

The Evolution of Asset Pricing: How Advanced Methodologies Addressed Static Factor Models Issues

Matteo Zerbin

19 September 2026

Word count: 5972

Project submitted in partial fulfilment of the MSc in Financial Economics

Department of Economics, City St George's University, London

Abstract

Asset pricing models provide expected return and cost of capital estimates, unveil risk-return trade-offs, and offer economic insights that enable market agents to make informed investment choices and proper portfolio allocations. Traditionally, static factor models have been the workhorse for these applications due to their minimal complexity and straightforward implementation. Despite their popularity, static factor models present several assumptions and mathematical limitations, which are not based on economic theory and are derived from OLS regressions, that hinder the empirical performance and explanatory power of those approaches. Such restrictions include the staticity and linearity of parameters, inability to handle and discern among large volumes of data, and fixed functional form. For these reasons, this literature review aims at explaining why it has been necessary and what has been the evolution of asset pricing. Various modelling choices are compared based on their empirical performance and interpretability, highlighting the benefits and disadvantages of each framework. Among time-varying models, Kalman filter methodologies provide the strongest empirical attainments, while IPCA and BMA prove to be the most effective among approaches dealing with the "factor zoo" and the curse of dimensionality. Feedforward neural networks offer the best out-of-sample prediction when accounting for non-linearities and complex variable interactions, while random forests result resilient to noisy datasets. Lastly, this review concisely examines the latest methodological advancements, which are able to tackle multiple issues of static factor models at once. Novel architectures feature and combine advanced neural networks to extract the maximum amount of information from the dataset, achieving strong empirical forecasting performances, while conceding economic interpretability and computational complexity to simpler methodologies.

Table of Contents

Abstract 2
Introduction 4
Chapter 1: Time-varying Models 5
1.1 Conditional Factor Models 5
1.2 Conditional Volatility Models 7
1.3 Latent Space Models 8
1.4 Performance Comparison 10
Chapter 2: "Factor Zoo" and Dimensionality Curse 12
2.1 Factor Selection and Shrinkage 12
2.2 Bayesian Model Averaging 13
2.3 Dimensionality Reduction 14
2.4 Performance and Interpretability Discussion 15
Chapter 3: Non-linear Models 17
3.1 Tree Ensemble Models 18
3.2 Feedforward Neural Networks 19
3.3 Performance and Interpretability Discussion 20
Chapter 4: Advanced Neural Network Architectures 22
Conclusions 24
Reference List 26

Introduction

Expected returns and risk, implicitly defined as uncertainty, affect the investment choices of individuals and institutional investors (Lintner, 1965). Asset pricing models are devices that link the two and provide to economic agents the tools to estimate the cost of capital, assess investment risks, and to create efficient portfolio allocations. For decades, static factor models have been the backbone of traditional asset pricing. Evolved from Sharpe's and Linter's CAPM (Sharpe, 1964; Lintner, 1965), static multifactor models, such as the Fama-French three-factor model (Fama and French, 1993), have provided economists with expected return forecasts and economic explanations. Despite their popularity, it soon became evident that the assumptions behind those models, which are mostly imposed by the ordinary least squares (OLS) regression, were not economically justifiable. Consequently, empirical applications proved to be not entirely satisfactory due to security mispricing and factor anomalies. For these reasons, practitioners and academics started to devise new methodologies addressing those issues.

This literature survey illustrates why it was necessary and what has been the methodological evolution of asset pricing models. This paper is divided into four chapters, each reviewing a branch of the literature handling specific issues of static factor models. Chapter 1 addresses the time-varying nature of factor loadings and risk premiums, Chapter 2 the "factor zoo" and dimensionality curse, Chapter 3 the fixed functional form and linearity assumption, and Chapter 4 briefly discusses the emerging architectures that handle several issues at once. Besides the theoretical improvements and drawbacks brought by new frameworks, each methodology is evaluated and compared based on its empirical performances and interpretability.

Chapter 1: Time-varying Models

Since static factor models usually rely on the Fama and MacBeth (1973) approach, which comprises multiple time series and cross-sectional OLS regressions on a fixed window of data, their estimated risk premiums and factor loadings are assumed to be constant over time. However, it is widely accepted that expected returns are time-varying and that they are affected by predictable changes in risk premiums due to the influence of macroeconomic variables (Keim and Stambaugh, 1986). Specifically, Ferson and Harvey (1991) showed that the major contributor of excess return variability is the time variation of factors' risk premia where, again, most of the predictability came from economic variables. Additionally, assets' factor loadings also change, usually reverting to one over time (Blume, 1975), as the company transforms and the business cycle advances. In fact, Jagannathan and Wang (1996) argue about the importance of dynamic betas, especially in case of recessions, as companies become riskier investments due to the increase in leverage. For these reasons, this chapter focuses on methodologies that address the time-varying nature of the betas and risk premiums to provide more accurate estimates.

1.1 Conditional Factor Models

Conditional factor models serve as a natural extension of static factor models, as they maintain most of their structure while allowing risk premia and factor loadings to vary. Under of conditional expectations (Jagannathan and Wang, 1996), factor loadings and risk premiums are conditioned with an information set that is supposed to contain the data that investors use to price securities (Ferson and Harvey, 1999).

Conditional models greatly improve the performance of static frameworks by explaining market anomalies and by separating non-diversifiable components that do not affect risk, such as sample selection biases (Lettau and Ludvigson, 2001). By allowing the CAPM to consider the business cycle and improving the market proxy selection by including human capital, Jagannathan and Wang (1996) found

the conditional CAPM to explain 55% of the variance, surpassing the CAPM performance that was merely at 1%. Despite the improvement, the conditional CAPM's ability to explain market anomalies is debated. Ferson and Harvey (1999) found the size factor of the competing Fama-French three factors model to be unable to explain firm's returns, while Avramov and Chordia (2006) state that the conditional CAPM is not able to explain either of the size, value, turnover, and past return effects.

As well as the CAPM, static multifactor models experience substantial improvement when conditioned with economic information. For example, Ferson and Harvey (1999) rejected the static three Fama-French model (FF3 hereafter) (Fama and French, 1993) and proved the performance improvements of the conditional version. Another benefit includes the detection of time-varying alphas, which vary consistently with the business cycle (Avramov and Chordia, 2006). In fact, Ferson and Schadt (1996) discovered that the negative alphas detected by static multifactor models were arising from model miss-specification, rather than manager inability.

The main issue with conditional models concerns the choice of the information set. In fact, such collection is unobservable, could be different for each investor, might cover only partially the real information set, and can be subject to estimation error (Ferson and Schadt, 1996). Lewellen and Nagel (2006) solve this issue by avoiding completely the use of proxies and by relying on short-term rolling window estimations with high-frequency data. They found betas and risk premium to co-vary less than anticipated, therefore cautioning on the use of conditional modelling with proxy instruments and directly opposing the findings in Jagannathan and Wang (1996) and Lettau and Ludvigson (2001). Building upon this, Ang and Kristensen (2012) proposed a non-parametric kernel smoothing methodology on short-term windows of data to eliminate the need for the conditioning of betas and risk premiums.

