volatility and market risk exposure under the time-varying concept, the conditional CAPM model was implemented. Recall that the beta coefficient is derived from the classical CAPM, where it reflects the stock’s exposure to market risk. The CAPM formula is presented below:

$$E(R_e) = R_f + \beta(R_m - R_f) \tag{5}$$

Where $E(R_e)$ is the expected return on the equity, R_f – risk-free rate, R_m – return on the market portfolio. β – beta coefficient, which can be calculated as it is written below:

$$\beta = \frac{cov(R_m, R_i)}{var(R_m)} \tag{6}$$

The CAPM is a remarkable framework, crafted by a diverse group of scholars. However, it has a significant drawback — it relies on static variables, which often lead to unrealistic outcomes. To address this issue, a subsequent version of the model was developed, giving rise to the Conditional CAPM. This version incorporates the dynamic nature of financial data, incorporating it into the formula. The model adopts a different perspective:

$$E_t(R_{i,t+1}) - R_f = \beta_{i,t}[E_t(R_{m,t+1}) - R_f] \tag{8}$$

Sometimes, the equation (7) can be expressed more conveniently in the following form:

$$E_t(\tilde{r}_{i,t+1}) = \beta_{i,t}E_t(\tilde{r}_{M,t+1}) \tag{9}$$

Where $E_t(\tilde{r}_{i,t+1}) = E_t(R_{i,t+1}) - R_f$, $E_t(\tilde{r}_{M,t+1}) = E_t(R_{m,t+1}) - R_f$, meaning the conditional expectation of the net excess return of asset i and of the market. Here, it is important to say that the new return is calculated as follows:

$$r = \ln\left(\frac{Price_{t+1}}{Price_t}\right) \tag{10}$$

Thus, the return is calculated as the log division of price change. Furthermore, in the context of a conditional CAPM model based on time-series data, the beta coefficient also becomes time-varying and can be expressed as follows:

$$\beta_{i,t} = \frac{cov_t(R_{m,t+1}, R_{i,t+1})}{var_t(R_{m,t+1})} \tag{11}$$

There are various methods for calculating the conditional variance and conditional beta, however, the current study will apply the GARCH(1,1) and Dynamic Conditional Correlation GARCH (DCC-GARCH) model respectively. The DCC-GARCH model was developed by Engle and Sheppard (2001) and Engle (2002) to estimate large, time-varying covariance matrices.

It combines dynamic correlation with the GARCH model, allowing it to handle heteroscedasticity as well as large, dynamic covariance matrices. Recall, that the GARCH(1,1) takes the following view:

$$\sigma_t^2 = \omega + a\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2 \quad (12)$$

Where σ_t^2 is the conditional variance at time t; ω is the constant term; a is the coefficient of the autoregressive term; ε_{t-1}^2 is the squared residual at time t-1; β is the coefficient of the moving average of squared shocks term; σ_{t-1}^2 is the conditional variance at time t-1. Regarding the DCC-GARCH formula, although it is quite complicated, the current study still aims to briefly explain it. There is the following set of formulas:

$$\begin{aligned} H_t &= D_t R_t D_t D_t = \text{diag}(\sqrt{h_{11t}}, \dots, \sqrt{h_{NNt}}) \\ Q_t &= (1 - a - \beta)Q + a\varepsilon_{t-1}\varepsilon'_{t-1} + \beta Q_{t-1} \end{aligned} \quad (13)$$

$$R_t = \text{diag}(Q_t)^{-1} Q_t \text{diag}(Q_t)^{-1}$$

Where H_t is the conditional covariance matrix; D_t is a diagonal matrix of time-varying standard deviations from univariate GARCH models; h_{NNt} is the conditional variance of the asset, which was calculated via GARCH before, R_t is the time-varying correlation matrix; Q_t is an intermediate correlation matrix; Q_{t-1} is the long-run unconditional correlation matrix of standardized residuals; a and β are parameters controlling the dynamics of correlation. The multivariate distribution is specified as Gaussian and estimation is conducted in R-studio using the `rmgarch` and `rugarch` packages, through which conditional covariance and variance are extracted from the estimated DCC model

3.2 Research Design

For this study, a U.S. AI-focused portfolio was constructed, consisting of 10 large-cap companies actively involved in artificial intelligence. Daily returns for each stock were collected from 10 December 2020 to 1 January 2025. The Nasdaq-100 index was used as the market benchmark and representative of the broader market. The dataset was divided into an in-sample (training) period from 10 December 2020 to 1 January 2024, used for model estimation, and an out-of-sample (testing) period from 1 January 2024 to 1 January 2025, used to evaluate model performance.

