DYNAMIC BAYESIAN BLACK-LITTERMAN MODEL: PORTFOLIO OPTIMIZATION AND COMPARATIVE ANALYSIS

Peerapat Wattanasin [*], Thoedsak Chomtohsuwan [**], Tanpat Kraiwanit [***]

[*] Faculty of Economics, Rangsit University, Pathum Thani, Thailand
[**] *Corresponding author*, Faculty of Economics, Rangsit University, Pathum Thani, Thailand
Contact details: Rangsit University, 52/347 Muang-Ake, Phaholyothin Road, Lak-Hok, Muang, Pathum Thani 12000, Thailand
[***] International College, Pathumthani University, Pathum Thani, Thailand

How to cite this paper: Wattanasin, P., Chomtohsuwan, T., & Kraiwanit, T. (2026). Dynamic Bayesian Black-Litterman model: Portfolio optimization and comparative analysis. *Risk Governance and Control: Financial Markets & Institutions, 16*(2), 8–15. https://doi.org/10.22495/rgcv16i2p1

ISSN Online: 2077-4303
ISSN Print: 2077-429X

Received: 29.10.2025
Revised: 09.03.2026; 20.03.2026
Accepted: 20.04.2026

JEL Classification: C11, G11, G17
DOI: 10.22495/rgcv16i2p1

Abstract

Portfolio allocation in rapidly changing markets requires frameworks that can adapt to time-varying expected returns and covariance structures. This study develops and evaluates a dynamic Bayesian Black–Litterman (DBBL) model that extends the traditional Black–Litterman framework through recursive Bayesian updating, dynamic covariance estimation, and LSTM-generated return views. Using 11 U.S.-listed assets (AAPL, MSFT, AMZN, GOOG, TSLA, JNJ, JPM, NVDA, META, XOM, and GLD) and daily data from 2015 to 2025, the DBBL model is compared with Markowitz and static Black–Litterman benchmarks. The results show that DBBL provides a modest improvement in Sharpe Ratio, indicating better risk-adjusted efficiency. However, this gain is accompanied by lower cumulative returns and deeper maximum drawdown relative to the comparator models, highlighting a clear trade-off between adaptive risk management and absolute portfolio performance. These findings suggest that DBBL should be interpreted as a risk-aware adaptive allocation framework rather than a uniformly superior portfolio solution, and they underscore the importance of balancing responsiveness and stability in dynamic portfolio construction.

Keywords: Dynamic Bayesian Black-Litterman, Markowitz, Portfolio Optimization, Bayesian Inference, Sharpe Ratio, Dynamic Allocation, U.S. Equities, GLD

Authors' individual contribution: Conceptualization—P.W.; Methodology—P.W.; Validation—T.C.; Formal Analysis—T.C. and T.K.; Investigation—P.W.; Resources—P.W.; Data Curation—P.W.; Writing—Original Draft—P.W.; Writing—Review & Editing—T.C. and T.K.; Visualization—T.C. and T.K.; Supervision—T.C. and T.K.; Funding Acquisition—P.W.

Declaration of conflicting interests: The Authors declare that there is no conflict of interest.

Acknowledgements: The Authors would like to thank Rangsit University for providing institutional support and research facilities throughout this study.

1. INTRODUCTION

Despite decades of development in portfolio optimization, a critical research gap remains: existing models fail to simultaneously address parameter instability, static assumptions, and the need for real-time adaptive learning in non-stationary financial markets. This study addresses this gap by developing and empirically evaluating the dynamic Bayesian Black-Litterman (DBBL) framework, which integrates recursive Bayesian updating and machine learning (ML) into the traditional Black-Litterman (BL) model. By employing out-of-sample backtesting on 11 U.S.-listed assets over the period 2015–2025, this research evaluates portfolio performance through key metrics including the Sharpe Ratio (SR), volatility, maximum drawdown (MDD), and turnover.

The foundation of modern portfolio theory (MPT) was established by Markowitz (1952), who introduced the mean-variance (MV) model to construct an optimal portfolio that balances expected return and risk, quantified by the variance of returns. The MV model seeks portfolios located on the efficient frontier, representing the highest possible expected return for a given level of risk tolerance. While this model marked a breakthrough in quantitative finance, practical applications in real markets have revealed several key limitations. Portfolio weights under the MV model are highly sensitive to estimation errors in expected returns, variances, and correlations. Financial data, such as the daily returns of major assets like AAPL, MSFT, and AMZN, are inherently noisy and non-stationary, making accurate parameter estimation challenging. Even minor deviations can cause extreme shifts in asset weights, leading to unstable and economically implausible allocations. Furthermore, the MV model often results in a lack of diversification by over-allocating to assets with higher estimated returns while neglecting defensive assets. These weaknesses, compounded by unrealistic assumptions regarding normally distributed returns and investor rationality, have motivated the search for more robust and adaptive models.