1.2 Conditional Volatility Models

Generalized Autoregressive Conditional Heteroskedasticity (GARCH) models are generally applied to forecast short-term asset volatility (Engle, 2002). Nevertheless, extensions such as the GARCH-in-mean (GARCH-M) and GARCH-DCC are applied for the forecasting of expected returns. GARCH methodologies can provide a direct link between time-varying non-diversifiable risk and expected returns because they are able to estimate a conditional version of the asset covariance matrix. In fact, expected returns (and risk premia) increase as unexpected volatility shocks take place (French, Schwert, and Stambaugh, 1987). Additionally, Bollerslev, Engle, and Wooldridge (1988) found the conditional covariance matrix to have improved explanatory power with respect to its static counterpart and that conditional covariances are highly volatile over time, which are an important determinant of time-varying risk premia.

GARCH-M models can be used in multivariate settings, therefore enabling the discovery of the factors that cause the time variability of expected returns (Dunne, 1999), but they also experience several limitations. Firstly, GARCH-M sustain only a limited amount of assets because the number of correlations leads to multicollinearity (Dunne, 1999), resulting in most studies to have no more than five assets (Engle, 2002). Secondly, results are dependent on portfolio sorting choices (Ng, 1991), and lastly, their estimation is more sensitive to model misspecification and measurement noise than other frameworks (French, Schwert, and Stambaugh, 1987).

GARCH-DCC provides a two-step estimation framework that models large time-varying covariance matrix, therefore enabling static models to be extended properly. Bali and Engle (2010) applied this framework to Merton's Intertemporal CAPM (Merton, 1973) and found it to be robust across market anomalies by detecting the value, investment-to-asset, and ROI factors, and estimating abnormal returns to be zero. An extension of GARCH-DCC models is DCB (dynamic conditional betas), which directly compute factor betas by exploiting the OLS

mathematical beta formulation and the conditional covariances estimated by DCC. DCB has two main advantages. It directly converts static factor models to conditionals and promptly detects beta fluctuations. For example, Engle (2016) studied the FF3 on a short-term timeframe, finding the value beta to be the most volatile, and assessed the risk profiles of European financial institutions during the 2008 and 2011 financial crisis.

DCC and DCB methodologies have two main limitations. Firstly, there is the risk of having biased estimates, as the conditional covariance matrix is sensible to misspecification due to the requirement of strict exogeneity of the independent variables (Engle, 2016). Secondly, DCC parameters might not be computable at all, due to the lack of convergence of the parameters in the maximum likelihood estimation process (Engle, 2016).

1.3 Latent Space Models

Latent space models address the time-variability of factor loadings and risk premiums by treating them as latent unobservable variables that drift and leap over time. Markov-switching models allow the capturing of structural brakes that otherwise would be averaged out in a rolling window setting (Timmermann, 2000), account for the increase in correlation between international markets during bear markets (Ang and Bekaert, 2002), capture asymmetric conditional return distributions of stocks during different latent states (Perez-Quiros and Timmermann, 2002), and can be applied to existing static multifactor models. Coggi and Manescu (2004) extended the FF3 to a version that allows for two distinct market regimes, one more persistent and the other attributable to financial distress, credit risk increase, and liquidity problems. Expectedly, the time-varying version of the FF3 outperformed the static counterpart in forecasting expected returns, confirming Timmermann's (2000) results on the importance of capturing structural breaks.

Regarding their limitations, Markov-switching models might not be implementable in real-time applications due to meaningful estimation errors, especially when more than two market regimes are considered (Guidolin and Timmermann, 2007). Additionally, within each regime, factor loadings and risk premia are constant similarly to static models.

Conversely, Kalman filters consider betas' variations to be continuous through a Bayesian learning approach, where factor loadings are recursively updated given the new data and past information (Adrian and Franzoni, 2009). This framework is justified by the assumption that economic agents update their expectations based on their previous beliefs and new available information (Cowan and Joutz, 2006). Additionally, conditional betas estimated through the Kalman filter are insensitive to trading noise because they are affected by low-frequency idiosyncratic variations, are subject to mean-reverting behaviours (Adrian and Franzoni, 2009), and use a recursive process that make estimates more efficient (He, Huh, and Lee, 2010). Furthermore, the estimated factors maintain their empirical interpretability, incorporate favourable statistical properties, such as mutual uncorrelation and minimum variance, and include all the historical information present in the information set (He, Huh, and Lee, 2010).

Kalman filter applications have also meaningful limitations. In fact, the model is sensitive to misspecification that can arise from the arbitrary choice of the functional form, beta dynamics, and probability distribution of the disturbance term (Choudhry and Wu, 2008). Additionally, since the framework assumes a continuous trend in the parameters, sudden discrete market shifts remain unaccounted for, resulting in unstable parameters and poor market dynamics approximation (Cowan and Joutz, 2006).

It is evident that the simultaneous application of Kalman filters and Markov-switching frameworks could solve each other problems. This approach was pioneered by Kim (1994), where he combined the Hamilton filter, which is a Markov-switching model, with the Kalman filter. This method is called the Kim

filter (or Markov-switching state space) and is able to detect both abrupt market changes and track continuous quantities within each state (Kim and Halbert, 1999). Despite the theoretical advantages proposed by the framework and its vast usage in other fields, such as market bubble identifications (Al-Anaswah and Wilfling, 2009), little empirical applications have been applied to expected return forecasting. The explanation might lay in the poor out-of-sample predictability performance of the model (Choi, Kim, and Park, 2017) and inability to consider individual assets due to the large number of computations required (Akay, Senyuz, and Yoldas, 2013). Nevertheless, newly developed methodologies are addressing the dimensionality problem. For example, Barigozzi and Massacci (2025) applied principal component analysis (PCA) to a large dimensional factor model, opening a new path for new empirical applications.

1.4 Performance Comparison

This chapter covered modelling solutions aimed at addressing the static assumption of unconditional factor models. The best performers are frameworks applying the Kalman filter, overperforming all other classes. Adrian and Franzoni (2009) measured the root mean square error (RMSE) of the OLS regression to be 1.076, whereas the learning algorithm reduced it to 0.814, both statistically significant at 5%. Additionally, by introducing the learning mechanism in the conditional CAPM, the aggregate pricing error was reduced by 45% and by 65% with respect to the conditional and unconditional versions. Similarly, He, Huh, and Lee (2010) found their dynamic factor pricing model to outperform both the FF3 and conditional FF3 by returning consistently lower and statistically different average mean squared forecasted errors (MSFE), especially out-of-sample. The Kalman filter also exceeds GARCH's performances. Brooks, Faff, and McKenzie (1998) calculated the mean square error (MSE) of the GARCH to be 0.00196 in-sample and 0.00041 out-of-sample, and the filter's MSE to be much lower, 0.00032 and a 0.00007 in and out-of-sample respectively.

Other models have less decisive empirical data on what of these methodologies perform best with respect to each other. Coggi and Manescu (2004) measured a Markov-switching models' average pricing error of 0.25 when applied to FF3 frameworks, against an average pricing error of 0.24 of the traditional FF3. Similarly, GARCH models and their application as DCBs perform comparably to the FF3 because they exhibit analogous unexplained returns (alphas), even though they are superior in reacting to regime shocks when compared to OLS estimations (Engle, 2016). Finally, when it comes to traditional conditional models, there are contrasting results. The conditional CAPM of Jagannathan and Wang (1996) improves the variance explained of the static CAPM from 1.35% to 55.21% but He, Huh, and Lee (2010) find the conditional FF3 to perform as good as the static FF3.

In conclusion, the Kalman filter seems to be the most reliable choice in empirical settings and there is no clear hierarchy between the other frameworks, prompting their use accordingly to specific theoretical necessities.

