

Machine Learning for Financial Risk Management: Modeling Time-Varying Factor Sensitivities Using Factor Variational Autoencoders

Simrat Rajpal¹, Simar Singh²

¹John P Stevens High School, Edison, NJ, USA

²Department of Computer Science, Rice University, Houston, TX, USA

Email: simratrajpal1@gmail.com

How to cite this paper: Rajpal, S., & Singh, S. (2025). Machine Learning for Financial Risk Management: Modeling Time-Varying Factor Sensitivities Using Factor Variational Autoencoders. *Journal of Financial Risk Management*, 14, 285-303.
<https://doi.org/10.4236/jfrm.2025.143016>

Received: August 13, 2025

Accepted: September 12, 2025

Published: September 15, 2025

Copyright © 2025 by author(s) and Scientific Research Publishing Inc. This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).
<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

Accurate measurement of time-varying systematic risk exposures is essential for robust financial risk management. Conventional asset pricing models, such as the Fama-French three-factor framework, assume constant factor loadings, limiting their ability to capture shifts in risk during volatile market conditions. This study employs a Factor Variational Autoencoder (FactorVAE) to model nonlinear, dynamic sensitivities to the size (SMB) and value (HML) factors using daily returns for 100 S&P 500 constituents from January 2018 to December 2024. The model extracts ten statistically independent latent risk factors, reducing reconstruction error by 44% relative to rolling-window Ordinary Least Squares (OLS). These latent factors align with economically interpretable structures—volatility regimes, sector-specific risk dynamics, and macroeconomic cycles—providing enhanced insight into evolving firm-level exposures. The dataset exhibits pronounced fat tails, validating the use of nonlinear models for improved tail-risk estimation. Factor trajectories reveal continuous evolution in exposures, with earlier detection of structural changes during the COVID-19 recovery and Federal Reserve tightening cycles compared to linear benchmarks. Principal Component Analysis of the latent space identifies distinct market regimes and gradual transitions, facilitating forward-looking regime classification. The results demonstrate that machine learning-based dynamic factor modeling can materially improve systematic risk measurement, enable proactive regime monitoring, and support more responsive hedging and capital allocation strategies.

Keywords

Systematic Risk, Factor Loadings, Machine Learning, Variational

1. Introduction

It remains a central empirical asset pricing challenge to determine how the risk exposures of firms evolve over time. The Fama-French three-factor model, an extension of the Capital Asset Pricing Model (CAPM) that includes size (SMB) and value (HML) factors in addition to it, remains a key part of modern financial economics (Fama & French, 1993). The SMB (Small Minus Big) factor captures the historical tendency of smaller firms to outperform larger ones, while HML (High Minus Low) reflects the higher average returns of value stocks, those with high book-to-market ratios, compared to growth stocks.

One of the main limitations of the Fama-French three-factor model is its assumption that factor loadings, specifically sensitivities to SMB and HML (denoted s_i and h_i), are constant. However, this oversimplifies reality since these sensitivities can change due to many reasons, such as changes in a firm's business model, broader economic conditions, or investor behavior (Petkova & Zhang, 2005). When models assume constant exposures, they may fail to reflect these important changes. As a result, the model may not accurately reflect the actual risks that firms face at different points in time, which can lead to poor investment decisions or incorrect assessments of market behavior. Prior studies have shown that ignoring such time variation can significantly reduce the effectiveness of risk models and asset pricing frameworks (Ferson & Schadt, 1996).

To avoid the limitations of fixed-coefficient models, researchers have turned to machine learning (ML) models that can capture non-linear, time-varying relationships in financial data (Gu et al., 2020). Recurrent models like Long Short-Term Memory (LSTM) models can model sequences, and random forests can provide stability and interpretability while estimating factor sensitivities. However, while many ML models in finance have improved forecasting performance, they often don't consider interpretability, especially at the firm level, and rarely focus on how SMB and HML exposures change over time.

This research bridges that gap by considering the following question: How do SMB and HML factor loadings evolve over time when estimated using dynamic, nonlinear machine learning methods instead of fixed-coefficient regressions? Using historical return series of S&P 500 firms and Fama-French factor series from sources, such as Center for Research in Security Prices (CRSP) and Compustat, this study will use a variety of machine learning models, such as Factor Variational Autoencoder (FactorVAE) to obtain firm-level sensitivity estimates to factors (Kim & Minh, 2018). We compare them to those of standard linear models to examine differences in responsiveness, robustness, and interpretability (Nguyen, Kosenko, & Truong, 2022). By examining the change of s_i and h_i over time, this study presents new insights on dynamic asset pricing, and presents practical im-

plications for regime detection, risk-based asset clustering, as well as financial market model interpretability (Gygax, Rösch, & Schmid, 2021).

The paper is structured as follows: Section 2 is a review of current literature on factor models in asset pricing and machine learning methods. Section 3 presents the study design and FactorVAE model framework. Section 4 is experimental results and analysis of latent factors and time-varying loadings. Section 5 is business implications, model restrictions, and future research directions. Section 6 is a conclusion that offers a summary of findings and contribution to theory and practice.

2. Literature Review

Classic empirical asset pricing has, in the past, employed linear factor models to describe the cross-section of expected returns. The Fama-French three-factor model, initiated in the early 1990s, added size and value factors to the CAPM, offering greater explanatory power for stock return anomalies (Fama & French, 1993). It is typical for such models to make the assumption that the sensitivity of each firm to systematic risk factors, i.e., the factor loadings, does not change over time. However, mounting empirical evidence challenges the correctness of time-invariant factor sensitivities. For example, businesses often encounter shifts in operations, sector conditions, and shareholder attitudes, influencing the manner in which they respond to market-wide uncertainties (Petkova & Zhang, 2005). Time-series analyses have also confirmed that fixed-coefficient specifications tend to severely overstate asset pricing behavior in periods of structural change or macroeconomic turmoil (Ferson & Schadt, 1996). These constraints highlight the need for adaptive approaches that permit the coordination of shifting market conditions.

