

Article

Asset Risk Pricing with Causal Representation Learning and Generative AI: Evidence from Complex Market Environments

Ao Zhang^{1,*}

¹ Norwich Business School, University of East Anglia, Norwich, NR4 7TJ, England

* Correspondence author: zhangao202605@163.com

Abstract: This study examines asset pricing in high-dimensional financial environments, where traditional factor models face limitations in interpretability and stability, and machine learning approaches remain largely correlation-driven. Existing methods often struggle to identify economically meaningful risk factors and exhibit limited robustness across changing market conditions. To address these limitations, this paper proposes a unified framework that integrates generative modeling with causal representation learning. A generative model is first employed to extract latent structures from high-dimensional financial data, and a causal learning module is subsequently used to identify invariant factors based on structural relationships. These factors are then incorporated into a standard asset pricing model using the Fama–MacBeth regression approach. Empirical results show that the proposed model achieves a cross-sectional R^2 of 0.38, compared to 0.26 for the Fama–French five-factor model and 0.31 for a deep learning benchmark. Pricing errors are reduced from 0.054 (FF5) and 0.049 (deep learning) to 0.036. In addition, the model maintains relatively stable performance across different market conditions, including bull, bear, and crisis periods, while benchmark models exhibit a more pronounced decline in explanatory power. These findings suggest that incorporating causal representation improves both explanatory power and robustness in asset pricing. By linking latent structure learning with causal invariance, the proposed framework provides a structured interpretation of risk premia and offers a more stable foundation for asset pricing in complex market environments.

Keywords: Asset Pricing; Causal Representation Learning; Generative Artificial Intelligence; Invariant Risk Factors; Financial Market Robustness

1. Introduction

Asset pricing remains a central topic in financial economics, with a long tradition of models aimed at explaining the cross-section of expected returns. Classical frameworks, such as the Capital Asset Pricing Model and multifactor extensions, provide parsimonious representations of systematic risk [1]. However, their empirical performance has been increasingly challenged in complex and evolving market environments. The proliferation of proposed factors, often referred to as the “factor zoo,” raises concerns about data mining, lack of economic interpretation, and limited out-of-sample stability. In response, machine learning techniques have been introduced to enhance predictive performance by exploiting high-dimensional data and nonlinear relationships [2]. While these approaches often improve statistical fit, they remain largely correlation-driven and provide limited insight into the underlying economic mechanisms [3]. As a result, their robustness across different market regimes and their interpretability for risk pricing remain open questions.

Recent advances in causal inference and causal representation learning offer a potential pathway to address these limitations. By explicitly modeling structural relationships and isolating invariant



components across environments, causal methods aim to move beyond prediction toward explanation [4]. In particular, causal representation learning seeks to extract latent factors that are stable under distributional shifts, which is particularly relevant in financial markets characterized by regime changes, policy interventions, and structural breaks [5]. At the same time, generative artificial intelligence has demonstrated strong capability in capturing complex data distributions and uncovering latent structures in high-dimensional settings [6]. These models provide a flexible framework for representing hidden factors that may not be directly observable in traditional financial datasets.

Motivated by these developments, this study explores how causal representation learning can be integrated into asset pricing, and how generative AI can facilitate the identification of latent structures underlying asset returns [7]. The central objective is to construct risk factors that are not only statistically significant but also structurally meaningful and stable across different market conditions. In doing so, the study seeks to bridge the gap between predictive modeling and economic interpretation, while addressing the need for robustness in empirical asset pricing.

This paper is guided by three main research questions. First, does the incorporation of causal representations improve the explanatory power of asset pricing models relative to traditional and machine learning-based approaches? Second, can generative AI enhance the discovery of latent structures that are relevant for risk pricing, beyond what can be achieved through conventional dimensionality reduction techniques? Third, do the resulting models exhibit robustness across heterogeneous market environments, including periods of heightened volatility and structural change?