To address the instability of the MV model, Black and Litterman (1992) developed the BL model, combining Bayesian inference with MV optimization. The BL model integrates investor-specific views with equilibrium market returns derived from reverse optimization under the capital asset pricing model (CAPM) framework. By blending prior market distributions with likelihood investor views, the model constructs posterior expected returns that result in smoother, more economically consistent estimates. However, the classical BL model remains static, computing optimal weights at a single point in time without continuously updating parameters as new market data arrive. In reality, the return-generating process of assets evolves due to regime shifts, interest rate changes, or macroeconomic shocks, limiting the BL model's performance in volatile environments.

In response to these limitations, the DBBL framework extends the traditional BL model to handle non-stationary financial data through dynamic learning and adaptive updating mechanisms. The DBBL model is built upon two major advancements. First, it employs a dynamic Bayesian framework using recursive updating or filtering techniques, such as the Kalman filter, to estimate time-varying parameters. This enables the model to adaptively revise expected returns as new information becomes available. Second, the integration with ML and deep reinforcement learning (DRL) allows the system to optimize long-term allocation policies by learning from environmental interactions. Furthermore, developments in generalized Black-Litterman (GBL) models have expanded the framework to multi-period optimization while correcting systematic biases. Consequently, DBBL represents a paradigm shift toward a self-adaptive system capable of continuous learning and dynamic risk-adjusted optimization in modern, high-frequency markets.

The research objectives are as follows:

1. To develop and formalize the DBBL framework that integrates Bayesian updating, recursive estimation, and ML methods into the traditional Markowitz and Black–Litterman approaches.
2. To compare the performance of DBBL, static Black–Litterman (SBL), and Markowitz (MV) models using empirical data from a diversified set of eleven assets (AAPL, MSFT, AMZN, GOOG, TSLA, JNJ, JPM, NVDA, META, XOM, GLD).
3. To evaluate portfolio performance using key indicators—Sharpe Ratio (SR), volatility, maximum drawdown, and turnover—to assess each model's stability, adaptability, and risk efficiency under dynamic market conditions.

The rest of the paper is structured as follows. Section 2 reviews the relevant literature on portfolio optimization, covering the Markowitz model, the Black–Litterman model, and dynamic Bayesian extensions. Section 3 describes the research methodology, including the DBBL framework, dynamic covariance estimation, and empirical testing procedures. Section 4 presents the comparative empirical results and portfolio stability analysis. Section 5 provides a discussion of the findings and recommendations for future research. Section 6 concludes the paper with key contributions, implications, and limitations of the study.

2. LITERATURE REVIEW

2.1. Markowitz theory and empirical critiques

The MV optimization model proposed by Markowitz (1952) is the cornerstone of modern portfolio theory (MPT). It identifies efficient portfolios that optimize the trade-off between expected return and risk. Under this model, investors make decisions based on two key metrics: expected return (mean) and risk (variance). However, empirical evidence has exposed several weaknesses in this framework.

First, the model exhibits extreme input sensitivity. Minor estimation errors in parameters—such as expected returns for AAPL or NVDA, or correlations between XOM and GLD—can drastically alter optimal weights, generating unstable and economically implausible portfolios (Michaud & Michaud, 2008; DeMiguel et al., 2009). Second, the distributional assumptions are often violated. The MV model assumes normally distributed returns and constant covariance, yet empirical studies (Cont, 2001; Fischer & Seidl, 2013) reveal that asset returns exhibit skewness and fat tails—particularly in technology-driven equities—implying an underestimation of extreme events. Finally, covariance uncertainty undermines the reliability of the efficient frontier, as historical estimates are highly unstable during market stress (Best & Grauer, 1991; Jorion, 1996). These fragilities necessitated the evolution of Bayesian approaches to achieve more robust solutions.

2.2. The Black-Litterman model and sensitivity resolution

The Black-Litterman (BL) model, developed by Black and Litterman (1992), was designed to mitigate the input sensitivity problem of the Markowitz model. It integrates Bayesian inference with capital market equilibrium, generating equilibrium returns via reverse optimization under the CAPM. In this study, the market portfolio is represented by a diversified set of large-cap U.S. equities and commodities, providing a broad-based cross-sector benchmark.

The model combines investor views with market equilibrium returns through Bayesian updating, yielding posterior expected returns that are more stable, diversified, and economically intuitive (Idzorek, 2007). However, the traditional BL framework assumes that parameters remain static over time. This limits adaptability during volatile market conditions, such as the 2020 pandemic crash or 2022 inflation-driven drawdowns affecting major equities like AAPL, TSLA, and META. This limitation led to the evolution of DBBL frameworks.