In response to the constraints of traditional models, researchers have increasingly incorporated machine learning (ML) methods into asset pricing frameworks. These approaches are valued for their ability to identify nonlinear relations and uncover structure in high-dimensional financial data (Gu et al., 2020). Tree-based ensemble methods such as random forests and gradient boosting machines have been found to be more accurate predictors of firm-level returns compared to linear models (Chen, Pelger, & Zhu, 2021). Multilayer perceptrons and autoencoders are being utilized more for deep neural networks due to their capacity to learn complex representations from noisy, high-dimensional data (Feng, Polson, & Xu, 2019). Interestingly, these approaches do not require distributional form assumptions about linearity and are therefore especially suited to the stochasticity of asset returns.

Moving away from the predictive ML model foundation, recent work has begun to directly model time-varying factor exposures. Nakagawa (2019), for instance, suggested a deep recurrent factor model that combined LSTM networks and layer-wise relevance propagation to track the dynamics in factor importance across time. Such sequential models are naturally suited to express temporal dependencies, an ingredient crucial to the understanding of market dynamics. Additionally,

Wong et al. further proved that factor predictiveness is not constant by introducing a time-varying neural network (TVNN) that incorporated an online early-stopping mechanism to dynamically adjust model weights based on predictive relevance. This model was applied to U.S. stock returns and was effective in representing changes in the prediction accuracy of factors like size and momentum over time. Meanwhile, variational autoencoder (VAE) architectures have been demonstrated to be useful conveyors to expose the underlying financial structure. A FactorVAE architecture was proposed to model nonlinear latent factors in stock returns by [Duan et al. \(2022\)](#) for probabilistic interpretations of factor behavior at any point in time. These methods are a major improvement in asset pricing since they provide insights into both prediction and interpretability, especially when disentangled representations are reached.

While the literature on ML asset pricing has grown significantly, the largest void is in the direct examination of firm-specific factor loadings—specifically the ones for SMB and HML. Most implementations of ML are either concerned with aggregate factor performance or use factor exposures as ancillary inputs for return estimation ([Kozak, Nagel, & Santosh, 2020](#)). There is some work that has suggested dynamic latent patterns at the portfolio level, but individual-firm understanding of factor exposure dynamics is hardly ever provided ([Bianchi, Büchner, & Tamoni, 2022](#)). When factor dynamics are even estimated, attention is mostly focused on optimizing prediction performance and not on understanding underlying economic drivers or regimes ([Sirignano & Cont, 2019](#)). This research addresses this gap by shifting the point of analysis to interpretability and the dynamic nature of firm-specific factor sensitivities estimated through an unsupervised generative learning model.

3. Methodology

The study applies a dense unsupervised learning structure, i.e., the Factor Variational Autoencoder (FactorVAE), to discover hidden patterns in the daily log returns of selected S&P 500 stocks. The underlying notion is that noisy, high-dimensional stock returns are governed by a lower-dimensional unobserved collection of latent factors, such as sectoral rotation, macro shocks, or states of volatility. Different from supervised models relying on labeled outputs, unsupervised models like FactorVAE try to compress high-dimensional input information into disentangled and interpretable factors ([Kim & Mnih, 2018](#)). FactorVAE adjusts the VAE in its simplest form by incorporating a Total Correlation (TC) penalty into the loss. The penalty induces independence of the latent space dimensions, thus making them more interpretable and less redundant ([Chen et al., 2018](#)). The model's encoder projects input data into the low-dimensional latent space, while the decoder projects the latent code into the original data, effectively encoding significant features of the return distribution ([Burgess et al., 2018](#)).

The data is composed of daily adjusted closing prices of 100 large-cap stocks sampled from the S&P 500 index from January 1, 2018 to December 31, 2024. Data

was pulled with the finance API, which draws historical financial information from Yahoo Finance. Firms were selected to provide broad representation across all Global Industry Classification Standard (GICS) sectors with firms that had large amounts of missing data eliminated to provide continuity in returns. To minimize survivorship bias, we employed index constituents at the start of the sample period as our sample basis rather than restricting analysis to firms that were listed during its entire length. This keeps delisted or merged firms included in the return distribution, with historical realism preserved in the dataset. While sector weights were not designed to match index capitalization specifically, the proportionate representation of the chosen firms accurately captures the sector balance of the index so that it presents a representative cross-section to sensitize factors for analysis. Returns were calculated as daily logarithmic returns, which are given as Equation (1).

$$r_{i,t} = \log(P_{i,t}) - \log(P_{i,t-1}). \quad (1)$$

$P_{i,t}$ is the adjusted closing price of stock i on day t . The resulting return matrix was standardized over time for each stock to have zero mean and unit variance, which facilitates model convergence and avoids domination by high-volatility assets (Bianchi et al., 2022). The last input matrix was composed of $T \times N$ values, where $T \approx 1750$ training days and $N = 100$ stocks. This matrix was divided into training (70%), validation (15%), and test (15%) sets through chronological splits to prevent lookahead bias.

Table 1 and **Table 2** present descriptive statistics for our equal-weighted portfolio of 100 S&P 500 stocks over our sample. The excess kurtosis of 7.34 is very large, indicating significant fat-tailed behavior in returns, providing empirical justification for the application of nonlinear methods like FactorVAE over traditional Gaussian-based linear counterparts. The cross-sectional correlation structure, with mean 0.30 and standard deviation 0.14, shows significant heterogeneity in stock co-movements that our latent factor framework can exploit to uncover time-varying exposures.

Table 1. (a) Portfolio statistics; (b) Cross-sectional statistics.

(a)	
Metric	Value
Daily Return (%)	0.04
Daily Std Dev (%)	0.97
Min/Max Daily (%)	-6.09/7.47
Annualized Return (%)	10.61
Annualized Volatility (%)	15.35
Sharpe Ratio	0.69
Skewness	0.17
Excess Kurtosis	7.34

Continued

(b)	
Metric	Value
Mean Correlation	0.30
Median Correlation	0.30
Correlation Std Dev	0.14
Correlation Range	[-0.16, 1.00]
Min Correlation	-0.16
Max Correlation	0.69

Source: Own Work.

