

**Libera Università Internazionale  
degli Studi Sociali Guido Carli**

**PREMIO TESI D'ECCELLENZA**

**Empirical Asset Pricing  
via Machine Learning  
and Fundamental Financial  
Metrics: evidence from  
the European Stock Market**

**Samuele Troiano**

**2023-2024**

Libera Università Internazionale  
degli Studi Sociali Guido Carli

Working Paper n. 8/2023-2024

Publication date: December 2025

*Empirical Asset Pricing via Machine Learning and Fundamental Financial Metrics: evidence from the European  
Stock Market*

© Samuele Troiano

ISBN 979-12-5596-344-8

This working paper is distributed for purposes of comment and discussion only.  
It may not be reproduced without permission of the copyright holder.

Luiss University Press is an imprint  
of Luiss Guido Carli  
Viale Pola 12, 00198 Roma  
Tel. 06 85225485  
E-mail [universitypress@luiss.it](mailto:universitypress@luiss.it)  
[www.luissuniversitypress.it](http://www.luissuniversitypress.it)

# Empirical Asset Pricing via Machine Learning and Fundamental Financial Metrics: evidence from the European Stock Market

By Samuele Troiano

## ABSTRACT

This research study investigates the predictive capabilities of various machine learning algorithms in forecasting European stock returns. It compares the performance of neural networks, tree-based models, and ensemble learning methods to assess whether robust economic and statistical outcomes can be achieved using a parsimonious set of predictors – primarily firm-specific characteristics and four momentum-based variables – applied to a relatively limited cross-section of 176 stocks. While numerous studies have reported strong predictive and economic results within the U.S. equity market, this research study aims to evaluate whether such findings can be extended to the European context, focusing on the main European equity indices. The empirical results reveal substantial heterogeneity in the models' statistical forecasting performance, leading to significant differences in economic profitability across approaches. Evaluation based on return and risk measures for long-only investment strategies indicates that portfolios constructed from machine learning model predictions can consistently outperform the market benchmark. Moreover, in long–short trading strategies, the models—except for the Stacking Regressor—achieve economically meaningful excess returns even when tasked with predicting downward price movements, once again outperforming the market portfolio. Overall, neural network architectures exhibit superior predictive accuracy and profitability compared to tree-based and ensemble approaches across all dimensions of analysis, highlighting their strong potential as effective tools for equity return prediction within European markets.

## INTRODUCTION

In recent times, machine learning has emerged as one of the most powerful tools in financial markets, challenging traditional methodologies for asset pricing and return forecasting. This research examines the potential applications of machine learning algorithms for stock returns prediction using fundamental financial indicators and momentum features within the context of the European stock market. Although numerous studies have demonstrated the efficacy of machine learning models in the U.S. market, there has been a notable dearth of research in the European context, despite significant differences in market dynamics and regulatory contexts. Given the nonlinear structure of these models, this study aims to test whether the success of machine learning in predicting U.S. stock returns can be replicated in Europe. To test whether machine learning models are able to explain and predict stock returns in the European context, this study employs neural networks, tree-based models, and ensemble techniques. The goal of the present study is to answer to the following questions:

1. Can machine learning models yield statistically and economical significant results in predicting stock returns within the European context?
2. Is it possible to achieve the performance sought in point 1 using a limited set of predictors consisting mainly of fundamental financial metrics of firms?
3. Is it possible to ascertain whether the aforementioned results could be achieved utilizing a relatively modest dataset of firms and a limited time span?

Previous studies have examined the capacity of machine learning models to study stocks movements in the European markets. The majority of these studies have employed predictors comprising hundreds (or, in some cases, thousands) of variables, resulting in only a limited number of predictors being of value to the models (see Fama (2008), Leippold, Q. Wang, and W. Zhou (2022), Martens and O. Penninga (2021)). Other studies, such as Drobetz and Otto (2021) and Marsi (2021), employed a limited number of predictors, but utilized a robust dataset comprising approximately 1,000 companies. Overall findings indicate that, in certain instances, machine learning models may offer superior performance compared to traditional statistical methods, like it happens when predicting US stock returns. This study aims to contribute to the existing body of literature on this subject by exploring these observations further.

The present research provides that machine learning models, especially neural networks, hold significant advantages in predicting stock returns within the European equity market. Aligned with the existing academic literature, i observe

that the level of complexity varies significantly between each model, yet they all identify similar predictors as relevant. The most influential predictors are based on current price movements, such as short-term reversal and stock momentum. Within the financial metrics predictors, the most relevant ones across models are instead the dividend price ratio, earning per share and the dividend yield. Despite these similarities, the forecast models considered in this analysis differ in statistical predictive effectiveness (as evidenced by predictive  $R^2$  metrics, MAE, MSE and Diebold and Mariano (1995) test statistics), resulting in significant variances in economic profitability. Neural network models, differently from the other models considered, assign greater weight to a more selective subset of predictors. The return and risk metrics of long-only investment strategies, which involve classifying stocks into decile portfolios based on predicted return estimations and buying the top decile portfolio, show that machine learning methods can provide predictive profits. Additionally, these models, through the construction of long-short investment strategies, demonstrate the ability to deliver economically relevant predictions even in declining markets, with all models —apart from the Voting Regressor— achieving notable returns during bear phases. Specifically, the top-performing long-short strategy originates from the neural model, which achieves a cumulative return of 3.63, compared to the market's 1.23 over the same period. The superiority of neural models in forecasting both the direction and magnitude of returns is further evidenced in the construction of decile portfolios; in contrast to other models, the forecasts generated by neural networks increase proportionally and monotonically with actual realized returns. Despite the high gains from using these models, it is presented that high levels of turnover, and consequently high transaction costs, can compromise and circumvent the superior performances achieved.

The rest of the work is organized as follows: in Chapter 1 the relevant academic literature with respect to stock return predictability and machine learning models is summarized. Special attention has been paid to the Efficient Market Hypothesis and how Machine Learning challenges traditional approaches to stock returns forecasting.

In Chapter 2, I'll describe the most relevant machine learning models used in the current thesis. Afterwards, detailed description of data handling and modeling is given, alongside the discussion on model selection and performance metrics.

Chapter 3 discusses the empirical results from this study. First, I illustrate the predictive accuracy and statistical significance of each machine learning model under analysis. At the end, portfolio performances both for the long-only and long-short strategies are discussed, while considering the economic significance of machine learning models in trading contexts, taking into consideration portfolio turnover and break-even trading costs.

## CHAPTER 1. LITERATURE REVIEW

For years, predictability in stock returns has been one of the debated topics in finance, standing against the Efficient Market Hypothesis. While traditional methods have considered either financial ratios or past prices for return forecasting, recent developments in machine learning techniques offer new paradigms. This section synthesizes academic research on these themes, focusing on the challenges on the EMH and the use of financial ratios in traditional forecasting models and ML techniques in order to predict stock

