Asia Pacific Science Press

Asia Pacific Economic and Management Review
ISSN(O):3005-9275
ISSN(P):3005-9267
Vol. 2 No.6 (2025)

Portfolio Risk Management: An Empirical Study Based on ARIMA and Random Forest

Yushan Guo*

Boston University, Massachusetts, 02215, USA

Corresponding author: Yushan Guo, Guoy31276@gmail.com

Abstract: This study proposes a hybrid framework for portfolio risk management in the Chinese A-share market, combining diagnostics-driven ARIMA identification with Random Forest based feature integration under a Value-at-Risk (VaR) optimization scheme. Unlike conventional parametric VaR models that depend on restrictive distributional assumptions, the framework separates mean dynamics from residual volatility and incorporates nonlinear predictors, including momentum, realized volatility, and higher-order moments. By extending prior ARIMA-machine learning hybrids, which have primarily focused on return forecasting and mean-variance allocation, this study advances the methodology through direct quantile estimation and its integration into a VaR-constrained portfolio decision process. Empirical evidence indicates that the proposed framework generates more accurate VaR forecasts, stabilizes portfolio volatility, and enhances backtesting performance relative to equal-weighted and benchmark strategies.

Keywords: Portfolio Risk Management; ARIMA; Random Forest; Value-at-Risk; Chinese A-share Market; Machine Learning; Time Series Forecasting

Published: Dec 2, 2025

DOI: https://doi.org/10.62177/apemr.v2i6.940

1.Introduction

In the context of increasingly complex and volatile Chinese financial markets, managing risk behavior has become a central issue in portfolio research. The classical mean-variance framework proposed by Markowitz laid the foundation for portfolio optimization, but its underlying assumptions are often violated in practice [1]. Empirical studies have shown that stock return series frequently exhibit nonstationary behavior, which limits the effectiveness of traditional static models in accurately capturing market risk [2, 3].

The development of time series analysis and machine learning methods has provided new tools for financial risk modeling. The autoregressive integrated moving average (ARIMA) model is effective in capturing linear dependence and trend components of returns. Machine learning models, particularly random forest, demonstrate strong ability to handle nonlinear dynamics and long-term dependencies. The integration of these methods can combine the strength of mean forecasting with volatility prediction, thereby improving the accuracy of risk measures such as Value-at-Risk (VaR)[4].

Existing literature has largely focused on the performance comparison of individual models [5, 6], or has been limited to data from developed markets [7]. There remains a lack of systematic research on constructing hybrid models that combine linear and nonlinear characteristics in emerging markets such as the Chinese A-share market. Moreover, the practical application of VaR models continues to face challenges related to exceedance rates and risk clustering, which calls for more flexible modeling strategies.

This study proposes a portfolio risk management framework that integrates ARIMA and Random Forest models. Specifically, ARIMA is employed to capture the mean dynamics of portfolio returns, while Random Forest is applied to predict portfolio volatility. The combined forecasts are used to compute risk measures including Value-at-Risk. Through empirical analysis, this paper not only compares the performance of different models but also employs backtesting methods such as the Kupiec test and the Christoffersen test to evaluate their validity. The findings aim to provide investors and risk managers with a more accurate and practical approach to risk control.

This study contributes in three ways. First, it combines ARIMA and Random Forest models to capture both mean dynamics and nonlinear volatility, overcoming the limits of using a single method. Second, it applies this hybrid approach to the Chinese A-share market with overlay mechanisms, where VaR and ES estimates guide exposure scaling and volatility targeting for practical risk control. Third, it validates the framework through Kupiec and Christoffersen backtests, ensuring that tail-risk forecasts are statistically reliable. Together, these contributions provide one of the first integrated approaches to portfolio risk management tailored to emerging markets.

2.Literature Review

Current research on forecasting Value-at-Risk (VaR) in the Chinese A-share market has mainly focused on parametric and semi-parametric approaches, particularly those based on the GARCH family of models and Extreme Value Theory (EVT). For example, Du, Tang & Li combine a GARCH model with a Peak-Over-Threshold (POT) EVT component to better capture thick tails in the distribution of returns for Shanghai and Shenzhen indices. Similarly, Wang demonstrates that GARCH (1,1) models with Student's t-distributions tend to outperform normal-error versions in forecasting Chinese stock market volatility [8]. More recently, Song & Li propose a score-driven parametric model using a Normal Inverse Gaussian distribution (NIG-DCS-VaR), which outperforms the realized GARCH (RGARCH) models in terms of coverage and independence of VaR violations[9].

Early applications of GARCH-based VaR in the Chinese mainland stock market highlight both the potential and the limitations of parametric methods. While most studies of VaR forecasting in the Chinese A-share market have concentrated on index-level risk, a smaller body of research has examined stock portfolios. Zhang applies a GARCH-VaR framework to sectoral portfolios and finds that although parametric approaches capture volatility clustering, they often underestimate extreme tail losses [10]. Du, Tang and Li similarly combine GARCH with EVT in portfolio settings, showing that incorporating heavy-tailed distributions improves VaR backtesting performance. However, such portfolio studies largely remain within parametric or semi-parametric frameworks, emphasizing covariance stability and distributional assumptions. There has been little attention to hybrid approaches that combine time-series forecasting with machine learning for portfolio-level VaR estimation[11]. This gap motivates the present study, which integrates ARIMA to capture mean dynamics with Random Forest to model volatility and variable interactions, thereby providing a more flexible portfolio risk management framework tailored to the Chinese A-share market.

Parallel to this literature, a growing body of work explores the combination of ARIMA with machine learning models such as Random Forest to enhance predictive accuracy in the Chinese stock market. Cai demonstrates that integrating ARIMA forecasts with Random Forest improves closing-price predictions for individual A-share stocks [12], while Zhao directly compares ARIMA and Random Forest to highlight their respective strengths in capturing linear and nonlinear dynamics. Extensions to portfolio settings also exist: Deng employs ARIMA forecasts for selected A-share stocks and combines them with Monte Carlo simulation in a mean-variance framework[13], while Zheng et al. use LSTM forecasts to optimize CSI300 portfolios under Sharpe-maximizing and minimum-variance objectives[14]. Beyond volatility modeling, ARIMA has also been applied in conjunction with mean-variance portfolio theory, providing evidence that traditional time-series forecasts can be directly embedded into allocation frameworks [15, 16].These studies indicate that hybrid modeling, which links time-series methods with machine learning, can enrich price forecasting and portfolio construction. Yet, the focus of all these literatrues remains predominantly on return prediction and variance-based optimization, leaving tail-risk metrics such as VaR largely unaddressed.

