An End-to-End Regime-Dependent Industry Rotation Strategy in China's A-Share Market

Qiuyu Shen, Hongsen Cheng*, Haifei Liu

School of Management and Engineering, Nanjing University, Nanjing 210023, China

*Corresponding author:502023150012@smail.nju.edu.cn

Abstract. This paper develops an integrated framework for industry rotation in China's A-share market by combining industry-level factor construction, deep-factor extraction, market-regime prediction, directional classification, and regime-dependent risk parity. Industry signals are first formed from Alpha-style predictors aggregated from representative constituent stocks and then enriched by a feed-forward neural network that learns latent deep factors. The market is subsequently classified into four interpretable states defined by volatility and rotation speed, and next-period regime probabilities are forecast with XGBoost. These probabilities affect both the directional signal layer and the covariance structure used in the allocation layer. Out-of-sample evidence for 2023-2024 shows that the full framework delivers stronger return stability, a higher Sharpe ratio, and better drawdown control than simpler equal-weight or non-regime alternatives. The empirical results indicate that regime information is valuable not only for return prediction, but also for dynamic portfolio risk allocation.

Keywords: industry rotation; market regime; deep factor; XGBoost; risk parity

1 Introduction

In recent years, the China A-share market has shown faster industry rotation, more frequent style switching, and a steady expansion of index-based investment products. These shifts have increased the importance of sector-level allocation because investors can increasingly express macro, style, and policy views through industry instruments rather than through concentrated single-stock positions. As a result, industry rotation has become both a research question in asset pricing and a practical question in portfolio management. [1-3]

At the same time, the return-prediction literature has moved from low-dimensional linear specifications toward high-dimensional and nonlinear learning systems. Machine-learning models are now widely used to process large sets of characteristics, capture nonlinear interactions, and improve out-of-sample forecasting performance. Yet many studies still treat prediction and portfolio construction as separate steps, so the final investment value of a signal is not evaluated jointly with the allocation rule through which it is implemented. [4-6]

T. G. Guan et al. (eds.), *Proceedings of the 2026 4th International Conference on Digital Economy and Management Science (CDEMS 2026)*, Advances in Economics, Business and Management Research 392,
https://doi.org/10.2991/978-94-6239-699-9_51

This separation is particularly limiting in industry rotation. The value of an industry score depends not only on whether it ranks industries correctly, but also on how the ranking is mapped into long-short baskets, how turnover is controlled, and how risk is distributed when market conditions change. Recent work on factor timing, ranking-oriented portfolio construction, and regime-aware allocation therefore suggests that signal generation and allocation design should be studied together rather than in isolation. [7-12]

Against this background, this paper studies whether industry rotation in China can be formulated as an integrated regime-dependent sequence in which factor mining, regime prediction, directional classification, and risk budgeting are all connected. The analysis is conducted at daily frequency on Shenwan Level-1 industries and focuses on the full empirical chain from feature construction to realized backtest performance. [2,8,13-15]

The paper contributes in two ways. First, it formulates industry rotation as a meso-level end-to-end allocation problem rather than as a disconnected set of forecasting tasks. Second, it extends regime information beyond signal interpretation and makes regime probabilities enter the covariance structure used for risk parity, thereby allowing the allocation layer to adapt to different market environments.

2 Recent Literature and Research Gap

The literature on industry rotation originates from the broader momentum and cross-sectional return-predictability literature. Industry components explain an important share of stock-level continuation, which implies that industry allocation can contain independent economic information. Recent studies have also linked industry predictability to business-cycle conditions, return spillovers, and volatility transmission; however, the profitability of simple business-cycle sector rotation is less robust than earlier narratives suggested once stricter out-of-sample standards are applied. [1-3,13]

A second stream of literature concerns factor construction and machine learning. Traditional factor models remain essential for economic interpretation, but machine-learning studies show that large sets of predictors can be reorganized into more informative nonlinear representations. This shift is especially relevant for industry rotation because industry signals aggregate information from many firms and therefore naturally benefit from representation learning. [4-6,16,17]

A third strand emphasizes that sorting and timing should be studied jointly. Factor momentum, portfolio-characteristic timing, and learning-to-rank approaches all suggest that a predictive signal should ultimately be judged by how it changes portfolio ordering, holding periods, turnover, and realized risk-adjusted returns. This idea aligns closely with the present paper, which evaluates model layers through their contribution to final portfolio performance rather than through prediction metrics alone. [7,8,18]

The final gap concerns state dependence. Regime-switching research has established that returns, volatility, and cross-asset dependence vary across market states. Recent tactical-allocation studies extend this insight by using interpretable regime probabilities or macro-financial clustering. Nevertheless, these developments are still seldom integrated into one unified meso-level framework for China's industry-rotation problem. This paper addresses that gap by connecting deep-factor extraction, regime prediction, directional classification, and regime-dependent risk parity in a single empirical design. [9-14]

