Vol. 19 (2026): Proceedings of 2026 International Conference on Financial Technology, Mathematical Statistics, and Business Analytics (ICFMB 2026)
E-ISSN: 3105-0522 P-ISSN: 3105-0514
DOI: https://doi.org/10.70267/icfmb.202619455463

The Interpretation-Prediction Trade-off and The Layered Integration Framework — A Three-Stage Evolution and Methodological Reconstruction of Asset Pricing Models

Yueqiao Ding*

Nanjing University of Finance & Economics, Nanjing, China

**Corresponding author: Yueqiao Ding.*

Abstract

The evolution of the asset pricing model has undergone a third paradigm transformation from the traditional linear factor period to the machine learning period, and then to the deep integration of the two. Based on the literature review, this article systematically sorts out this evolutionary vein and puts forward the "three-stage evolution model" as a historical framework for understanding the changes in asset pricing theory. On this basis, the "interpretation-prediction trade-off curve" is constructed as an analysis tool to describe the three types of models in "economic interpretation" and "prediction accuracy" with dimensional positioning differences; then, the "layered integration framework" is proposed as a methodological path to meet the weighing challenges, and the stochastic discount factor theory, machine learning methods and interpretability tools are organically integrated to form an operable empirical research system. The theoretical contribution of this article is: first, the system puts forward the evolutionary logic of asset pricing research, revealing the intrinsic motives of the transformation from "interpretation-oriented" to "interpretation-prediction coordination"; second, deepen the understanding of "interpretation-prediction trade-off" from the perspective of model uncertainty and causal inference; third, build a five-layer integrated framework and provide methodological guidance for the systematic construction of the fusion model. This article provides a clear theoretical basis and methodological path for subsequent empirical research, and promotes the evolution of asset pricing research in a more interpretable and predictive direction.

Keywords

asset pricing, interpretation-predicted trade-off, layered integration framework, random discount factor, machine learning

1. Introduction

The evolution of the asset pricing model is undergoing the third paradigm transformation: from the traditional linear factor period [1-3] to the machine learning period [8-11], and then to the new stage of deep integration of the two [12-17]. The traditional factor model faces the dilemma of "factor zoo" [5] and the bottleneck of prediction accuracy [6-7]. Although the machine learning model is significantly better than the linear model [8-9] in yield prediction, its "black box" nature [12] leads to insufficient economic explanatory power, which has become a new bottleneck. Against this background, the "fusion model" [13-14], which integrates economic theory and machine learning methods, has become a research hotspot.

However, there are still three shortcomings in the existing fusion research: first, the integration path is scattered, and there is a lack of a unified theoretical framework to integrate factor learning, generation modelling, interpretability analysis and other different directions; second, there is a lack of systematic analysis of the essence of "interpretation-prediction trade-off", and most studies focus on method innovation and ignore the inherent of trade-offs mechanism; Third, the construction method of the fusion model presents fragmented characteristics and lacks operable and replicable methodological guidance. How to systematically build a fusion model and how to achieve an effective balance between explanatory power and predictive power are still methodological problems that need to be solved urgently.

In response to the above problems, this article puts forward three interrelated core concepts based on a literature review: first, "three-stage evolution model", as a historical framework for understanding the changes in asset pricing theory; second, "interpretation-prediction trade-off curve", as an analysis of the three types of models in "economic interpretation" and "prediction accuracy" The theoretical tools for positioning differences in individual dimensions, and deepen the understanding of the essence of trade-offs from the perspective of model uncertainty [18] and causal inference [19]; the third is the "layered integration framework", as a methodological path to meet the challenges of trade-offs, random discount factor theory, machine learning methods and interpretability tools [24-27]. Organic integration forms a five-layer progressive empirical research system.

This article is positioned as methodological construction research based on literature review, without empirical verification, aiming to provide a clear theoretical basis and methodological guidance for subsequent research.

2. Theoretical Background

This article constructs a trinity theoretical framework of "evolutionary logic-analysis tool-methodological path", and systematically explains the inherent logic and realisation path of asset pricing research from interpretation-oriented to interpretation-prediction collaborative transformation.