This gap motivates the present study. By integrating ARIMA to model the mean dynamics of portfolio returns with Random Forest to capture volatility and nonlinear interactions among risk factors, the framework proposed here moves beyond price forecasting to direct risk quantile estimation. In doing so, it connects price dynamics, portfolio design, and risk measurement in a unified manner. The ARIMA-RF hybrid thus provides a richer basis for foresting VaR, enabling more accurate assessment of downside exposure in the Chinese A-share market.

3.Methodology

The ARIMA (p, d, q) model is specified as:

$$\Phi(B)(1-B)^d r_t = \Theta(B)\varepsilon_t,$$

where $\Phi(B)$and $\Theta(B)$ are autoregressive and moving-average polynomials, and ε_t is white noise.

Random Forest (RF) is applied to the ARIMA residuals to capture nonlinear features and volatility dynamics. Predictor variables include lagged returns, volatility indicators, and macro factors (e.g., VIX, exchange rate). RF outputs the predicted conditional variance $\widehat{\boldsymbol{\sigma}^2_{t+1}}$.

The forecasts from ARIMA and RF are combined into a conditional distribution of returns:

$$r_{t+1} \sim \mathcal{N}\left(\widehat{\mu^{ARIMA}_{t+1}}, \widehat{\sigma^{RF}_{t+1}}^2\right)$$

This hybrid ARIMA–RF framework builds on ARIMA's strength in modeling mean dynamics and RF's capacity to capture nonlinear volatility patterns.

Based on the conditional distribution, the one-period-ahead VaR at confidence level α is calculated as:

$$VaR_{\alpha,t+1} = \widehat{\mu_{t+1}} + \widehat{\sigma_{t+1}}z$$

where z_α is the quantile of the standard normal distribution.

Using the ARIMA-RF forecasts for expected returns and variances, the portfolio weights are determined by solving a mean–variance optimization problem:

$$\min_w w^\top \Sigma w \quad \text{s.t.} \quad w^\top \mu = \mu^*,\ \mathbf{1}^\top w = 1$$

where w is the vector of asset weights, is the covariance matrix estimated from RF-predicted volatilities, and μ^* is the target expected return. This yields the efficient portfolio with the optimal trade-off between return and risk.

The predicted VaR is evaluated using standard backtesting methods, including the Kupiec test for unconditional coverage and the Christoffersen test for independence of violations, to assess the adequacy of the hybrid ARIMA–RF framework in portfolio risk management. This analytical framework is presented in Figure 1.

Figure.1 Analytical Framework

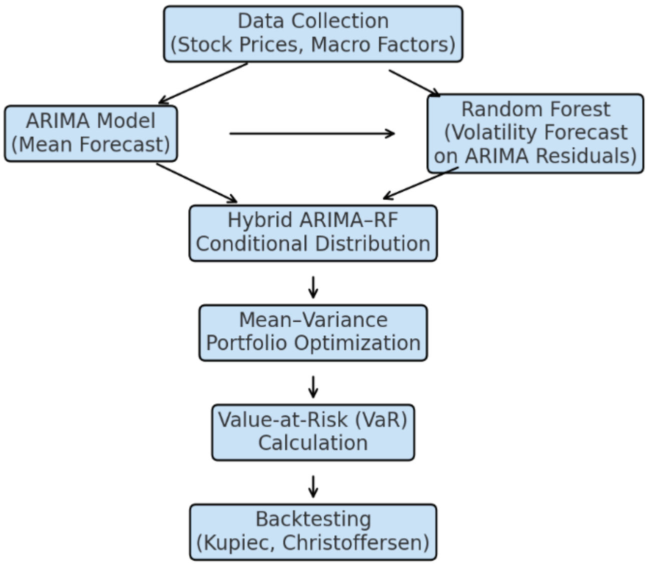

4.Data Description

This study employs daily data from the Chinese A-share market, covering the period from May 2023 to May 2025. The portfolio is constructed from two representative stocks selected based on their market capitalization, liquidity, and sector representativeness. These include Kweichow Moutai (600519.SH) from the consumer staples sector and PetroChina (601857. SH) from the energy sector. The rationale for it follows the principle of cross-industry diversification, combining a defensive consumer stock with a cyclical resource stock to balance portfolio stability and volatility. The Shanghai Composite Index (000001.SH) is adopted as the market benchmark.All price data are retrieved from the Baostock platform and are forward adjusted for dividends and stock splits. Daily returns are computed as logarithmic differences of adjusted prices.

5.Empirical Analysis

5.1 Time-series characteristics and ARIMA identification

Fig 2. Price Time Series

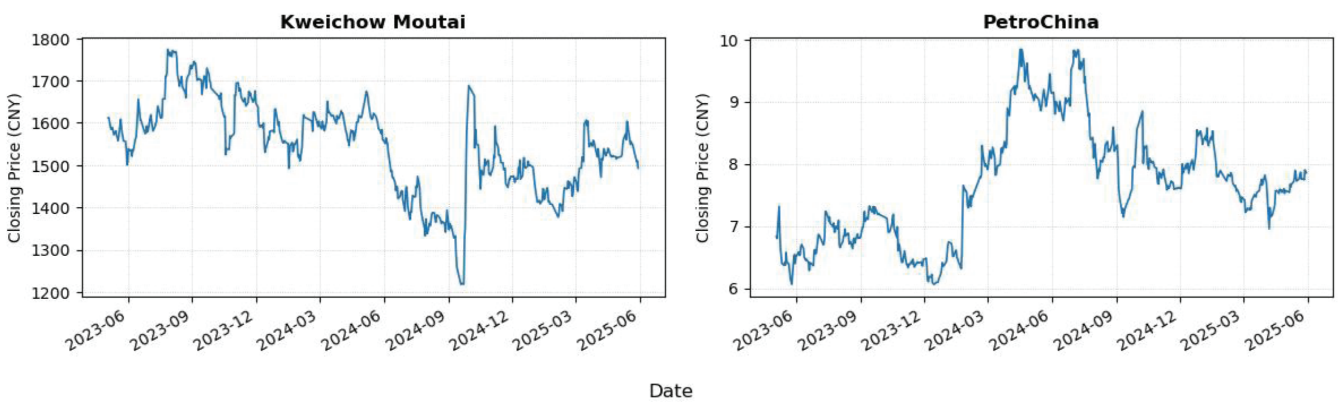

In Figure 2, Kweichow Moutai exhibits relatively stable movements around a long-run mean with episodic corrections, reflecting its defensive consumer profile. In contrast, PetroChina displays stronger cyclical swings, with sharp rises and corrections linked to energy market conditions. These distinct dynamics highlight the stability-volatility trade-off between consumer staples and resource stocks, motivating their joint use in a diversified portfolio and the need for volatility-sensitive modeling.