3 Methodology

3.1 Industry-Level Factor Construction and Deep-Factor Extraction

The empirical universe consists of Shenwan Level-1 industries observed at daily frequency. Because many established predictors are defined at the stock level, the framework follows a bottom-up aggregation logic. Representative constituent-level Alpha-style signals are first computed, and the resulting information is aggregated to construct industry-level factor features. This design preserves micro-level information while ensuring that the predictive signals are aligned with the actual industry-level trading universe. [4,5,16]

On top of the handcrafted factor panel, the paper introduces a feed-forward neural network that learns latent deep factors. The purpose of this module is not only to reduce dimensionality, but also to capture nonlinear interactions and common structures that may not be visible in the original signals. The deep-factor layer is evaluated through rolling-window diagnostics, factor-return trajectories, and cross-factor correlation checks, which together support its use in the downstream ranking and allocation stages. [5,6,17]

3.2 Regime Definition and Two-Layer XGBoost Prediction

To keep regime construction interpretable, the market state is defined using two variables: 20-day market volatility and 20-day industry rotation speed. Each day is assigned to one of four regimes according to whether volatility and rotation speed are above or below their respective medians. This approach preserves economic meaning and is better suited to empirical interpretation than a fully latent classification that offers little direct insight into the mechanics of rotation. [9,10,13]

The forecasting structure contains two XGBoost layers. The first layer predicts next-period regime probabilities from market-level indicators such as return, volatility, cross-sectional dispersion, rank-based speed, and liquidity-related variables. The second layer predicts industry return directions conditional on both industry features and regime information. Using probabilities rather than hard labels allows the downstream allocation rule to react smoothly to transitional states instead of mechanically switching between discrete categories. [5,13-15]

3.3 Regime-Dependent Risk Parity and Experimental Design

Regime information enters the portfolio stage a second time through regime-dependent risk parity. Rather than allocating capital directly from predicted scores, the model first forms candidate long and short baskets and then computes a regime-conditioned covariance estimate. Weights are determined so that risk contributions are distributed more evenly than under naive equal weighting, which helps prevent the portfolio from becoming excessively exposed to a small number of volatile industries.

Training follows a temporally consistent design. The deep-factor module uses rolling windows, while the classification system applies a fixed chronological split aligned with the original implementation: training ends on 2021-12-31, validation covers 2022, and testing or backtesting begins in 2023. This arrangement reduces look-ahead risk and makes the different model layers directly comparable within one unified evaluation framework. [5,8]

4 Empirical Results

4.1 Experimental Settings and Predictive Results

Table 1 summarizes the core experimental design. The sample begins in 2015 so that the framework can combine rolling deep-factor learning with a later clean out-of-sample window. This setting preserves the full logical chain from factor mining to regime-aware allocation and allows the main empirical mechanism to be evaluated under a temporally consistent design. [5,11]

Table 1. Core experimental settings.

Item	Setting
Universe	Shenwan Level-1 industry indices
Base factors	10 Alpha-style industry factors
Deep factors	5 factors from a feed-forward network
Regime definition	4 states from volatility and rotation speed
Training split	Date ≤ 2021-12-31
Validation split	2022-01-01 to 2022-12-31
Test / backtest split	Date ≥ 2023-01-01
Allocation layer	Regime-dependent risk parity

The market-regime classifier performs strongly. As shown in Table 2, validation-set accuracy, balanced accuracy, and macro-F1 all exceed 0.91, while the corresponding test-set values are even higher. These metrics suggest that the volatility-speed definition is sufficiently regular to be learned with stable out-of-sample performance, yet still flexible enough to differentiate economically distinct market environments. [13-15]

Table 2. Performance of the market-regime classifier.

Metric	Validation	Test
Accuracy	0.9218	0.9688
Balanced accuracy	0.9176	0.9605
Macro-F1	0.9184	0.9645

4.2 Backtest Results

The full portfolio results are reported in Table 3. During the 2023-2024 out-of-sample period, the complete regime-dependent system achieves a cumulative return of 23.20%, an annualized return of 11.45%, and a Sharpe ratio of 1.399. Its maximum drawdown is contained at -7.42%, which is markedly better than the equal-weight factor baseline and still superior to the ordinary risk-parity alternative that ignores regime-conditioned covariance information. [11,12]

These metrics matter because they show that the improvement is not driven by a simple risk-seeking shift. The annualized volatility of the full model remains moderate, yet the gain in return is large enough to improve both Sharpe and Calmar ratios. In other words, the regime layer does not merely forecast direction; it also improves the way risk is organized across industries.