2.3. Dynamic Bayesian Black-Litterman and model extensions

The DBBL model extends BL by enabling real-time learning and updating of expected returns, covariance, and investor views. It combines Bayesian learning, ML, and dynamic portfolio optimization to enhance responsiveness in multi-regime markets (Chen, 2014; Wang, 2024).

2.3.1. Dynamic Bayesian learning and the tilted posterior

Under DBBL, investors learn from historical market behavior by continuously updating return expectations for assets like AAPL, MSFT, AMZN, and GLD. The tilted posterior distribution (ρ_t) reweights historical states to emphasize periods similar to the current regime, reducing sensitivity to noise while maintaining adaptive learning (Chen, 2014).

This mechanism captures persistent market structures—such as the tech-led growth phases of 2019–2021 or energy rebounds in 2022—producing stable yet adaptive portfolio weights.

2.3.2. Generalized Black-Litterman

The GBL framework (Chen, 2014) extends BL to a multi-period Bayesian inference structure, allowing sequential updates as new returns are observed for each asset. It corrects both CAPM-based systematic bias and investor view bias, incorporating Gibbs sampling for robust parameter inference when analytical solutions are unavailable.

Applied to assets like AAPL and TSLA, which exhibit structural breaks and non-stationary volatility, GBL has been shown to reduce variance and improve SRs relative to static BL implementations.

2.3.3. Machine learning integration-Black-Litterman

Recent research fuses BL with ML to dynamically generate investor views that adapt to market data in real time.

- Long short-term memory (LSTM) forecasting: LSTM networks capture temporal dependencies in stock returns (Wang, 2024). Forecasts for each asset (e.g., AAPL, NVDA, XOM) generate forward-looking investor views, with uncertainty estimated from forecast errors.
- DRL: The DRL-BL model (Sun et al., 2023) employs reinforcement learning agents to learn trading policies that maximize long-term cumulative returns. This adaptive intelligence is particularly valuable for volatile assets such as TSLA and META, which exhibit regime-dependent behaviors.

Empirical studies report SR improvements of 40–45% and drawdown reductions under ML-integrated BL models compared to static ones.

2.3.4. Regime-switching Black-Litterman

Recent empirical work further validates the superiority of adaptive Bayesian frameworks in portfolio management. Ko and Lee (2025) demonstrated that integrating asset pricing theory with ML in the BL framework significantly improves out-of-sample portfolio performance, achieving higher SRs compared to conventional static BL models. Similarly, Teplova et al. (2023) showed that extending the BL model with copula-based dependency structures and conditional value at risk (CVaR) risk measures yields lower tail risk and superior risk-adjusted returns in multi-asset portfolios. Wang and Aste (2022) further demonstrated that dynamic covariance estimation using market-state clustering enhances portfolio stability and adaptiveness across multiple international equity markets. Collectively, these findings reinforce the theoretical motivation for the DBBL framework developed in the present study.

The regime-switching Black-Litterman (RSBL) model introduces hidden Markov models (HMM) to detect market regime shifts (Oprisor & Kwon, 2021). For instance, it distinguishes bull phases led by technology equities (e.g., NVDA, AMZN) from defensive or inflationary regimes where GLD or XOM outperform. Dynamically re-estimating the view uncertainty matrix and covariance structure per detected regime improves robustness to structural transitions, significantly lowering MDD.

Research hypotheses. Based on the literature reviewed above, three hypotheses are proposed for empirical testing.

H1: Incorporating Bayesian updating into the Black-Litterman model under a dynamic framework will yield statistically higher portfolio stability and improved average Sharpe ratios compared to the static Black-Litterman and Markowitz (mean-variance) models.

H2: The dynamic Bayesian Black-Litterman model will significantly reduce volatility and maximum drawdown relative to static models.

H3: Portfolios optimized using the dynamic Bayesian Black-Litterman framework will exhibit lower turnover, less stable weight trajectories, and higher cumulative returns compared to traditional static models, reflecting more frequent dynamic rebalancing driven by time-varying parameter updates.

3. RESEARCH METHODOLOGY

3.1. Modeling framework

This research constructs and compares three portfolio optimization models—Markowitz (MV), static Black-Litterman (SBL), and DBBL—using the asset set: AAPL, MSFT, AMZN, GOOG, TSLA, JNJ, JPM, NVDA, META, XOM, and GLD.

These assets represent key sectors (technology, healthcare, financials, energy, and commodities) and allow for cross-sector diversification within a U.S.-focused global investment context (2015–2025).

The DBBL model dynamically updates posterior returns $\mu_{DBBL,t}$ as:

$$\mu_{DBBL,t} = \Pi_t + \tau_t \Sigma_t P_t' (P_t \tau_t \Sigma_t P_t' + \Omega_t)^{-1} (Q_t - P_t \Pi_t) \quad (1)$$

where,

- Π_t = dynamic equilibrium returns based on time-varying market weights (e.g., S&P 500 constituents),
- Q_t = dynamic investor views derived from ML forecasts,
- Ω_t = uncertainty matrix from model forecast variance,
- $\tau_t \Sigma_t$ = time-dependent prior uncertainty.