### *1.1 Predictability of stock returns and Efficient Market Hypothesis (EMH)*

Economists have long been fascinated by the nature and sources of variations in the stock market. By the early 1970s, a consensus had emerged among financial economists suggesting that stock prices were well proxied by a random walk model and that changes in stock returns were basically unpredictable. Fama (1970) provides an early, definitive statement of this position. Theoretically, the random walk theory of stock prices is historically preceded by theories relating movements in the financial markets to the business cycle. One notable example concerns the interest of Keynes in the variation in stock returns over the business cycle. The efficient market hypothesis began to take shape in the 1960s from the random walk theory of asset prices advanced by Samuelson (1973). Samuelson (1973) showed that, in an informationally efficient market, the changes in prices have to be unforecastable. Indeed, statistical evidence on the random character of the changes in equity prices had already been provided by Kendall (1953), Cowles (1960), Osborne (1959), and many others. However, it was Samuelson's contribution that had given academic respectability for the hypothesis; although the random walk model has been around for many years, having been originally discovered by Louis Bachelier, a French statistician, back in 1900. Whereas a number of studies did find some statistical evidence against the random walk hypothesis, these were dismissed as economically unimportant—trading rules based upon them could not generate profits in the presence of transaction costs—and statistically suspect—they might be due to data mining. For instance, Fama (1965) concluded that “there is no evidence of important dependence from either an investment or a statistical point of view.” Despite its apparent empirical success, the random walk model remained a statistical statement, rather than a coherent theory of asset prices. For example, it need not hold in markets with risk-averse traders, even under market efficiency. The term ‘market efficiency’ first received formal attention in the seminal review by Fama (1970); it has since typically been known as the informational efficiency of financial markets and focuses on information as the key to price setting. The efficient markets hypothesis (EMH) more narrowly defines an efficient market as one in which a new piece of information is quickly and accurately reflected in the current security price. In the seminal review, Fama (1970) surveys the empirical evidence for the weak-form, semi-strong-form, and strong-form EMH. Being

from the first category, he gives relatively wider coverage. Most of the empirical studies before 1970 generally used serial correlation tests and technical trading rules, and their findings strongly indicated that stock markets were weak-form efficient. Exploring the predictability of equity returns then allows an investigation of the degree of efficiency in the equity market. Two decades later, Fama (1991) conducts a second review of the market efficiency literature. Instead of anchoring on past returns, he expands the coverage of weak-form EMH to tests of return predictability using other variables such as the dividend-price ratio, earnings-price ratio, book-to-market ratio, and various measures of the interest rates. Event studies and tests for private information are merely the new names for tests of the respective semi-strong-form and strong-form EMH. His review indicates that evidence of return predictability from past returns, dividend yields, and a number of term structure variables is mounting, but he argues this may be spurious and should be treated skeptically.

Fama and French (1993) extended work by identifying five common risk factors in the returns on stocks, three of which are stock market factors based on financial information of firms: an overall market factor, and factors related to firm size and book-to-market equity. They use the time-series regression approach of Black, Jensen, and Scholes (1972), regressing monthly returns on stocks and bonds on the returns to a market portfolio of stocks and mimicking portfolios for size, book-to-market equity, and term-structure risk factors in returns. The time-series regression slopes are factor loadings that, unlike size or BE/ME, have an unequivocal interpretation as risk-factor sensitivities for bonds as well as for stocks. Their findings indicate that portfolios designed to depict the risk factors of size and book-to-market equity reflect strong co-movement in common stock returns, supporting the hypothesis that size and book-to-market equity proxy for sensitivity to common risk factors in stock returns. Their multifactor approach documents that, while the market factor captures a significant portion of variability in stock returns, additional size and book-to-market equity factors provide crucial explanatory power regarding cross-sectional variation in average returns. This evidence challenges the strict interpretation of the EMH, suggesting that stock returns may be somewhat predictable, particularly when considering the multifactor dimensions of risk. In addition to the canonical Fama and French papers, other widely cited studies like Rozeff (1984), Campbell and Shiller (1988), Cochrane (1991), Goetzmann et al. (1993), Hodrick (1992), and Lewellen (2004) report the main predictors of stock returns to be dividend yield, dividend-price ratio, dividend payout, earnings-price ratio, and book-to-market ratio.

Earnings-based measures have also proven useful in predicting stock returns. Basu (1977) finds that stocks with low price-to-earnings (P/E) ratios earn higher returns than those with high P/E ratios, suggesting that earnings information is not fully captured by market prices. This substantiates the idea that financial and

accounting information is instrumental in acquiring a better forecast of stock returns, providing investors with valuable insights beyond the scope of market efficiency as traditionally defined. Indeed, it is the integration of financial ratios and accounting information that constitutes the fundamental analysis of a stock, in contrast to trading strategies based on technical indicators like Moving Average or traded volume. The estimate in fundamental analysis is the intrinsic value of a company, which investors try to assess using information provided by the firm through public accounting disclosures and by observing market prices to evaluate whether the market price reflects the firm's financial performance. The outcome of this evaluation is to understand whether a firm is undervalued or overvalued in the market, and then to speculate on the stock's return over the following weeks, months, or years, depending on the variables used for evaluation and the investor's time horizon. Intuitively, if a stock is considered undervalued, investors will expect an increase in the stock price and predict a positive—and sometimes abnormal—excess return. This implies that if investors can identify the crucial variables for assessing firm value, those variables can be input into models used to predict and forecast future stock returns.

Academic literature shows that stock returns can indeed be predicted from historical accounting data, particularly financial ratios (Chattopadhyay et al (2016), Kheradyar et al (2011), Chen et al (2010)) Fundamental factors such as profitability, solvency, liquidity, and operating efficiency are crucial in developing successful fundamental investment strategies, and there is an optimistic correlation between these cumulative fundamental indicators and high-performance companies (Venkates et al (2016), Agustin. (2019)). Among these fundamentals, various variables have been identified as influencing stock returns. Khan (2012) found that EPS, ROE, cash flow ratio, and D/E had a significant positive impact, while the net profit margin ratio had a significant negative impact on stock returns. Lev and Thiagarajan (1993) predicted that earnings-related fundamental signals like sales growth and gross margin are significantly related to stock returns and future earnings predictions. Except for the debt-to-equity ratio (D/E), all other fundamental variables have shown a significant positive relationship with stock returns (Akbar (2023), Dempsey (2010)). Hatta (2012) used price-earnings ratio, EPS, D/E ratio, dividend per share, net profit margin, and return on assets as proxies for financial factors in determining stock returns. His findings confirmed that EPS and the price-earnings ratio had a significantly positive relationship, while D/E ratio and net profit margin showed a significantly negative relationship with stock returns. Martani et al. (2009) studied the role of cash flow, firm size, market-based ratios, liquidity ratios, profitability ratios, and leverage ratios on stock returns finding negative relationships except for profitability, market-based ratios, and turnover ratios on stock returns.

## *1.2 Predictability of stock returns and Machine Learning techniques*

As pointed out, the ability to predict equity returns has long been a fundamental question in the empirical finance literature. Despite the extensive work on empirical asset pricing, there remains a divide on whether equity returns are predictable. Goyal and Welch (2008) favor a negative answer to this question, while others, reviewed later and mentioned above, suggest that variables with predictive power can indeed be found. Our work contributes to the many attempts made over the years to add evidence to the prediction of stock returns within the finance and economics literature.

Recent developments in the machine learning (ML) literature allow for the study of equity return predictability using a vast set of candidate predictors simultaneously, overcoming limitations faced by conventional statistical methods. In some cases, machine learning techniques more than double the performance of leading regression-based strategies from the literature, with predictive gains linked to the non-linear predictor interactions that traditional methods miss (Wang et al. (2018)). Machine learning techniques are designed to build predictive models even when the number of predictors is extremely large, sometimes exceeding the number of available observations. These algorithms are highly appealing in finance, both for practitioners and researchers, given the abundance of data and the typical challenges associated with predicting stock returns.