First of all, the article would review the "three-stage evolution model" as historical framework divides asset pricing research into three stages: the traditional linear factor period [1-3], the machine learning period [8-11] and the fusion period [12-17]. This framework reveals that the traditional model is interpretation-oriented but lacks predictability [4-5], and the machine learning model is prediction-oriented but lacks explanatory power. The two form tension in the dimensions of "economic interpretation" and "prediction accuracy", which constitutes the fundamental motive of paradigm transformation.

Secondly, as an analysis tool, the "interpretation-prediction trade-off curve" portrays the positioning differences of the three types of models in two-dimensional space. From the perspective of model uncertainty [18-19], the non-sparseness and model equivalence characteristics rooted in the asset pricing data structure are weighed; from the perspective of causal inference [20-21], prediction depends on statistical correlation and interpretation requires causal understanding. The essential difference between the two determines the inevitability of trade-offs. This analysis reveals the theoretical necessity of the rise of the fusion model.

Finally, the "layered integration framework" takes the stochastic discount factor theory as the core [12], and decomposes the construction system of the fusion model into five progressive levels: economic theory foundation -- management combination construction -- SDF coefficient estimation -- use of interpretability tools -- empirical test. The theoretical innovation of this framework is to integrate the distributed fusion path into a systematic methodology, so that researchers can introduce nonlinear modelling capabilities while maintaining economic explanatory power, and transform the tension of "explanation-prediction" into the driving force of collaborative improvement.

3. Literature Review

3.1 Traditional Linear Factor Period: The Theoretical Foundation and Limitations of Interpretation-oriented

The traditional linear factor model takes CAPM as the foundation theory to establish the linear relationship between systemic risk and expected benefits. However, Roll pointed out that because the real market combination is unobservable, it is difficult for CAPM to be effectively verified in reality [1]. In order to make up for this shortcoming, a multi-factor model was born: Fama and French incorporated scale and value factors to build a three-factor model [2], and Carhart further introduced momentum factors to form a four-factor model [3]. The core contribution of these models is to establish the dominant position of the multi-factor framework by introducing additional factors to explain the visions that CAPM cannot capture.

However, the limitations of traditional linear models are gradually emerging. The main limitations are reflected in three aspects: first, the problem of missing variables in the factor set [4]; second, the "factor zoo" leads to model redundancy and inefficiency [5-6]; third, the static hypothesis is difficult to capture the risk-benefit relationship in the dynamic market environment [7]. These limitations provide a theoretical motive for the introduction of machine learning methods.

3.2 Machine Learning Period: Prediction-oriented Method Breakthroughs and Challenges

The introduction of machine learning methods marks the expansion of asset pricing research from explanation-oriented to prediction-oriented. The core feature lies in nonlinear modelling ability and data-driven characteristics; representative methods include neural networks, gradient promotion, extreme random trees, etc.; the main breakthroughs are reflected in two aspects: in the dimension of prediction accuracy, machine learning is significantly better than the traditional linear model in the prediction of stock yield and bond yield [9]; in Theoretical dimension, the discovery of the "double decline" phenomenon [10-11] challenges the traditional deviation-variance trade-off and reveals the potential advantages of complex models in yield prediction.

However, the "black box nature" of the machine learning model [12] makes the lack of interpretation a new bottleneck, limiting the application value of the model in theoretical construction and policy formulation.

3.3 Integration Period: Path Exploration of Interpretation and Prediction Coordination

In the face of the dual dilemma of strong explanatory power but weak predictive power of traditional models and strong predictive power but weak explanatory power of machine learning models, the "fusion model" that integrates economic theory and machine learning methods has become the forefront of research. The random discount factor (SDF) framework provides a unified theoretical basis for fusion: in the linear model, SDF is a linear combination of factors, and in the nonlinear model, it can be an arbitrary function form of factors [13]. Shi proposes a unified framework for SDF on this basis, and introduces nonlinear flexibility while retaining economic interpretation through function mapping [14].