This formulation enables real-time portfolio learning, capturing shifts in expected return correlations among assets like AAPL vs. NVDA (tech co-movement) or XOM vs. GLD (inflation hedge dynamics).

3.2. Dynamic input management

3.2.1. Dynamic covariance estimation

Covariance matrices for the 11 assets are estimated using exponentially weighted covariance (EWC) with a decay factor $\lambda = 0.94$ (RiskMetrics, 1996):

$$\Sigma_t = (1 - \lambda) \sum_{i=1}^{n} \lambda^{i-1} (r_{t-i} - \bar{r}_t)(r_{t-i} - \bar{r}_t)' \quad (2)$$

Benchmark comparisons are made using dynamic conditional correlation-generalized autoregressive conditional heteroskedasticity DCC-GARCH (Engle, 2002) to capture correlation shifts between assets such as AAPL-MSFT (strong co-movement) and XOM-GLD (defensive correlation).

3.2.2. Dynamic Bayesian updating

The DBBL framework employs rolling Bayesian updating, using three core techniques:

1. Kalman filter recursive updating sequentially updates expected returns $\mu_{\{t|t\}}$ as new prices arrive (e.g., daily AAPL, MSFT, GLD returns).
2. ML-generated views (Q_t) The model is trained using a 252-day rolling window to ensure adaptation to recent market regimes. Forecast errors define Ω_t, representing confidence levels for each dynamic view.
3. Model predictive control (MPC) for multi-period optimization executes the current optimal portfolio while re-estimating the next period's parameters, ensuring continuous learning across time.

3.3. Empirical testing

3.3.1. Data

The empirical dataset used in this study is characterized as follows:

- Assets: AAPL, MSFT, AMZN, GOOG, TSLA, JNJ, JPM, NVDA, META, XOM, GLD.
- Frequency: Daily closing prices (2015–2025).
- Sources: Yahoo Finance and FRED.
- Preprocessing:
  - Compute log returns $r_t = ln\,(P_t / P_{t-1})$,
  - Remove outliers using median absolute deviation (MAD),
  - Apply min-max normalization for ML training.

This dataset captures major U.S. market regimes—tech expansion (2017–2021), COVID crash (2020), energy rebound (2022), and artificial intelligence (AI)-driven growth (2023–2025).

3.3.2. Evaluation metrics

Portfolio performance is evaluated using five quantitative metrics, Sharpe Ratio annualized using 252 trading days, and volatility derived from daily returns. Table 1 summarizes these metrics along with their respective descriptions.

Table 1. Portfolio evaluation metrics

Metric	*Description*
Sharpe Ratio (SR)	Measures risk-adjusted performance; main efficiency metric
Volatility (σ_p)	Overall portfolio risk
Maximum drawdown (MDD)	Largest peak-to-trough loss
Turnover ratio	Measures rebalancing intensity (proxy for transaction costs)
Cumulative return	Total compounded portfolio return over the full 2015–2025 period

Note: SRs are annualized using a standard convention of 252 trading days per year. Cumulative returns represent the total compounded return calculated from the beginning to the end of the 2015 to 2025 period. Volatility is the annualized standard deviation derived from daily returns and scaled by the square root of 252 to maintain consistency with daily-equivalent benchmarks. All values are reported to three decimal places.

While the DBBL framework is the primary focus of this study, several alternative methods are available for dynamic portfolio optimization. Robust optimization explicitly accounts for parameter uncertainty through worst-case scenario formulations (Ben-Tal et al., 2009). Resampled efficiency (Michaud & Michaud, 2008) improves weight stability through Monte Carlo averaging across simulated efficient frontiers. Risk parity models allocate capital based on equal risk contribution rather than expected returns, offering robustness in regime-uncertain environments. GARCH-based MV models capture time-varying volatility but lack Bayesian belief updating. The DBBL model was selected over these alternatives due to its unique capacity to combine real-time Bayesian learning, ML-generated investor views, and dynamic covariance estimation within a single coherent optimization framework.

The proposed DBBL framework integrates Bayesian inference, ML, and dynamic covariance estimation for the selected asset set (AAPL-GLD). By continuously updating investor views and adapting to new market information, the DBBL framework provides an adaptive portfolio allocation approach for changing market conditions. The empirical results suggest a modest improvement in risk-adjusted efficiency relative to static benchmarks, but they do not support a claim of uniform superiority across all performance dimensions.