Machine learning algorithms provide the ability to study equity return predictability in an agnostic manner. Unlike traditional approaches, machine learning algorithms do not require pre-specifying which predictors to use or the functional form of the predictive model, allowing for both linear and non-linear relationships. This flexibility enables the data to speak for themselves without imposing theoretical constraints on how future returns are determined. However, this comes at the cost of being unable to draw theoretical implications from the model, which is a limitation of this approach. Recently, Gu et al. (2020) synthesized empirical asset pricing literature with machine learning methods in a comparative analysis for stock return prediction. They present evidence clearly rejecting the ordinary least squares (OLS) benchmark in favor of machine learning methods in terms of both statistical performance and investor economic performance. The OLS benchmark represents a typical approach in one of two main strands of empirical literature on stock return prediction. Specifically, in cross-sectional stock return predictability research (e.g., Fama and French (200), Drobetz and Otto (2021)), future stock returns are regressed on a few lagged stock characteristics. The second strand, time-series stock return predictability research, focuses on predicting the time-series of returns, often forecasting stock indices using macroeconomic predictors. Due to the underperformance of the linear benchmark in their study, Gu et al. (2020) recommend using machine

learning techniques to overcome the serious limitations of commonly applied methods.

Gu et al. is one of the first paper to comprehensively use ML methods in cross-sectional predictability research. Many concurrent papers have followed this seminal work, each contributing to the literature by investigating the generality of the conclusions derived by the researchers. This has been done by applying their research design to other stock markets (Tobek (2021), Drobetz and Otto (2021), Leippold et al. (2022), Chopra (2021), Rubesam (2022), Lalwani (2022)) for international or regional applications, or to other asset classes (Martens (2021) for bond returns). More precisely, I focus on two separate issues: (1) overfitting and irrelevant predictors, and (2) a U.S. bias. The following sections discuss these concerns and how they can be addressed.

The aforementioned literature typically relies on large predictor sets. For example, the predictor set in Gu et al. (2020) includes 94 firm characteristics and interactions of each firm characteristic with eight macroeconomic time-series from Welch (2008) and 74 industry sector dummy variables. This results in a set of 900+ predictors. Following Kelly (2020), Leippold et al. (2022) even extend this predictor set (i.e., 1160 predictors) while relying on a smaller cross-section (i.e., 3,900 stocks) and time-series (i.e., a study period from 2000 to 2020). Tobek and Hronec (2021) and Martens et al. (2021) also use more than 100 predictors when investigating stocks internationally and bonds, respectively. Such large predictor sets and the enhanced flexibility of machine learning methods over more traditional prediction techniques such as the OLS benchmark come at the risk of overfitting the data, which may put machine learning methods at a disadvantage. My set of predictors, 28 in total, as they'll be discussed in the section 'Data', are constituted by the main drivers of value firm. The goal of the current thesis is to test whether machine learning approaches work even if the research is conducted without utilizing a large number of predictors from which the model gets knowledge, and the universe of stocks is not sufficiently enough to allow the model to "learn too much".

Another issue is the U.S. bias and lack of external validity in economics Kalimipalli (2013), and finance Karolyi (2016) research. Karolyi (2016) finds that only 16% to 23% of empirical studies published in top finance journals focus on non-U.S. markets, which is disproportionate to their economic significance. This presents two concerns: generalizing conclusions from U.S. data alone can be misleading, as research shows these findings don't always hold internationally (see Goyal (2015), Woodhouse (2017), Jacobs (2020)). Every replication adds value by extending studies out-of-sample. Harvey et al. (2017) suggest three ways to address multiple testing bias: (1) out-of-sample validation, (2) using statistical frameworks for multiple testing, and (3) looking across multiple asset classes. Like Fieberg et al. (2023), I use the first approach, conducting empirical analysis based on European markets. It is interesting to see if a market with a

shorter history and less research shows different predictability patterns. Replicating the main findings on stock return predictability outside the U.S. is key to validating these studies' conclusions. Furthermore, in opposed to Fieberg et al. (2023), who used 12,000 stocks across multiple markets in the MSCI Developed Europe Index<sup>1</sup>, so using both developed and emerging markets. I deploy 176 companies from major European indexes (EURO STOXX 50, DAX, CAC40, and AEX). Firms listed in more than one index will only be counted once. The goal of this work is to test whether the results from the discussed academic literature hold with a different number of predictors, only firm related, and fewer stocks. Other recent studies have examined the forecasting performance of machine learning algorithms on European stock returns, building on GKX's U.S. data analysis. Drobetz and Otto (2021) report strong predictive performance using neural networks and support vector machines, employing 22 predictors but focusing on a larger sample of 800 firms. Antonio Marsì (2021) conducted a similar analysis, using 45 stock-level characteristics drawn from the EURO STOXX 50 index

The remainder of the current work is organized as follows: Chapter 2 will discuss the machine learning methodology that will be deployed and will illustrate the organization of the data and predictors used. Chapter 3, on the other hand, will show the results of the empirical research conducted. Specifically, the latter will be divided in classification analysis, statistical analysis, variable importance and portfolios performances based on the machine learning methodologies discussed. The last part, the Conclusions, will draw summaries of the previously conducted analysis.

---

<sup>1</sup> That includes: Austria, Belgium, Denmark, Finland, France, Germany, Ireland, Italy, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, and the United Kingdom.

## CHAPTER 2. EMPIRICAL ANALYSIS

## 2.1 Methodology

In this empirical study, I am dealing with a panel of stocks, where months are indexed as  $t = 1, \dots, T$  and stocks are indexed as  $i = 1, \dots, N$ . Accordingly, the future stock return  $r$  of asset  $i$  at month  $t + 1$  can be defined in general terms as

$$r_{i,t+1} = E_t(r_{i,t+1}) + \epsilon_{i,t+1},$$

With

$$E_t(r_{i,t+1}) = g(x_{i,t}),$$

Current expectations about future returns are expressed as a function  $\hat{g}(\cdot)$  of a vector of predictor variables  $x_{i,t}$  (i.e., firm characteristics). The vectors of predictor variables at time  $t$  are defined as a  $P$ -dimensional vector  $x_{i,t} = (x_{i,t,1}, \dots, x_{i,t,P})$ . Thus, the following methods aim to provide an estimate  $\hat{g}(\cdot)$ . Despite its flexibility, this framework imposes some important restrictions. The function  $\hat{g}(\cdot)$  does not depend on  $i$  or  $t$ . Since the same functional form is adopted for all time periods and stocks, the model uses all the information in the panel and stabilizes estimates of risk premiums for an individual asset. The objective of most Machine Learning models is to find the optimal functional representation for the cross-sectional return  $y_t$  based on the information (characteristics) available at previous time point  $x_{i,t-1}$  as:

$$\hat{g}(\cdot) = \operatorname{argmin}_{g \in \mathcal{G}} \sum_{i=1}^N \sum_{t=1}^T (y_{i,t} - g(x_{i,t-1}))^2$$

All the models used are designed to predict the true returns by minimizing the out-of-sample mean squared forecast error:

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2$$

## 2.1.1 Regression trees and Random Forest

Regression Trees are a special class of decision trees used for predicting continuous outcomes. As a fundamental tool in statistical modeling and machine learning, they are widely applied due to their intuitive interpretability and ability to capture nonlinear relationships. Regression trees provide a systematic method for partitioning the feature space and modeling complex relationships among predictors. In our application, Regression Trees group firms by their characteristics, so that observations within each group are similar with respect to their future returns. A Tree algorithm splits the predictor space into  $K$  non-

overlapping regions  $R_1, R_2, R_3 \dots R_K$ . In each region, the fitted value of the target variable is its mean.

$$\hat{y}_k = \frac{\sum_{i=1}^n y_i \mathbb{1}(x_i \in R_k)}{\sum_{i=1}^n \mathbb{1}(x_i \in R_k)}$$