The contributions of this study are threefold. From a methodological perspective, it proposes a unified framework that integrates generative modeling with causal representation learning to identify invariant risk factors. From an empirical perspective, it evaluates the performance of the proposed approach across different market conditions, with a focus on explanatory power and pricing errors. From a theoretical perspective, it provides a causal interpretation of risk premia, emphasizing the role of invariant mechanisms in determining asset returns. Together, these contributions aim to advance the understanding of asset pricing in complex environments and to provide a foundation for more robust and interpretable financial models.

2. Literature Review and Theoretical Framework

Modern asset pricing is grounded in the Capital Asset Pricing Model (CAPM), which links expected returns to systematic market risk. While CAPM provides a parsimonious benchmark, its empirical limitations have motivated multifactor models, particularly the Fama-French three- and five-factor frameworks [8]. These models incorporate size, value, profitability, and investment factors, improving cross-sectional explanatory power. However, they face important challenges. The rapid expansion of factors raises concerns about data mining and multiple testing, often referred to as the “factor zoo.” Many factors also lack clear economic interpretation, weakening their theoretical foundation [9], and their performance is often unstable across time and market conditions. These limitations highlight the need for more principled approaches to identifying economically meaningful and stable risk factors.

Machine learning methods have been widely applied to asset pricing to improve predictive accuracy. Techniques such as Lasso, random forests, and deep neural networks enable high-dimensional analysis and capture nonlinear relationships between characteristics and returns [10], often improving out-of-sample performance. However, these models are largely correlation-based and do not distinguish between predictive associations and causal relationships [11]. As a result, they may capture spurious patterns that fail to generalize across environments. In addition, limited interpretability, especially in deep learning, restricts their usefulness for understanding economic mechanisms. While feature selection can mitigate overfitting, it does not address the identification of structurally relevant risk factors.

Causal inference distinguishes correlation from causation by modeling structural relationships among variables. Structural causal models (SCMs) formalize these relationships through directed acyclic graphs, enabling identification of causal effects under specific assumptions. Recent advances in causal representation learning extend these ideas to high-dimensional settings by recovering latent variables associated with underlying causal mechanisms [12]. A key concept is invariance, referring to the stability of causal relationships across environments. Methods such as invariant risk minimization (IRM) identify representations that remain predictive under distributional shifts [13], which is particularly relevant in financial markets with regime changes. However, applying causal methods in asset pricing remains challenging due to latent confounders, measurement error, and the absence of controlled experiments, and causal assumptions are often difficult to verify empirically.

Generative artificial intelligence models, including variational autoencoders (VAE), generative adversarial networks (GAN), and diffusion models, model complex data distributions and learn latent representations in high-dimensional settings [14]. In asset pricing, they can extract latent factors from

noisy and heterogeneous data and capture nonlinear relationships more flexibly than traditional methods such as principal component analysis [15]. However, their application in finance remains limited. Generative models are typically designed to reproduce data distributions rather than identify economically meaningful structures, and without additional constraints, the learned latent variables may lack interpretability.

Building on these strands, this study proposes an integrated framework combining generative modeling with causal representation learning for asset pricing. The approach first recovers latent structures from high-dimensional data and then applies causal constraints to identify invariant components relevant for risk pricing. This integration addresses key limitations of existing methods: generative models handle complex data, while causal learning focuses on invariant relationships. Together, they provide a pathway to identifying risk factors that are both robust and economically interpretable. However, the effectiveness of this framework depends on the validity of its assumptions and the extent to which causal invariance can be established, making empirical validation essential.

3. Model Construction and Methodology

3.1. Overall Framework

This study proposes a unified asset pricing framework that integrates generative modeling with causal representation learning. The framework consists of three layers: data representation, causal structure, and pricing, jointly aimed at identifying structurally meaningful and robust risk factors.

As illustrated in Figure 1, high-dimensional financial data, including firm characteristics, macroeconomic variables, and market indicators, are first processed through a generative model to obtain latent representations that capture underlying structures while reducing noise. These latent variables are then input into a causal learning module, where structural dependencies are identified and invariant components are extracted. The resulting causal factors are subsequently incorporated into an asset pricing model to explain the cross-section of expected returns.