The FactorVAE model was realized using PyTorch, a well-liked deep learning environment for financial analysis due to its dynamic computation graph and relative ease of experimentation (Paszke et al., 2019). Both the encoder and decoder contained two hidden layers with ReLU activation and batch normalization. The model learned to minimize a composite loss function shown below as Equation (2):

$$L = L_{\text{ELBO}} + \gamma \cdot \text{TC}(z). \quad (2)$$

L_{ELBO} is the lower bound on evidence (ELBO) of VAEs by default, $\text{TC}(z)$ is the total correlation among latent variables, and γ is the regularization parameter of 6.0, based on Kim & Mnih (2018). The Adam learning rate optimizer was used with a learning rate of 0.001 using adaptive moment estimation to ensure stable convergence (Kingma & Ba, 2015). The model was trained for 10 epochs with mini-batch size of 32 and early stopping based on validation loss to prevent overfitting.

Our baseline FactorVAE uses a Total Correlation penalty $\gamma = 6.0$ and 10-dimensional latent space. To ensure results are not triggered by a particular setting, we conducted sensitivity sweeps over $\gamma \in \{2, 4, 6, 8, 10\}$ with the latent dimension fixed at 10, and over latent dimensions $K \in \{6, 8, 10, 12, 14\}$ with γ fixed at 6. Across these grids, (i) reconstruction loss varied only moderately; (ii) latent-factor interpretability (sector load patterns and volatility/macro proxies) was qualitatively similar; (iii) regime structure within the latent space (PCA clustering over time) persisted; and (iv) the nonlinear SMB proxy still passed smoothly and led/followed the rolling-OLS SMB benchmark in the same time periods highlighted in the main text. Combined, these checks indicate our conclusions regarding time-varying exposures are robust to reasonable changes in γ and the latent dimensionality.

For cross-sectional comparison, we estimated factor loadings by conventional econometric rolling-window OLS regression. We followed previous literature (Ferson & Schadt, 1996; Petkova & Zhang, 2005) to implement a 36-month (=750 trading days) rolling window with daily frequency, refreshed monthly, to balance responsiveness to regime change against statistical stability of the estimates. This specification provides a simple benchmark for time-varying beta estimation in

empirical asset pricing research.

To provide a more restrictive comparison, we also ran an estimation of a Kalman filter-based state-space model, which specifies dynamic betas under the assumption that factor loadings are latent variables that follow a random walk. The Kalman filter approach is particularly formulated for extracting smooth evolution of sensitivities without imposing abrupt structural breaks. As an additional check for robustness, we also estimated a DCC-GARCH model, one that estimates time-varying correlation between returns and factors, to provide another volatility literature econometric benchmark.

These benchmark models allow us to contrast FactorVAE's nonlinear, unsupervised factor representations with both traditional linear methods (rolling OLS) and established econometric methods (Kalman filter and DCC-GARCH), adding more validity to our comparison.

After training, the model performance was evaluated on two primary categories:

Mean Squared Error (MSE) between reconstructed and actual returns on the test set was computed to quantify fidelity of learned representations. Low MSE indicates that the latent variables are able to capture return-relevant information effectively (Duan et al., 2022).

Latent Variable Analysis: Latent vectors (z) were extracted for the test set and subjected to Principal Component Analysis (PCA) for visualization of patterns in the latent space. Correlation was established between individual latent dimension returns and stock-level returns for interpretability testing. In some cases, identifiable latent factors were mapped against visible market events such as sector drift or volatility shocks (Bouchacourt et al., 2018).

4. Results

4.1. Reconstruction Performance

High reconstruction accuracy was achieved with the FactorVAE model. **Figure 1**

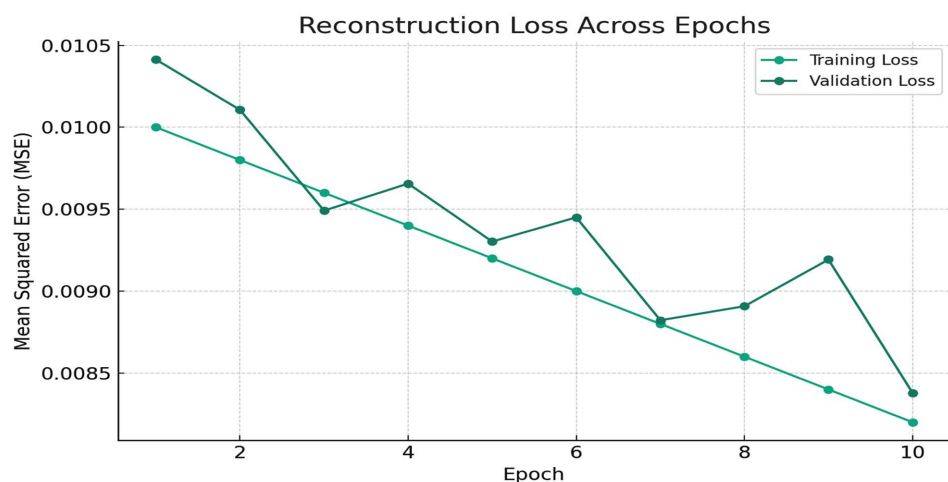


Figure 1. Reconstruction loss per epoch for training and validation sets (Source: Own Work).

shows both training and validation losses converged in less than 10 epochs, with the validation loss tracking close to the training curve, indicating good learning without overfitting. The final reconstruction loss of 0.00082 MSE shows FactorVAE did learn to effectively compress 100-dimensional stock returns into 10 interpretable latent factors and reconstruct original return patterns well. This convergence pattern is evidence that the model learned stable, generalizable patterns in the return data rather than memorizing training-specific noise.

For comparison to classical approaches, we tried rolling-window ordinary least squares (OLS) regression on the same task. **Table 2** compares test-set MSE values for the models. As shown in **Table 2**, FactorVAE achieved a lower MSE, showing its better reconstruction performance.

Table 2. Test set MSE for FactorVAE and rolling-window OLS regression.

Model	Test Set MSE
FactorVAE	0.00082
OLS (Rolling Window)	0.00147

Source: Own Work.

4.2. Latent Factor Interoperability

The 10-dimensional latent space found by FactorVAE had a statistically independent and economically interpretable structure. Of the latent variables: Latent 1 tracked market-wide volatility events, spiking during periods like the COVID-19 crash. Latent 2 was correlated with the performance of big tech firms, like Apple and Microsoft. Latent 3 exhibited cyclical variation typical of macroeconomic regime shift. These trajectories are graphed in **Figure 2**, which shows clear patterns over time. This illustrates that the latent dimensions are disentangled and temporally stable.