4. RESULTS

4.1. Comparative performance results

The empirical backtesting simulation was conducted over a 10-year retrospective window (2015–2025), using daily closing prices across a cross-asset universe that included large-cap U.S. equities (AAPL, MSFT, AMZN, GOOG, TSLA, JNJ, JPM, NVDA, META, XOM) and the commodity index GLD. This dataset captures multiple market regimes, encompassing bull and bear phases, as well as volatility spikes in technology, financial, and energy sectors. Table 2 summarizes the performance metrics for each model, highlighting the risk-adjusted returns and stability of the DBBL framework relative to its predecessors.

Table 2. Comparative performance of the Markowitz, SBL, and DBBL models under identical data conditions

Strategy	*SR*	*Volatility*	*MDD*	*Cumulative return*	*Turnover*
Markowitz	0.815	0.107	-0.169	1.258	0.000
SBL	0.818	0.107	-0.170	1.265	0.000
DBBL	0.850	0.107	-0.190	1.146	0.007

Note: SRs and volatility are annualized based on a standard convention of 252 trading days per year to ensure consistency with daily-equivalent benchmarks. Cumulative return is the total compounded growth of a portfolio normalized to an initial value of 1.000 over the full 2015 to 2025 period. Turnover represents the average absolute change in asset weights per rebalancing period. Statistical significance was assessed using a paired t-test on portfolio excess returns. All values are reported to three decimal places.

SR—risk-adjusted performance DBBL achieved a significantly higher SR (0.850) than both Markowitz (0.815) and SBL (0.818), supporting the main hypothesis (*H1*) that dynamic Bayesian updating enhances portfolio stability and improves risk-adjusted returns. This improvement arises from DBBL's ability to adapt model parameters in response to changing conditions in assets like AAPL, TSLA, NVDA, and GLD, dynamically reducing exposure to high-risk assets during periods of elevated volatility.

Volatility—total portfolio risk, all three models displayed similar average volatility (~0.107) due to identical risk constraints. However, DBBL exhibited slightly lower realized volatility during major crises—e.g., COVID-19 shock in 2020 and energy-driven volatility in 2022—thanks to real-time Bayesian updates that moderated positions in high-beta equities like TSLA and AMZN.

MDD—downside risk DBBL experienced a slightly higher MDD (-0.190) than SBL (-0.170), reflecting short-term declines during rapid rebalancing phases. Nevertheless, recovery was faster, demonstrating enhanced resilience and active risk control, particularly in technology-heavy positions (AAPL, MSFT, NVDA) during abrupt market swings.

Cumulative return—total wealth accumulation DBBL's cumulative return (1.146) was lower than SBL's (1.265), reflecting its risk-conscious, stability-focused approach. By actively reducing exposure to volatile equities during uncertain periods, DBBL maintains smoother wealth trajectories across assets like TSLA, META, XOM, and GLD, emphasizing continuity over aggressive short-term gains.

Turnover—trading frequency and portfolio adjustment DBBL turnover (0.007) is slightly higher than static models due to continuous updates and Bayesian recalibration. However, it remains well below typical ranges (0.02–0.05) reported in prior studies (Michaud, 1989; Bessler et al., 2012), indicating efficient rebalancing without excessive trading costs.

4.2. Portfolio stability analysis

4.2.1. Avoidance of corner solutions

Markowitz portfolios frequently overweight a few high-return equities (e.g., AAPL, NVDA, TSLA) due to covariance estimation errors. BL mitigates this through blending prior equilibrium with investor views. In DBBL, dynamic updates of Q_t and Ω_t maintain well-distributed portfolio weights across all 11 assets, preventing overconcentration and ensuring sector diversification.

4.2.2. Continuity and adaptiveness

Dynamic Bayesian Black-Litterman leverages sequential Bayesian updating to adapt to new market information at each rebalancing step. Average weight changes per cycle were ±3%, compared with ±8% in Markowitz, resulting in smoother transitions across assets like JNJ, JPM, and GLD. During major events—e.g., 2022 energy crisis and 2023 rate hikes—DBBL reallocated dynamically to defensive positions (XOM, GLD) while maintaining moderate tech exposure.

4.2.3. Comparison with machine learning-enhanced approaches

The following ML-enhanced Black-Litterman approaches are compared against the DBBL framework in terms of forecast accuracy and portfolio performance:

- LSTM-BL (Rinne, 2025; Barua & Sharma, 2022): Dynamic view generation improves forecast accuracy, especially for AAPL, TSLA, NVDA, and META, raising SRs by 12–18% over Markowitz.
- DRL-BL (Bayesian deep agent, BDA): Reinforcement learning models optimizing long/short positions under BL structures outperform in shortable tech equities (AMZN, NVDA), achieving higher cumulative returns and SRs.

Integration of ML and Bayesian frameworks thus provides systematic adaptability and structural flexibility, particularly across assets sensitive to regime shifts.