Chapter 2: “Factor Zoo” and Dimensionality Curse

The number of newly discovered factors, which currently amounts to more than 300 (Harvey, Liu, and Zhu, 2016), creates an issue named “factor zoo”, which questions whether all discovered factors add real explanatory power in pricing securities (Cochrane, 2011), as most of them are a scaled version of price (Fama and French, 1992). In fact, it is difficult to distinguish between real determinants of asset returns and temporary market anomalies, as different factor combinations, constructions, and omissions might change the relevance of factors themselves (Feng, Giglio, and Xiu, 2020). For example, Fama and French five-factor model (FF5 hereafter) found the value factor (HML), previously deemed in Fama and French (1993) as a true price determinant, to lose its explanatory power (Fama and French, 2015). For this reason, it is plausible that most of the discovered factors and research findings over this piece of literature are incorrect (Harvey, Liu, and Zhu, 2016). The factor zoo induces another problem, the curse of dimensionality. The dimensionality problem is a mathematical issue that arises when the elevated number of predictors makes the model’s estimates to be meaningless or computationally unfeasible (Feng, Giglio, and Xiu, 2020). Consequently, this chapter focuses on architectures addressing both obstacles.

2.1 Factor Selection and Shrinkage

Ordinary least squares regressions are doomed to fail when the number of regressors approaches the number of observations and are proven to be unstable when some independent variables are highly correlated among themselves (Gu, Kelly, and Xiu, 2020). Penalized regressions, such as LASSO, Ridge, and elastic net, solve this problem by selecting and/or shrinking the coefficients of variables that have low explanatory power. LASSO is quite popular because it offers a straightforward interpretation of the estimated coefficients and provides better out-sample-performance than standard OLS regressions (Messmer et al., 2022). Despite its popularity, the recent academic literature suggests using caution. In fact, by

aggressively setting parameters to zero, LASSO is subject to omitted variable bias, therefore hindering the interpretability of the estimated coefficients and the identification of the true price drivers (Feng, Giglio, and Xiu, 2020). Furthermore, LASSO assumes the regressors' vector to be sparse, meaning that only a small subset of the characteristics considered are priced. Kozak, Nagel, and Santhosh (2020) consider this assumption "futile" because the cross-section of returns is so complex and the redundance among predictors so low that models with only few predictors cannot perform well out-of-sample. For this reasons, Kozak, Nagel, and Santhosh (2020) advocate for Ridge, which shrinks highly correlated parameters without dropping them. Alternatively, Gu, Kelly, and Xiu (2020) propose elastic net frameworks that select and shrink loadings simultaneously. Interestingly, the parameters selection ability of elastic nets results to be superior to other penalized regression methodologies, as they tend to choose less parameters without sacrificing predictive performance (Gu, Kelly, and Xiu, 2020).

2.2 Bayesian Model Averaging

Bayesian model averaging (BMA) is a mathematical framework that incorporates specification uncertainty by simultaneously considering the entire possible space of models given a set of factors (Avramov, 2002). BMA assigns a weight to each model based on its posterior probability, which favours high explanatory power, to create a composite weighted model. Since this methodology considers all highly informative models, it is possible to robustly gauge the probability of a factor to be part of the pricing kernel (Bryzgalova, Huang, and Julliard, 2023). Furthermore, advanced BMA methodologies offer protection from factor proliferation and spurious correlations by accounting for the risk premia of factor themselves. By empirically implementing this methodology, Avramov et al. (2023) were able to confirm that unconditional models perform poorly and that characteristics such as post-earnings announcement drift, quality minus junk, and intermediary capital might be important pricing factors.

BMA's limitation is its inability to address the dimensionality problem of the factor zoo directly, as it structurally needs to compute all possible model specifications (Harvey, Liu, and Zhu, 2016). Despite the possibility of circumventing this issue by applying spike-and-slab priors (Bryzgalova, Huang, and Julliard, 2023), priors selection must not be solely dictated by computational hurdles because it heavily affects posterior probability distributions. In fact, the weights assigned to each model are quite sensitive to the prior assumptions, an effect that does not fade away as the sample size increases (Steel, 2020).

2.3 Dimensionality Reduction

Principal component analysis (PCA) is a statistical methodology that aims to capture the covariance matrix of factors with an approximated lower dimensional object (Chamberlain and Rothschild, 1983). This methodology is justified in asset pricing by the asymptotic equivalence of factor analysis when a strict factor structure is assumed (Chamberlain and Rothschild, 1983). The main advantage of using PCA is the creation of few highly explicative statistical factors that can be used in conjunction to other modelling techniques to solve the dimensionality problem (Barigozzi and Massacci, 2025).

PCA presents several limitations. Firstly, PCA does not account for firm-specific characteristics changes, therefore making it an unconditional framework (Lettau, 2024). Secondly, by being purely statistical, its estimation needs large datasets (Connor and Korajczyk, 1986) and compromises its economic interpretability (Giglio, Kelly, and Xiu, 2022). Thirdly, PCA is prone to misspecification due to omitted risk factors and incorrect assumptions on conditional expected returns' statistical distributions (Zaffaroni, 2025). Lastly, PCA is an unsupervised method that computes each principal component so that it explains the maximum amount of variance among the predictors, leading factors with low variance but high importance for price determination to be unaccounted for (Lettau and Pelger, 2020).

Partial least squares (PLS) addresses PCA's tendence of seeking unexplanatory variances by acknowledging the covariance of predictors with the forecasted target, therefore reducing significantly the amount of noise captured (Kelly and Pruitt, 2013). In all fairness, PCA can occasionally outperform PLS out-of-sample (Gu, Kelly, Xiu, 2020), but the variance problem makes it unreliable. On the other hand, PLS's reliability led Huang et al. (2015) to create an investor's sentiment index able to satisfactorily predict expected returns on portfolios sorted by industry, size, value, and momentum, and to infer on investor's beliefs.

Despite these improvements, PCA and PLS assume factor loadings to be constant over time and do not provide almost any economic interpretation on return drivers. Instrumented principal component analysis (IPCA) addresses both issues by considering firm characteristics as instruments that partially influence factor loadings (Kelly, Pruitt, and Su, 2019). By doing so, factor loadings become conditional and, since IPCA tracks how stock characteristics influence each beta, the estimated latent factors can be economically interpreted. For example, Kelly, Pruitt, and Su (2019) found their first factor to be comparable to the value factor of the FF3 model, as it was mainly driven by the book to market ratio and market capitalization, and Kelly, Moskowitz and Pruitt (2021) were able to study the stock momentum and long-term reversal characteristics and their ability to predict future betas. When it comes to performance, IPCA greatly outperforms static models and PCA, while remaining parsimonious (Kelly, Pruitt, and Su, 2019). A limitation of IPCA is its reliance to the quality of the information disclosed, in fact, in developing markets IPCA's advantage persists but shrinks (Wang and Zhang, 2025).

2.4 Performance and Interpretability Discussion

This chapter discussed the advantages and drawbacks of the methodologies developed to solve either the factor zoo, the dimensionality problem that comes with it, or both. The models considered can be organized in three performance brackets. The best performing one comprises IPCA and PLS, as both frameworks

provide strong and stable out-of-sample R^2 and forecast performances (Kelly, Pruitt, and Su, 2019; Kelly and Pruitt, 2015). The second bracket features BMA and PCA but with an important caveat. Despite the two methodologies having similar forecasting abilities, their out-of-sample output being strongly correlated, and their mean square forecasted errors being comparable (Ouysse, 2011), BMA results to be way more robust. In fact, especially when considering individual stock returns, PCA can perform very poorly out-of-sample by yielding negative R^2 (Kelly, Pruitt, and Su, 2019). The last grouping includes penalized regressions. Elastic nets are generally better performing than LASSO because they do not overly exclude necessary parameters (Kozak, Nagel, and Santhosh, 2020) but even with respect to PCA, their out-of-sample R^2 is sensibly lower (Gu, Kelly, and Xiu, 2020). Importantly, these methodologies overperform standard OLS regressions regardless of their proposed ordering, therefore validly solving the dimensionality problem.