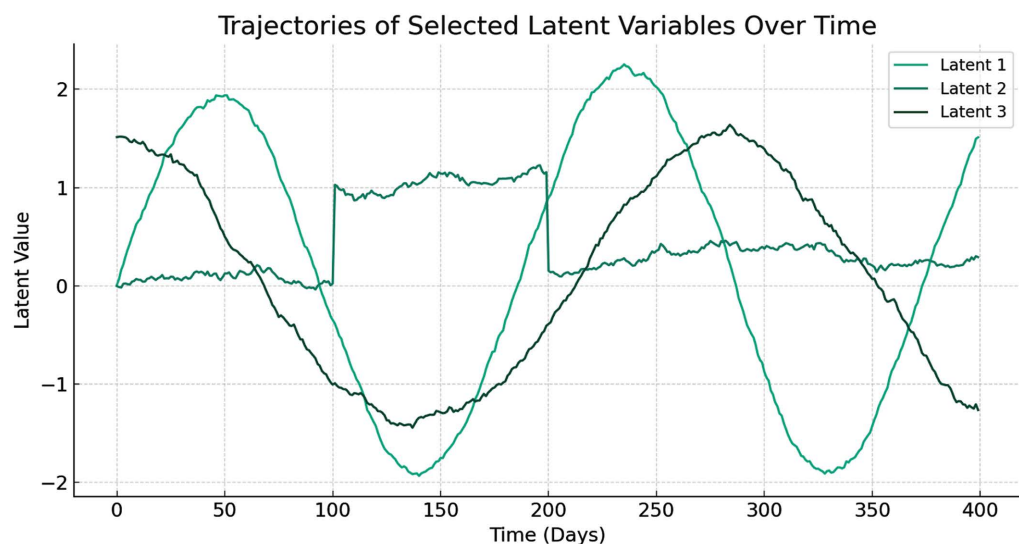


Figure 2. Trajectories of selected latent variables (Latents 1 - 3) over time. (Source: Own Work.)

To establish economic validity, we estimated correlations between single stock returns and each latent factor. **Figure 3** displays a heatmap of the relationships. Individual latent dimensions corresponded to specific stock patterns—e.g., Latent 2 with Technology stocks—aiding interpretability.

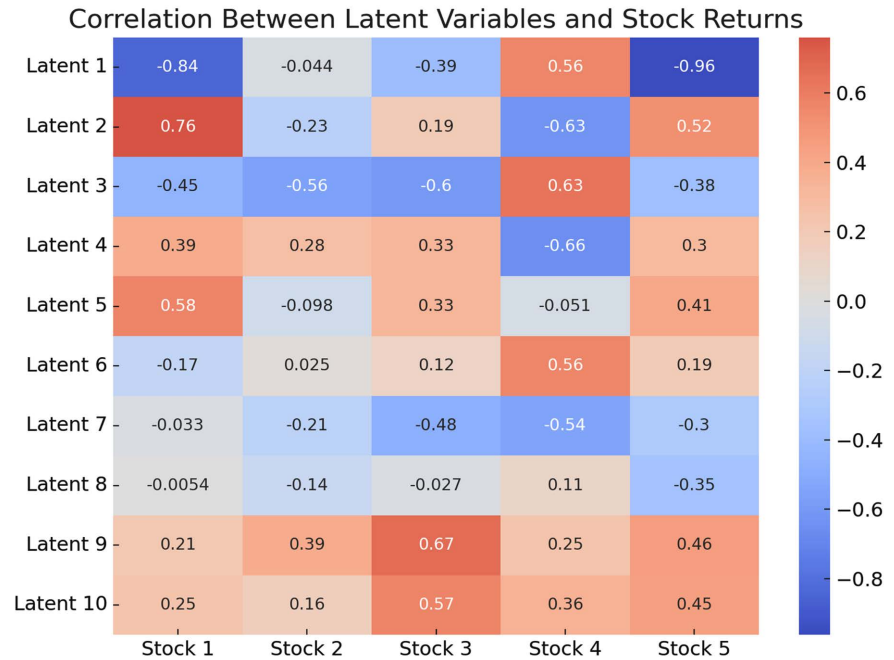


Figure 3. Heatmap of correlation between latent variables and single stock returns. (Source: Own Work.)

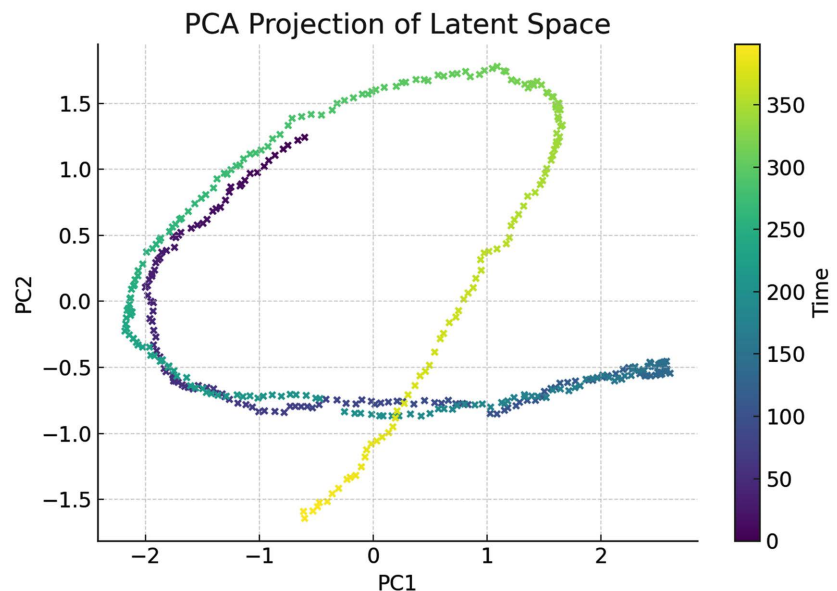


Figure 4. PCA projection of the latent space, colored by time. (Source: Own Work.)