A tree 'grows' by completing a succession of binary splits based on cutoffs for concrete properties at each branching point. As a result, each split adds another layer of depth, allowing for more complicated interaction effects. Beginning with all observations, the tree divides the feature space into two rectangular partitions. The firm characteristic and cutoff value are selected to provide the best fit in terms of predicting accuracy. The generated rectangular patches in the predictor space approach the unknown function  $\hat{g}(\cdot)$  as the partition's mean. The splitting method continues on one or both sections, resulting in increasingly smaller rectangles until a stopping requirement is satisfied, such as tree depth (J) or no benefit from additional splits. Regions are defined by finding critical values of the predictors such that by fitting  $y$  with its mean over the region the MSE is minimized. To sum up the algorithm works following these 3 main steps:

1. Choose a predictor  $x_j \in X$  and a value  $x_j^*$  such that two regions are defined:  $R_1: x_j \leq x_j^*, R_2: x_j > x_j^*$ . The predictor and its critical value are chosen such that the MSE is minimized:

$$\text{MSE} = \frac{1}{N} \left( \sum_{i \in R_1(s)} (y_i - \hat{y}_{R_1})^2 + \sum_{i \in R_2(s)} (y_i - \hat{y}_{R_2})^2 \right)$$

$$(X_j^*, s^*) = \underset{j,s}{\text{argmin}} \text{MSE}(X_j, s)$$

Where  $\hat{y}_{R_1}$  and  $\hat{y}_{R_2}$  are the forecasted mean responses for the regions  $R_1(s)$  and  $R_2(s)$ , respectively, and  $N$  is the total number of observations. The critical value  $x^*$  is not chosen at random: the model takes each predictor and uses all the values that the predictor  $j$  takes in the dataset to identify the optimal one that minimizes the MSE.

2. Choose a region  $R_k$ , a predictor  $x_j$  and a value  $x_j^*$ , split the region  $R_k$  into 2 subregions accordingly and fit the target variable mean to each region. Again, the split is made to minimize prediction error.
3. Repeat step 2 recursively until a stopping criterion is met.

Figure 2.1. illustrates a regression tree of depth  $J = 2$  based on exemplary firm characteristics  $\text{char1}$  and  $\text{char2}$  (2 predictors) and 3 regions formed – the leaves of the tree  $K$  – These three parameters are critical in tree formation, and their optimal values will be found in the context of Random Forest using grid search,

which entails trying every combination of values and selecting the one that results in the lowest RMSE. Figure 2a depicts the tree architecture and its representation in the rectangular feature space. Initially, all data are divided by the firm characteristic  $char1$  using a cutoff value of 0.7. Firm-month observations with values less than that threshold are assigned to the left branch, whilst all other observations are assigned to the right branch.




Figure 2.1. Regression Trees structure. This figure shows the structure of an example regression tree (a) with two characteristics ( $char1$ ,  $char2$ ) and its representation as divided partitions of the predictor space (b)

Formally, the unknown function  $\hat{g}(\cdot)$  is approximated by a regression tree as:

$$\hat{g}^{\text{tree}}(x_{i,t}; K) = \sum_{k=1}^K c_k \mathbb{1}(x_{i,t} \in R_k),$$

Regression trees are commonly used in the ML literature because (as shown in figure 2.1.), the succession of binary decision rules allows for a very straightforward and intuitive explanation. Furthermore, the underlying structure is perfect for depicting multiway interactions and nonlinearities. However, regression trees are rarely utilized on their own since they are unstable in response to changes in the input data and are prone to significant overfitting difficulties. Unlike OLS, a regression tree can substantially overfit even with few variables, and with enough depth, a perfect fit is possible. In the simplest instance, each leaf node has simply one observation. This is why in machine learning modeling; Random Forests are spread out instead of regression trees.

Random Forests go a step beyond regression trees in using ensemble learning techniques that will enhance predictive accuracy and robustness. They mitigate problems of single trees and aggregate their multiple-tree outputs, each built from a different boot-strapped<sup>2</sup> sample of the data. Random Forests follow the

<sup>2</sup> Statistical method used to create multiple subsets of the original dataset by sampling with replacement.

Bagging-Bootstrap Aggregating approach: several regression trees are independently trained on different data subsets, generated via an i.i.d. bootstrap sampling. First, it will create  $B$  different bootstrap samples from the original dataset. It builds, for each observation, a regression tree by randomly selecting, at each node, a different subset  $p < P$  of predictors. This ensures diversity among the trees and reduces both the correlation as well as the variance of the estimator (cf. Breiman 2001). It then constructs each tree without pruning<sup>3</sup>, allowing maximum complexity to capture diverse patterns. This approach is quite important because the ensemble nature of random forests inherently reduces overfitting due to averaging. The final output of the forest will be the average of each bootstrap sample chosen by the model:

$$\hat{g}(x_{i,t}; B, K) = \frac{1}{B} \sum_{b=1}^B \hat{g}_b^{\text{tree}}(x_{i,t}; K).$$

I will apply grid search algorithm for choosing the best hyperparameters in the analysis. Grid search approach is used in order to understand how performance of a model measured by RMSE varies in case of change of a certain parameter. Correspondingly, the methodology allows finding the best combination of parameters that minimize the RMSE and provide more accurate predictions. The optimal parameters searched by the GridSearch algorithm are the following:

- **n\_estimators (B)**: this parameter specifies the number of decision trees that the Random Forest algorithm will create during training. In the formulas it corresponds to  $B$  or the so-called bootstrap sample.
- **max\_features (K)**: controls the number of features that are randomly chosen to evaluate for each split at a decision node. In the literature, the value chosen is the square root of  $P$  (the number of predictors). In this study it's allowed for variation using grid search, setting also the choice between other values.
- **Max\_depth (J)**: refers to the maximum number of levels (or depth) from the root node to the deepest leaf node in a decision tree. It essentially limits how many splits can occur along any path from the root to a leaf.

The value of the previously listed parameters used in GridSearch, and the one chosen by the optimization procedure are presented in figure 2.2.

---

<sup>3</sup> It represents the removal of superfluous branches or nodes to make the tree less complex, hence less expensive, simpler in its logic, and more general. The ultimate objective of pruning is to improve the generalizing performance of the decision tree by avoiding overfitting to the training data

### 2.1.2 Gradient Boosting regression trees

While trees are fit independently in the random forest algorithm, GBRTs are estimated in a manner that is adaptive so as to reduce bias. Thus, B trees are estimated sequentially and combined additively to form an ensemble prediction. The iterative procedure goes as: In step, the GBRT approach computes a first shallow tree to fit the realized returns. This oversimplified tree exhibits a high forecast error. Second, it calculates another shallow tree that fits the forecast residuals of the first tree. Its forecasts are summed over to arrive at an ensemble prediction. This forecast component given by this second tree is shrunk-a factor  $\nu \in (0,1)$  to prevent overfitting of forecast residuals. Every additional shallow tree is fitted to the forecast residual from the previous ensemble prediction; its shrunk forecast component gets added into the ensemble forecast. In other words, on each iteration, b, a new decision tree  $\hat{g}_b$  is fitted on the previous residuals added to the ensemble. The contribution of every tree, however, is controlled according to the learning rate shrinkage factor  $0 < \nu < 1$ , in order for the residuals not to overfit. Finally, given by the GBRT estimation function:

$$\hat{g}(x_{i,t}; B, K, \nu) = \sum_{b=1}^B \nu \hat{g}_b^{\text{tree}}(x_{i,t}; K).$$

In the context of Gradient Boosted regression trees, the following hyperparameters are set using grid search:

- **n\_estimators (B):** this parameter specifies the number of decision trees that the Gradient Boosted regression trees algorithm will create during training. In the formulas it corresponds to  $B$  or the so called bootstrap sample
- **Max\_depth (J):** refers to the maximum number of levels (or depth) from the root node to the deepest leaf node in a decision tree. It essentially limits how many splits can occur along any path from the root to a leaf.
- **learning\_rate ( $\nu$ ):** it regulates the learning rate shrinkage factor in order to avoid overfitting.