When it comes to interpretability, BMA is a strong choice. By computing the probability of each factor being present in the SDF at multiple prior Sharpe ratio multiples, each factor can be assessed more thoroughly. For example, Avramov et al. (2023) found the market factor, post-earnings announcement drift, and quality minus junk to have a probability of 1 (100%) at all prior Sharpe ratios multiples, and Bryzgalova, Huang, and Julliard (2023) were able to confirm the proposition hypothesized by Kozak, Nagel, and Santhosh (2020) of a dense factor space. Penalized regressions are also highly interpretable and, as previously discussed, elastic nets should be preferred over LASSO to avoid factor misinterpretation. Lastly, by tracking how observable characteristics affect each factor loading, IPCA functions as a bridge between economical understanding and statistical black box algorithms like PCA and PLS.

In conclusion, BMA and IPCA are the strongest choices among the models examined. BMA prevalently addresses the factor zoo while maintaining strong predictive performance, whereas IPCA shines in out-of-samples applications and provides meaningful economic interpretation.

Chapter 3: Non-linear Models

Static multifactor models have a fixed functional form because they use ordinary least squares regressions (OLS) to estimate their parameters (Fama and MacBeth, 1973). This modelling choice imposes linearity in the parameters and supposes that the contribution of each term is additive, therefore overlooking complex non-linear relationships (Freyberger, Neuhierl, and Weber, 2020). Fama and French (2008) empirically reject both assumptions. In fact, they confirmed the non-linear relationships on the size factor by detecting its relevance mostly on microcaps and identified conditional interactions among momentum and asset growth with size itself, therefore rejecting the additivity assumption. Additionally, as a consequence of linearity, static multifactor models are also unable to capture asymmetric risks (Harvey and Siddique, 2000).

To improve the performance of the pricing kernel, early solutions employed polynomials of higher order to capture coskwedness (Harvey and Siddique, 2000) and cokurtosis (Dittmar, 2002). For example, Borri et al., (2026) found the addition of squared factors to common static multifactor models to improve out-of-sample explanatory power and Chung, Johnson, and Schill (2006) were able to explain the size and value factors with high order comoments. Nevertheless, since OLS remains the framework used, these corrections do not directly address the fixed functional form problem and create other issues. The dimensionality problem of static multifactor models worsens (Martellini and Ziemann, 2010), estimation errors might be amplified (Boudt, Lu, and Peeters, 2015), and factors to the order as high as the tenth are needed to make traditional factors lose their explanatory power (Chung, Johnson, and Schill, 2006). For these reasons, the modern literature prevalently leverages advanced machine learning methodologies to devise true non-linear models that overcome the aforementioned hurdles.

3.1 Tree Ensemble Models

Tree ensemble models are a natural candidate to solve the fixed functional form problem of static multifactor models because they are able to capture highly complex non-linear relationships (James et al., 2021). Unlike traditional methodologies, trees are fully nonparametric and incorporate multiway predictor interactions (Gu, Kelly, and Xiu, 2020). Asset pricing applications from the tree ensemble family stem from the use of either bagging or boosting frameworks. Both methodologies are regularization techniques that rely on the aggregation of numerous weak learners (James et al., 2021) because single trees are prone to overfitting (Gu, Kelly, and Xiu, 2020).

Random forests (RF) for expected return estimation, which are a special application of bagging, are a collection of regression trees where each tree uses a subsample of the training data and a random selection of predictors (James et al., 2021). This setting decorrelates the trees in the forest, making the whole model resilient against noisy data and robust to overfitting (Hastie et al., 2009). In particular, increasing the number of trees does not overfit the model and improves the estimation by decreasing its variance. Furthermore, random forests require little tuning of their hyperparameters (Rasekhschaffe and Jones, 2019), therefore making the result less dependent from researcher choices. Lastly, since the characteristics with the highest predictive power are the ones that are selected to split the branches, the random selection feature of predictors for each potential node prevents factor dominance of the strongest predictors, as it allows other tree configurations and therefore more pattern discoveries (Gu, Kelly, and Xiu, 2020). A specific limitation of random forests is the dilution of weak predictors. It occurs because trees become deep without accounting for previous nodes' mistakes, therefore making splitting choices on noisy characteristics that do not add explanatory power, rather than on important factors with small variance (Coqueret and Guida, 2020).

Gradient boosting ensembles (GB) address this issue by creating one strong learner that sequentially connects multiple shallow trees, where each tree's goal is to decrease the prediction error of the previous one (Gu, Kelly, and Xiu, 2020). This

approach forces the model to split on factors that are the most explicative at each specific node regardless of their absolute strength (Rossi, 2018), therefore enabling the next tree to discover patterns that are difficult to detect. Accordingly, Gu, Kelly, and Xiu (2020) found gradient boosting to provides better estimations than random forests, respectively scoring an out-of-sample R^2 of 0.34% against a 0.33%. On the other hand, gradient boosting requires finer parameter tuning than RFs (Hastie et al., 2009), which consist in controlling the shrinkage and depth of each tree (Gu, Kelly, and Xiu, 2020). Additionally, differently from random forests, boosting algorithms can be sensitive to signal noise (Oza, 2004).

Finally, a general limitation of tree ensemble models is their inability to extrapolate forecasts for data that exceed the boundaries of the training sample, essentially creating bounded outputs from extreme inputs (James et al. 2021).

3.2 Feedforward Neural Networks

Neural networks are machine learning methodologies formed by multiple layers of "neurons" that are activated according to a specific non-linear activation function (James et al., 2021). Despite feedforward neural networks (FNN) being the simplest framework of the family, as the information flows strictly in one direction without any loops or memory, their application is justified as they "are capable of approximating any measurable function to any desired degree of accuracy" (Hornik, Stinchombe, and White, 1989, p. 363). Their versatility stems from the hidden layers of neurons that can capture highly non-linear relationships and interactions while performing a task of data reduction (Feng, He, and Polson, 2018). This dimensionality reduction enables the model to estimate few highly explicative latent factors from lots of firm characteristics and economic data without the need for theoretical assumptions on how they interact with each other. Consequently, FNNs empirically outperform out-of-sample OLS and penalized regression models by almost 50% on out-of-sample R^2 terms, as they do not present a fixed functional form (Gu, Kelly, and Xiu, 2020). Despite deeper networks being generally able to

capture more complex patterns than shallower versions (Eldan and Shamir, 2016), Gu, Kelly, and Xiu (2020) found shallower models to perform better. An explanation might arise from the low signal-to-noise ratio of financial data, which could have made deeper models to fit "false" patterns (Hornik, Stinchombe, and White, 1989). Another reason for the efficacy of FNNs is the synchronous computation of the parameters (Feng, He, and Polson, 2018), which promotes system error minimization. There are multiple costs for their heightened predictive accuracy and flexibility. Firstly, FNNs require ad-hoc optimization algorithms, as brute force methodologies are computationally demanding and do not ensure convergence (Gu, Kelly, and Xiu, 2020). Secondly, model interpretation is limited (Gu, Kelly, and Xiu, 2020) and the patterns discovered cannot infer causality (Feng, He, and Polson, 2018). Lastly, if applied to trading strategies, portfolio solutions tend to be unstable, therefore generating high turnovers (Avramov, Chen, and Metzker, 2023).