Table 3.1 summarizes the key characteristics of the stocks included in the AI portfolio. Data were obtained in mid-May 2025 via TradingView and Tiingo (free access). Notably, daily returns are calculated from raw closing prices and are not adjusted for dividends or stock splits. While this may affect absolute return levels, it does not influence the comparative analysis across the portfolio optimization models.

Table 3.1 The structure of the US AI Index

Company	Capitalization	Main AI Involvement
MSFT	3.42T USD	Investments in OpenAI, Azure AI
NVDA	3.3T USD	GPUs and architecture for AI training
AMZN	2.19T USD	AWS AI/ML services, Alexa
GOOGL	2.1T USD	Leaders in machine learning (DeepMind, Gemini, Bard)
META	1.61T USD	
ORCL	454.03B USD	AI integration into cloud products
PLTR	291.19B USD	Big data analytics and military AI
IBM	244.65B USD	Watson, enterprise AI solutions
AMD	185.75B USD	NVIDIA alternative: chips for AI
AI	3.09B USD	Pure-play AI company: enterprise AI solutions

Source: TradingView

As the primary objective of this study is to compare alternative Bayesian portfolio optimization approaches applied to an AI portfolio, five Bayesian model specifications (B1–B5) are developed alongside a classical Markowitz mean–variance optimization (MVO) benchmark. Table 3.2 provides a summary of their key characteristics. Below, there is a formal specification of each Bayesian model.

Table 3.2 All models’ description

Generation	Description
Bayesian 1	The initial model with a mean equal sample mean, sigma is 0.01, and weights follow the Dirichlet distribution. All weights almost equally distributed.
Bayesian 2	The developed Bayesian 1 model with increased sigma to 0.05.
Bayesian 3	A model with sigma equals 0.02, the single asset limitation to 0.2 and implementation a penalty multiplied by 1000.
Bayesian 4	A developed Bayesian 3 model with the same sigma, but with asset limitation 0.25 and a penalty multiplied by 10000.
Bayesian 5	A grid search model, with weight limitation = 0.25; sigma = 0.01, 0.02, 0.05; penalty multiplied by 1000, 10000, 100000.
Markowitz	The Markowitz Mean-Variance Optimization model with Sharpe ratio maximization.

Source: Calculated by the Author

Bayesian 1 (B1) represents a baseline Bayesian portfolio optimization framework that incorporates parameter uncertainty in expected returns while enforcing fully invested, long-only portfolio weights. Expected returns are assigned independent normal priors centered at their empirical sample means with a sigma equal to 0.01. Portfolio weights follow the Dirichlet distribution. Rather than specifying a likelihood for returns, B1 defines a utility-based objective proportional to the Sharpe ratio. Posterior inference is conducted using the No-U-Turn Sampler (NUTS), an adaptive variant of Hamiltonian Monte Carlo. Sampling is performed with 2000 posterior draws following

1000 tuning iterations. Final portfolio weights are obtained as posterior means of the sampled weight distributions.

Bayesian 2 (B2) extends the baseline specification by increasing prior uncertainty over expected returns and allowing for greater dispersion in portfolio weights. Expected asset returns are assigned independent normal priors centered at their empirical sample means, with a larger prior sigma of 0.05. Portfolio weights follow a Dirichlet distribution with concentration parameters below one, which permits more extreme long-only allocations while maintaining fully invested portfolios without short selling. Portfolio return, volatility, the Sharpe-ratio-based utility function, and inference are defined identically to Bayesian Model 1.