The value of the previously listed parameters used in GridSearch, and the one chosen by the optimization procedure are presented in table 2.1.

### 2.1.3. Neural Networks

Neural networks are the most complex method in this empirical analysis. They are highly parameterized – thus suitable to solve very complicated machine learning problems – but they are opaque and hard to interpret. In general, they map inputs – predictors – to outputs – realized returns. Inspired by the

functioning of the human brain, a network is composed of many interconnected computational units, called "neurons". But every single neuron on its own contributes very little, whereas a network of neurons works cohesively to improve the performance in prediction.

My analysis focuses on traditional "feed-forward" networks. These consist of an "input layer" of raw predictors, one or more "hidden layers" that interact and nonlinearly transform the predictors, and an "output layer" that aggregates hidden layers into an ultimate outcome prediction. The principle of this approach is to make the predictions through a linear combination of non-linear combinations of predictors, passed through a hidden layer. The structure of a neural network with 1 hidden layer, 4 predictors, 5 nodes is represented in the figure 2.2. below<sup>4</sup>. The input layer includes predictor variables, and each neuron in this layer represents one of the characteristics of a firm, such as price-to-earnings ratio, market capitalization, and past returns among others, while the output layer contains an estimation of the dependent variable-realized returns. Without any hidden layer, the simplest neural network equals the OLS regression model. Hidden layers added allow us to move from shallow to deep architectures and capture interactions and nonlinear effects. The number of units in the input layer is equal to the dimension of the predictors, for which i choose four in this example and denote  $z_1, \dots, z_4$ . The left panel shows the simplest possible network without hidden layers (OLS case). Each of the predictor signals is amplified or attenuated according to a 5-dimensional parameter vector,  $\omega$  that includes an intercept and one weighing parameter per predictor. The output layer sums up the weighted signals into the forecast:

$$\theta_0 + \sum_{k=1}^4 z_k \theta_k$$

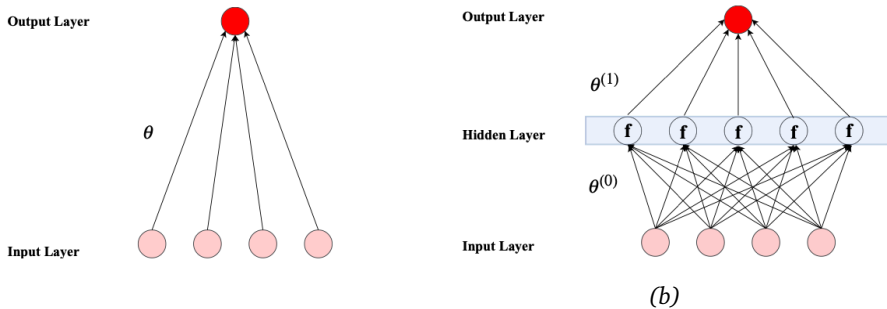
However, like in the simplest OLS instance, the intercept represents the model's bias or unexplained component, and the weights—or omegas—are the coefficients of each predictor, indicating how much influence each has on the prediction outcome. It incorporates more flexible prediction correlations by introducing hidden layers between inputs and outputs. The right panel of figure 2.2. shows an example of a hidden layer with five neurons. Each neuron receives information linearly from all of the input units, just like in the simple network on the left. Next, each neuron applies a nonlinear "activation function"  $f(\cdot)$  to its aggregated signal before sending it to the next layer. For example, the second neuron in the buried layer converts inputs into outputs as:

---

<sup>4</sup> S. Gu, B. Kelly, and D. Xiu. "Empirical asset pricing via machine learning". In: *The Review of Financial Studies* 33.5 (2020), pp. 2223–2273

$$x_2^{(1)} = f \left( \theta_{2,0}^{(0)} + \sum_{j=1}^4 v_j^{(0)} z_j \theta_{j,2}^{(0)} \right)$$

The second neuron in the first hidden layer, indicated as  $x_2^{(1)}$ , analyzes the input data by first performing a linear combination of the input features  $z_j$ , the weights  $\theta_{j,2}^{(0)}$ , and the bias term  $\theta_{2,0}^{(0)}$ . The linear combination occurs at the input level, whereas the activation function  $f(\cdot)$  occurs at the next layer, in this example, the first hidden layer.



(a) Neural network with no hidden layer

(c) Neural network with one hidden layer

Figure 2.2. Neural Network structure. This figure provides diagrams of two simple neural networks with (right) or without (left) a hidden layer. Pink circles denote the input layer, and dark red circles denote the output layer. Each arrow is associated with a weight parameter. In the network with a hidden layer, a nonlinear activation function  $f$  transforms the inputs before passing them on to the output

In this example, there are a total of  $31 = (4 + 1) \times 5 + 6$  parameters (five parameters for each neuron and six weights to aggregate them into a single output). The three most frequent activation functions are the sigmoid (or logistic) function, the hyperbolic tangent (tanh) function, and the Rectified Linear Unit (ReLU) function. Their expressions are given below:

$$\text{sigmoid: } f(x) = \frac{1}{1 + \exp^{-x}}$$

$$\text{tanh: } f(x) = \frac{\exp^x - \exp^{-x}}{\exp^x + \exp^{-x}}$$

$$\text{ReLU: } f(x) = \max(0, x)$$

Since the sigmoid function has a range from 0 to 1 and the tanh between -1 and 1, they usually are adopted in the hidden layer to mimic the behavior of biological neurons. More precisely, ReLU function is more used in the output layer for classification problems, while linear function,  $f(x) = x$ , is preferred for regression problems. The choice of the activation functions plays an important role in the architecture of an NN. To see this, notice that if all the activation functions were linear, the NN would boil down to a linear regression model, whereas nonlinear functions such as the sigmoid and the tanh function allow neural networks to discover complex functional relationships that might exist between outputs and inputs.

In order to implement the model used Keras,<sup>5</sup> an open-source library that provides a Python interface for constructing ANNs via TensorFlow. The activation function used in the hidden layer is 'relu', while a linear function was used for the output layer. As the optimizer, I choose Adam (Adaptive Moment Estimation), a form of mini-batch gradient descent that adjusts the learning rate for each model parameter at each iteration. It is an optimizer that has been shown to produce better results in terms of accuracy, but it requires more time and computational power.

Neural networks predict the output  $y$  by applying learned weights to the input features  $x$ . In the simplest linear model, such as ordinary least squares (OLS) regression, the coefficients are analogous to weights, but neural networks extend this by using multiple layers and non-linear transformations. As for the usual the weights in a neural network are optimized by minimizing a loss function. In this case, I used Root Mean Squared Error (RMSE) as the loss function to guide the optimization process In line with Jacobs (2020) and Fieberg et al. (2023), due to the complexity of Neural Networks, I choose to deploy two separate models, reducing computational efforts. The appropriate number of neurons in hidden layers is determined using the geometric pyramid rule from Masters (1993):

- **NN1:** with 1 hidden layer with 32 neurons ( $HL = 1; N = \{32\}$ )
- **NN2:** with 2 hidden layers with 64 and 32 neurons each ( $HL = 2; N = \{64,32\}$ )

One of the main parameters set by the search grid is the **batch\_size**: it determines how many training examples are used to compute the gradient and update the model's weights in one iteration. In other words, it controls how much data the model "sees" before making an update. Figure 2.3 depicts the minimization procedure for selecting the batch size, with each batch size having

---

<sup>5</sup> Keras is an open-source library that provides a Python interface for developing ANNs through another library for machine learning called TensorFlow

a matching RMSE. The grid search algorithm will select the smallest RMSE value that corresponds to a given batch size. For NN1 the optimal batch size is 32, for NN2 is 128.