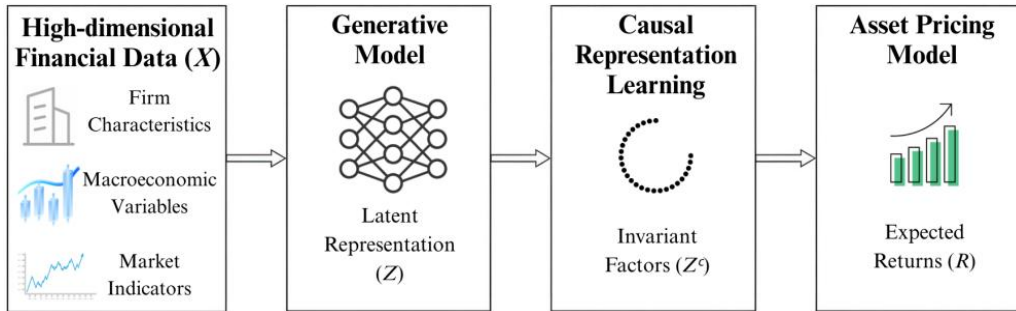


Figure 1. Framework of Causal Representation Learning with Generative AI for Asset Pricing

Formally, let $X \in \mathbb{R}^d$ denote the observed high-dimensional feature space, and $Z \in \mathbb{R}^k$ denote the latent representation with $k \ll d$. The objective is to construct a transformation from X to a subset of causally relevant latent factors $Z^c \subseteq Z$, which are subsequently used in the pricing equation. This sequential structure ensures that representation learning and causal identification are jointly aligned with the ultimate goal of explaining asset returns. The latent space serves as a structured representation that mitigates noise and enables more reliable identification of causal relationships.

3.2. Generative AI for Latent Factor Representation

To model the complex and noisy structure of financial data, this study employs generative models to learn latent representations. Specifically, a variational autoencoder (VAE)-type architecture is used to map observed variables X into a lower-dimensional latent space Z . The generative process can be expressed as:

$$X = g(Z, \delta) \quad (1)$$

where $g(\cdot)$ is a nonlinear function parameterized by neural networks, and δ represents stochastic noise.

The encoder network approximates the posterior distribution $q(Z|X)$, while the decoder

reconstructs X from Z . The training objective balances reconstruction accuracy and regularization of the latent space, ensuring that Z captures essential structural information while mitigating overfitting. Compared to traditional dimensionality reduction techniques, generative models offer two advantages. First, they can capture nonlinear dependencies among variables, which are prevalent in financial data. Second, they provide a probabilistic framework that explicitly models uncertainty, which is important for noisy and incomplete datasets. However, it is important to note that the latent variables obtained at this stage are not guaranteed to have economic or causal interpretation, necessitating further structural refinement.

3.3. Causal Representation Learning

Building on the latent space obtained from the generative model, the next step is to identify causally meaningful factors. This is achieved by modeling the structural relationships among latent variables using a causal framework. Specifically, we assume that the latent variables follow a structural causal model (SCM):

$$Z_i = f_i(pa(Z_i), U_i) \quad (2)$$

where $pa(Z_i)$ denotes the set of parent variables of Z_i in a directed acyclic graph, and U_i represents exogenous noise.

A key objective is to identify a subset of latent variables Z^c that exhibit invariance across different environments. To this end, the framework incorporates principles from invariant risk minimization (IRM), which seeks representations that maintain predictive consistency under distributional shifts. Formally, the invariant factors are those that satisfy:

$$\arg \min_{Z^c} \sum_{e \in \mathcal{E}} L^e(Z^c) \quad (3)$$

subject to the constraint that the optimal predictor remains invariant across environments e .

This approach allows the model to distinguish between spurious correlations, which may vary across market regimes, and stable causal relationships, which are more likely to reflect underlying economic mechanisms. Nevertheless, the identification of causal structures relies on assumptions regarding the data-generating process, and thus requires careful empirical validation.