Bayesian Model 3 (B3) extends the previous specifications by introducing soft portfolio concentration constraints and a more agnostic prior over expected returns, where the mean is 0.001, and the sigma is 0.02. The model assumes that the maximum desired weight for the single asset in the portfolio is 0.2, while the penalty term is scaled by a coefficient of 1000, controlling the strength of the constraint without imposing hard bounds. Portfolio weights are parameterized via a softmax transformation of unconstrained latent variables, ensuring long-only and fully invested portfolios without short selling. Portfolio performance is evaluated using a Sharpe-ratio-based utility function computed from posterior draws of expected returns and an empirical covariance matrix. Inference is conducted using the NUTS, employing four chains with 5000 posterior draws following 2000 tuning iterations. Final portfolio weights are computed as averages over the top 5% of posterior draws ranked by the Sharpe ratio, emphasizing high-performing portfolios.

Bayesian 4 (B4) builds upon the penalized Sharpe-ratio framework of Bayesian Model 3 by adjusting the portfolio concentration parameters. Expected returns and portfolio weights are modeled identically to B3. The soft concentration threshold is increased to 0.25. At the same time, the strength of the penalty applied to constraint violations is increased to 10000.

Bayesian 5 (B5) extends the penalized Sharpe-ratio framework by conducting a systematic grid search over key hyperparameters governing prior uncertainty and constraint enforcement. The objective is to assess the robustness of portfolio outcomes to alternative Bayesian specifications and penalty magnitudes. Expected asset returns are assigned independent normal priors with a constant mean of 0.001 and prior standard deviations of 0.01, 0.02, and 0.05. Portfolio weights are parameterized via a softmax transformation to ensure long-only and fully invested portfolios. A soft concentration constraint is imposed by penalizing weights exceeding 0.25, with penalty scaling factors of 1000, 10000, and 100000. Portfolio performance is evaluated using a Sharpe-ratio-based utility function computed

with an empirical covariance matrix For each combination of prior variance and penalty strength, posterior inference is conducted using the NUTS with identical sampling settings. Across all model configurations, the specification achieving the highest Sharpe ratio is selected, and the corresponding portfolio weights are reported as the B5 outcome.

Given the use of daily return data and a short investment horizon, the risk-free rate is set to zero across all portfolio optimization models. This choice is equivalent to working with daily excess returns under the assumption that the one-day risk-free rate is negligible relative to asset return variability. Moreover, because the risk-free rate enters the Sharpe ratio as a constant shift, its omission does not affect portfolio rankings and has no material impact on the comparative performance of the Bayesian models. As a robustness check, the Sharpe ratio is additionally computed using the 3-month U.S. Treasury bill yield as a proxy for the risk-free rate. To ensure consistency with the log-transformed daily index returns, the Treasury yield is converted to a daily rate and expressed in logarithmic form before calculating excess returns. This approach allows the Sharpe ratio to fully reflect movements in short-term interest rates while preserving comparability with the baseline specification.

All Bayesian models are implemented in Python using the PyMC library. Posterior inference is conducted via Hamiltonian Monte Carlo using the NUTS. Convergence is assessed through multiple chains, tuning phases, and inspection of posterior summaries. Portfolio weights are constrained to be long-only and fully invested through either Dirichlet priors or softmax transformations. Final portfolio allocations are computed as posterior averages or performance-weighted averages of high-Sharpe posterior draws.

To provide clearer empirical structure, the study formulates testable hypotheses regarding portfolio performance and risk characteristics. First, it is hypothesized that Bayesian models incorporating concentration penalties will outperform unconstrained Bayesian specifications in terms of out-of-sample risk-adjusted performance, while simultaneously reducing tail risk and conditional market beta. Second, it is expected that the classical Markowitz mean–variance optimization will achieve the highest Sharpe ratio but at the cost of substantially higher portfolio concentration and lower diversification. These hypotheses reflect the theoretical trade-off between pure Sharpe maximization and robustness to estimation error and concentration risk, and they guide the comparative evaluation of the proposed models.

3.3 Descriptive Statistics

As the main objective of the study is to compare different Bayesian portfolio optimization models, it is essential to calculate the optimal weighting for securities in the AI portfolio based on the training period. Table 3.3 shows the share of each stock in the portfolio according to each generation decision.

Interestingly, while the first generation used an equally weighted portfolio, subsequent generations preferred to increase the shares of certain companies due to their high profitability. However, the greatest concentration can be seen in the Markowitz (MVO) model, which allocates only three out of the ten companies. Regarding Bayesian optimization models, some generation reduce concentration (B3 and B4), while others (B5) can also become highly concentrated depending on grid settings and objective. While this may be mathematically sound, the lack of diversity is still a risky decision.