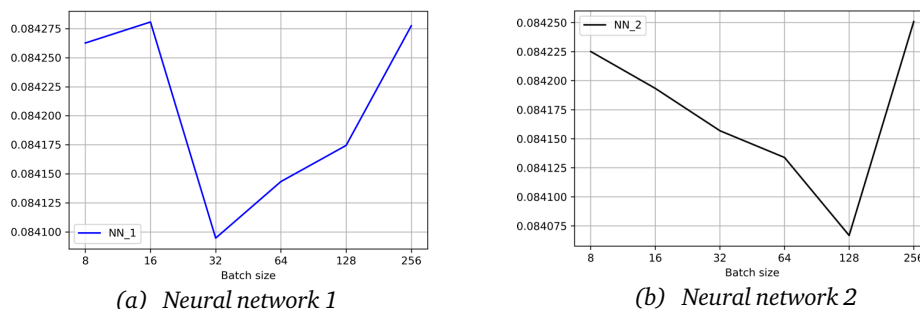


Figure 2.3. RMSE for varied batch sizes. The figure depicts the grid search algorithm's minimization phase. Since  $cv=3$  was chosen, each batch size has three different RMSEs. The graph depicts the minimal values of them

Yet another crucial tuning parameter for the learning algorithm is the **learning rate lr**, which controls the step size of every parameter update. In the following, the **learning\_rate** is set to 'adaptive', which means the learning rate will be changed dynamically across the course of training, based on an initial learning rate, **learning\_rate\_init**. The search grid also tunes another hyperparameter to avoid overfitting: **alpha**, alias the **L2 regularization parameter**, penalizes large weights by preventing overfitting. L2 regularization adds an additional term to the loss function discouraging the weights from getting too large, which can occur when the model starts memorizing the training data. The smaller the alpha value, the less regularization is applied, and the better the fit of the model.

I also use **early\_stopping**, a regularization technique that stops the training process when the model performance on the validation set doesn't show further improvements to prevent overfitting. If set to True, it enables early stopping to stop training when the model stops improving performance on the validation set, which saves time and prevents overfitting to the training data. Another important parameter set by the search grid is **max\_iter**, or **epochs**, meaning the maximum number of iterations the optimization process can perform. The optimization process is show in figure 2.4 where different values of RMSE are plotted in correspondence of the epochs values. An iteration is one complete pass through the training data. If the training converges early (i.e., the loss function stabilizes), **early\_stopping** will stop the process before reaching the iteration limits.




Figure 2.4. RMSE for varied epochs. The figure depicts the grid search algorithm's minimization phase. Since  $cv=3$  was chosen, each value of epochs has three different RMSEs. The graph depicts the minimal values of them

#### 2.1.4. Stacking Ensemble Regressor

Stacking is a sophisticated ensemble learning technique designed to enhance predictive performance by strategically combining multiple models. The core idea of stacking is to exploit the benefit of different machine learning models. One method for leveraging the benefits of these distinct models is to train them independently and give their predictions as input to a high-level model, also known as a **meta-model** or **level-1 model**. This meta-model aims to correct any systematic errors that the base models (**level-0 models**) might make, leading to a more accurate and robust final prediction. Mathematically speaking let  $\mathbf{X}$  represent the feature matrix and  $\mathbf{y}$  the target variable. Suppose you are working with  $M$  base models denoted by  $h_1(\mathbf{X}), h_2(\mathbf{X}), \dots, h_M(\mathbf{X})$ . The stacking procedure will deploy through the following steps:

1. **Base Model Training:** Each base model  $h_m(\mathbf{X})$  is trained on the dataset  $(\mathbf{X}, \mathbf{y})$ , resulting in predictions  $\hat{y}_m = h_m(\mathbf{X})$  for each model  $m$ .
2. **Meta-Model Training:** The predictions from all base models are compiled into a new dataset

$$\mathbf{X}' = [\hat{y}_1, \hat{y}_2, \dots, \hat{y}_M].$$

The meta-model  $h_{\text{meta}}(\mathbf{X}')$  is then trained on this dataset to predict the target variable  $\mathbf{y}$ .

This essentially encapsulates the original notion of two-layered prediction, with diversity in the base models at the first layer, followed by a meta-model that learns how to optimally combine these base models at the second layer. The meta-model itself might be as relatively simple as a linear regressor, but also might be more complex – if the problem complexity requires so – like a neural

network, depending on the chain of relationships between the base model's outputs.

The main theoretical benefit of stacking is the reduction of generalization error as a result of a balancing effect of individual model biases. Recent literature, such as those by Van der Laan, Polley, and Hubbard (2007) suggest that stacking does indeed combine diverse perspectives of a set of competing models to make impressive improvements upon the baseline predictive accuracy. This is because each of the base models is capable of learning to represent different patterns or aspects of the data, thus requiring the meta-model to be quite significant for synthesizing these perspectives into an integrated final forecast.

In the context of analyzing financial markets, stacking has gained more importance because of its power in model combination, which captures varied market phenomena, such as momentum, mean reversion, or volatility clustering. For example, Gu, Kelly, and Xiu (2020) show that the accuracy of forecasting stock returns is substantially improved by integrating various signals extracted from financial data by stacking a sequence of machine learning models, including gradient boosting machines and random forests. While the performance of stacking generally develops on the quality of the base models and the complexity of the meta-model. As Lessmann et al. (2015) note, an ensemble stacking can outperform an individual model only when a set of selected base models is adequately diverse and the meta-model parameters are suitably optimized. This is especially true for high-dimensional financial data, where different models would capture different dimensions, and some kind of stacking approach synthesizes them into something more reliable to make a prediction.

#### 2.1.5. Voting Ensemble Regressor

The Voting Ensemble Regressor is one of the simpler yet equally powerful ensemble techniques that, at its core, combines predictions by implemented voting mechanism. In regression tasks, voting generally means averaging the predictions generated by all base models in order to generate a final output. The intuitive idea underlying voting is the reduction of variance in the predictions, using diversity in the base models in order to yield a more robust and stable result. Suppose you have  $M$  base models  $h_1(\mathbf{X}), h_2(\mathbf{X}), \dots, h_M(\mathbf{X})$ . The prediction of the Voting Regressor,  $\hat{y}_{\text{voting}}$ , is computed as the mean of the predictions from these models:

$$\hat{y}_{\text{voting}} = \frac{1}{M} \sum_{m=1}^M h_m(\mathbf{X})$$

This approach finds its rationale in the law of large numbers, stating that an average of multiple independent estimators should outperform and yield a

prediction closer to the true value, under the assumption these estimators are unbiased and their errors are not highly uncorrelated (Kuncheva (2004)). The performance of a voting ensemble strongly depends on diversity among base models: the less correlated the errors of base models, the larger the potential for improvement by averaging (Rokach (2010)). Voting ensembles are of particular appeal in financial markets, where forecasts generally have a high degree of uncertainty owing to the stochastic nature of asset prices. A voting ensemble can reduce the impacts of outliers and noisy predictions by averaging the outputs of different models, thus making more reliable forecasts. Recent works like Patent et al. (2015), show that the performance of voting ensembles is superior to that of single models in the prediction of stock returns by appropriately aggregating several models like support vector machines, decision trees, and neural networks. Another important strength of regressor Voting is simplicity and interpretability, as the final prediction is a mean of the base models' predictions; therefore, it is simple to understand and explain, which can often be a critical choice in financial applications. Another advantage of voting, which has been highlighted by Dong (2020), is that such a combination is straightforward, so this technique is especially useful for tasks where model interpretability is valued about as much as predictive accuracy. On the other hand, voting does not explicitly train a meta-model to combine the base models; it relies entirely on the assumption that their combined prediction will be better than the predictions of its individual constituent models. This may turn out to be beneficial in scenarios where the relationships between the models are very rich and could not effectively be modeled by a meta-model. Furthermore, voting ensembles may be less computationally intensive than stacking because they do not require an additional layer of model training.