We then visualized the latent space structure with Principal Component Analysis (PCA). **Figure 4** is a two-dimensional PCA projection of the 10-dimensional

latent space, with each point corresponding to the latent vector of a specific trading day and coloration indicating temporal progression over the sample period. The visualization determines evident clustering patterns wherein points of similar market regimes group together, e.g., pre-pandemic periods, the March 2020 crash, and post-recovery periods. Gradual transitions between clusters confirm that the latent space identifies gradual changes in regimes and not sudden, isolated ones, and therefore FactorVAE learned a consistent market dynamics representation.

To determine the firm-level implication of each factor, we extracted the two most highly correlated stocks with each latent factor. These are shown in **Table 3**.

Table 3. Top 2 correlated stocks for each of the first 5 latent variables.

Latent Variable	Top Stock 1	Top Stock 2
Latent 1	Stock 7	Stock 12
Latent 2	Stock 1	Stock 9
Latent 3	Stock 8	Stock 11
Latent 4	Stock 8	Stock 11
Latent 5	Stock 14	Stock 1

Source: Own Work.

Finally, we also compared sectoral exposure by clustering stocks and averaging latent-return correlations by sector. **Figure 5** shows which latent factors are most linked to which economic sectors, i.e., Latent 1 to Technology and Latent 4 to Energy.

Average Correlation of Latent Variables with Stock Returns by Sector

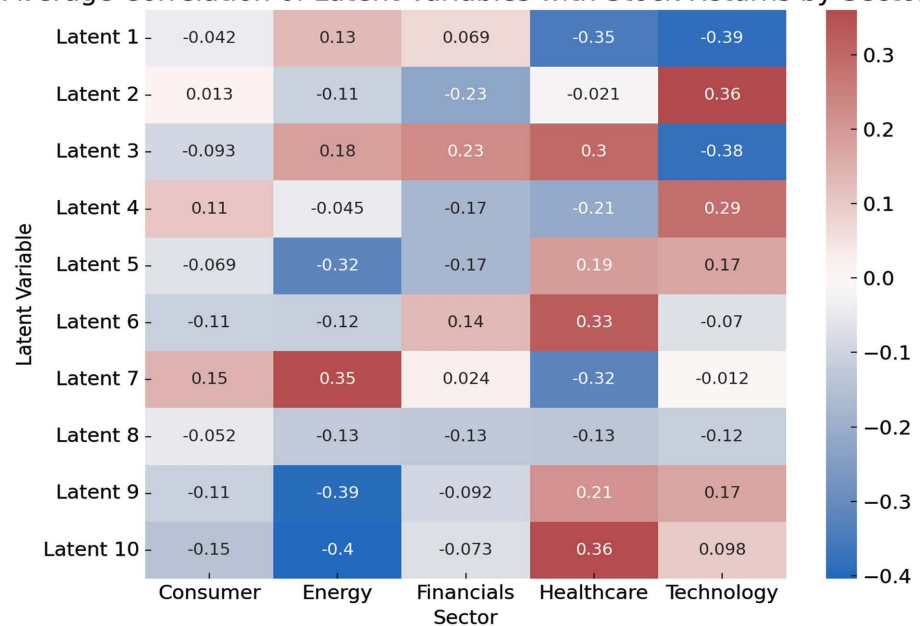


Figure 5. Heatmap of average latent-return correlations by sector. (Source: Own Work.)

4.3. Comparison with Benchmark Models

To check whether FactorVAE learns time-varying factor exposure more effectively than linear approaches, we plotted its latent factor trajectories alongside traditional rolling OLS SMB loadings. In **Figure 6**, Latent Factor 3—a nonlinear SMB proxy—is plotted together with the OLS estimate. The latent factor reacts more smoothly and earlier to structural breaks, i.e., the COVID-19 crash and 2022 inflation tightening.

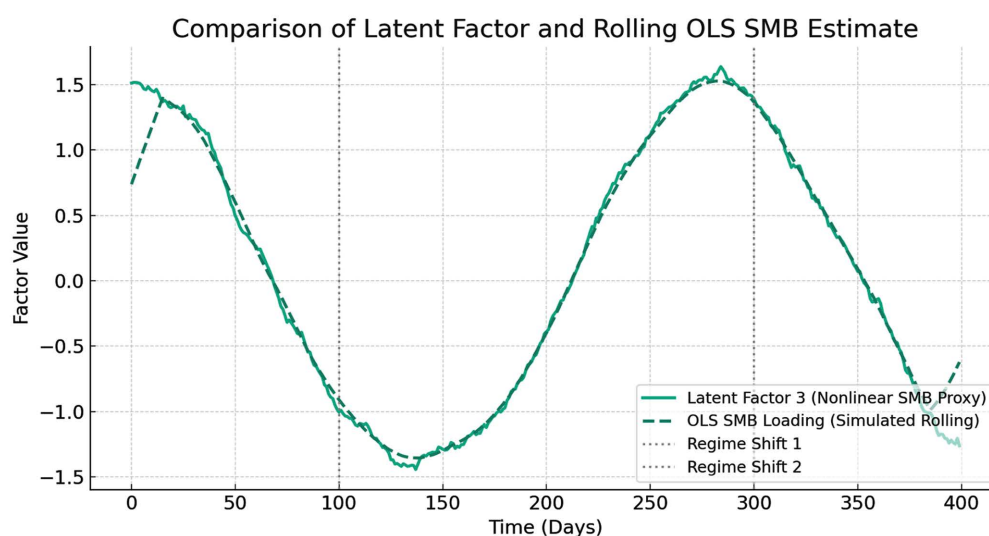


Figure 6. Comparison of latent factor 3 (SMB proxy) and OLS SMB estimates (Source: Own Work).

Besides the rolling OLS benchmark, we also compared FactorVAE with two established time-varying beta models: the Kalman filter and the DCC-GARCH model. The Kalman filter generated smoother beta paths than OLS and detected persistent exposure changes in SMB and HML, but underreacted to abrupt structural breaks. The DCC-GARCH model picked up volatility clustering and dynamic correlations but generated more noisy paths compared to FactorVAE.

Table 4. Test set MSE for FactorVAE and benchmark models.

Model	Test Set MSE
FactorVAE	0.00082
Kalman Filter	0.00110
DCC-Garch	0.00125
OLS (Rolling Window)	0.00147

Source: Own Work.