The existing integration research can be summarised into three core paths:

1. The Nonlinearisation Path of Factor Learning. Extend the Traditional Factor Model into a Nonlinear Time-varying Form to Introduce Machine Learning Flexibility While Maintaining Factor Interpretation. Ipca [16] Realises the Dynamics of Factor Exposure, and Cvae [17] Realises Nonlinear Dimension Reduction Through Self-encoding. the Core Contribution is to Expose the Factor from the Expansion of the Static Constant to the Characterised Function.

2. Generate the Path of the Model. Introduce Generative Machine Learning Methods to Construct Nonlinear Factors. GAN-SDF [14] Enhances the Robustness of Sdf Through Confrontation Training, and Risk-premium PCA [18] Embeds Risk Premium Information to Identify "Weak Factors" and Expands the Boundaries of Factor Identification.

3. Estimate the Structured Constraint Path of the Process. Introduce Structured Constraints to Alleviate Over-fitting in the Sdf Coefficient Estimation Link. the Study of Bayesian Contraction [23] and AP-Trees [24]

Shows That the Application of Theoretically Driven Structured Constraints is a Key Mechanism to Achieve the Balance Between Interpretation and Prediction.

In addition, the introduction of interpretability tools provides help for the fusion model. The wide application of methods such as SHAP [25] and LIME [26, 27] reveals that traditional factors such as momentum and mobility are still at the core of complex models [28]. This shows that the fusion model is not a subversion of traditional theory, but a deepening and expansion on its basis.

Looking at the above research, although the exploration of the integration field has made progress in multiple dimensions, the overall situation shows the characteristics of scattered paths and lack of system integration: different studies cut in from their own perspectives and have not yet formed a unified construction methodology. This provides a problem-oriented approach for this article to propose layered integration framework.

4. Key Finding

4.1 The Intrinsic Logic of "Three-Stage Evolution Model"— the "Interpretation -Prediction Trade-off" Curve

To deepen the understanding of the internal logic underlying the three-stage evolution model, this study introduces the "interpretation-prediction trade-off" as a unified analytical perspective. The fundamental goal difference between econometrics and machine learning -- the former focuses on causal inference, while the latter focuses on predictive accuracy -- forms a core tension in the field of asset pricing. In order to systematically analyse this tension, this study examines the asset pricing model in the two dimensions of "economic interpretation" and "prediction accuracy", and constructs an " interpretation — prediction trade-off curve", as presented in Figure 1, to visually describe the relative position of each stage in the "three-stage evolution model": the traditional linear factor models have high explanation but low predictive power; complex machine learning models are the opposite; while hybrid models such as IPCA try to seek a better balance.