Thus, both stacking and voting ensemble regressors have their merits regarding the stock return prediction for the European market. While one can combine several models together by a stacked ensemble regressor in a more sophisticated way to potentially capture complex relationships in the data, a voting-based ensemble regressor is far simpler to produce, more interpretable, and yet sometimes quite accurate since a combination of voters reduces the variance in the predictions. Specific dataset characteristics, single model performance, and the trade-off between model complexity, interpretability, and computational load are some of the factors guiding the choice among these methods.

## *2.2. Data and Predictors*

Previous study revealed that there is a U.S. bias in the field of research of economics (Das et al. (2013)) and finance (Karolyi (2016)). In this regard, recent research, like Harvet et al. (2016) calls for a higher statistical threshold (t-statistic of 3 rather than 2) in empirical studies on the US stock market, as well as a greater emphasis on other asset classes and equity markets. I follow this demand by investigating the European stock market. More specifically the

sample of stocks deployed in the current work includes companies contained in the EURO STOXX 50 index, DAX, CAC40, e AEX. The target variable in my predictive analysis is the monthly return. Monthly price data was sourced from Refinitiv DataStream. As highlighted by Kelly (2020), using higher-frequency data like monthly returns can significantly improve the predictive power of machine learning models in asset pricing and return forecasting. Summary statistics on returns are provided in Table 2.1. After removing companies that were listed in multiple indices, the sample was narrowed down to 216 firms. The data spans from June 2004 to June 2023, covering a period of 19 years and providing 229 monthly observations per firm. Firms that lacked data for five or more years were excluded from the sample, leaving a total of 176 companies. The accounting variables used to predict the target variable were also downloaded from Refinitiv DataStream. The starting predictors are 28 and are derived from the extensive empirical asset pricing literature previously discussed, as well as from the CFI website. The complete list and descriptions and summary statistics of these variables, are provided in Table 2.1.

*Table 2.1. Summary statistics*

| <b>Predictor</b> | <b>Avg</b> | <b>Std</b> | <b>Description</b>            |
|------------------|------------|------------|-------------------------------|
| p/e              | 27.0       | 135.7      | Price-to-earnings ratio       |
| p/s              | 2.3        | 4.8        | Price-to-sales ratio          |
| p/cf             | 10.5       | 89.8       | Price-to-cash flow ratio      |
| p/b              | 5.1        | 70.4       | Price-to-book ratio           |
| g_profmargi      | 0.4        | 0.2        | Gross profit margin           |
| o_profmargi      | 0.2        | 0.3        | Operating profit margin       |
| n_profmargi      | 0.1        | 1.1        | Net profit margin             |
| eps              | 34.7       | 70.9       | Earnings per share            |
| eps_fully        | 2.1        | 4.2        | Fully diluted EPS             |
| bookval/n        | 18.8       | 28.6       | Book value per share          |
| tan_bv/n         | 7.5        | 23.3       | Tangible book value per share |
| cf/n             | 4.4        | 8.2        | Cash flow per share           |
| roe              | 0.2        | 0.8        | Return on equity              |
| roa              | 0.1        | 0.1        | Return on assets              |
| eps_5y           | 8.3        | 19.5       | 5-year EPS growth             |
| salgro_5y        | 0.1        | 0.2        | 5-year sales growth           |
| quick            | 1.0        | 0.8        | Quick ratio                   |
| d/e              | 1.2        | 6.6        | Debt to equity ratio          |
| d/v              | 0.4        | 1.7        | Debt to total capital         |
| w_cap/t_cap      | 0.1        | 0.3        | Working cap to total cap      |
| inv_turn         | 0.34       | 10.8       | Inventory turnover            |

Table 2.1. Summary statistics

| Predictor       | Avg  | Std  | Description              |
|-----------------|------|------|--------------------------|
| asset_turnr     | 0.7  | 0.5  | Asset turnover ratio     |
| dy              | 0.03 | 0.03 | Dividend yield           |
| $roa^2$         | 0.4  | 0.3  | Square of ROA            |
| $dy^2$          | 0.03 | 0.0  | Square of dividend yield |
| $roe^2$         | 0.7  | 14.6 | Square of ROE            |
| roa*roe         | 0.1  | 1.8  | ROA multiplied by ROE    |
| mom_1m          | 0.1  | 0.1  | Momentum over 1 month    |
| mom_6m          | 0.03 | 0.2  | Momentum over 6 months   |
| mom_9m          | 0.04 | 0.3  | Momentum over 9 months   |
| mom_36m         | 0.2  | 0.5  | Momentum over 36 months  |
| monthly_returns | 0.0  | 0.1  | Monthly returns          |

Following Otto (2020), I also consider two-way interactions between firm characteristics as predictors, as well as second and third-order polynomials of firm characteristics to capture nonlinear relationship. Due to computational limitations, I only selected some specific predictors for interaction and the integration of higher-order polynomials. The selection is based on the simple correlation with the target variable: the predictors that covary most strongly (positively or negatively) with the target variable were chosen. The list of correlations between predictors and returns is shown in the APPENDIX A in table A.2. Additionally, I constructed momentum indicators as predictors to capture the importance that investors place on past stock movements:

- **mom\_1m**: rolling 1-month return – short term reversal
- **mom\_6m**: rolling 6-month return
- **mom\_9m**: rolling 9-month return
- **mom\_36m**: rolling 36-month return – long term reversal

Incorporating rolling window returns into stock return prediction models exploits the momentum effect, where security prices tend to continue moving in the same direction. This enhances the viewpoint and boosts the predictive capability of the models. Various empirical studies by Jegadeesh and Titman (1993), Carhart (1997), recent works by Han, Lee, and Worah (2018), and Nordvig and Firoozye (2019) affirm the significance of momentum indicators in machine learning frameworks and validate their utility in financial prediction. Xia, Hu, and Sargent (2020) demonstrate that machine learning models incorporating momentum are likely to outperform those that do not.

To ensure accurate computation of firm momentum features without introducing bias, calculations are performed separately within a data subset of each firm. This approach maintains the integrity of predictive features by preventing mixing of calculations across firms. However, NaN values are inevitably generated in the momentum features' calculation, particularly at the initial stages of the time series due to insufficient data points for computing rolling returns. It is crucial to address this issue to prevent errors during model training, which necessitates a complete dataset.

In order to handle these NaN values, two imputation methods are deployed: the Spline Interpolation and K-Nearest Neighbors. Spline interpolation is a mathematical process based on fitting a smooth curve through available data points. In particular, cubic splines will be used for interpolating those missing values with the goal of imputed values to keep continuity and, meanwhile, the first and second derivatives smooth. This is a very effective method in financial time series data, where the underlying process is often a smooth trend, therefore making spline interpolation a good choice. Recent studies like Falk et al. (2019) support this application of the spline interpolation in financial contexts where it has performed well in maintaining the integrity of the time-series data.

Complementary, I used the K-Nearest Neighbors (KNN) Imputation method in addition to spline interpolation. KNN imputation works by finding the  $k$ -nearest neighbors, based on Euclidean distance or some other suitable metric, to the observation with the missing value and imputes the missing value with the mean of these nearest neighbors. KNN proves particularly useful when there is, possibly on the local scale, a pattern that might be lost with mere interpolation. For example, if certain stocks tend to exhibit similar momentum, KNN can use that to interpolate among stocks that do not have explicit values available on a specific day or time period. Its robustness concerning the handling of diverse data structures has been tested in many studies, including the recent comprehensive work from Beretta (2016), which underlined its efficiency in financial datasets.