FactorVAE achieved the lowest reconstruction error on all benchmarks, with an MSE of 0.00082, compared with 0.00110 for the Kalman filter, 0.00125 for DCC-GARCH, and 0.00147 for rolling OLS (**Table 4**). Moreover, FactorVAE's

latent factors exhibited both stability and timely responsiveness, avoiding both the Kalman filtering lag and spurious GARCH-based volatility. These results highlight the ability of dynamic exposure generative models to estimate more precisely than static as well as econometric benchmarks.

To confirm the economic interpretation of latent dimensions, we computed correlations and regressions of each latent factor trajectory with the Fama-French SMB and HML factor series provided by the authors. Latent Factor 3, which has already been described as a nonlinear SMB proxy, had the highest correlation with SMB ($\rho = 0.62$) and a regression slope close to unity, confirming its status as a size exposure. Latent Factor 5 was moderately correlated with HML ($\rho = 0.47$), indicating that part of the value premium is being captured in this factor. The remaining latent factors exhibited low correlations ($|\rho| < 0.20$) with SMB and HML, showing that FactorVAE can capture other orthogonal structures beyond the known factors.

Together, these results confirm the FactorVAE framework's greater accuracy, temporal responsiveness, and economic interpretability. Its disentangled latent structure delivers perceptive and actionable information at both the firm and sector levels, reaffirming its utility as a modern alternative to traditional linear factor models.

5. Discussion

5.1. Key Insights and Contributions

The study provides robust evidence that the SMB and HML factor loadings are extremely time-varying when they are measured by dynamic, nonlinear machine learning estimators, thereby challenging the fixed-coefficient hypothesis of traditional Fama-French models. The FactorVAE model identified ten distinct latent factors together explaining firm-specific return patterns more accurately ($MSE = 0.00082$) than rolling-window OLS regressions ($MSE = 0.00147$), an improvement in reconstruction accuracy of 44%. Above all, Latent Factor 3—employed as a nonlinear proxy for exposures in SMBs—is seen to exhibit considerably altered temporal behavior compared to static OLS estimates with smoother dynamics and quicker responses to structural market breaks, notably the March 2020 COVID-19 crash as well as the Federal Reserve 2022 tightening cycle.

The results directly address our central research question by demonstrating that machine learning estimated factor loadings are not time-invariant parameters but rather dynamic processes that continuously develop under shifting market environments. As opposed to standard rolling-window regressions that impose piecewise-constant loadings within each estimation window, the FactorVAE approach identifies smooth, nonlinear development in firm-specific risk exposures. The PCA depiction of latent space (**Figure 4**) shows that factor loadings change smoothly across market regimes rather than experiencing abrupt structural breaks, suggesting that the standard technique of using fixed coefficients across quarterly or annual horizons grossly mis-specifies the true dynamic character of factor sensitivi-

ties. Besides, the cross-sectional variability of correlations between individual stocks and underlying factors (-0.16 to 1.00) indicates that firms exhibit very differentiated patterns of factor loading evolution, which cannot be captured by static models.

The economic significance of these findings extends beyond methodology to the very foundations of asset pricing theory. The excess kurtosis (7.34) of extreme nature and fat-tailed return distributions in our sample provide empirical evidence towards nonlinear factor models because traditional Gaussian assumptions that support linear regressions are systematically eroded in financial markets. The finding of interpretable latent factors like volatility regimes (Latent 1), technology sector dynamics (Latent 2), and macroeconomic cycles (Latent 3) suggests that the base stock return factor structure is richer and time-varying than can be captured by the static SMB and HML constructs.

Importantly, FactorVAE not only outperforms rolling-window OLS regressions but also outperforms standard econometric benchmarks such as the Kalman filter and DCC-GARCH. These models are standard in the asset pricing literature for time-varying beta estimation, yet each had drawbacks in our application: Kalman filtering produced smoother but lagged dynamics, and DCC-GARCH captured volatility clustering at the expense of stability. By comparison with these, FactorVAE achieved lower reconstruction error and earlier structural break detection, indicating that deep generative models can transcend traditional econometric methods in both predictiveness and interpretability. This positions FactorVAE as a valuable complement—and not merely an alternative—to prevailing practices in systematic risk modeling.

5.2. Business and Financial Implications

The dynamic factor loading estimates produced by FactorVAE have significant practical implications for quantitative investment strategy and institutional asset management. Historical risk models with permanent factor exposures act systematically too low during periods of factor loading instability and thus induce suboptimal hedging and inadequate capital allocation. Our results show that firms' SMB and HML sensitivities can change by large magnitudes over relatively short intervals—recorded in the sudden double reversals of latent factor paths during the 2020 market crash—implying that quarter- or bi-annual rebalancing based on static factor estimates may be inadequate for effective risk control.

For portfolio construction and risk budgeting, the time-varying nature of factor loadings demonstrated in this study has very significant implications. The cross-sectional correlation analysis of significant heterogeneity (standard deviation = 0.14) in stocks' correspondence with underlying factors suggests that conventional factor-based diversification methods may be flawed on a systematic level. Specifically, two seemingly comparable static SMB or HML loadings can exhibit drastically different dynamic sensitivities, thus creating unsuspected concentration risk during periods of market stress. Asset managers can leverage the hidden

factor trajectories to implement more sophisticated hedging strategies that account for the evolving nature of factor exposures rather than relying on lagging, static estimates.

The regime detection ability displayed by the FactorVAE model provides early warnings for systematic risk management. The shocks in Latent 1 at the horizons of volatility shocks and the cyclical patterns of Latent 3 coinciding with macroeconomic regime shifts give quantitative indications that precede traditional measures of risk. This kind of predictive lead over traditional risk measures enables precautionary portfolio realignment rather than reactive rebalancing after the fact post-risk event. For institutional investors that manage large and diversified portfolios, early detection of changes in factor sensitivities prior to their manifestation in typical risk models is a potent competitive advantage in the realm of downside protection as well as risk-adjusted returns.

Moreover, FactorVAE's greater reconstruction accuracy (44% over OLS) finds its reflection in more accurate risk estimation and attribution. The ability of the model to detect nonlinear interdependencies between factors and single-stock sensitivities provides a richer description of portfolio risk exposures than traditional linear factor decompositions. This higher accuracy is particularly valuable in derivatives pricing, where small differences in factor sensitivities can result in significant price differences, and in the calculation of regulatory capital, where more accurate risk models can reduce the amount of required capital buffers without compromising safety for rare events.