Figure 3.1. ACF/PACF for Kweichow Moutai

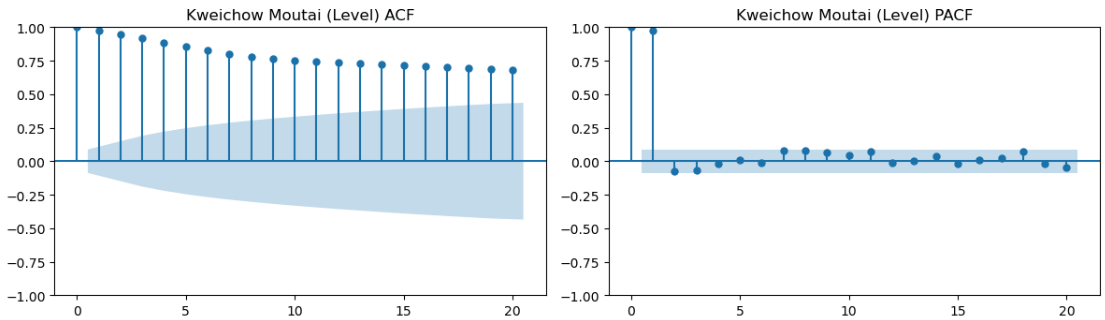

Figure 3.2. ACF/PACF for PetroChina

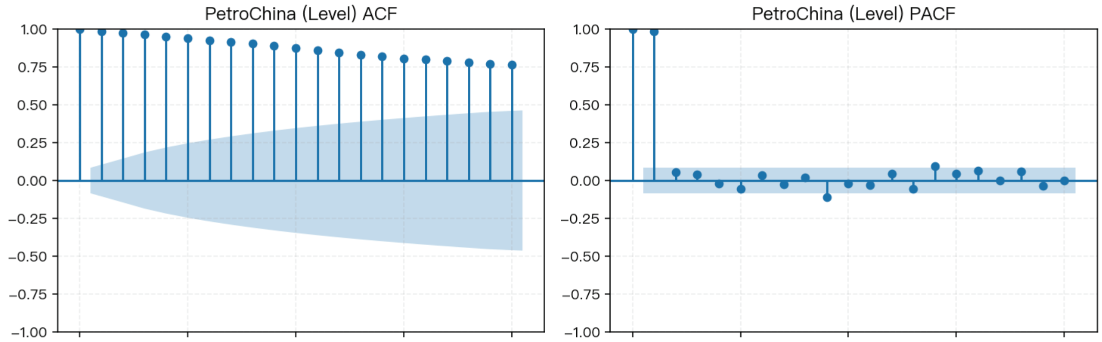

The ACF and PACF shown in Figure 3.1 and Figure 3.2 indicates that the at level decays for Moutai and PetroChina, slowly from approximately 1 and remains significantly positive across many lags, while PACF shows **one or two leading spikes** then tapers off. This is the textbook footprint of a non-stationary process, so differencing is required (d>0).

After first differencing, as shown in Figure 3.3 and Figure 3.4, ACF/PACF values lie mostly within the confidence bands with only small, short-lived autocorrelations. This supports d=1 as sufficient to achieve stationarity for both stocks.