Table 3. Out-of-sample portfolio performance of the full model.

Metric	Full model
Cumulative return	23.20%
Annualized return	11.45%
Annualized volatility	8.19%
Sharpe ratio	1.399
Maximum drawdown	-7.42%
Calmar ratio	1.543

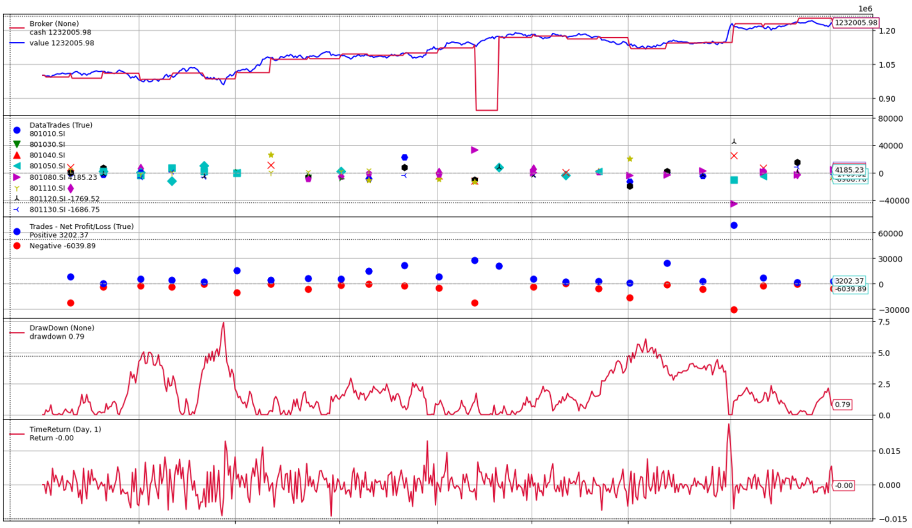

Fig. 1. Backtest curve of the regime-dependent risk-parity strategy.

Figure 1 reports the main backtest curve of the regime-dependent risk-parity strategy. The net-asset-value path rises in a stepwise but persistent manner, while the drawdown panel remains contained relative to the realized return path. The figure is informative because it indicates that the full-model advantage is not the result of one isolated jump; instead, performance is accumulated through a sequence of moderate gains realized under controlled downside episodes. [11,12]

4.3 Ablation and SHAP-Based Interpretation

Ablation results provide a direct view of how each module contributes to final performance. Table 4 compares the full framework with three simplified settings: an equal-weight factor strategy, an equal-weight XGBoost signal strategy, and a regime-dependent ordinary risk-parity strategy without the complete end-to-end signal structure. Performance improves monotonically as richer predictive structure and regime-aware allocation are introduced. [5,8,11]

Table 4. Ablation results.

Metric	Equal-weight factors	Equal-weight XGBoost	Regime-dependent ordinary risk parity
Cumulative return	-7.66%	10.53%	14.01%
Annualized return	-3.90%	5.13%	6.77%
Annualized volatility	6.52%	6.48%	6.57%
Sharpe ratio	-0.601	0.834	1.067
Maximum drawdown	-10.64%	-4.59%	-4.22%
Calmar ratio	-0.367	1.119	1.604

To interpret the regime classifier, the paper follows the explainability logic proposed by Lundberg and Lee. The most influential variables are the lagged market-volatility measure, the lagged rotation-speed indicator, and related cross-sectional dispersion variables. This result clarifies why the regime module improves the allocation layer rather than merely improving a stand-alone classification score.

Economically, this result implies that regimes in the present framework are not broad labels for market stress alone. They summarize the environment in which leadership changes across industries become either smooth and persistent or abrupt and unstable. That interpretation is fully consistent with the paper's design: the same variables that define and forecast regimes also influence the covariance environment in which risk parity operates. [10-15]

4.4 Discussion

The empirical evidence supports an important methodological point. A stronger ranking model helps, but the gain becomes more durable when the allocation layer is allowed to react to regime-dependent covariance conditions. This is why the full framework outperforms not only the equal-weight factor baseline but also the regime-dependent ordinary risk-parity benchmark. The paper therefore argues that signal quality and weight construction should be viewed as complementary sources of performance rather than as substitutable design choices.

The results also have a practical interpretation. When market volatility and rotation speed change together, the concentration risk of an industry portfolio can increase sharply. A regime-conditioned risk-parity mechanism helps redistribute risk before such concentration becomes large enough to dominate the portfolio. This role is particularly valuable in the China A-share market, where style switching and policy-sensitive leadership changes can accelerate cross-industry re-pricing. [2,10-14]

Overall, the empirical presentation concentrates on the main backtest evidence, portfolio metrics, ablation results, and the textual interpretation of the regime model. The combined evidence shows that performance gains arise from the interaction between signal generation and allocation design rather than from any single isolated modeling choice.