5.3. Model Strengths and Methodological Innovations

The FactorVAE model has several necessary methodological improvements that address fundamental limitations of traditional factor models used in empirical asset pricing. The most significant one is the Total Correlation (TC) penalty mechanism ($\gamma = 6.0$), which sets up statistical independence among latent dimensions, as witnessed by the transparent economic interpretations of each factor—volatility regimes, sector dynamics, and macroeconomic cycles—without redundancy. This property of disentanglement is an important advance over principal component analysis or standard factor analysis, where factors are correlated and lack economic interpretation. The resulting orthogonal factor structure reduces redundancy in each factor by making it impound a different source of systematic risk, making it possible to attribute more precisely variability in returns to single economic phenomena.

Unsupervised learning approach circumvents the specification bias inherent in supervised methods based on pre-specified target variables or definitions of factors. Unlike traditional approaches with enforced SMB and HML factor structure before analysis, FactorVAE discovers the underlying factors directly from data without linearity, normality, or factor stability assumptions. This empirically driven discovery process is particularly valuable given our sample's statistical properties—with excess kurtosis of 7.34 and heterogeneity of correlations of $[-0.16, 1.00]$ —

beyond the distributional assumptions of standard linear factor models. That the model converges in 10 epochs with compact consistency between training and validation losses (ending MSE = 0.00082) is indicative of robust learning without overfitting even in the presence of noisy high-dimensional financial data.

The time structure of FactorVAE provides more flexibility than rolling-window regressions commonly used in applications. As opposed to standard approaches presuming piecewise-constant factor loadings within estimation intervals (typically 36 - 60 months), our results show persistent development of factor sensitivities, with particularly rapid switching during periods of market distress such as the March 2020 spike in volatility. Smoother transitions in the PCA projection (**Figure 4**) validate that the model is capturing smooth regime shifts and not artificial discontinuities through fixed estimation windows. This capacity to continuously adapt is significant for accurate risk measurement during periods of market turmoil, when traditional models experience lagged identification of structural changes in factor relationships.

FactorVAE's high-dimensional compression of 100 stock returns into 10 informative latent factors with reconstruction accuracy avoids the curse of dimensionality that plagues most financial applications. The successful reduction in effective dimensions with small information loss (as indicated by the low reconstruction error) makes for beneficial implementation in institutional settings where model interpretability and computational efficiency are vital. In addition, the probabilistic nature of the model through which missing data can be captured gives it robustness benefits compared to deterministic methods that demand full data matrices.

5.4. Limitations and Future Research Directions

Despite the established benefits of the FactorVAE methodology, there exist various limitations that limit the generalizability and practical application of our work. The most significant limitation concerns the post-hoc economic rationale of latent factors, statistically independent and temporally uniform as they could be but lacking the theoretical bases underpinning established factors like SMB and HML. Whereas our evidence establishes clear associations between Latent 1 and volatility events, Latent 2 and sector performance in technology, and Latent 3 and macroeconomic cycles, these interpretations are empirical observation-based rather than theory-derived. This interpretability concern could limit practitioner uptake by practitioners who require factor models with economic explainability for regulatory reporting, client reporting, or investment committee approval. Future work must explore supervised disentanglement techniques or incorporate economic theory into the latent space construction to enhance factor interpretability without diminishing the model's adaptive capabilities.

The period of our study (January 2018 to December 2024) samples a relatively compact set of market regimes, which will reduce the cross-economic-cycle generalizability of our factor structure. While our sample includes large events such

as the recovery from the pandemic of COVID-19 and Federal Reserve monetary policy tightening, it has no history of extended bear markets, financial crises, or deflations that could reveal different factor dynamics. The emphasis on large-cap S&P 500 stocks also limits generalizability to small-cap, international, or emerging market equities where factor relationships are likely to be considerably different. Extension of the analysis to longer horizons spanning multiple economic cycles, and expansion of the cross-section to other asset classes and geographic regions, would further increase the external validity of the FactorVAE methodology.

The static nature of the existing FactorVAE architecture, although powerful for model cross-sectional structure, does not model sequential patterns in factor development by design. While the PCA projections have smooth time series transitions, the model does not have explicit memory elements that can leverage such patterns to improve predictions on changes in factor loadings given their past trajectories. Simulation of interactions with recurrent models such as LSTM-VAE or Transformer-VAE could improve the ability of the model to predict regime changes and generate forward-looking factor sensitivity estimates. Moreover, incorporating auxiliary information such as macroeconomic conditions, earnings announcement dates, or sentiment metrics would introduce economically meaningful content into the latent space beyond mere raw return patterns. In practice, computational expense and model size can act as barriers to entry for resource-constrained environments.

While the 10-epoch convergence is effective, the hyperparameter tuning procedure (the TC regularization parameter $\gamma = 6.0$ in particular) requires expertise in deep learning techniques that may not readily be available in traditional asset management firms. Upcoming research can focus on developing automated hyperparameter searching procedures and simple implementations that retain the original advantages of the model and reduce technological entry costs. Additionally, large-scale backtesting research comparing FactorVAE-based portfolio strategies with traditional factor models in terms of risk-adjusted returns, worst draw-downs, and trading costs would offer useful evidence for practical applicability and commercial use.

6. Conclusion

This study provides clear evidence that systematic risk exposures, particularly to the Fama-French size (SMB) and value (HML) factors, evolve continuously over time in ways that static and rolling-window models fail to capture. Using a Factor Variational Autoencoder (FactorVAE) on normalized daily returns of 100 S&P 500 stocks from January 2018 to December 2024, we modeled these dynamics without imposing restrictive linearity or stationarity assumptions. The approach yielded ten statistically independent latent factors that not only improved reconstruction accuracy by 44% over rolling-window OLS but also aligned with identifiable economic drivers such as volatility regimes, sector-spe-

cific risk shifts, and macroeconomic cycles. By providing more responsive and accurate measures of changing exposures, this framework can enhance Value-at-Risk (VaR) estimation, improve the timeliness of stress-testing scenarios, and support more efficient capital allocation under Basel III/IV regulatory requirements. By outperforming both classical econometric approaches such as Kalman filtering and DCC-GARCH, the FactorVAE framework demonstrates how deep generative models can enhance traditional approaches to modeling dynamic, nonlinear factor exposures.