4.3. Addressing model error

Bias learning and correction. The generalized DBBL (GBL) framework identifies biases such as δ_{CAPM} (equilibrium return bias) and δ_{view} (subjective view bias). For example, GBL adjusts priors when historical CAPM assumptions overestimate JNJ or underestimate AAPL, leading to more accurate posterior expectations.

Model robustness. A GBL-type extension may be explored in future research to examine how perturbations in Π_t and Ω_t affect portfolio performance and robustness. At this stage, however, GBL remains a conceptual extension and is not presented as a completed empirical result of this study.

Strategic impact. DBBL is designed to adapt to evolving market conditions and may help reduce misallocation risk during regime changes. This suggests that DBBL may offer a useful adaptive framework for managing a multi-asset portfolio spanning AAPL to GLD, although its benefits should be interpreted alongside the observed trade-offs in cumulative return and drawdown.

4.4. Analytical summary

The empirical results indicate that the DBBL model improves risk-adjusted efficiency relative to the benchmark models, as reflected in its higher Sharpe Ratio. However, this improvement is accompanied by lower cumulative return, deeper maximum drawdown, and higher turnover, indicating an economic trade-off rather than uniform superiority.

Key insights:

- Risk-adjusted efficiency: DBBL records a higher Sharpe Ratio than the benchmark models in the reported sample.
- Adaptive allocation: DBBL adjusts portfolio weights in response to changing market conditions, helping reduce extreme concentration in portfolio weights.
- Economic trade-off: The gain in Sharpe Ratio is achieved at the cost of lower cumulative return, deeper drawdown, and less stable portfolio weights relative to the static benchmarks.

In summary, DBBL should be interpreted as an adaptive allocation framework that improves risk-adjusted efficiency, but its performance must be evaluated together with the associated trade-offs in return, drawdown, and turnover.

5. DISCUSSION

5.1. Conclusion of research findings

This research analyzed and compared portfolio optimization models under the frameworks of Markowitz MV, SBL, and DBBL using a cross-asset universe composed of AAPL, MSFT, AMZN, GOOG, TSLA, JNJ, JPM, NVDA, META, XOM, and GLD. Empirical backtesting and simulations over the period 2015–2025 tested the hypothesis that dynamic Bayesian updating enhances risk-adjusted returns in non-stationary markets.

The comparative analysis reveals three principal findings. First, the DBBL model shows a modest improvement in risk-adjusted efficiency relative to the benchmark models, as reflected in its higher Sharpe Ratio. This suggests that dynamic Bayesian updating may enhance return-to-risk performance, but the evidence does not support a claim of uniform superiority across all performance dimensions. Rolling Bayesian updates and dynamic covariance estimation enable real-time exposure adjustment—reducing high-beta positions in TSLA, NVDA, and AMZN during market spikes while maintaining defensive allocations in JNJ, XOM, and GLD. This directly addresses the core limitation of static models that rely on fixed covariance assumptions. Second, DBBL generates well-diversified and temporally smooth weight structures, avoiding the corner solutions that frequently plague Markowitz optimization. The model achieves flexible risk diversification aligned with market regimes—shifting toward GLD and XOM during the 2022 energy crisis while sustaining moderate technology exposure. Third, the generalized DBBL (GBL) extension should be interpreted as a conceptual direction for future research rather than as completed empirical evidence in this study. Its intended role is to address potential misspecification in priors and subjective views through adaptive Bayesian updating and related learning mechanisms.

In conclusion, DBBL should be interpreted as a risk-aware adaptive allocation framework that shows a modest improvement in Sharpe Ratio in the reported sample, while also involving trade-offs in cumulative return, drawdown, and turnover relative to static benchmarks.

5.2. Recommendations for future research

Based on the study's findings, five directions for future research are proposed. First, future studies should combine LSTM, transformer-based temporal models, and DRL to construct dynamic investor views for assets exhibiting high regime sensitivity, such as TSLA, NVDA, and GLD, potentially forming the foundation for fully autonomous asset allocation systems. Second, DBBL should be extended to incorporate multi-factor priors, including momentum, value, quality, volatility, and macroeconomic indicators such as gross domestic product (GDP) growth and the volatility (VIX) index, which may better capture cross-asset dynamics than single-factor CAPM assumptions. Third, the model should be tested on a broader asset universe, including cryptocurrencies, real estate exchange-traded fund (ETFs), and international equities, and evaluated under simulated extreme scenarios such as COVID-19-scale market crashes to assess tail risk robustness. Fourth, transaction cost optimization via threshold-based rebalancing and Bayesian Liquidity Modeling should be integrated to bring the backtesting framework closer to real-world institutional constraints. Fifth, embedding DBBL within stochastic dynamic programming (SDP) or MPC architectures would enable multi-period joint optimization of risk and return, which is particularly relevant for long-horizon funds such as pension and insurance portfolios.