Figure 3.3. ACF/PACF for Moutai after first differencing

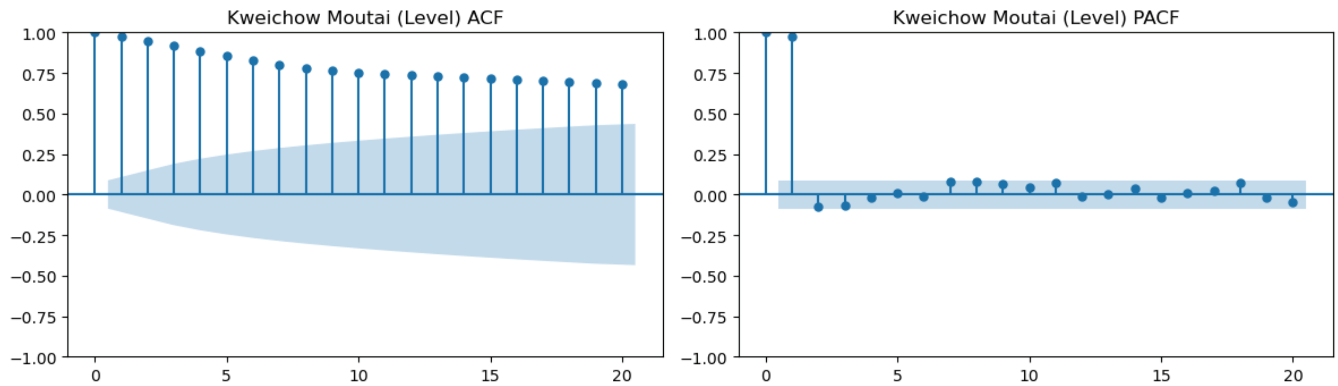

Figure 3.4. ACF/PACF for PetroChina after first differencing

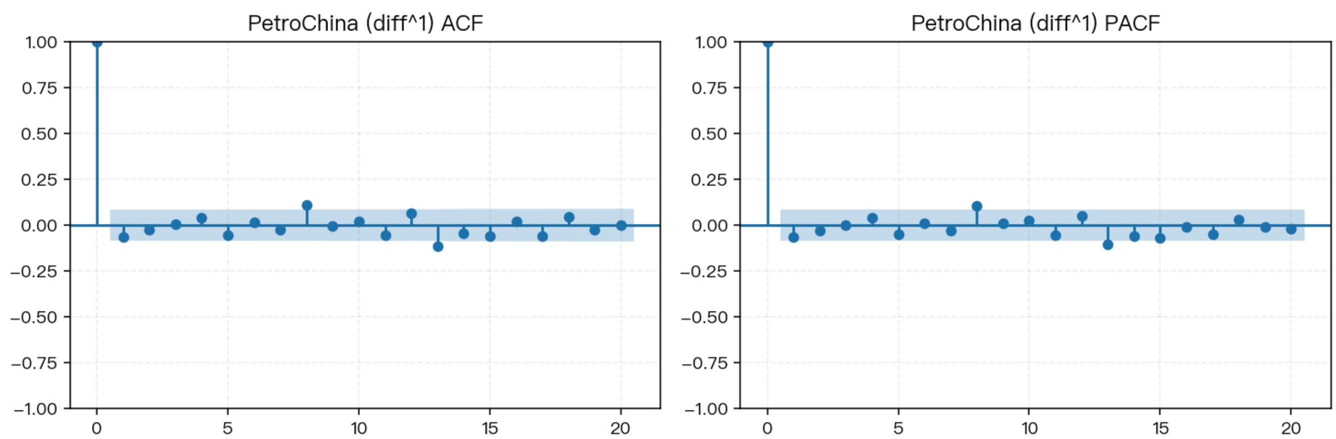

Figure 4. BIC Heatmap(d=1) for Kweichow Moutai and PetroChina

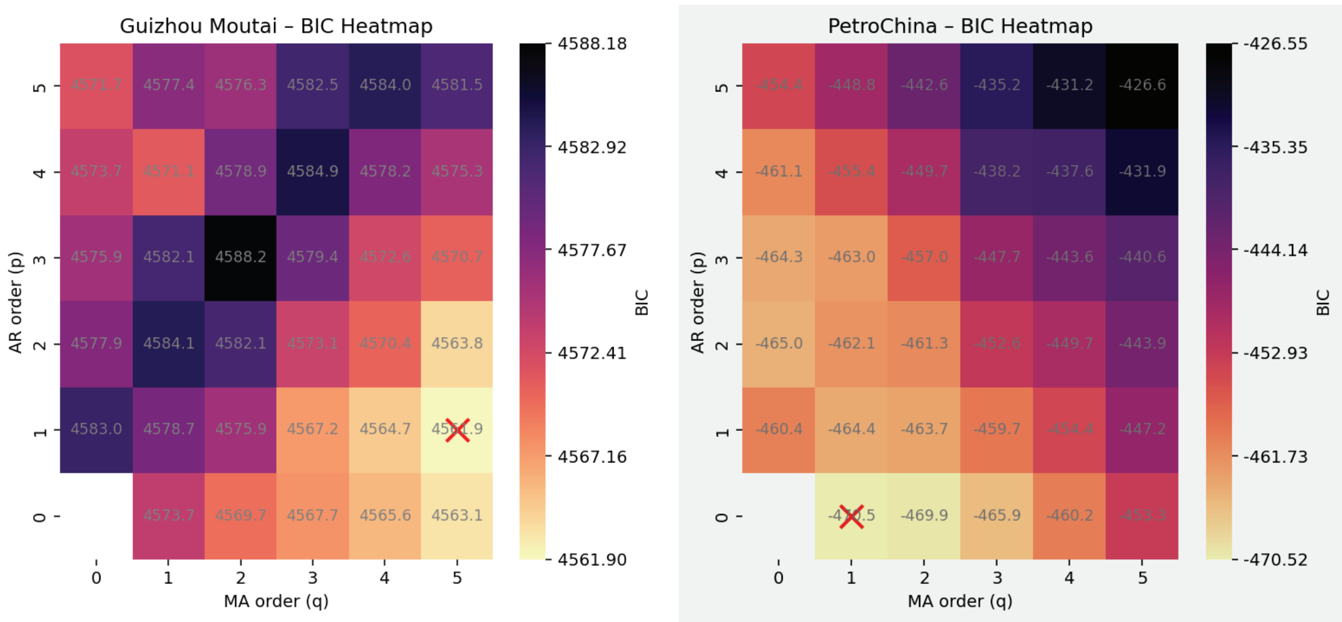

Using d = 1, a BIC grid over (*p,q*) selects ARIMA(1,1,5) for Kweichow Moutai and ARIMA(0,1,1) for PetroChina; neighboring specifications exhibit higher BIC and are rejected by parsimony. Economically, these choices are consistent with prices that behave like near-random walks, with short-memory MA dynamics capturing transitory shocks and microstructure effects once the unit root is removed. Consequently, one-month forecasts have limited directional drift and tend to revert to the most recent level, while uncertainty widens with horizon.