3.3 Performance and Interpretability Discussion

The technological leap from OLS to non-linear machine learning methods is considerable. In fact, solely due to overfitting, OLS yields negative out-of-sample R^2 (-3.46%), while both tree ensembles (0.33%-0.34%) and feedforward neural networks (0.39%-0.40%) dominate linear factor selection (0.11%) and dimensionality reduction devices (0.26%-0.27%) covered in Chapter 2 (Gu, Kelly, and Xiu, 2020). In general, machine learning's accuracy is attributed to its ability to discover non-linear relationships and to its robustness to multicollinearity in the predictors (Rasekhschaffe and Jones, 1999). FNNs tend to overperform tree ensembles due to their global parameter optimization procedures (Gu, Kelly, and Xiu, 2020) and ability to capture tail behaviour (Feng, He, and Polson, 2018), enabling deep learning architectures to predict more accurately market crashes than RF and GB (Feng, He, and Polson, 2018). On the other hand, neural networks require careful tuning and are computationally demanding (Krauss, Do, and Huck, 2017). When it comes to tree ensembles, gradient boosting algorithms tend to

outperform random forests (Gu, Kelly, and Xiu, 2020). Exception is made when the data is highly noisy. In such circumstances, the resilience to overfitting and to noisy signals makes random forest preferrable to GB and even to deep learning if not adequately tuned (Krauss, Do, and Huck, 2017). A common drawback of machine learning is the black box problem, where the researcher in not able to interpret the patterns discovered by the model. Despite the existence of methodologies to better understand the processes inside machine learning models (Moritz and Zimmermann, 2016; Bracke et al. 2019), OLS remains better at providing clear economic explanations.

Chapter 4: Advanced Neural Network Architectures

Each of the previous chapters have been discoursing on how different issues of static factor models are addressed by newer frameworks. This chapter briefly reviews the frontier of the most modern architectures, which entail various types of advanced neural network configurations that are able to address multiple modelling limitations at once.

The family of autoencoders is a deep learning methodology that natively aims to perform dimensionality reduction tasks. They are constituted by three parts: the encoder, which progressively compresses the information, the bottleneck, which distils the latent factors, and the decoder, which reconstructs the input (Gu, Kelly, and Xiu, 2021). Similarly to PCA, autoencoders retrieve the statistical information from the bottleneck, therefore ensuring a parsimonious representation of the dataset, but it also allows for non-linear relationships and interactions in its mapping (Gu, Kelly, and Xiu, 2021). Another advantage comes from conditional autoencoders (CAE) and conditional variational autoencoders (CVAE) where, as the names suggest, they enable the conditioning of factor loadings and latent factors with firm characteristics (Nechvátalová, 2025; Yang et al., 2024). Additionally, CVAE consider the temporal dependency of factor loadings with previous values (Chen, Wang, and Huang, 2026), therefore allowing the parameters to be dynamic. Autoencoders are not just pure statistical machines but can also provide economic explanations and operate under real-world constraints. In fact, Chen et al. (2026) found their implied pricing kernel to be in line with the implications of the intertemporal CAPM and Gu, Kelly, Xiu (2021) were able to enforce the no arbitrage condition. When it comes to predictive performance, autoencoders greatly outperform other methodologies. Gu, Kelly, Xiu (2021) saw the conditional autoencoder with three hidden layers to outperform IPCA with a predictive out-of-sample R^2 of 0.57% against a 0.30%, while both FF3 and PCA did not provide any forecasting capacity. Unfortunately, performances drop significantly when

economic constraints, such as liquidity and transaction costs, are considered. Nechvátalová (2025) measured a drop ranging from 60% to 85% in portfolio profitability when economic constraints were implemented. Regardless, liquid strategies continued to beat the benchmark. Lastly, autoencoders cannot often be perfectly optimized due to their nonlinearities, nonconvexity, and high parametrization, therefore requiring approximation algorithms to make tuning computationally efficient (Gu, Kelly, and Xiu, 2021).

Other modelling choices combine multiple neural network architectures. For example, Chen, Pelger, and Zhu (2019) implemented an FNN to estimate the functional form of the stochastic discount factor, a Long-Short Term-Memory (LSTM) network to allow time-varying parameters, and a generative adversarial network (GAN) to identify the portfolios and states with the most inexplicable pricing information. Their model addresses the problematics found in penalized regressions, as it detects the factors with high-risk premia but low variability, and improve the economic space selection of Markov-switching models. The model provides a Sharpe ratio three times larger than the FF5 and, similarly to other machine learning methodologies, those performances come at the cost of interpretability, computational hurdles, and difficulty in understanding its statistical properties (Chen, Pelger, and Zhu, 2019; Fallahgoul, Franstianto, and Lin, 2024).

Conclusions

This literature survey covered the evolution of asset pricing models. By grouping together methodologies devised to address common issues concerning traditional static factor models, this paper highlights the necessity for more advanced frameworks and the respective outcomes of those applications. Chapter 1 addressed the time variability of parameters, proposing Kalman Filters as the best alternative, despite being unable to capture abrupt regime shifts. Chapter 2 discussed the factor zoo and dimensionality curse, presenting IPCA and BMA as preferred methodologies due to their strong performance and interpretability. Chapter 3 tackled non-linear relationships and the fixed functional form problem, highlighting how the strong return forecasting abilities of feedforward neural networks require fine tuning, are sensitive to market noise, and lack economic interpretability. Chapter 4 illustrated how asset pricing is evolving towards advanced neural networks such as autoencoders, LSTMs, and GANs. They provide strong theoretical out-of-sample predictions because they enable feature extraction, conditioning with firm-specific characteristics and macroeconomic variables, time variability, and non-linearities at once. Nevertheless, such models do not provide substantial information on return drivers, which is vital for sound economic decision making, are computationally intensive, and, when real-world constraints are applied, their effectiveness drops significantly. For these reasons, the application of simpler methodologies remains worthwhile. Furthermore, Chen, Pelger, and Zhu (2019) found that most firm characteristics affect linearly the stochastic discount factor, partially rehabilitating linear models. However, the empirical performances of static factor models might depend on the involuntary capturing of latent features (He, Huh, and Lee, 2010), making assessments of performance and factor determination unreliable.

This literature survey provides a comprehensive, cohesive, and accessible review on the vast academic landscape of asset pricing models, highlighting strengths, flaws, and trade-offs among comparable models. It also analysed conflicting pieces of literature, such as the validity of conditional models using a conditioning

information set, and unveiled gaps in the academic research, like the lack of implementation of the Kim filter. The limitations of this work include the restricted performance comparisons between each chapter's frameworks and the reliance of empirical results on the specific data used, model implementation, and parameter tuning.