Table 3.3 Assets' allocation in the portfolio

Generation	AI	ORCL	AMZN	MSFT	AMD	NVDA	IBM	GOOGL	META	PLTR
Bayesian 1	9.89%	9.92%	9.95%	10.13%	9.95%	10.01%	10.14%	10.04%	9.93%	10.05%
Bayesian 2	8.69%	11.50%	7.68%	11.47%	9.37%	8.63%	16.18%	7.64%	9.70%	9.12%
Bayesian 3	2.84%	22.12%	2.25%	21.32%	5.75%	2.91%	23.41%	2.10%	13.41%	3.90%
Bayesian 4	2.27%	24.51%	1.80%	24.09%	4.61%	2.47%	25.08%	1.72%	10.21%	3.24%
Bayesian 5	0.15%	0.32%	0.20%	33.46%	0.17%	0.19%	64.96%	0.15%	0.22%	0.19%
Markowitz	0.00%	32.04%	0.00%	47.38%	0.00%	0.00%	20.58%	0.00%	0.00%	0.00%

Source: Calculated by the Author

Table 3.4 provides a summary of the descriptive statistics for six custom-built AI indexes (B1 to B5 and M), each created using different Bayesian portfolio optimization models, compared to the Nasdaq-100 (NDX). The data spans from 2021 to 2024.

Table 3.4 Descriptive statistics of models' performance

2021-2024	Mean	Median	Sd	Min	Max	Sharpe	Sharpe*	Skew.	Kurt.
Index B1	-0.07%	0.09%	2.57%	-30.11%	9.04%	-0.026	-0.030	-4.15	48.80
Index B2	-0.04%	0.09%	2.24%	-23.03%	8.38%	-0.019	-0.024	-2.86	30.43
Index B3	0.03%	0.06%	1.52%	-6.69%	6.38%	0.017	0.011	-0.30	4.78
Index B4	0.03%	0.05%	1.44%	-5.70%	5.96%	0.022	0.016	-0.22	4.42
Index B5	0.05%	0.06%	1.18%	-6.59%	4.58%	0.038	0.031	-0.44	6.01
Index M	0.07%	0.10%	1.33%	-5.81%	6.06%	0.049	0.043	-0.03	4.58
NDX	0.04%	0.09%	1.50%	-5.70%	7.22%	0.027	0.021	-0.18	4.28

Sharpe* - risk-free adjusted ratio
 Source: Calculated by the Author

Indexes B1 and B2 have negative mean returns (-0.07% and -0.04%, respectively) and pronounced negative skewness (-4.15 and -2.86) indicates exposure to asymmetric downside events, while the high kurtosis values (48.80 and 30.43) reflect heavy-tailed return realizations typical of equity markets during turbulent periods, while the B1 and B2 made almost equal allocation in weights among assets. Indexes B3 through B5 demonstrate improved performance, with increasing average returns (up to 0.05%), decreasing standard deviations, and lower skewness and kurtosis values, indicating more stable and symmetric return behavior. The index M, which was created using the traditional Markowitz optimization approach, exhibits the highest average return (0.07%) and the best Sharpe ratio (0.049) among all

custom indices. Notably, since the study uses the daily returns, so the Sharpe ratio is also presented on daily scale. The index M also demonstrates relatively low-risk characteristics, with a standard deviation of 1.33% and a minimal skewness of -0.03. In contrast, the Nasdaq-100 (NDX) achieved a mean return of 0.04%, with a Sharpe ratio of 0.027 - higher than most Bayesian models, except for B5 and M. This makes B5 and M the most promising models during the training period, with M standing out in terms of overall risk-adjusted performance.

Figure 3.1 illustrates the dynamics of all Bayesian model indices, based on the USA AI companies' portfolio during the training period.

Figure 3.1 Index performances during the training period



Source: Calculated by the Author

Table 3.5 assesses the out-of-sample performance of the AI-powered indices over the timeframe of 2024 to 2025, using fixed portfolio weights that were determined during the preceding training period. This stage serves to test the models' practicality and their ability to generalize beyond the data they were trained on.