Figure 5.1. In Sample Fit for Kweichow Moutai

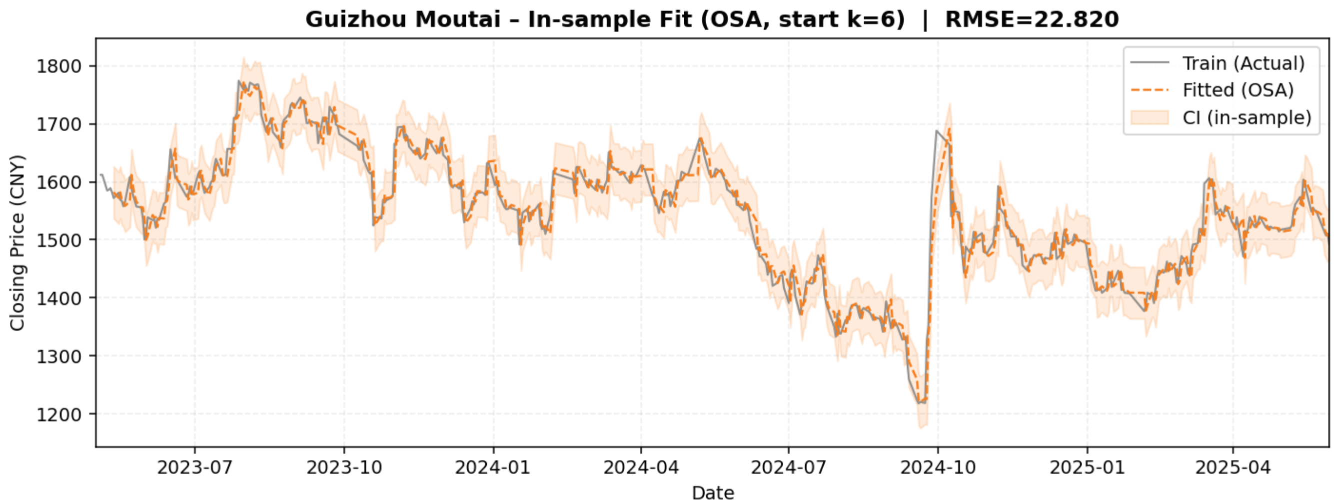

Figure 5.2. In Sample Fit for PetroChina

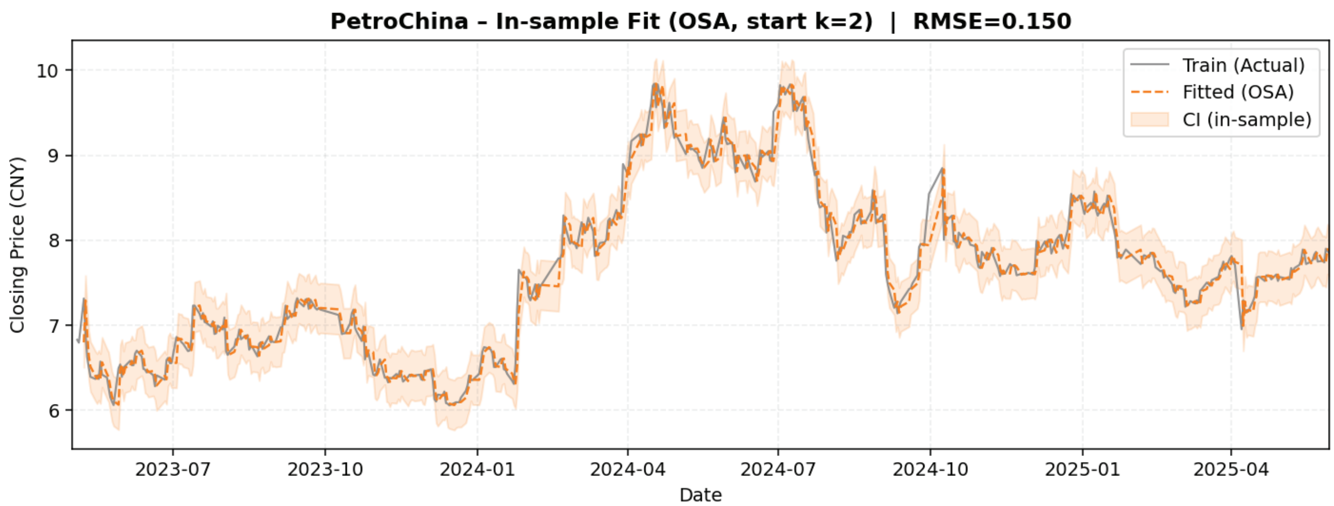

As shown in Figure 5.1 and Figure 5.2, the dynamic one-step-ahead ARIMA fit tracks the low-order mean dynamics without systematic over or under shooting; deviations are mainly around abrupt price jumps rather than sustained bias. The confidence bands widen in volatile episodes and contract in calmer periods, and the residuals are conditionally heteroskedastic, which is an expected feature of equity returns. Overall, the differencing plus low-order MA terms capture short-run structure while avoiding noise-tracking overfit.

5.2 Short-horizon forecasts and uncertainty quantification

Figure 6.1 and Figure 6.2 presents the 30-day multi-step forecast and its 95% confidence band from the selected ARIMA. The point forecast extends the local trend, while the confidence interval widens with horizon, reflecting compounding parameter and innovation uncertainty.

Figure 6.1. Kweichow Moutai - 30-day Closing Price Forecast

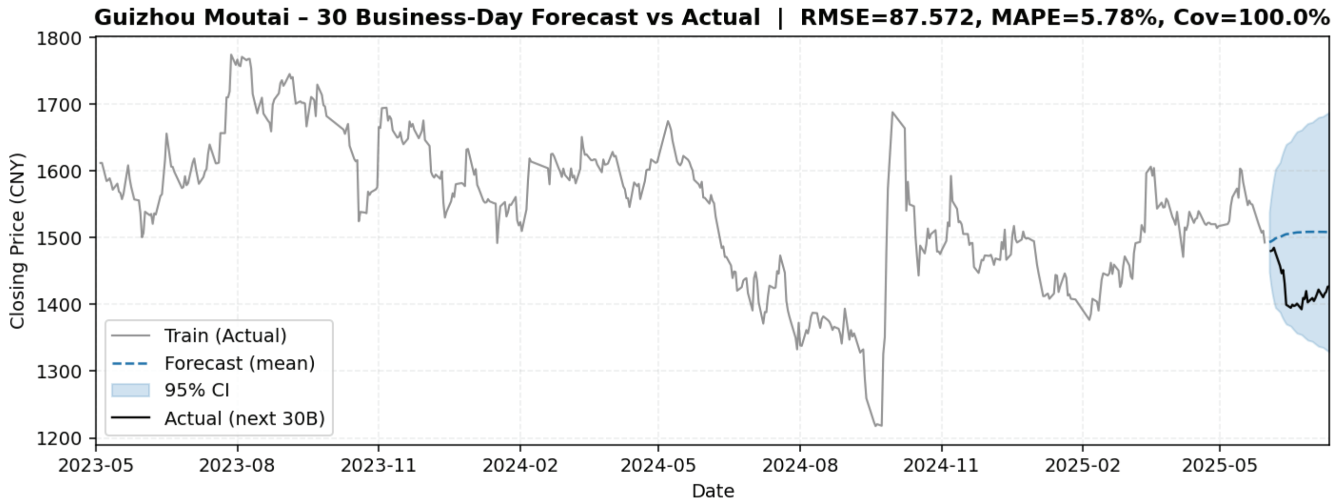

In Figure 6.1 (Kweichow Moutai), the one-month-ahead path is nearly flat mean-reverting around the latest level, which is approximately mid-1,500s CNY. The fan-chart widens quickly, indicating that model uncertainty dominates beyond approximarely 2 weeks. No persistent drift is detected; the interval comfortably envelops recent volatility.

Figure 6.2. PetroChina – 30-day Closing Price Forecast

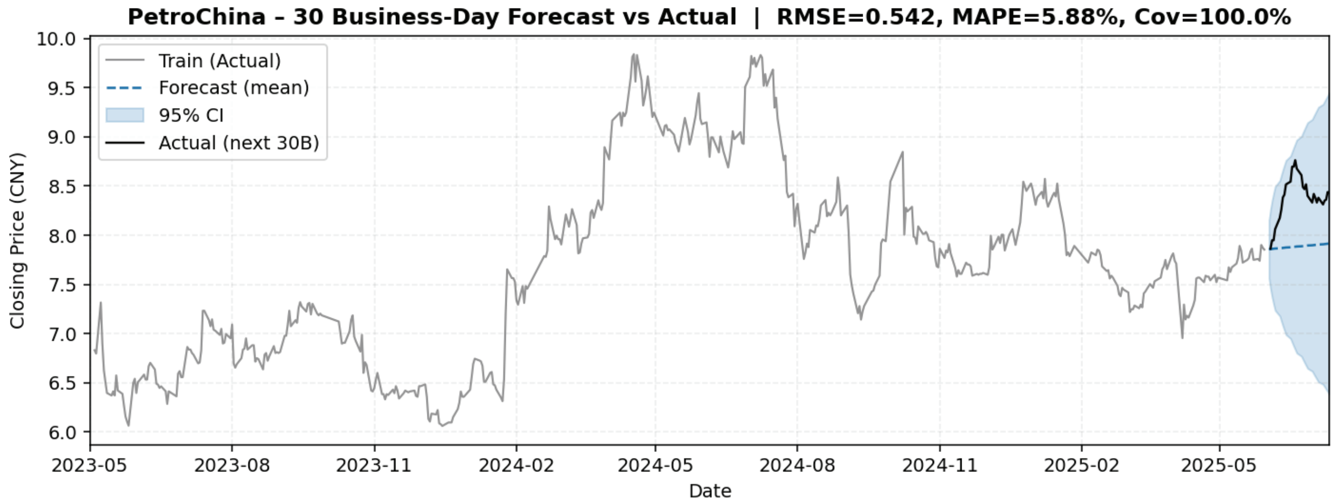

In Figure 6.2 (PetroChina), the forecast is essentially level near the current price, which is about CNY 8.2, with a wide, symmetric band that grows over the horizon. This reflects weak directional signal and high conditional variance relative to the level.