Figure 1: Interpretation–prediction trade-off curve

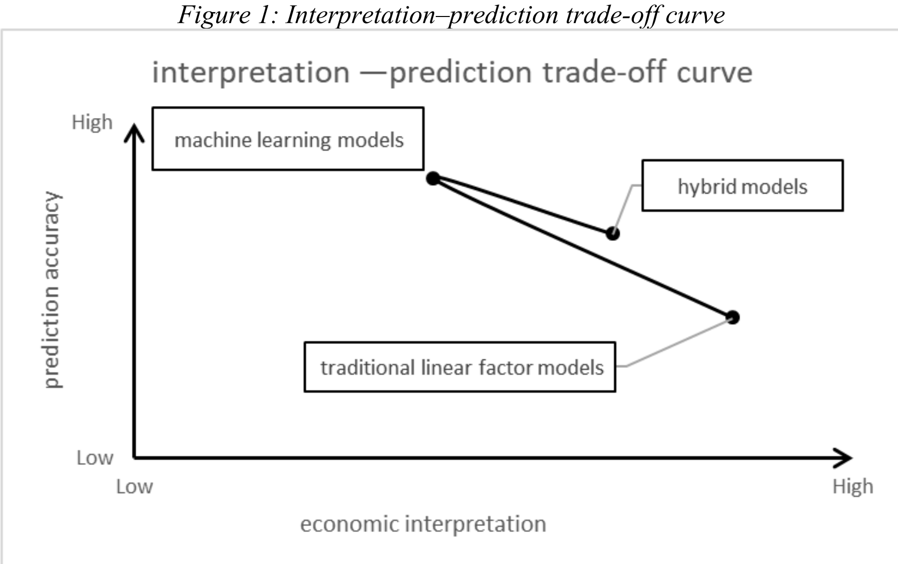

The model uncertainty provides key evidence for the inevitability of "interpretation —prediction trade-off". The study of Giannone et al. shows that there is no sparse structure in asset pricing problems - almost all variables contain certain predictive information, and a sufficient pricing model cannot be built with a few variables [19]. Bryzgalova et al. analysed more than 20 trillion factor combinations and found that many combinations produce similar empirical results, and there is no single optimal model [20]. The model uncertainty reveals that the "interpretation —prediction trade-off" is not a temporary technical problem, but a fundamental feature rooted in the asset pricing data structure, which must be dealt with through a systematic methodology.

And the perspective of causal inference reveals the root cause of the trade-off. Lopez de Prado and Zoonekynd proposed the "causal factor investment", emphasising the importance of distinguishing between correlation and causality [21]. Machine learning realises prediction by capturing statistical correlation, while economics requires causal understanding. Hünermund and Bareinboim systematically sorted out the application of causal inference and data integration in econometrics [22], providing a methodological basis for asset pricing. Even if the model predicts accurately, its theoretical value is still limited if the economic mechanism behind it cannot be identified. Therefore, the future research direction should be to develop models with both predictive ability and causal interpretation. Therefore, the future research direction should be committed to developing models that can not only give full play to the advantages of machine learning prediction, but also embed economic theoretical constraints, and strive to transform the tension of "interpretation-prediction" into the driving force of collaborative improvement.

4.2 Method of Constructing Hybrid Model Based on "Interpretation-Prediction Trade-off Curve"—Layered Integration Framework

In view of the current situation of scattered research paths and lack of systematic construction methods, this paper proposes a "layered integration framework" based on the SDF theory. The framework decomposes the construction process system of the fusion model into five progressive levels, each of which corresponds to a specific functional goal and method selection space, forming a complete logical chain from theoretical constraints to nonlinear modelling, to interpretability analysis and empirical testing.

The first layer: the foundation of economic theory -- embedding theoretical constraints to ensure the economic consistency of the model. With SDF as the core, it embeds the assumption of no arbitrage constraints and rational expectations to ensure that the model meets the basic requirements of economic theory. SDF unified framework [14] formalises this idea: mapping the original features to the factor space through management combination, which not only retains the economic explanation, but also lays the foundation for subsequent nonlinear modelling.

The second layer: management combination construction -- introduce nonlinear mapping to capture complex relationships. Map the original features to the management combination through functions, and introduce nonlinear modelling capabilities while maintaining interpretability. Researchers can choose methods such as IPCA [16], CVAE [17] or GAN-SDF [15] according to the research questions. The core function of this layer is to realise the characteristic nonlinear transformation under theoretical constraints.

The third layer: SDF coefficient estimation - applying structured constraints to mitigate the risk of over-fitting. After the construction of the management combination is completed, the SDF coefficient needs to be estimated. In order to balance the fit excellence and generalisation ability, it is recommended to adopt methods such as Bayesian contraction [23] or AP-Trees [24] to introduce theoretically driven structural constraints in the estimation link. This is a key mechanism to realise the balance between "predictive accuracy" and "economic interpretation".

The fourth layer: interpretability tools -- reveal the logic of decision-making and improve the transparency of the model. Use SHAP [25], LIME [26] and other tools to analyse the internal mechanism of the model and identify the contribution of core variables. Explainable tools not only serve model diagnosis, but also assume the bridge function of connecting "black box prediction" and "theoretical explanation" -- for example, Gu et al. found that traditional factors such as momentum and mobility are still at the core in complex models [28], which verifies the explanatory power of traditional theories.