Table 3.5 Descriptive statistics of models' performance

2024-2025	Mean	Median	Sd	Min	Max	Sharpe	Sharpe*	Skew.	Kurt.
Index B1	0.10%	0.16%	2.11%	-22.70%	4.65%	0.046	0.039	-5.03	56.22
Index B2	0.10%	0.13%	1.92%	-19.59%	4.43%	0.053	0.046	-4.28	45.77
Index B3	0.12%	0.18%	1.35%	-6.54%	4.45%	0.091	0.082	-0.49	5.66
Index B4	0.12%	0.16%	1.30%	-5.57%	4.69%	0.093	0.084	-0.32	5.05
Index B5	0.10%	0.16%	1.17%	-6.45%	6.08%	0.084	0.074	-0.52	9.42
Index M	0.11%	0.16%	1.23%	-4.03%	5.44%	0.089	0.080	0.13	6.11
NDX	0.10%	0.15%	1.15%	-3.72%	3.01%	0.083	0.072	-0.45	3.95

Sharpe* - risk-free adjusted ratio
 Source: Calculated by the Author

Remarkably, all six indices (B1 through B5 and M) achieved positive average returns ranging from 0.10% to 0.12%, outperforming their performance during the training period. Notably, Indexes B3, B4, and M continued to exhibit strong Sharpe ratios (0.091, 0.093, and 0.089, respectively), indicating attractive returns adjusted for risk. These findings reinforce the earlier conclusion that Bayesian models with lower variance (such as B3 and B4) and the Markowitz-based Index M remain resilient when applied to new data. Despite the positive mean returns, indexes B1 and B2 continue to exhibit extreme skewness (-5.03 and -4.28) and high kurtosis (56.22 and 45.77), indicating persistent left-tail risk and non-normal distributions. Their Sharpe ratios (0.046 and 0.053) remain the lowest, suggesting that these models may not be suitable for real-world application, despite their improved returns. The Nasdaq-100 (NDX) achieved a mean return of 0.10% and a Sharpe ratio of 0.083, performing similarly to indexes B5 and slightly below M, B3, and B4. This suggests that some of the custom Bayesian and Markowitz-weighted indexes may outperform the benchmark, particularly in terms of risk-adjusted returns.

Figure 3.2 represents the dynamics of all indexes on the AI company's portfolio during the out-of-sample period, with respect to different optimal weight allocations.

Figure 3.2 Index performances during the out-of-sample period



3.4 Time-varying beta

This part aims to analyze the market risk exposure for the out-of-sample period. While all generations already made their allocation decision, it is essential to examine how risky their strategies are on another data sample. The following Table 3.6 presents the results of conditional volatility and market beta estimates for each AI index, calculated using the DCC-GARCH model under the conditional CAPM theory. This approach allows us to analyze the time-varying exposure to market risk (beta) and dynamically estimate portfolio risk (volatility), providing a more accurate picture of portfolio behavior in changing market conditions.

Table 3.6 Descriptive statistics of models' volatility

	Mean	Median	Sd	Min	Max	Skewness	Kurtosis
Vol_B1	2.02%	2.02%	0.0461%	1.9422%	2.1011%	0.06	1.80
Vol_B2	1.85%	1.84%	0.0389%	1.7806%	1.9147%	0.06	1.80
Vol_B3	1.31%	1.31%	0.0169%	1.2849%	1.3429%	0.07	1.81
Vol_B4	1.28%	1.28%	0.0137%	1.2462%	1.2999%	0.08	1.81
Vol_B5	1.13%	1.13%	0.0231%	1.0884%	1.1678%	0.06	1.80
Vol_M	1.22%	1.22%	0.0002%	1.2238%	1.2243%	-0.09	1.81
Vol_NDX	1.14%	1.10%	0.1419%	0.9624%	1.6743%	1.41	4.84
Beta_B1	1.31	1.35	0.35	-0.85	1.85	-2.10	11.24
Beta_B2	1.22	1.27	0.32	-0.77	1.70	-2.13	11.33
Beta_B3	0.94	0.96	0.13	0.08	1.16	-1.81	10.62