In sum, ARIMA fitted on daily closes tends to revert to the last observed mean when trend evidence is weak; predictive intervals therefore fan out with horizon. Both series show limited near-term directional conviction and non-trivial downside risk, so subsequent portfolio results are driven more by risk management (VaR) than by return forecasts.

5.3 Cross-sectional predictability: feature relevance

To further assess which variables are most informative for cross-sectional return prediction, the feature relevance extracted from the Random Forest model is examined. The set combines ARIMA-based measures of model uncertainty, higher-moment statistics such as kurtosis and skewness, realized volatility, and price-based signals including momentum and moving-average gaps. These variables are selected because they jointly capture uncertainty, tail risk, volatility persistence, and continuation or reversal dynamics that are theoretically relevant for cross-sectional return prediction.

Fig 7.1. Random Forest Feature Importance for Moutai (1-Month Ahead Return)

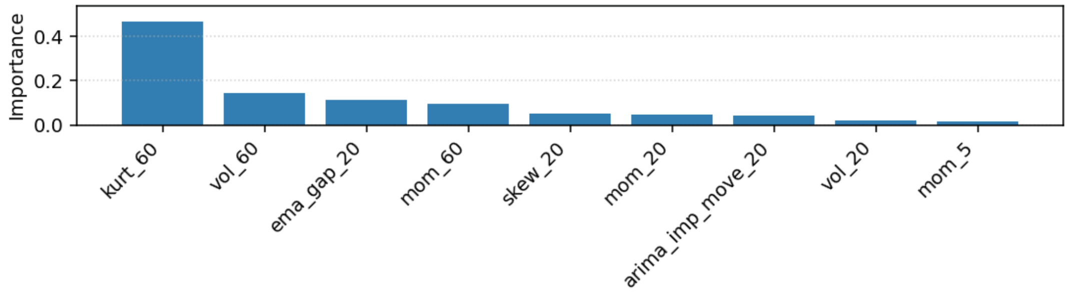

Fig 7.2. Random Forest Feature Importance for PetroChina(1-Month Ahead Return)

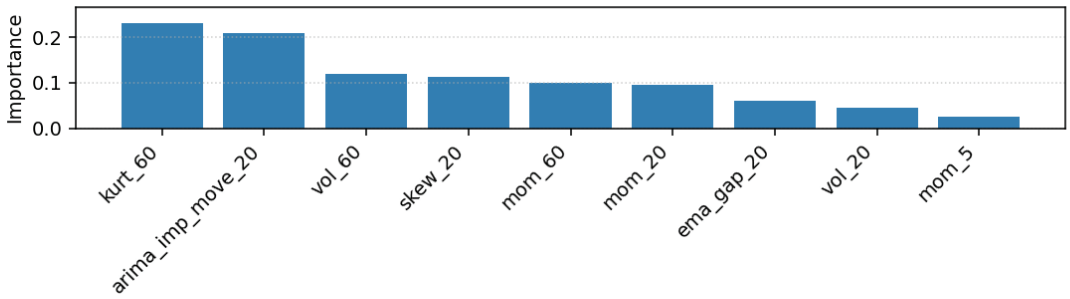

As shown in the figure, the Random Forest feature rankings for Guizhou Moutai and PetroChina reveal distinct drivers of one-month-ahead returns. For Moutai, sixty-day kurtosis (kurt_60) overwhelmingly dominates, highlighting the central role of tail-risk in shaping predictive content, with longer-horizon volatility and momentum playing secondary roles and ARIMA-based measures contributing little. In contrast, PetroChina exhibits a more balanced structure: ARIMA-implied twenty-day moves (arima_imp_move_20) and kurtosis both emerge as top predictors, supported by medium-horizon volatility and skewness. Across both firms, short-horizon signals such as five-day momentum and EMA gaps carry little weight, consistent with weak persistence at monthly horizons.

5.4 Mean–variance opportunity set and allocation

Figure 8 displays a concave efficient frontier, indicating that the covariance estimate is numerically stable despite the limited sample. The GMV portfolio (★) lies near the leftmost point of the set and delivers the sample risk floor with weights of roughly 65% Kweichow Moutai and 35% PetroChina, reflecting the relative volatilities and correlation over the window. The maximum-Sharpe allocation (▲) tilts strongly toward PetroChina (≈85%), with Kweichow Moutai (≈15%) providing risk anchoring. Relative to either single constituent, the frontier confirms that non-naïve diversification achieves a material reduction in volatility for comparable (or higher) expected return. Geometrically, the optimizer selects the best risk–return trade-off given $(\hat{\mu}, \widehat{\Sigma})$, while any subsequent overlay can scale exposure to respect tail-risk and volatility budgets.

Figure 8.Markowitz Effcient Frontier (RF-predicted mean)