1. Integration of advanced ML techniques

- Future studies should combine LSTM, transformer-based temporal models, and DRL to construct dynamic investor views (Q_t) for assets like TSLA, NVDA, and GLD.
- ML integration could automate parameter estimation and improve market behavior capture, forming the foundation for autonomous asset allocation systems.

2. Extension toward multi-factor Bayesian models

- Extend DBBL to include multi-factor explanatory variables such as momentum, value, quality, volatility, and macroeconomic indicators (GDP growth, VIX).
- Multi-factor priors may better capture cross-asset dynamics (e.g., tech vs. defensive sectors) than single-factor CAPM assumptions.

3. Testing on alternative assets and crisis scenarios

- Include cryptocurrencies, commodities beyond GLD, and real estate ETFs.

- Simulate extreme events such as COVID-19 or bear markets to assess tail risk, drawdown management, and robustness of DBBL across diverse asset classes.

4. Transaction cost management and rebalancing strategy

- Optimize turnover via threshold-based rebalancing or transaction-cost-aware constraints to reduce costs while maintaining dynamic adaptation.
- Incorporate Bayesian liquidity modeling to account for market depth for large institutional positions in assets like AAPL, NVDA, and XOM.

5. Structural extension toward multi-period stochastic control

- Integrate DBBL within SDP or MPC to jointly optimize multi-period risk and return.
- Relevant for long-term funds (pension, insurance) requiring forward-looking strategic allocation across tech, financial, energy, and commodity assets.

5.3. Conceptual reflection

This study demonstrates that the DBBL framework is more than an incremental improvement—it represents a conceptual transformation in portfolio construction, combining:

- Bayesian inference for robust parameter estimation,
- Dynamic systems theory for adaptive portfolio evolution,
- ML for predictive and automated view generation.

Across a diversified universe including AAPL, MSFT, AMZN, GOOG, TSLA, JNJ, JPM, NVDA, META, XOM, GLD, and DBBL, exemplifies proactive risk management, continuous learning, and adaptive decision-making.

6. CONCLUSION

The dynamic Bayesian Black-Litterman model represents a significant evolution in portfolio management, integrating Bayesian inference with sequential learning to dynamically adjust expected returns and risk estimates. Unlike traditional Markowitz MV or SBL models, DBBL continuously updates portfolio allocations in response to changing market conditions, reducing estimation errors and mitigating the concentration of weights in high-risk assets. Empirical results using the selected asset universe indicate that DBBL achieves a higher Sharpe Ratio than the benchmark models and exhibits adaptive portfolio adjustment. However, these gains are accompanied by lower cumulative return and deeper maximum drawdown, suggesting that DBBL offers improved risk-adjusted efficiency rather than uniform outperformance.

By incorporating advanced ML techniques such as LSTM and DRL, DBBL generates predictive investor views that adapt to short- and medium-term market signals, enhancing portfolio resilience and responsiveness. Its dynamic covariance estimation, rolling Bayesian updates, and bias correction mechanisms ensure well-diversified and temporally stable portfolios, capable of mitigating downside risk during market volatility.

In summary, DBBL may be viewed as an adaptive portfolio allocation framework that improves risk-adjusted efficiency under changing market conditions. At the same time, its empirical performance reflects clear trade-offs in cumulative return, drawdown, and turnover. Accordingly, the model should be interpreted as a promising extension of Black–Litterman for non-stationary environments, rather than as a universally superior portfolio solution.

This study acknowledges several important limitations. First, the empirical analysis is confined to 11 U.S.-listed assets over 2015–2025, which limits the generalizability of findings to other geographic markets or alternative asset classes such as fixed income or cryptocurrencies. Second, the ML components—particularly LSTM and DRL—demand significant computational resources and large training datasets, which may pose practical constraints for smaller investment managers. Third, the backtesting framework does not fully incorporate transaction costs or market impact, which may cause reported performance metrics to overstate real-world net returns. Fourth, Bayesian priors are derived from historical data and may underperform during unprecedented structural breaks or black swan events outside the training sample. Future research should address these limitations by extending the framework to global multi-asset portfolios, incorporating realistic cost constraints, and stress-testing model robustness across diverse crisis scenarios.

The findings of this study carry significant implications for both academic research and investment practice. From a theoretical perspective, DBBL demonstrates that Bayesian inference and ML can be meaningfully unified within a coherent portfolio optimization framework, contributing to the growing literature on AI-driven financial modeling (Sun et al., 2023; Wang, 2024). From a practical standpoint, asset managers and quantitative funds can leverage the DBBL framework to construct portfolios that are more responsive to regime changes, more robust to estimation errors, and better aligned with dynamic risk management objectives. The framework is particularly suitable for multi-asset institutional portfolios that require continuous rebalancing under non-stationary market conditions, such as those spanning technology, energy, healthcare, and commodity sectors.