Reference List

Adrian, T., Franzoni, F. (2009) 'Learning about beta: Time-varying factor loadings, expected returns, and the conditional CAPM', *Journal of Empirical Finance*, 16(4), pp. 537-556. Available at: https://doi.org/10.1016/j.jempfin.2009.02.003

Akay, O., Senyuz, Z., Yoldas, E. (2013) 'Hedge fund contagion and risk-adjusted returns: A Markov-switching dynamic factor approach', *Journal of Empirical Finance*, 22, pp. 16-29. Available at: https://doi.org/10.1016/j.jempfin.2013.02.005

Al-Anaswah, N., Wilfling, B. (2009) *Identification of speculative bubbles using state-space models with Markov-switching*, CQE Working Papers 0309, Center for Quantitative Economics (CQE), University of Muenster. Available at: https://www.wiwi.uni-muenster.de/oeew/sites/oeew/files/CQE_paper/cqe_wp_3_2009.pdf (Accessed on: 13/07/2026)

Ang, A., Bekaert, G. (2002) 'International Asset Allocation with Regime Shifts', *The Review of Financial Studies*, 15(4), pp. 1137-1187. Available at: https://www.jstor.org/stable/1262694 (Accessed on: 05/07/2026)

Ang, A., Kristensen, D. (2012) 'Testing conditional factor models', *Journal of Financial Economics*, 106, pp. 132-156. Available at: http://dx.doi.org/10.1016/j.jfineco.2012.04.008

Avramov, D. (2002) 'Stock return predictability and model uncertainty', *Journal of Financial Economics*, 64(3), pp. 423-458. Available at: https://doi.org/10.1016/S0304-405X(02)00131-9

Avramov, D., Cheng, S., Metzker, L., (2023) 'Machine learning vs. Economic restrictions: Evidence from stock return predictability', *Management Science*, 69(5), pp. 2587–2619. Available at: http://dx.doi.org/10.2139/ssrn.3450322

Avramov, D., Cheng, S., Metzker, L., Voigt, S. (2023) 'Integrating Factor Models', *The Journal of Finance*, 78(3), pp. 1593-1646. Available at: https://doi.org/10.1111/jofi.13226

Avramov, D., Chordia, T. (2006) 'Asset Pricing Models and Financial Market Anomalies', *The Review of Financial Studies*, 19(3), pp. 1001-1040. Available at: https://doi.org/10.1093/rfs/hhj025

Bali, T. G., Engle, R. F. (2010) 'The intertemporal capital asset pricing model with dynamic conditional correlations', *Journal of Monetary Economics*, 57, pp. 377–390. Available at: https://doi.org/10.1016/j.jmoneco.2010.03.002

Barigozzi, M., Massacci, D. (2025) 'Modelling large dimensional datasets with Markov switching factor models', *Journal of Econometrics*, 247, Article 105919. Available at: https://doi.org/10.1016/j.jeconom.2024.105919

Blume, M. E. (1975) 'Betas and Their Regression Tendencies', *The Journal of Finance*, 30(3), pp. 785-795. Available at: https://www.jstor.org/stable/2326858 (Accessed: 03/06/2026)

Bollerslev, T., Engle, R. B., Wooldridge, J. M. (1988) 'A Capital Asset Pricing Model with Time-Varying Covariances', *Journal of Political Economy,* 96(1), pp. 116-131. Available at: https://www.jstor.org/stable/1830713 (Accessed: 08/07/2026)

Borri, N., Chetverikov, D., Liu, Y., Tsyvinski, A. (2026) 'Higher-Order Asset Pricing Factors via Forward Selection Fama-MacBeth Regression', ArXiv. Available at: https://doi.org/10.48550/arXiv.2503.23501

Boudt, K., Lu, W., Peeters, B. (2015) 'Higher order comoments of multifactor models and asset allocation', *Finance Research Letters*, 13, pp. 225-233. Available at: https://doi.org/10.1016/j.frl.2014.12.008

Bracke, P., Datta, A., Jung, C., Sen, S. (2019) 'Machine learning explainability in finance: an application to default risk analysis', *SSRN*. Available at: http://dx.doi.org/10.2139/ssrn.3435104

Brooks, R. D., Faff, R. W., McKenzie, M. D. (1998) 'Time-Varying Beta Risk of Australian Industry Portfolios: A Comparison of Modelling Techniques', *Australian Journal of Management*, 23(1), pp. 1-22. Available at: https://journals.sagepub.com/doi/pdf/10.1177/031289629802300101 (Accessed on: 16/07/2026)

Bryzgalova, S., Huang, J., Julliard, C. (2023) 'Bayesian Solutions for the Factor Zoo: We Just Ran Two Quadrillion Models', *Journal of Finance*, 78 (1). pp. 487-557. Available at: https://doi.org/10.1111/jofi.13197

Chamberlain, G. Rothschild, M. (1983) 'Arbitrage, Factor Structure, and Mean-Variance Analysis on Large Asset Markets', *Econometrica*, 51(5), pp. 1281-1304. Available at: https://www.jstor.org/stable/1912275 (Accessed on: 26/07/2026)

Chen, J., Han, Y., Tang, G., Zhu, Y. (2026) 'Taming the global factor zoo', *Journal of International Money and Finance*, 160, Article 103466. Available at: https://doi.org/10.1016/j.jimonfin.2025.103466

Chen, L., Pelger, M., Zhu, J. (2019) ,Deep Learning in Asset Pricing', *SSRN*. Available at: http://dx.doi.org/10.2139/ssrn.3350138

Chen, X., Wang, J., Huang, R. (2026) 'Variational Autoencoder Asset Pricing models with economic restrictions', *International Review of Economics & Finance*, 109, Article 105402. Available at: https://doi.org/10.1016/j.iref.2026.105402

Choi, K. H., Kim, C.-J., Park, C. (2017) 'Regime Shifts in Price-Dividend Ratios and Expected Stock Returns: A Present-Value Approach', *Journal of Money, Credit and Banking*, 49(2-3), pp. 417-441. Available at: https://doi.org/10.1111/jmcb.12384

Choudhry, T., Wu, H. (2008) 'Forecasting Ability of GARCH vs Kalman Filter Method: Evidence from Daily UK Time-Varying Beta', *Journal of Forecasting*, 27(8), pp. 670-689. Available at: https://doi.org/10.1002/for.1096

Chung, Y. P., Johnson, H., Schill, M. J. (2006) 'Asset Pricing When Returns Are Nonnormal: Fama-French Factors versus Higher-Order Systematic Comoments', *The Journal of Business,* 79(2), pp. 923-940. Available at: https://doi.org/10.1086/499143

Cochrane, J. H. (2011) 'Presidential Address: Discount Rates', *The Journal of Finance*, 66, pp. 1047-1108. Available at: https://doi.org/10.1111/j.1540-6261.2011.01671.x

Coggi, P., Manescu, B. (2004) 'A multifactor model of stock returns with endogenous regime switching', Discussion paper no. 2004-01, *University of St. Gallen: Department of Economics*. Available at: www.fgn.unisg.ch/public/public.htm (Accessed on 11/07/2026)

Connor, G., Korajczyk, R. A. (1986) 'Performance measurement with the arbitrage pricing theory: A new framework for analysis', *Journal of Financial Economics*, 15(3), pp. 373-394. Available at: https://doi.org/10.1016/0304-405X(86)90027-9

Coqueret, G., Guida, T. (2020) 'Training trees on tails with applications to portfolio choice', *Annals of Operations Research*, 288, pp. 181-221. Available at: https://doi.org/10.1007/s10479-020-03539-2

Cowan, A. M., Joutz, F. L. (2006) 'An unobserved component model of asset pricing across financial markets', *International Review of Financial Analysis*, 15(1), pp. 86-107. Available at: https://doi.org/10.1016/j.irfa.2004.10.004