The fifth layer: empirical test -- verify the applicability of the model and consider the practical feasibility. The model is tested in a multi-asset category [29, 30] and multi-market environment, and included in transaction cost considerations [31] to ensure that the model is not only statistically significant, but also has practical application value.

The theoretical innovation of this framework is to integrate the dispersed fusion path into a systematic methodology, so that researchers can introduce nonlinear modelling ability along a clear hierarchy while maintaining economic interpretation, and transform the tension of "explanation-prediction" into the driving force of collaborative improvement.

5. Discussion

This study systematically reviews the three-stage evolution of asset pricing models and constructs a tripartite theoretical system comprising the three-stage evolution model, an "interpretation —prediction trade-off curve", and a "layered integration framework". This system clarifies the core logic underlying the paradigm shift in asset pricing research. Based on the core research findings, this section extends the theoretical value and practical inspiration of the research from the four dimensions of practical application, cross-field integration, limited reflection and future trends.

5.1 Practical Application

This study has guiding significance for multiple subjects in the financial market.

For institutional investors, the three-stage evolution model clarifies the applicable scenarios of different models: the traditional linear model adapts to the interpretation of asset income cross-section and the construction of basic factor combinations; the machine learning model has significant advantages in dynamic yield prediction and short-term time selection; the hierarchical fusion model takes into account interpretation and prediction accuracy, which is suitable for Medium- and long-term investment and complex asset pricing. Interpretation -- The predictive trade-off curve can be used as a model selection tool to help investors avoid the deviation of a single pursuit of explanatory or predictive power. At the same time, the cross-asset comparison verified the adaptability of the layered integration framework: the performance of the two traditional models in the stock market is balanced [8, 28], the predictive power of the integrated neural network in the bond market is outstanding [9], and the option market [29] is more dependent on the fusion model due to its nonlinear characteristics. The framework can achieve multi-asset refinement through differentiated constraints.

For policymakers, the layered integration framework provides methodological support to improve market pricing efficiency and prevent systemic risks. The framework emphasises that economic theory constraints are embedded in nonlinear modelling. Combined with SHAP, LIME and other tools, it can accurately identify core risk factors and provide a basis for targeted supervision. At the same time, it incorporates the requirements of transaction costs and multi-market adaptability to make the research conclusions more suitable for market reality.

For academic researchers, the system integrates the dispersion direction of fusion model research, disassembles complex modelling processes, solves the problem of method fragmentation, and promotes theoretical accumulation and reuse of results.

5.2 Cross-field Theory Integration

This study realises the organic integration of financial metrology, machine learning and causal inference theory. As a core analysis tool, the interpretation-prediction trade-off curve deeply explores the intrinsic mechanism of trade-off from the dual perspective of model uncertainty and causal inference. The non-sparseness and multi-model equivalence of the asset pricing data structure lead to the inevitability of model uncertainty [19, 20], which is the objective basis for weighing existence; machine learning focuses on statistical correlation and economics focuses on the essential difference in causal cognition, which is the theoretical root of its existence [21, 22]. The introduction of causal inference theory has made asset pricing research break through the limitation of focusing only on correlation and move forward in the direction of exploring the causal mechanism.

With SDF theory as the core, the hierarchical fusion framework integrates the nonlinear modelling ability and interpretable tools of machine learning, which not only ensures the economic logical consistency of the model, but also builds a bridge between "black box" prediction and theoretical interpretation, and provides a new model of interdisciplinary research for asset pricing and other financial measurement fields.

5.3 Reflection on Research Limitations

There are still three limitations in this research: first, the calculation complexity of the hierarchical fusion framework is high. The application threshold is high. Mainwhile, the integration of emerging theories such as behavioural finance is insufficient, and the explanatory power in complex markets needs to be verified. Second, as a methodological construction research, empirical testing has not been carried out, and the framework is in different assets, markets and extremes. The effect and robustness of the environment lack data support. Third, the interpretable analysis relies on general tools and does not fully adapt to the characteristics of asset pricing data, and the interpretation results may be biased.