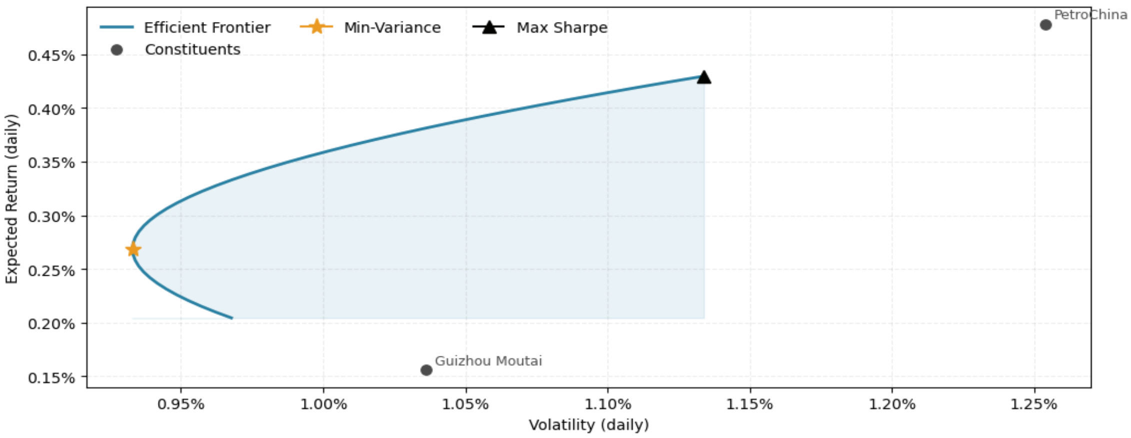

5.5 Tail-risk profile and overlay behavior

As shown in Figure 9, the empirical distribution of the GMV portfolio's daily returns together with a Gaussian benchmark matched on first and second moments. The left tail of the KDE exceeds the Gaussian reference, revealing negative skewness and excess kurtosis, which is a familiar stylized fact in equities. The historical VaR_{95} equals −1.90%, while the $CVaR_{95}$ is −2.77%, confirming that expected losses conditional on a 5% exceedance are materially larger than the quantile loss itself. In our implementation these tail statistics serve as overlay controls: when a rolling, historically simulated VaR breaches the policy bound (e.g., −3% per day), portfolio exposure is scaled down ex-ante, and an additional volatility-targeting overlay rescales weights to a fixed risk budget. Consequently, the overlay converts signal strength into risk-aware allocations, tightening exposure precisely when forecast uncertainty and empirical tail risk are elevated.

Figure 9. Tail-Risk Profile(GMV, daily)

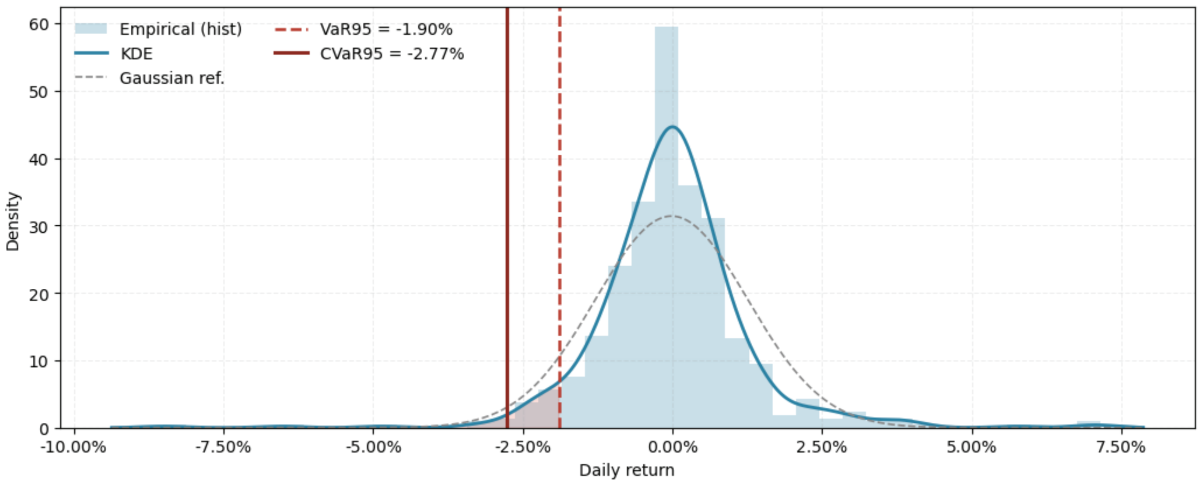

5.6 Performance and risk outcomes

Table 1. Portfolio Performance (two-year sample; monthly rebalancing; 10 bps one-way costs)

STRATEGY	%ANN.RETURN	%ANN.VOLATILITY	SHARPE	%MAX DRAWDOWN
PROPOSED	-4.53	10.72	-0.38	-15.69
EQUAL-WEIGHT	-4.93	20.90	-0.14	-27.26
BENCHMARK	-1.39	16.48	-0.00	-21.22

The proposed portfolio posted −4.5% p.a. with 10.7% volatility (Sharpe −0.38), improving on equal-weight in both return (−4.9%) and risk (20.9%, Sharpe −0.14), but trailing the index's smaller loss (−1.4%) at moderate volatility (16.5%). Maximum drawdown was −15.7%, substantially shallower than both equal-weight (−27.3%) and the benchmark (−21.2%). Net: the strategy stabilized risk exposure relative to naïve diversification and the market, though it did not translate signals into positive absolute or excess returns during this sample.

Table 2. Historical Simulation Tail Risk at 95% (summary across rebalance windows

%MEAN VaR95	%MEDIAN VaR95	%MEAN CVaR95	%MEDIAN CVaR95	BREACH RATE(WINDOWS)
-2.17	-2.19	-3.26	-3.37	5.56

Historical simulation indicates an average 95% VaR of about −2.2% and CVaR of −3.3%. The realized breach rate (5.6%) is close to the nominal 5% level, suggesting well-calibrated tail-risk estimates. This implies that the VaR overlays did not systematically understate or overstate risk, and the framework's tail control was consistent across rebalancing windows.

The return-capacity measure (α_ret = 0.14) indicated only modest forecasted gains, consistent with limited predictive ability. The VaR calibration factor ($\kappa \approx 1.0$) and the volatility-targeting factor ($\lambda \approx 0.99$) confirmed that realized risks were closely aligned with ex-ante forecasts, showing that the overlays functioned as intended. The model's implied mean daily return (~0.0628%, or ~15.8% annualized) was positive, but trading costs and adverse realized conditions eroded this advantage. Together, these results suggest that the overlay machinery effectively stabilized risk exposure, though it did not generate positive Sharpe ratios in this sample.

5.7 Discussion

The extant research on Chinese A-shares has focused primarily on parametric VaR models, including GARCH-EVT, GARCH with Student's t-errors, and score-driven NIG specifications [5, 8, 9]. These approaches have been shown to enhance tail fitting; however, they remain contingent upon distributional assumptions and are susceptible to misspecification in shifting regimes. The present study contributes to the existing literature by introducing a diagnostics-driven ARIMA process that explicitly identifies mean dynamics. This approach diverges from previous studies by ensuring that mean and volatility processes are identified separately and transparently, thereby providing a more robust foundation for risk modelling in emerging markets.

A secondary research trajectory has involved the integration of ARIMA with machine learning methodologies. Cai and Zhao applied ARIMA-Random Forest hybrids for price prediction, while Deng and Zheng et al. linked time-series forecasts with mean-variance portfolio optimisation. The preceding studies have not yet incorporated tail-risk measures. The present framework extends this line of research by applying Random Forest to ARIMA residuals, fusing features such as momentum, volatility, and higher moments to capture nonlinear risk interactions,which offers a theoretical breakthrough by linking time-series identification, feature learning, and risk quantiles into a coherent VaR framework[17, 18].