Beta_B4	0.91	0.93	0.12	0.21	1.10	-1.50	7.87
Beta_B5	0.58	0.59	0.07	0.38	0.68	-0.82	3.02
Beta_M	0.82	0.84	0.09	0.58	0.95	-0.90	3.12

Source: Calculated by the Author

The levels of conditional volatility reveal a clear trend: indexes B3, B4, and B5 consistently exhibit the lowest risk levels (1.31%, 1.28%, and 1.13%, respectively), while indexes B1 and B2 display higher volatility (2.02% and 1.85%). The Markowitz-optimized portfolio (Index M) achieves both low volatility (1.22%) and extremely low dispersion ($SD \approx 0.00\%$), indicating exceptional stability. Interestingly, while the Nasdaq-100 has a relatively low mean volatility (1.14%), it also has the highest variability ($SD = 0.14\%$) and skewness (1.41), suggesting occasional spikes in risk exposure. Examining the conditional betas, the study notices that B1 and B2 have the highest average market exposure (1.31 and 1.22), accompanied by strong negative skewness and extreme kurtosis, indicating their susceptibility to asymmetric market shocks. In contrast, B5 and M have the lowest average betas (0.58 and 0.82), suggesting limited market dependence and greater potential for diversification benefits. Moreover, the low volatility–low beta profiles of B3, B4, B5, and M support the notion that these indexes can provide more stable performance with reduced exposure to systemic risk.

Overall, the DCC-GARCH results reinforce the findings from previous observations: indexes derived from Bayesian models with lower variance (B3–B5) and the Markowitz-optimized index (M) not only offer favorable return-risk ratios but also exhibit desirable conditional characteristics, such as low volatility and reduced sensitivity to market-wide fluctuations.

Table 3.7 reports the diagnostics of the DCC-GARCH models, including parameter estimates for both the univariate GARCH components of the Index and Market series, as well as the DCC correlation dynamics. The “Sum” column represents volatility persistence for univariate GARCH (Index/Market) or correlation persistence for DCC. All parameters are positive and satisfy the stationarity condition ($\alpha + \beta < 1$), and convergence was achieved without numerical issues. These diagnostics confirm that the estimated models are well-behaved, and the results are robust.

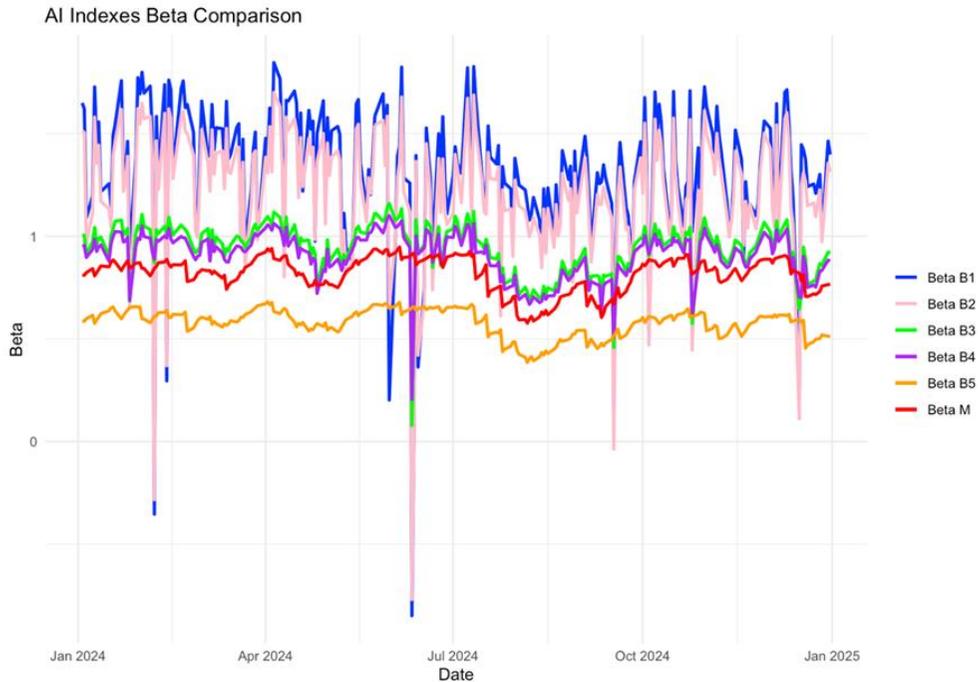
Table 3.7 Descriptive statistics of models’ volatility