5.4 Future Research Trends

Based on research findings and limitations, future research can be promoted from three aspects: at the theoretical level, further integrate behavioural finance, complex networks and other theories, and develop low-complexity modelling methods to reduce the threshold of framework application; at the method level, build exclusive interpretability tools for asset pricing, and extract non-structure in combination with large language models. Information integrate into the causal inference framework. At the practical level, carry out multi-asset and multi-market empirical testing, incorporate actual factors such as transaction costs, and provide realistic support for investment strategies and regulatory policies.

In general, the coordination of economic explanation and forecasting ability is the inevitable direction of asset pricing research. Future research will take the layered integration framework as the core, develop in the direction of multi-theory integration, method innovation and practical implementation, and provide solid support for the efficient operation of the financial market.

6. Conclusion

This study proposes the “three-stage evolution model” as a theoretical framework for understanding the transformation of asset pricing, and constructs the “interpretation -prediction trade-off curve” to describe the relative positioning of the traditional linear factor model, machine learning model and hybrid model in the two dimensions of “economic interpretation” and “prediction accuracy”, thus revealing the main logical line of evolution from traditional models to hybrid models. Under this framework, this study further proposes a “layered integration framework” based on the SDF theory - a methodological system that integrates economic theory, machine learning and explainable tool systems. The framework transforms abstract theoretical integration into specific modelling steps through five-layer decomposition so that researchers can introduce nonlinear modelling ability to achieve prediction accuracy and reason while maintaining economic interpretation. This framework not only systematises the construction process of the hybrid model, but also provides operational methodological guidelines for subsequent research. Of course, this study also has certain limitations: the computational complexity of the framework is relatively high, and its robustness in the extreme market environment still needs to be tested. Future research can be deepened from the following directions: developing structured nonlinear modelling methods to enhance theoretical constraints; improving interpretation tools to improve model transparency; further incorporating transaction cost and cross-asset feasibility considerations into empirical testing, promoting the hybrid model from theory to practice, and providing academic support for improving financial market efficiency.

References

[1] Roll, R. (1977). A critique of the asset pricing theory’s tests Part I: On past and potential testability of the theory. Journal of Financial Economics, 4(2), 129–176. https://doi.org/10.1016/0304-405x(77)90009-5

[2] Fama, E. F., & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. Journal of Financial Economics, 33(1), 3–56. https://doi.org/10.1016/0304-405x(93)90023-5

[3] Carhart, M. M. (1997). On persistence in mutual fund performance. The Journal of Finance, 52(1), 57–82. https://doi.org/10.1111/j.1540-6261.1997.tb03808.x

[4] Ang, A., Hodrick, R. J., Xing, Y., & Zhang, X. (2006). The Cross-Section of volatility and expected returns. The Journal of Finance, 61(1), 259–299. https://doi.org/10.1111/j.1540-6261.2006.00836.x

[5] Cochrane, J. H. (2011b). Presidential address: Discount rates. The Journal of Finance, 66(4), 1047–1108. https://doi.org/10.1111/j.1540-6261.2011.01671.x

[6] Harvey, C. R., Liu, Y., & Zhu, H. (2015). . . . and the Cross-Section of expected returns. Review of Financial Studies, 29(1), 5–68. https://doi.org/10.1093/rfs/hhv059

[7] Fama, E. F., & French, K. R. (2019). Comparing Cross-Section and Time-Series factor models. Review of Financial Studies, 33(5), 1891–1926. https://doi.org/10.1093/rfs/hhz089

[8] 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

[9] Bianchi, D., Büchner, M., & Tamoni, A. (2020). Bond Risk Premiums with Machine Learning. Review of Financial Studies, 34(2), 1046–1089. https://doi.org/10.1093/rfs/hhaa062

[10] Belkin, M., Hsu, D., Ma, S., & Mandal, S. (2019). Reconciling modern machine-learning practice and the classical bias–variance trade-off. Proceedings of the National Academy of Sciences, 116(32), 15849–15854. https://doi.org/10.1073/pnas.1903070116

[11] Kelly, B., Malamud, S., & Zhou, K. (2023). The virtue of complexity in return prediction. The Journal of Finance, 79(1), 459–503. https://doi.org/10.1111/jofi.13298

[12] Coqueret, G., & Guida, T. (2020). Machine learning for factor investing: R version. Chapman and Hall/CRC. https://doi.org/10.1201/9781003034858

[13] Cochrane, J. H. (2005). Asset pricing (Rev. ed.). Princeton University Press.