Dittmar, R. F. (2002) 'Nonlinear Pricing Kernels, Kurtosis Preference, and Evidence from the Cross Section of Equity Returns', *The Journal of Finance*, 57(1), pp. 369-403. https://doi.org/10.1111/1540-6261.00425

Dunne, P. G. (1999) 'Size and book-to-market factors in a multivariate GARCH-in-mean asset pricing application', *International Review of Financial Analysis*, 8(1), pp. 35-52. Available at: https://doi.org/10.1016/S1057-5219(99)00004-6

Eldan, R., Shamir, O. (2016) 'The power of depth for feedforward neural networks', in Feldman, V., Rakhlin, A., and Shamir, O. 29th Annual Conference on Learning Theory, Brookline, MA: Microtome Publishing, pp. 907-940. Available at: https://proceedings.mlr.press/v49/eldan16.pdf (Accessed: 11/08/2026)

Engle, R. F. (2002) 'Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models', *Journal of Business & Economic Statistics*, 20(3), pp. 339-350. Available at: https://www.jstor.org/stable/1392121 (Accessed: 19/07/2026)

Engle, R. F. (2016) 'Dynamic Conditional Beta', *Journal of Financial Econometrics*, 14(4), pp. 643–667. Available at: https://doi.org/10.1093/jjfinec/nbw006

Fallahgoul, H., Franstianto, V., Lin, X. (2024) 'Asset pricing with neural networks: Significance tests', *Journal of Econometrics*, 238(1), Article 105574. Available at: https://doi.org/10.1016/j.jeconom.2023.105574

Fama, E. F., French, K. R. (1992) 'The Cross-Section of Expected Stock Returns', *The Journal of Finance,* 47(2), pp. 427-465. Available at: https://doi.org/10.2307/2329112

Fama, E. F., French, K. R. (1993) 'Common Risk Factors in the Returns on Stocks and Bonds.', *Journal of Financial Economics,* 33(1), pp. 3–56. Available at: https://www.bauer.uh.edu/rsusmel/phd/Fama-French_JFE93.pdf (Accessed: 02/06/2026)

Fama, E. F., French, K. R. (2008) 'Dissecting Anomalies', *The Journal of Finance*, 63(4), pp. 1653-1678. Available at: https://www.jstor.org/stable/25094486 (Accessed: 05/08/2026)

Fama, E. F., French, K. R. (2015) 'A five-factor asset pricing mode', *Journal of Financial Economics*, 116, pp. 1-22. Available at: http://dx.doi.org/10.1016/j.jfineco.2014.10.010

Fama, E. F., MacBeth, J. D. (1973) 'Risk, Return, and Equilibrium: Empirical Tests', *Journal of Political Economy*, 81(3), pp. 607–636. Available at: http://www.jstor.org/stable/1831028 (Accessed: 09/07/2026)

Feng, G., Giglio, S., Xiu, D. (2020) 'Taming the Factor Zoo: A Test of New Factors', *The Journal of Finance*, 75(3), pp. 1327-1370. Available at: https://doi.org/10.1111/jofi.12883

Feng, G., He, J., Polson, N.G. (2018) 'Deep Learning for Predicting Asset Returns', ArXiv. Available at: https://arxiv.org/pdf/1804.09314 (Accessed: 11/08/2026)

Ferson, W. E., Harvey, C. R. (1991) 'The Variation of Economic Risk Premiums', *Journal of Political Economy*, 99(2), pp. 385-415. Available at: https://www.jstor.org/stable/2937686 (Accessed: 30/03/2026)

Ferson, W. E., Schadt, R. W. (1996) 'Measuring Fund Strategy and Performance in Changing Economic Conditions', *The Journal of Finance*, 51(2), 425-461. Available at: https://www.jstor.org/stable/2329367 (Accessed: 04/07/2026)

Ferson, W. E., Harvey, C. R. (1999) 'Conditioning Variables and the Cross Section of Stock Returns', *The Journal of Finance*, *54*(4), pp. 1325–1360. Available at: http://www.jstor.org/stable/798007 (Accessed: 04/07/2026)

French, R. K., Schwert, G. W., Stambaugh, R. F. (1987) 'EXPECTED STOCK RETURNS AND VOLATILITY', *Journal of Financial Economics*, 19, pp. 3-29.

Available at: https://finance.martinsewell.com/stylized-facts/volatility/FrenchSchwertStambaugh1987.pdf (Accessed: 07/07/2026)

Freyberger, J., Neuhierl, A., Weber, M. (2020) ‚Dissecting Characteristics Nonparametrically', *The Review of Financial Studies*, *33*(5), pp. 2326–2377. Available at: https://www.jstor.org/stable/48574462 (Accessed: 11/08/2026)

Giglio, S., Kelly, B., Xiu, D. (2022) 'Factor Models, Machine Learning, and Asset Pricing', Annual *Review of Financial Economics*, 14, pp. 337-368. Available at: https://doi.org/10.1146/annurev-financial-101521-104735

Gu, S., Kelly, B. T., Xiu, D. (2020) 'Empirical Asset Pricing via Machine Learning', *The Review of Financial Studies*, 33(5), Pages 2223–2273. Available at: https://doi.org/10.1093/rfs/hhaa009

Gu, S., Kelly, B. T., Xiu, D. (2021) 'Autoencoder asset pricing models', *Journal of Econometrics*, 222(1), pp. 429-450. Available at: https://doi.org/10.1016/j.jeconom.2020.07.009

Guidolin, M., Timmermann, A. (2007) 'Asset allocation under multivariate regime switching', *Journal of Economic Dynamics and Control*, 31(11), pp. 3503-3544. Available at: https://doi.org/10.1016/j.jedc.2006.12.004

Harvey, C. R., Liu, Y., Zhu, H. (2016) '… and the Cross-Section of Expected Returns', *The Review of Financial Studies*, 29(1), pp. 5–68. Available at: https://doi.org/10.1093/rfs/hhv059

Harvey, C. R., Siddique, A. (2000) Conditional Skewness in Asset Pricing Tests, *'The Journal of Finance'*, 55(3), pp. 1263-1295. https://doi.org/10.1111/0022-1082.00247

Hastie, T., Tibshirani, R., Friedman, J. (2009) *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd edn, New York, NY: Springer New York. Available from: ProQuest Ebook Central. (Accessed: 08/08/2026)

He, Z. (Lawrence), Huh, S.-W. and Lee, B.-S. (2010) 'Dynamic Factors and Asset Pricing', *Journal of Financial and Quantitative Analysis*, 45(3), pp. 707–737. Available at: https://doi.org/10.1017/S0022109010000207

Hornik, K., Stinchcombe, M., White, H. (1989) 'Multilayer Feedforward Networks are Universal Approximators', *Neural Networks,* 2(5), pp. 359-366, Available at: https://doi.org/10.1016/0893-6080(89)90020-8

Huang, D., Jiang, F., Tu, J., Zhou, G. (2015) 'Investor Sentiment Aligned: A Powerful Predictor of Stock Returns', *The Review of Financial Studies*, 28(3), pp. 791-837. Available at: https://www.jstor.org/stable/24465727 (Accessed: 26/06/2026)

Jagannathan, R., Wang, Z. (1996) 'The Conditional CAPM and the Cross-Section of Expected Returns', *The Journal of Finance*, 51(1), pp. 3-53. Available at: https://www.jstor.org/stable/2329301 (Accessed: 04/07/2026)

James, G., Witten, D., Hastie, T., Tibshirani, R. (2021) *An Introduction to Statistical Learning: with Applications in R*, 2nd Edition, Springer.