3.4. Asset Pricing Model with Causal Factors

The causally identified latent factors are incorporated into a standard asset pricing framework. Specifically, the excess return of asset i at time t is modeled as:

$$R_{it} = \alpha_i + \beta_i Z_t^c + \delta_{it} \quad (4)$$

where Z_t^c denotes the vector of causal factors, β_i represents factor loadings, and δ_{it} is an idiosyncratic error term.

Estimation is conducted using cross-sectional regression methods, such as the Fama-MacBeth procedure, to evaluate the pricing ability of the proposed factors. The performance of the model is assessed in terms of explanatory power, pricing errors, and stability across different market conditions. By replacing traditional observable factors with causally informed latent variables, the model aims to provide a more structurally grounded explanation of expected returns. In particular, the use of invariant factors is expected to improve robustness under regime changes and reduce sensitivity to sample-specific patterns.

3.5. Methodological Contributions

The proposed framework contributes to the asset pricing literature by integrating generative modeling and causal representation learning into a unified approach. Unlike conventional methods that rely either on predefined factors or purely predictive models, this framework combines data-driven representation learning with structural constraints. The key innovation lies in the sequential identification of latent structures and causal invariance. Generative models provide a flexible mechanism for capturing complex patterns in financial data, while causal learning imposes discipline by focusing on stable relationships. This combination addresses both the high dimensionality of modern financial datasets and the need for interpretability and robustness in asset pricing. Overall, the methodology offers a novel perspective on risk factor identification, emphasizing the importance of

structural stability rather than purely statistical significance.

4. Data and Empirical Design

4.1. Data Sources

This study employs a combination of stock market data and macroeconomic variables to construct a comprehensive dataset for empirical analysis. The stock-level data are obtained from standard financial databases (e.g., CRSP/Compustat or CSMAR/WIND, depending on the market), and include monthly returns, firm characteristics, and trading-related variables. Firm characteristics typically cover size, book-to-market ratio, profitability, investment, momentum, and liquidity measures, which are widely used in the asset pricing literature. These variables serve as the primary inputs for both traditional models and the proposed framework.

In addition, macroeconomic variables are incorporated to capture broader economic conditions that may influence asset returns. These variables, including interest rates, inflation, industrial production growth, and market volatility indicators, are collected from publicly available macroeconomic databases (e.g., FRED or national statistical sources). The inclusion of macro-level information allows the model to account for systematic changes in the economic environment and provides a richer context for identifying latent structures.

The sample period spans multiple market cycles, ensuring sufficient variation in economic conditions. To evaluate robustness, the dataset is further divided into subperiods corresponding to different market regimes, such as expansion, recession, and periods of financial stress.

4.2. Variable Construction

The empirical analysis relies on three categories of variables: observed features, latent factors, and causal factors. Observed features are constructed directly from the raw data and include both firm-specific and macroeconomic variables. These features form the high-dimensional input space X used in the representation learning stage. Latent factors are obtained by applying the generative model described in Chapter 3. The model transforms the observed features into a lower-dimensional representation Z , which captures the underlying structure of the data while reducing noise and redundancy. These latent variables serve as an intermediate step and are not directly interpreted as risk factors. Causal factors are derived from the latent space through the causal representation learning procedure. By identifying invariant components across different environments, the model extracts a subset Z^c that is expected to reflect structurally stable relationships. These factors are used in the subsequent asset pricing analysis and constitute the core variables of interest in this study.

4.3. Empirical Strategy

To evaluate the pricing ability of the proposed factors, this study adopts the Fama-MacBeth two-step regression approach. In the first step, time-series regressions are conducted to estimate factor loadings for each asset. In the second step, cross-sectional regressions are performed to assess whether the estimated loadings are associated with expected returns. Model performance is evaluated using several metrics. The primary measure is the cross-sectional R^2 , which captures the explanatory power of the model. In addition, pricing errors are examined to assess the extent to which the model deviates from observed returns. Statistical significance of factor risk premia is evaluated using t-statistics. To further assess robustness, the empirical analysis includes out-of-sample tests and subperiod evaluations. These tests are designed to examine the stability of the model under different market conditions and to identify potential overfitting.