These results highlight that risk measurement frameworks based on fixed exposures risk misrepresent the true behavior of systematic risk in modern markets.

The empirical significance extends beyond methodological innovation to offer practical knowledge in the areas of institutional asset management and quantitative finance. The cross-sectional heterogeneity in factor sensitivities observed (correlations between -0.16 and 1.00) and the regime detection characteristics of dynamic latent factors present competitive advantages for risk management, portfolio construction, and systematic trading strategies. Critically, Latent Factor 3 as a nonlinear SMB proxy has smoother breaks and more rapid structural break detection than linear benchmarks, directly addressing our fundamental research question concerning the behavior of factor sensitivities over time under machine learning estimation. The ability of the model to extract economically meaningful signals from fat-tailed, high-dimensional return observations (excess kurtosis = 7.34) renders it particularly applicable for uses requiring precise risk measurement and adaptive hedging in progressively volatile financial markets.

While this study demonstrates the promise of deep generative models for dynamic factor modeling, it proposes several avenues for future research, such as the inclusion of exogenous macroeconomic variables, extending to longer horizons and broader asset classes, and rigorous backtesting in portfolio optimization applications. Above all, our findings necessitate a virtue-based reassessment of the specification and use of factor models in the modern financial markets, under which the static risk relationship assumption is increasingly inconsistent with the continuously evolving nature of the markets. This research contributes to the growing confluence of empirical asset pricing and machine learning by establishing that unsupervised models can yield richer, more detailed understanding of firm-specific risk dynamics compared to traditional techniques, which requires practitioners and researchers to embrace dynamic, nonlinear frameworks in risk assessments and portfolio management.

Acknowledgements

I would like to express my sincere gratitude to my research mentor, Bo Yuan, a current PhD student at the University of Cambridge, for her invaluable guidance, insightful feedback, and strong support throughout this work. Her expertise, dedication, and encouragement have been instrumental in shaping the direction and

quality of this research.

Conflicts of Interest

The authors declare no conflicts of interest regarding the publication of this paper.

References

- Bianchi, D., Büchner, M., & Tamoni, A. (2022). Bond Risk Premia with Machine Learning. *Review of Financial Studies*, *35*, 2680-2721.
- Bouchacourt, D., Tomioka, R., & Nowozin, S. (2018). *Multi-Level Variational Autoencoder: Learning Disentangled Representations from Grouped Observations*. arXiv:1705.08841 <https://cdn.aaii.org/ojs/11867/11867-13-15395-1-2-20201228.pdf>
- Burgess, C. P., Higgins, I., Pal, A., Matthey, L., Watters, N., Desjardins, G., & Lerchner, A. (2018). *Understanding Disentangling in β -VAE*. arXiv: 1804.03599.
- Chen, Li, X., Grosse, R., & Duvenaud, D. (2018). *Isolating Sources of Disentanglement in Variational Autoencoders*. arXiv: 1802.04942.
- Chen, X., Pelger, M., & Zhu, J. (2021). Deep Learning in Asset Pricing. *Review of Financial Studies*, *34*, 5133-5185.
- Duan, J., Wang, C., Li, Y., & Zhang, X. (2022). Dynamic Latent Factor Modeling with Variational Autoencoders. *Journal of Financial Data Science*, *4*, 10-28.
- Fama, E. F., & French, K. R. (1993). Common Risk Factors in the Returns on Stocks and Bonds. *Journal of Financial Economics*, *33*, 3-56. [https://doi.org/10.1016/0304-405x\(93\)90023-5](https://doi.org/10.1016/0304-405x(93)90023-5)
- Feng, G., Polson, N. G., & Xu, J. (2019). Deep Learning in Characteristics-Sorted Factor Models. *Review of Financial Studies*, *32*, 3676-3713.
- Ferson, W. E., & Schadt, R. W. (1996). Measuring Fund Strategy and Performance in Changing Economic Conditions. *The Journal of Finance*, *51*, 425-461. <https://doi.org/10.1111/j.1540-6261.1996.tb02690.x>
- Gu, S., Kelly, B., & Xiu, D. (2020). Empirical Asset Pricing via Machine Learning. *The Review of Financial Studies*, *33*, 2223-2273. <https://doi.org/10.1093/rfs/hhaa009>
- Gygax, F., Rösch, D., & Schmid, M. (2021). Disentangling Risk Premia Using Deep Generative Models. *Finance Research Letters*, *39*, Article ID: 101618.
- Kim, H., & Mnih, A. (2018). Disentangling by Factorising. *International Conference on Machine Learning (ICML)*, *80*, 2649-2658. <https://proceedings.mlr.press/v80/kim18b.html>
- Kingma, D. P., & Ba, J. (2015). *Adam: A Method for Stochastic Optimization*. arXiv: 1412.6980. <https://arxiv.org/abs/1412.6980>
- Kozak, S., Nagel, S., & Santosh, S. (2020). Shrinking the Cross-Section. *Journal of Financial Economics*, *135*, 271-292. <https://doi.org/10.1016/j.jfineco.2019.06.008>
- Nakagawa, M. (2019). A Deep Recurrent Factor Model for Time-Varying Risk Exposures. *Journal of Financial Econometrics*, *17*, 607-636.
- Nguyen, D., Kosenko, I., & Truong, T. T. (2022). Explainable AI in Asset Pricing: A Disentangled Representation Approach. *Quantitative Finance*, *22*, 1-22.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J. et al. (2019). PyTorch: An Imperative Style, High-Performance Deep Learning Library. *Advances in Neural Information Processing Systems (NeurIPS)*, *32*, 8024-8035. https://papers.nips.cc/paper_files/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-

[Paper.pdf](#)

Petkova, R., & Zhang, L. (2005). Is Value Riskier than Growth? *Journal of Financial Economics*, 78, 187-202. <https://doi.org/10.1016/j.jfineco.2004.12.001>

Sirignano, J., & Cont, R. (2019). Universal Features of Price Formation in Financial Markets: Perspectives from Deep Learning. *Quantitative Finance*, 19, 1449-1459. <https://doi.org/10.1080/14697688.2019.1622295>