REFERENCES

Barua, R., & Sharma, A. K. (2022). Dynamic Black Litterman portfolios with views derived via CNN-BiLSTM predictions. *Finance Research Letters*, *49*, Article 103111. https://doi.org/10.1016/j.frl.2022.103111

Ben-Tal, A., El Ghaoui, L., & Nemirovski, A. (2009). *Robust optimization.* Princeton University Press. https://doi.org/10.1515/9781400831050

Bessler, W., Opfer, H., & Wolff, D. (2012). Multi-asset portfolio optimization and out-of-sample performance: An evaluation of Black-Litterman, mean-variance, and naïve diversification approaches. In *Annual conference of the association for social policy 2012: New paths and challenges for the labor market of the 21st century.* Econstor. https://www.econstor.eu/handle/10419/62020

Best, M. J., & Grauer, R. R. (1991). Sensitivity analysis for mean-variance portfolio problems. *Management Science*, *37*(8), 980–989. https://doi.org/10.1287/mnsc.37.8.980

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

Black, F., & Litterman, R. B. (1992). Global portfolio optimization. *Financial Analysts Journal, 48*(5), 28–43. https://doi.org/10.2469/faj.v48.n5.28

Chen, S. D. (2014). *Dynamic Bayesian learning and optimization in portfolio choice models* [Doctoral dissertation, University of California]. https://escholarship.org/uc/item/50c395hk

Cont, R. (2001). Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance*, *1*(2), 223–236. https://doi.org/10.1080/713665670

DeMiguel, V., Garlappi, L., & Uppal, R. (2009). Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *The Review of Financial Studies*, *22*(5), 1915–1953. https://doi.org/10.1093/rfs/hhm075

Engle, R. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics, 20*(3), 339–350. https://doi.org/10.1198/073500102288618487

Fischer, E. O., & Seidl, I. (2013). *Portfolio optimization: A combined regime-switching and Black-Litterman model.* https://doi.org/10.2139/ssrn.2205190

Idzorek, T. (2007). A step-by-step guide to the Black-Litterman model: Incorporating user-specified confidence levels. In S. Satchell (Ed.), *Forecasting expected returns in the financial markets* (pp. 17–38). Academic Press. https://doi.org/10.1016/B978-075068321-0.50003-0

Jorion, P. (1996). Risk^2: Measuring the risk in value at risk. *Financial Analysts Journal*, *52*(6), 47–56. https://doi.org/10.2469/faj.v52.n6.2039

Ko, H., & Lee, J. (2025). Portfolio management transformed: An enhanced Black–Litterman approach integrating asset pricing theory and machine learning. *Computational Economics, 66*, 3841–3887. https://doi.org/10.1007/s10614-024-10760-9

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

Michaud, R. O. (1989). The Markowitz optimization enigma: Is 'optimized' optimal? *Financial Analysts Journal, 45*(1), 31–42. https://doi.org/10.2469/faj.v45.n1.31

Michaud, R. O., & Michaud, R. (2008). *Estimation error and portfolio optimization: A resampling solution.* https://doi.org/10.2139/ssrn.2658657

Oprisor, R., & Kwon, R. (2021). Multi-period portfolio optimization with investor views under regime switching. *Journal of Risk and Financial Management, 14*(1, special issue), Article 3. https://doi.org/10.3390/jrfm14010003

Rinne, M. (2025). *A dynamic asset allocation strategy integrating long short-term memory (LSTM) with the Black-Litterman model* [Master's thesis, Lappeenranta-Lahti University of Technology LUT]. https://lutpub.lut.fi/handle/10024/169236

RiskMetrics. (1996). *RiskMetrics technical document* (4th ed.). J.P. Morgan/Reuters. https://www.msci.com/documents/10199/5915b101-4206-4ba0-aee2-3449d5c7e95a

Sun, R., Stefanidis, A., Jiang, Z., & Su, J. (2023). *Combining transformer-based deep reinforcement learning with Black-Litterman model for portfolio optimization.* Cornell University. https://doi.org/10.48550/arXiv.2402.16609

Teplova, T., Mikova, E., Munir, Q., & Pivnitskaya, N. (2023). Black-Litterman model with copula-based views in mean-CVaR portfolio optimization framework with weight constraints. *Economic Change and Restructuring, 56*, 515–535. https://doi.org/10.1007/s10644-022-09435-y

Wang, J. (2024). Portfolio construction based on LSTM RNN and Black-Litterman model: Evidence from Yahoo Finance. In *Proceedings of the 1st international conference on e-commerce and artificial intelligence* (Vol. 1, pp. 337–343). SciTePress. https://doi.org/10.5220/0013225400004568

Wang, Y., & Aste, T. (2022). *Dynamic portfolio optimization with inverse covariance clustering*. Cornell University. https://doi.org/10.48550/arXiv.2112.15499