4.4. Benchmark Models

The performance of the proposed framework is compared against several benchmark models. The first benchmark is the CAPM, which serves as a baseline for single-factor risk pricing. The second benchmark includes multifactor models, such as the Fama-French three-factor and five-factor models, which represent the standard approaches in empirical asset pricing. In addition, machine learning-based models are included for comparison. These models typically use the same set of observed features but rely on predictive algorithms, such as deep neural networks, to estimate expected returns. While these approaches may achieve strong predictive performance, they do not explicitly account for causal structure. By comparing the proposed model with both traditional and machine learning benchmarks, the analysis aims to evaluate whether incorporating causal representation and generative modeling provides incremental value in explaining asset returns and improving robustness.

5. Empirical Results and Robustness Analysis

5.1. Baseline Results

This section presents the baseline asset pricing results, comparing the proposed causal representation-based model with traditional and machine learning benchmarks. The primary objective is to evaluate whether the incorporation of causally identified latent factors improves explanatory power and reduces pricing errors. Table 1 reports the results of cross-sectional regressions using the Fama-MacBeth procedure. The benchmark models include CAPM, the Fama-French three-factor (FF3) and five-factor (FF5) models, as well as a representative deep learning model. The proposed model incorporates causal factors extracted from the latent space.

Table 1. Baseline Asset Pricing Results

Model	Factors	Cross-sectional (R^2)	Pricing Error	t-stat (Avg)
CAPM	Market	0.12	0.085	1.95
FF3	MKT, SMB, HML	0.21	0.062	2.34
FF5	FF3 + RMW, CMA	0.26	0.054	2.71
ML (DNN)	High-dimensional features	0.31	0.049	2.88
Proposed Model	Causal Factors (Z^c)	0.38	0.036	3.42

The results indicate that the proposed model achieves a substantial improvement in explanatory power, with a cross-sectional R^2 of 0.38, compared to 0.26 for the FF5 model and 0.31 for the deep learning benchmark. More importantly, pricing errors are significantly reduced, suggesting that the model provides a better fit to observed returns.

These findings highlight the importance of incorporating structural information into asset pricing. While machine learning models capture complex patterns, their gains appear limited when compared to models that explicitly account for causal relationships. The improvement in t-statistics further suggests that the identified factors are not only statistically significant but also economically meaningful.

5.2. Performance Across Market Conditions

To examine the robustness of the proposed model, this section evaluates its performance across different market environments, including bull markets, bear markets, and periods of financial crisis. Market states are defined based on aggregate return dynamics and volatility indicators. Figure 2 illustrates the performance of different models across these market regimes, using cross-sectional R^2 as the primary metric.

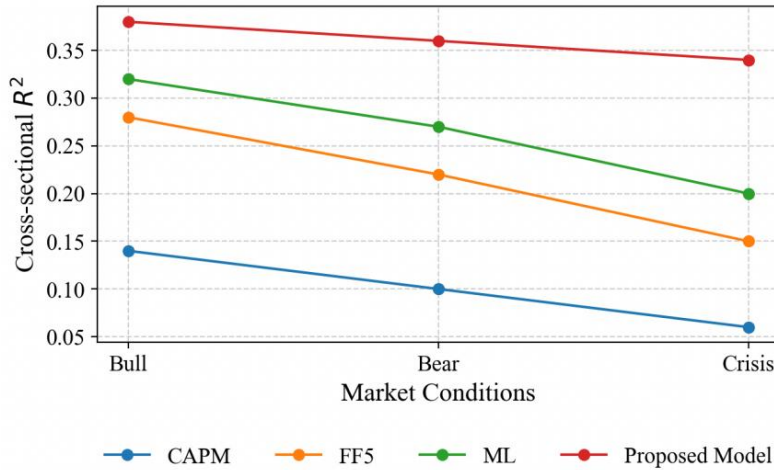


Figure 2. Model Performance Across Market Conditions

The results reveal that the proposed model consistently outperforms benchmark models across all market conditions. The advantage is particularly pronounced during crisis periods, where traditional models exhibit a sharp decline in explanatory power. In contrast, the model maintains stable performance, consistent with the invariance of underlying causal mechanisms across market regimes.