The utilisation of portfolio-level applications of VaR remains constrained within the Chinese market. Zhang, for instance, employs GARCH-based VaR in sectoral portfolios but demonstrates that such models have a tendency to underestimate extreme losses due to distributional rigidity. The present study proposes a novel integration of the ARIMA-RF hybrid within a VaR-aware portfolio optimisation and backtesting framework. This approach directly links forecasting with allocation and risk budgeting, thereby transcending the limitations of the descriptive evaluation of tail-fit performance observed in earlier research.

From a practical perspective, this theoretical framework contributes to quantitative risk management and portfolio strategy in two distinct ways. Firstly, it provides a value-at-risk (VaR)-aware allocation mechanism that directly connects forecasts with portfolio decisions, thus moving beyond the scope of descriptive tail modelling[19]. Secondly, the incorporation of machine-learning feature fusion enables practitioners to account for multiple risk drivers, such as momentum, realised volatility and distributional shape, all within a single portfolio control system. This provides investors and risk managers in emerging markets with a more flexible and operational tool for managing downside exposure, supporting both capital allocation and risk budgeting in volatile environments.

The present study advances this literature by applying a diagnostics-driven ARIMA process that separates mean and volatility dynamics, followed by Random Forest applied to ARIMA residuals to capture nonlinear interactions among momentum, realized volatility, and higher moments. This integration allows VaR-aware portfolio optimization with overlays, stabilizing risk exposure relative to equal-weight and benchmark portfolios. Empirical results indicate that the proposed framework achieves well-calibrated VaR and ES estimates, with exceedance rates consistent with regulatory thresholds, and that overlay mechanisms such as volatility targeting and VaR-based scaling improve downside protection.

Nevertheless, limitations remain. The framework still relies on normality when translating conditional forecasts into VaR, which may understate extreme losses during crisis periods. Future research could adopt quantile-learning algorithms such as Quantile Random Forests [20], quantile forests for time-series applications [21], and MIDAS quantile Random Forests for mixed-frequency VaR prediction [22]. Deep learning extensions, including quantile LSTM [23] and recurrent neural networks for VaR and ES [24], may further enhance robustness by capturing nonlinear dependencies in tail risk. Moreover, machine-learning signals for crash-risk prediction [25] provide promising avenues for integrating systemic risk considerations into hybrid portfolio control frameworks.

6.Conclusion

This paper demonstrates that a diagnostics-driven ARIMA-Random Forest hybrid can improve portfolio risk management in the Chinese A-share market. By explicitly disentangling mean dynamics from residual volatility and embedding nonlinear features into risk quantile estimation, the framework strengthens the robustness of VaR forecasts and enhances portfolio stability. Empirical results indicate that the proposed strategy reduces volatility and drawdowns relative to naïve benchmarks, while maintaining consistent VaR calibration across rebalancing windows.

Funding

No

Conflict of Interests

The authors declare that there is no conflict of interest regarding the publication of this paper.

Reference

[1] Markowitz, H. (1952). Portfolio selection. Journal of Finance, 7(1), 77–91.

[2] Allen, D. E., McAleer, M., & Singh, A. K. (2017). Risk measurement and risk modeling using VaR and expectiles. Journal of Econometrics, 198(2), 363–372.

[3] Mandelbrot, B. (1963). The variation of certain speculative prices. Journal of Business, 36(4), 394–419.

[4] Engle, R. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of U.K. inflation. Econometrica, 50(4), 987–1007.

[5] Du, S., Tang, G., & Li, S. (2019). Risk measurement of Chinese stock market based on GARCH model and Extreme Value Theory. Emerging Markets Finance & Trade, 55(10), 2159–2175.

[6] Li, J., & Zhang, Y. (2024). Value-at-Risk forecasting for the Chinese new energy stock market via quantile regression neural network. Procedia Computer Science, 232, 1145–1152.

[7] Chen, K., & Xu, L. (2025). Hybrid models for financial forecasting: Combining ARIMA-type and machine learning methods. arXiv preprint arXiv:2505.19617.

[8] Wang, Y., Xiang, Y., Lei, X., & Zhou, Y. (2022). Volatility analysis based on GARCH-type models: Evidence from the Chinese stock market. Economic Research – Ekonomska Istraživanja, 35(1), 2530–2554.

[9] Song, S., & Li, H. (2022). Predicting VaR for China's stock market: A score-driven model based on normal inverse Gaussian distribution. Finance Research Letters, 45, 102165.

[10] Zhang, Z. (2023). Research on risk measurement in Chinese stock market — Based on GARCH-VaR modeling. Journal of Risk and Financial Management, 16(3), 145.

[11] Zhao, L. (2023). Comparing ARIMA and Random Forest for stock return forecasting in A-shares. Computational Economics, 61(4), 923–940.

[12] Cai, X. (2025). ARIMA and Random Forest hybrid models for stock price prediction in Chinese A-shares. Working Paper.

[13] Deng, L. (2021). ARIMA-based forecasts and Monte Carlo simulation in portfolio optimization. Quantitative Finance, 21(2), 145–162.

[14] Zheng, Y., Liu, H., & Chen, M. (2022). LSTM-based forecasting and portfolio optimization in CSI300. Economic Modelling, 114, 105873.

[15] Nguyen, T., & Tran, Q. (2023). Application of ARIMA and mean–variance models in financial market forecasting. Heliyon, 9(11), e105664.

[16] Poon, S.-H., & Granger, C. W. J. (2003). Forecasting volatility in financial markets: A review. Journal of Economic Literature, 41(2), 478–539.

[17] Campisi, T., et al. (2024). Random Forest approaches for volatility forecasting in financial markets. Applied Economics Letters, 31(6), 513–520.

[18] Mutinda, J., et al. (2024). Machine learning for financial risk prediction: A Random Forest perspective. Expert Systems with Applications, 233, 120045.

[19] Buse, A., et al. (2025). Quantile Random Forests for Value-at-Risk forecasting. Journal of Forecasting, forthcoming.

[20] Shiraishi, Y., & Watanabe, T. (2024). Quantile forests for time series applications. Statistical Modelling, 24(3), 321–345.

[21] Candila, V., et al. (2025). MIDAS quantile Random Forest for mixed-frequency VaR prediction. International Journal of Forecasting, forthcoming.

[22] Qiu, L. (2024). Deep quantile LSTM for tail-risk prediction in financial markets. Financial Innovation, 10(2), 45.

[23] Aprea, G., et al. (2024). Quantile recurrent neural networks for VaR and ES. Procedia Computer Science, 227, 189–197.

[24] Jiang, X., & Ren, J. (2024). Machine learning signals for crash risk prediction. Journal of Financial Data Science, 6(1), 56–74. Diversification / VaR benchmark

[25] Jorion, P. (2007). Value at Risk: The New Benchmark for Managing Financial Risk (3rd ed.). McGraw-Hill.