Ind.	Conv.	Index			Market			DCC		
		alpha	beta	sum	alpha	beta	sum	alpha	beta	sum
B1	0	3,93E-11	9,99E-01	0,999	5,82E-02	8,65E-01	0,923	0,73	0,08	0,806
B2	0	2,56E-10	9,99E-01	0,999	5,82E-02	8,65E-01	0,923	0,73	0,08	0,808
B3	0	2,80E-09	9,99E-01	0,999	5,82E-02	8,65E-01	0,923	0,29	0,12	0,416
B4	0	9,53E-11	9,99E-01	0,999	5,82E-02	8,65E-01	0,923	0,23	0,18	0,405
B5	0	5,33E-10	9,99E-01	0,999	5,82E-02	8,65E-01	0,923	0,00	0,90	0,900
M	0	2,24E-08	9,99E-01	0,999	5,82E-02	8,65E-01	0,923	0,04	0,67	0,708

Source: Calculated by the Author

Figure 3.3 illustrates the dynamics of the conditional betas (market risk exposure) of all indexes during the out-of-sample period.

Figure 3.3 Index betas during the out-of-sample period



Source: Calculated by the Author

Conclusion

This study examined the performance of alternative Bayesian portfolio optimization models using a constructed AI index of the United States composed of ten large-cap companies over the period from December 2020 to January 2025. The empirical analysis was conducted within a three-year training period and a one-year out-of-sample evaluation, with portfolio weights held constant during the testing phase. Five Bayesian model specifications were compared against a classical Markowitz mean-variance optimization benchmark, with particular attention paid to risk-adjusted performance, concentration, and time-varying market risk exposure.

The results indicate that within this narrow AI-focused investment universe, the Markowitz MVO model achieved the highest Sharpe ratio both in-sample and out-of-sample. This “win” is likely influenced by the small number of assets, overlapping factor exposures, and the use of Sharpe ratio maximization with a zero risk-free rate, which can magnify the apparent efficiency of unconstrained allocations. However, this performance came at the cost of substantial portfolio concentration, increasing exposure to idiosyncratic risk.

Bayesian models incorporating soft concentration penalties (B3–B4) demonstrated a more balanced profile, while others (B5) can also become highly concentrated depending on grid settings and objective. These models achieved competitive Sharpe ratios while maintaining lower conditional volatility and reduced market beta, as evidenced by DCC-GARCH and conditional CAPM results. The trade-off between marginally lower Sharpe ratios and greater diversification suggests that such Bayesian specifications may be preferable for investors prioritizing risk management, robustness to estimation error, or adherence to diversification constraints.

It is important to note potential overfitting risks, particularly for the B5 model, which was selected through a grid search on Sharpe ratios. Without a nested validation set or rolling-window evaluation, performance estimates may be upward-biased. Future studies could adopt a three-way split or rolling-window frameworks to mitigate this risk.

Finally, the analysis is subject to several limitations. The small, sector-specific universe of ten AI stocks restricts generalizability to broader equity markets or other asset classes. Portfolio weights were held constant out-of-sample, with no consideration of transaction costs or dynamic rebalancing. Results also reflect a specific market environment characterized by rapid technological growth and unique volatility patterns, which may influence model performance.

Overall, the evidence suggests that Bayesian portfolio optimization provides a flexible framework for addressing parameter uncertainty and concentration risk in narrowly defined, high-growth sectors. For investors willing to accept a slight reduction in Sharpe ratio, models with explicit regularization (B3–B5) offer a practical balance between performance, diversification, and time-varying market risk exposure.

Acknowledgements

I would like to express my sincere gratitude and appreciation to everyone who has supported me during this challenging period. I would like to extend my sincere gratitude to my supervisor, Dr. Jie Zhu. His invaluable guidance, wisdom, and unwavering support have been crucial throughout this research process. His profound insights and encouragement not only shaped this thesis but also contributed to my academic and personal growth.

Conflict of Interest: The author reported no conflict of interest.

Data Availability: All data are included in the content of the paper.

Funding Statement: The author did not obtain any funding for this research.