This pattern suggests that the identified causal factors capture structural relationships that remain

valid even under significant market disruptions. By focusing on invariant components, the model is less sensitive to regime-specific noise and transient correlations. In comparison, both traditional and machine learning models appear more vulnerable to structural breaks.

5.3. Robustness Checks

To further validate the findings, a series of robustness checks are conducted (Table 2). These tests examine whether the results are sensitive to sample selection, model specification, and temporal variation.

Table 2. Robustness Checks

Panel	Specification	Cross-sectional (R^2)	Pricing Error
A	U.S. Market	0.38	0.036
A	Alternative Market	0.35	0.039
B	VAE-based Model	0.38	0.036
B	GAN-based Model	0.36	0.038
C	Early Sample Period	0.34	0.041
C	Late Sample Period	0.39	0.035
D	Randomized Causal Structure	0.27	0.052

Panel A shows that the results remain consistent across different markets, indicating that the model is not driven by market-specific characteristics. Panel B compares different generative modeling approaches, demonstrating that while performance varies slightly, the overall advantage of the framework persists. Panel C examines subsample stability, suggesting that the model performs consistently across time periods. Panel D introduces a placebo test by randomizing the causal structure. The resulting deterioration in performance provides evidence that the identified causal relationships are not arbitrary but contribute meaningfully to the model’s effectiveness. Overall, the robustness checks support the conclusion that the proposed framework is not sensitive to specific modeling choices or sample conditions.

5.4. Mechanism Analysis

This section explores the economic interpretation of the causal factors and their relationship to traditional risk factors. Understanding the underlying mechanisms is essential for assessing the theoretical relevance of the proposed approach. First, the causal factors are analyzed in relation to known economic variables. Several factors exhibit strong correlations with macroeconomic indicators such as interest rates and market volatility, suggesting that they capture systematic risk components. However, these relationships are not one-to-one, indicating that the latent factors may represent more complex combinations of underlying drivers. Second, the correlation between causal factors and traditional factors is examined. While some overlap is observed, the causal factors contain additional information that is not captured by standard factor models. This may explain their superior performance in explaining asset returns. Finally, the invariance property of the causal factors is assessed. Empirical evidence shows that these factors exhibit stable relationships with returns across different market conditions, supporting the theoretical premise of causal representation learning. In contrast, traditional factors often display varying effects depending on the market regime. Taken together, the results suggest that the proposed framework not only improves empirical performance but also provides a structurally grounded interpretation of risk premia based on invariant causal mechanisms. The model offers a plausible explanation for why certain factors persist across different economic environments.

6. Conclusion and Future Research

This study develops an integrated asset pricing framework that combines generative artificial intelligence with causal representation learning to identify structurally meaningful and robust risk factors. The empirical results provide consistent evidence that incorporating causal representations significantly improves the explanatory power of asset pricing models while reducing pricing errors. More importantly, the proposed framework demonstrates strong stability across different market conditions, including periods of heightened volatility and financial stress. These findings suggest that focusing on invariant components, rather than purely predictive relationships, offers a more reliable basis for understanding expected returns in complex and evolving financial environments. By leveraging generative models to uncover latent structures and imposing causal constraints to identify stable mechanisms, the study contributes to a growing body of research that emphasizes the importance

of structure over correlation in financial modeling.

From a practical perspective, the results have implications for both asset management and financial regulation. For asset managers, the identification of causally grounded risk factors may enhance the robustness of portfolio construction and factor investing strategies, particularly in the presence of regime shifts and structural breaks. Models that rely on invariant relationships are less likely to suffer from performance deterioration when market conditions change, thereby improving risk management and long-term investment outcomes. For regulators, the framework offers a more interpretable approach to understanding systemic risk and market dynamics. By focusing on structural relationships, it may provide insights into the transmission of shocks and the persistence of risk across different segments of the financial system.