[14] Shi, C. (2025). From Econometrics to Machine Learning: Transforming Empirical asset pricing. Journal of Economic Surveys, 40(1), 528–548. https://doi.org/10.1111/joes.70002

[15] Chen, L., Pelger, M., & Zhu, J. (2023). Deep learning in asset pricing. Management Science, 70(2), 714–750. https://doi.org/10.1287/mnsc.2023.4695

[16] Kelly, B. T., Pruitt, S., & Su, Y. (2019). Characteristics are covariances: A unified model of risk and return. Journal of Financial Economics, 134(3), 501–524. https://doi.org/10.1016/j.jfineco.2019.05.001

[17] Gu, S., Kelly, B., & Xiu, D. (2021). Autoencoder asset pricing models. Journal of Econometrics, 222(1), 429–450. https://doi.org/10.1016/j.jeconom.2020.07.009

[18] Lettau, M., & Pelger, M. (2020). Estimating latent asset-pricing factors. Journal of Econometrics, 218(1), 1–31. https://doi.org/10.1016/j.jeconom.2020.01.006

[19] Giannone, D., Lenza, M., & Primiceri, G. E. (2021). Economic predictions with big data: the illusion of sparsity. Econometrica, 89(5), 2409–2437. https://doi.org/10.3982/ecta17842

[20] Bryzgalova, S., Huang, J., & Julliard, C. (2022). 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

[21] Lopez de Prado, M., & Zoonekynd, V. (2024). Correcting the factor mirage: A research protocol for causal factor investing. SSRN.

[22] Hünermund, P., & Bareinboim, E. (2025). Causal inference and data fusion in econometrics. Econometrics Journal, 28(1), 41–82. https://doi.org/10.1093/ectj/utad008

[23] Kozak, S., Nagel, S., & Santosh, S. (2020). Shrinking the cross-section. Journal of Financial Economics, 135(2), 271–292. https://doi.org/10.1016/j.jfineco.2019.06.008

[24] Bryzgalova, S., Pelger, M., & Zhu, J. (Forthcoming). Forest through the trees: Building cross-sections of stock returns. Journal of Finance.

[25] Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30, 4765–4774. https://doi.org/10.48550/arXiv.1705.07874

[26] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). “Why should I trust you?”: Explaining the predictions of any classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 1135–1144). Association for Computing Machinery. https://arxiv.org/abs/1606.05386

[27] Coqueret, G., & Guida, T. (2020). Machine learning for factor investing: R version. Chapman and Hall/CRC. https://doi.org/10.1201/9781003034858

[28] Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. The Review of Financial Studies, 33(5), 2223–2273. https://doi.org/10.1093/rfs/hhaa009

[29] Büchner, M., & Kelly, B. (2022). A factor model for option returns. Journal of Financial Economics, 143(3), 1140–1161. https://doi.org/10.1016/j.jfineco.2021.12.007

[30] Giglio, S., & Xiu, D. (2021). Asset pricing with omitted factors. Journal of Political Economy, 129(7), 1945–1990. https://doi.org/10.1086/714090

[31] Avramov, D., Cheng, S., & Metzker, L. (2023). Machine learning vs. economic restrictions: Evidence from stock return predictability. Management Science, 69(5), 2587–2619. https://doi.org/10.1287/mnsc.2022.4449

Funding

This research received no external funding.

Conflicts of Interest

The authors declare no conflict of interest.

Acknowledgment

This paper is an output of the science project.

Open Access