5 Conclusion

This paper develops an integrated framework for industry rotation in China's A-share market by connecting industry-level factor construction, deep-factor extraction, regime prediction, directional classification, and regime-dependent risk parity in one sequential design.

The main empirical message is clear. In the 2023-2024 out-of-sample period, the full model delivers attractive risk-adjusted performance, moderate drawdowns, and consistent improvements over simpler baselines. The ablation evidence indicates that stronger signal generation and better allocation design reinforce each other rather than acting as substitutes.

Future work can extend the framework by incorporating macro, sentiment, option-implied, or liquidity-based state variables; by replacing the two-stage ranking process with more fully differentiable portfolio layers; and by bringing trading frictions into the optimization loop at the industry level.

References

1. Moskowitz, T. J., & Grinblatt, M. (1999). Do industries explain momentum? The Journal of Finance, 54(4), 1249-1290. https://doi.org/10.1111/0022-1082.00146.
2. Chava, S., Hsu, A., & Zeng, L. (2020). Does history repeat itself? Business cycle and industry returns. Journal of Monetary Economics, 116, 201-218. https://doi.org/10.1016/j.jmoneco.2019.10.005.
3. Molchanov, A., & Stangl, J. (2024). The myth of business cycle sector rotation. International Journal of Finance & Economics, 29(4), 4419-4442. https://doi.org/10.1002/ijfe.2882.
4. Fama, E. F., & French, K. R. (2015). A five-factor asset pricing model. Journal of Financial Economics, 116(1), 1-22. https://doi.org/10.1016/j.jfineco.2014.10.010.
5. Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. Review of Financial Studies, 33(5), 2223-2273. https://doi.org/10.1093/rfs/hhaa009.
6. Bryzgalova, S., Huang, J., & Julliard, C. (2023). Bayesian solutions for the factor zoo: We just ran two quadrillion models. The Journal of Finance, 78(1), 487-557. https://doi.org/10.1111/jofi.13197.

7. Ehsani, S., & Linnainmaa, J. T. (2022). Factor momentum and the momentum factor. The Journal of Finance, 77(3), 1877-1919. https://doi.org/10.1111/jofi.13131.
8. Kagkadis, A., Vasilas, N., Nolte, I., & Nolte, S. (2024). Factor timing with portfolio characteristics. Review of Asset Pricing Studies, 14(1), 84-118. https://doi.org/10.1093/rapstu/raad010.
9. Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. Econometrica, 57(2), 357-384. https://doi.org/10.2307/1912559.
10. Ang, A., & Bekaert, G. (2002). International asset allocation with regime shifts. Review of Financial Studies, 15(4), 1137-1187. https://doi.org/10.1093/rfs/15.4.1137.
11. Costa, G., & Kwon, R. H. (2019). Risk parity portfolio optimization under a Markov regime-switching framework. Quantitative Finance, 19(3), 453-471. https://doi.org/10.1080/14697688.2018.1486036.
12. Maillard, S., Roncalli, T., & Teiletche, J. (2010). The properties of equally weighted risk contribution portfolios. The Journal of Portfolio Management, 36(4), 60-70. https://doi.org/10.3905/JPM.2010.36.4.060.
13. He, M., Bai, Y., & Ren, X. (2023). Forecasting aggregate stock market volatility with industry volatilities: The role of spillover index. Research in International Business and Finance, 65, 101958. https://doi.org/10.1016/j.ribaf.2023.101958.
14. Oliveira, D. C., Sandfelder, D., Fujita, A., Dong, X., & Cucuringu, M. (2025). Tactical asset allocation with macroeconomic regime detection. arXiv. https://arxiv.org/abs/2503.11499.
15. Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30, 4765-4774. https://proceedings.neurips.cc/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html.
16. Tang, G., Zhu, L., Liao, C., & Jiang, F. (2024). Asset pricing based on autoencoding machine learning: Evidence from Chinese stock-market financial big data. Journal of Management Sciences in China, 27(9), 82-97. http://www.jmsc.tju.edu.cn/CN/Y2024/V27/I9/82.
17. Li, X., Li, Z., Li, Q., Liu, Y., & Tang, W. (2025). A survey of machine-learning-based stock return prediction. Chinese Journal of Management Science, 33(1), 311-322. http://www.zgglkx.com/CN/Y2025/V33/I1/311.
18. Wu, M.-C., Huang, S.-H., & Chen, A.-P. (2024). Momentum portfolio selection based on learning-to-rank algorithms with heterogeneous knowledge graphs. Applied Intelligence, 54(5), 4189-4209. https://doi.org/10.1007/s10489-023-05106-3.