Keim, D. B., Stambaugh, R. F. (1986) 'Predicting returns in the stock and bond markets, *Journal of Financial Economics*, 17(2), pp. 357-390. Available at: https://doi.org/10.1016/0304-405X(86)90070-X

Kelly, B., Moskowitz, T. J., Pruitt, S. (2021) 'Understanding momentum and reversal', *Journal of Financial Economics*, 140, pp. 726-743. Available at: https://doi.org/10.1016/j.jfineco.2020.06.024

Kelly, B., Pruitt, S. (2013) 'Market Expectations in the Cross-Section of Present Values', *The Journal of Finance*, 68, pp. 1721-1756. Available at: https://doi.org/10.1111/jofi.12060

Kelly, B., Pruitt, S. (2015) 'The three-pass regression filter: A new approach to forecasting using many predictors', *Journal of Econometrics*, 186, pp. 294-316. Available at: http://dx.doi.org/10.1016/j.jeconom.2015.02.011

Kelly, B., Pruitt, S., Su, Y. (2019) 'Characteristics are covariances: A unified model of risk and return', *Journal of Financial Economics*, 134, pp. 501-524. Available at: https://doi.org/10.1016/j.jfineco.2019.05.001

Kim, C.-J. (1994) 'Dynamic linear models with Markov-switching', *Journal of Econometrics*, 60(1-2), pp. 1-22. Available at: https://doi.org/10.1016/0304-4076(94)90036-1

Kim, C.-J., Halbert, D. C. R. (1999) *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*, Cambridge, The MIT Press. Available at: https://doi.org/10.7551/mitpress/6444.003.0007

Kozak, S., Nagel, S., Santosh, S. (2020) 'Shrinking the cross-section', *Journal of Financial Economics*, 135(2), pp. 271-292. Available at: https://doi.org/10.1016/j.jfineco.2019.06.008

Krauss, C., Do, X. A., Huck, N. (2017) 'Deep neural networks, gradient-boosted trees, random forests: Statistical arbitrage on the S&P 500, *European Journal of Operational Research*, 259(2), pp. 689-702. Available at: https://doi.org/10.1016/j.ejor.2016.10.031

Lettau, M. (2024) *High-Dimensional Factor Models and the Factor Zoo*, NBER Working Paper No.31719, Cambridge, National Bureau of Economic Research. Available at: http://www.nber.org/papers/w31719 (Accessed on: 13/07/2026)

Lettau, M., Ludvigson, S. (2001) 'Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying', *Journal of Political Economy*, 109(6), pp. 1238-1287. Available at: https://www.jstor.org/stable/10.1086/323282 (Accessed: 10/07/2026)

Lettau, M. Pelger, M. (2020) 'Estimating latent asset-pricing factors', *Journal of Econometrics*, 218, pp. 1-31. Available at: https://doi.org/10.1016/j.jeconom.2019.08.012

Lewellen, J., Nagel, S. (2006) 'The conditional CAPM does not explain asset-pricing anomalies', *Journal of Financial Economics*, 82(2), pp. 289-314. Available at: https://doi.org/10.1016/j.jfineco.2005.05.012

Lintner, J. (1965) 'The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets', *The Review of Economics and Statistics*, 47(1), pp. 13-37. Available at: https://doi.org/10.2307/1924119

Martellini, L., Ziemann, V. (2010) 'Improved Estimates of Higher-Order Comoments and Implications for Portfolio Selection', *The Review of Financial Studies*, 23(4), pp. 1467–1502. Available at: https://doi.org/10.1093/rfs/hhp099

Merton, R. C. (1973) 'An Intertemporal Capital Asset Pricing Model', *Econometrica*, *41*(5), pp. 867–887. Available at: https://doi.org/10.2307/1913811

Messmer, M., Audrino, F. (2022) 'The Lasso and the Factor Zoo-Predicting Expected Returns in the Cross-Section', *Forecasting*, 4, pp. 969–1003. Available at: https://doi.org/10.3390/forecast4040053

Moritz, B., Zimmermann, T. (2016) 'Tree-Based Conditional Portfolio Sorts: The Relation between Past and Future Stock Returns', *SSRN*. Available at: http://dx.doi.org/10.2139/ssrn.2740751

Nechvátalová, L. (2025) 'Autoencoder asset pricing models and economic restrictions —international evidence', *International Review of Financial Analysis*, 107, Article 104642. Available at: https://doi.org/10.1016/j.irfa.2025.104642

Ng, L. (1991) 'Tests of the CAPM with Time-Varying Covariances: A Multivariate GARCH Approach', *The Journal of Finance*, 46(4), pp. 1507-1521. Available at: https://www.jstor.org/stable/2328869 (06/06/2026)

Ouysse, R. (2011) 'Comparison of Bayesian moving Average and Principal Component Forecast for Large Dimensional Factor Models', Discussion Papers 2012-03, School of Economics, The University of New South Wales. Available

at: http://research.economics.unsw.edu.au/RePEc/papers/2012-03.pdf (Accessed: 06/08/2026)

Oza, N. C. (2004) 'AveBoost2: Boosting for Noisy Data', in Roli, F., Kittler, J. and Windeatt, T. (eds) Multiple Classifier Systems: MCS 2004. Lecture Notes in Computer Science, vol. 3077. Berlin and Heidelberg: Springer, pp. 31-40. Available at: https://doi.org/10.1007/978-3-540-25966-4_3

Perez-Quiros, G., Timmermann, A. (2002) 'Firm Size and Cyclical Variations in Stock Returns', *The Journal of Finance*, 55(3), pp. 1229-1262. Available at: https://doi.org/10.1111/0022-1082.00246

Rasekhschaffe, K. C., Jones, R. C. (2019) ,Machine Learning for Stock Selection', *Financial Analysts Journal*, *75*(3), pp. 70–88. Available at: https://doi.org/10.1080/0015198X.2019.1596678

Rossi, A. G. (2018) 'Predicting Stock Market Returns with Machine Learning', University of Maryland. Available at: https://mendoza.nd.edu/wp-content/uploads/2019/07/2018-Alberto-Rossi-Fall-Seminar-Paper-1-Stock-Market-Returns.pdf (Accessed: 07/08/2026)

Sharpe, W. F. (1964) 'Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk', *The Journal of Finance*, 19(3), pp. 425–442. Available at: https://doi.org/10.2307/2977928

Steel, M. F. J. (2020) 'Model Averaging and Its Use in Economics', *Journal of Economic Literature*, 58(3), pp. 644–719. Available at: https://doi.org/10.1257/jel.20191385

Timmermann, A. (2000) 'Moments of Markov switching models', *Journal of Econometrics*, 96, pp- 75-111. Available at: https://doi.org/10.1016/S0304-4076(99)00051-2

Wang, C., Zhang, G. (2025) ‘In the shadows of opacity: Firm information quality and latent factor model performance’, *International Review of Financial Analysis*, 100, Article 103970. Available at: https://doi.org/10.1016/j.irfa.2025.103970

Yang, X., Zhu, Z., Li, D., Zhu, K. (2024) ‘Asset Pricing via the Conditional Quantile Variational Autoencoder’, *Journal of Business & Economic Statistics*, 42(2), pp. 681-694. Available at: https://doi.org/10.1080/07350015.2023.2223683

Zaffaroni, P. (2025) ‘Factor Models for Conditional Asset Pricing’, *Journal of Political Economy*, vol. 133(8), pp. 2615-2642. Available at: http://dx.doi.org/10.1086/735513