References:

1. Aloy M, Laly F, Laurent S, Lecourt C. Modeling Time-Varying Conditional Betas. A Comparison of Methods with Application for REITs. In: *Dynamic Modeling and Econometrics in Economics and Finance.* ; 2020:229-264. doi:10.1007/978-3-030-54252-8_9
2. Bade A, Frahm G, Jaekel U. A general approach to Bayesian portfolio optimization. *Mathematical Methods of Operations Research.* 2008;70(2):337-356. doi:10.1007/s00186-008-0271-4
3. Bayes Mr, Price Mr. LII. An essay towards solving a problem in the doctrine of chances. By the late Rev. Mr. Bayes, F. R. S. communicated by Mr. Price, in a letter to John Canton, A. M. F. R. S. *Philosophical Transactions of the Royal Society of London.* 1763;53:370-418. doi:10.1098/rstl.1763.0053
4. Black F. Capital Market Equilibrium with Restricted Borrowing. *The Journal of Business.* 1972;45(3):444. doi:10.1086/295472
5. Black, F., & Litterman, R. (1992). Global Portfolio Optimization. *Financial Analysts Journal,* 48(5), 28–43. <https://doi.org/10.2469/faj.v48.n5.28>
6. Corradi, V., Distaso, W., & Fernandes, M. (2013). Conditional alphas and realized betas. Paper presented at the Moscow Financial Economics Conference, Moscow, Russia.
7. Laplace PS. *Théorie analytique des probabilités.* Published online June 11, 2018. doi:10.14711/spcol/b718972
8. Engle RF. Dynamic Conditional beta. *Journal of Financial Econometrics.* 2016;14(4):643-667. doi:10.1093/jjfinec/nbw006
9. Engle R, Sheppard K. Theoretical and Empirical Properties of Dynamic Conditional Correlation Multivariate GARCH.; 2001. doi:10.3386/w8554
10. Fabozzi FJ, Francis JC. Beta as a random coefficient. *Journal of Financial and Quantitative Analysis.* 1978;13(1):101. doi:10.2307/2330525
11. Garlappi, L., Uppal, R., & Wang, T. (2006). Portfolio Selection with Parameter and Model Uncertainty: A Multi-Prior Approach. *Review of Financial Studies,* 20(1), 41–81. <https://doi.org/10.1093/rfs/hhl003>
12. Gonzalez J, Lezmi E, Roncalli T, Xu J. Financial applications of Gaussian processes and Bayesian optimization. arXiv (Cornell University). Published online March 12, 2019. doi:10.48550/arxiv.1903.04841
13. Homan MD, Gelman A. The No-U-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. arXiv (Cornell University). 2014;15(1):1593-1623. <https://arxiv.org/pdf/1111.4246.pdf>

14. Jain K. Time-Varying Beta: The heterogeneous Autoregressive Beta model. 2011. <http://hdl.handle.net/10161/3402>
15. Ledoit, O., & Wolf, M. (2003). Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of Empirical Finance*, 10(5), 603–621. [https://doi.org/10.1016/s0927-5398\(03\)00007-0](https://doi.org/10.1016/s0927-5398(03)00007-0)
16. Lintner J. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *The Review of Economics and Statistics*. 1965;47(1):13. doi:10.2307/1924119
17. Markowitz H. Portfolio selection. *The Journal of Finance*. 1952;7(1):77. doi:10.2307/2975974
18. Pfarrhofer M, Stelzer A. Scenario Analysis with Multivariate Bayesian Machine Learning Models. arXiv (Cornell University). Published online February 12, 2025. doi:10.48550/arxiv.2502.08440
19. Sharpe WF. Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk. *The Journal of Finance*. 1964;19(3):425. doi:10.2307/2977928
20. Zellner A, Chetty VK. Prediction and Decision Problems in Regression Models from the Bayesian Point of View. *Journal of the American Statistical Association*. 1965;60(310):608-616. doi:10.1080/01621459.1965.10480817
21. Zhang Y, Choudhry T. Forecasting the daily Time-Varying Beta of European banks during the crisis period: Comparison between GARCH models and the Kalman Filter. *Journal of Forecasting*. 2016;36(8):956-973. doi:10.1002/for.2442