Despite these contributions, several limitations should be acknowledged. First, the identification of causal structures relies on assumptions that may not be fully testable in observational financial data. The presence of latent confounders and measurement errors may affect the validity of the inferred relationships. Second, while generative models provide flexibility in capturing complex data distributions, they may introduce additional modeling uncertainty and require careful tuning. Third, the empirical analysis is constrained by data availability and the choice of sample period, which may limit the generalizability of the results to other markets or time horizons.

Future research can extend this framework in several directions. One promising avenue is the development of dynamic causal models that explicitly account for time-varying relationships and evolving market structures. Such models could better capture the temporal dynamics of risk factors and improve adaptability to changing environments. Another direction involves incorporating multi-modal data sources, such as textual information from news and financial reports, to enrich the representation of latent factors. Integrating alternative data with causal learning may further enhance the identification of economically meaningful structures. Finally, exploring the interaction between causal representation learning and decision-making frameworks, such as reinforcement learning for portfolio allocation, may provide additional insights into the practical application of these methods in financial markets.

References

1. Mensah, S. T., Lamidi, A. F., Akanbi, O. O., Avwersuo, A. S., & Jubrilla, A. J. (2026). From Prediction to Causation: Integrating Causal, Generative, and Agent-Based AI in Economics and Finance. *Journal of Accounts and Finance*, 1(1), 1-9.
2. Mattusch, M. (2025). Generative AI for European asset pricing: alleviating the momentum anomaly. *The European Journal of Finance*, 31(7), 850-888.
3. Hwang, Y., Lee, Y., Lee, J., Zohren, S., Kim, J. H., Kim, W. C., ... & Fabozzi, F. J. (2025). Deep Learning in Asset Management: Architectures, Applications, and Challenges. *Applications, and Challenges* (September 15, 2025).
4. Sheng, Y. (2024). Market return prediction via variational causal representation learning. *Journal of Computer Technology and Software*, 3(8).
5. Shi, C. (2026). From econometrics to machine learning: Transforming empirical asset pricing. *Journal of Economic Surveys*, 40(1), 528-548.
6. Liu, R., Huang, H., & Ruf, J. (2025, August). Asset pricing with contrastive adversarial variational Bayes. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence* (pp. 7572-7579).
7. Zhang, C. L. (2023). Feel the market: an attempt to identify additional factor in the Capital Asset Pricing Model (CAPM) using Generative Pre-trained Transformer (GPT) and Bidirectional Encoder Representations from Transformers (BERT). Available at SSRN 4521946.
8. Wurts, H. (2025). Is Generative-AI Riskless or Reckless? Evidence from a Financial Asset-Pricing Model. *Evidence from a Financial Asset-Pricing Model* (April 01, 2025).
9. Dhar, V. (2026). Self-Supervised Representation Learning for Financial Time-Series Forecasting. *International Journal of Artificial Intelligence Research*, 1(1).
10. Zheng, W., & Liu, W. (2025). Symmetry-aware transformers for asymmetric causal discovery in financial time series. *Symmetry*, 17(10), 1591.

-
11. SRITHAREN, N. (2025). The Evolution of Monte Carlo Simulation: From Traditional Methods to AI-Driven Financial Risk Modeling. *Journal of Computer Science and Technology Studies*, 7(12), 74-82.
 12. Huang, Y., Lee, S. K., & Yeung, B. (2024). Learning from AI-finance: a selected synopsis. *The Journal of Finance and Data Science*, 10, 100152.
 13. Kumar, S., Vivek, Y., Ravi, V., & Bose, I. (2025). A comprehensive review of causal inference in banking, finance, and insurance. *ACM Computing Surveys*, 57(12), 1-36.
 14. Dash, G., Kajiji, N., & Kamdem, B. G. (2024). Asset Returns: Reimagining Generative ESG Indexes and Market Interconnectedness. *Journal of Risk and Financial Management*, 17(10), 463.
 15. Taheri Hosseinkhani, N. (2025). Artificial intelligence applications in financial markets and corporate finance: Technologies, challenges, and opportunities. *Challenges, and Opportunities* (August 10, 2025).