By using Spline Interpolation along with KNN imputation, momentum features were imputed in a manner such that the temporal and cross-sectional characteristics of the data were preserved to retain the predictive integrity of the features.

During the course of the cleaning phase of the dataset, two principal actions were undertaken. firstly, in cases where a predictor is missing for a particular firm  $i$ , the median value of the cross-section for the corresponding month  $t$  is imputed, following the approach commonly seen in academic literature such as Fieberg et al. (2023) and Marsì (2021). Secondly, too correlated predictors were dropped. According to Lewellen (2015), the high correlations resulting from the similarity in construction, issuance metrics, and incorporation of comparable firm information, such as profitability measures, do not pose a significant issue in this

empirical analysis. The focus of the researcher lies in the collective predictive capability of the machine learning models rather than the individual effects of each predictor. Machine learning models offer a potential solution to address multicollinearity by choosing/reducing variables, thereby enhancing predictive accuracy. Despite this, I construct a correlation matrix to identify the most correlated one, following Otto (2021). The correlation matrix is shown in the APPENDIX A in figure A.1. The aim is to drop variables that exhibit correlation equal or higher of 80%. In my dataset the predictors are **current\_ratio** and **eps\_basic**, positively correlated respectively with **quick\_ratio** and **EPS\_diluted**. To enhance and confirm the results of my variable elimination, I performed a variable importance factor (VIF) analysis, that basically measures the correlation of one variable with all the others. Standard practices makes that a variable exhibiting values higher than 5 should be treated accordingly. Fortunate the highest value in my sample is 4.95 as it is shown in table A.3 in the APPENDIX A, meaning that no other variable should be dropped. I refined the initial dataset through two additional steps: firstly, I winsorized all monthly firm characteristics at the 1% and 99% levels to address outliers. Unlike Otto (2021), I followed Gu et al. (2020) in not winsorizing returns. Secondly, I ranked all firm characteristics on a monthly basis cross-sectionally and normalized the ranks to the  $(-1, +1)$  interval, in line with the approach of Kelly et al. (2019) and Fieberg et al. (2023), instead of a simple linear transformation:

$$\text{normalized\_rank} = \frac{2 \times \text{rank}_i}{n + 1} - 1$$

Due to the fact that machine learning methods allow us to overcome the high dimensionality problem, which arises when the number of predictors tends to be very large compared with the number of observations, and at the same time be useful in solving the problem of multicollinearity, they tend to overfit, so i will need to control for the degree of model complexity by tuning respective model hyperparameters. Examples of such hyperparameters include the number and/or depth of trees in boosted regression trees or random forests. In order to avoid overfitting and ensure maximum predictive power out-of-sample, hyperparameters cannot be set a priori, but they have to be adaptively determined from the sample data. When I'm dealing with time series data, the temporal order needs to be preserved. I therefore use a chronological split instead of random splitting. It is important to preserve the temporal order of the data in order to prevent leakage, which can result in overly optimistic evaluations of model performance.

In many circumstances, there is no theoretical guidance on how to tune hyperparameters for new data. To address this, i take a typical technique from current research literature by dynamically picking tuning parameters from a validation sample using GridSearch. This method allows us to construct a pseudo out-of-sample scenario by giving the model a variety of options and ranges to

choose from and then determining the optimal combinations. By splitting the data into two in-sample sets (testing and validation) and one out-of-sample set (test set), I can successfully optimize hyperparameters by minimizing a loss function, such as RMSE, on a validation sample rather than the training set. This strategy helps to prevent the creation of an overfit model during training, hence enhancing its performance

The values contained in the GridSearch and the optimal parameters chosen after the optimization procedure are summarized in Table 2.2.

The training sample is used to estimate the model for multiple parameter specifications, while the validation sample is used to tune the parameters. Based on the models estimated from the training sample, the RMSE within the validation sample is calculated for each parameter specification. The model with the parameter specification that minimizes the RMSE is used for out-of-sample testing. Note that, because the tuning parameters are chosen from the validation sample, it is not truly out-of-sample. The test sample, however, is used for neither model estimation nor parameter tuning. This is why it is truly out-of-sample and appropriate for evaluating a model's out-of-sample predictive power.

My sample data are then split into three different sets using a fixed window approach:

3. **Training set:** 8.5 years of data, or 100 monthly observations per firm, for a total of 17.600 observations.
4. **Validation set:** 5 years of data, or 60 monthly observations per firm, for a total of 10.560 observations.
5. **Test set:** 5.5 years of data, or 68 monthly observations per firm, for a total of 11.968 observations.

The target objective of the analysis is to predict the magnitude (and the sign) of the next month's return. To do so, the column `monthly_returns` is shifted, and a binary column based on the latter is created for classification purposes. The column target will take the value 1 when the return in the next period is  $> 0$ , and 0 otherwise.

*Table 2.2. Hyperparameters, search grid values and optimal parameter. This table display the hyperparameters search by grid search for each of the machine learning models deployed and described in the section Methodology. The algorithm searches for the best combination and gives as an output the optimal value. Since  $vc = 3$  is chosen, for each combination, the grid search algorithm will perform the task 3 times and will display the lower RMSE.*

| Model | Hyperparameter           | Grid search             | Optimal value |
|-------|--------------------------|-------------------------|---------------|
| RF    | n_estimators (B)         | [300,400,600,800]       | 800           |
|       | max_features (K)         | [sqrt, log2, 0.6]       | 0.6           |
|       | max_depth (J)            | [25,30,35,45]           | 25            |
| GBR   | n_estimators (B)         | [350,500,600,750]       | 750           |
|       | max_depth (J)            | [10,15,20]              | 10            |
|       | shrinkage_rate ( $\nu$ ) | [0.001,0.003,0.005]     | 0.003         |
| NN1   | batch_size               | [8,16,32,64,128,256]    | 32            |
|       | neurons                  | (32,)                   | (32,)         |
|       | epochs                   | [250,500,750,1000,1250] | 500           |
|       | learn_rate_init          | [0.001,0.003,0.005]     | 0.001         |
|       | l2_reg                   | [0.001,0.003,0.005]     | 0.001         |
| NN2   | batch_size               | [8,16,32,64,128,256]    | 128           |
|       | neurons                  | (64,32)                 | (64,32)       |
|       | epochs                   | [250,500,750,1000,1250] | 1000          |
|       | learn_rate_init          | [0.001,0.003,0.005]     | 0.001         |
|       | l2_reg                   | [0.001,0.003,0.005]     | 0.001         |
|       |                          |                         |               |

## CHAPTER 3. EMPIRICAL RESULTS

*3.1 Predictive Accuracy*

To assess the predictive performance of our models, the analysis of the predictions generated by our model will be based using two different approaches: regression and classification-based. As underlined by of Gu, Kelly, and Xiu (2020), the benefit of a dual approach lies in how combining regression and classification metrics can lead to a more holistic evaluation of model performance. While the regression metrics provide a detailed view of the magnitude of the prediction errors and the model's ability to capture variance, the classification metrics provide insight into the practical applicability of the predictions in real-world trading where the focus is often on the direction of returns.

Initially, i analyze the predictions by treating the analysis as a classification problem. It is feasible to do so by setting the target variable and prediction to 0 or 1 depending on their sign, which is analogous to the challenge of forecasting whether a stock's price will rise or fall, with the target being 1 when the return is positive and 0 when it is negative. Thus, i can now use classic classification metrics to evaluate our models, such as accuracy, precision, recall, and F1 score, which are defined as follows:

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}}$$

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}$$

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}$$

$$\text{F1 Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

This type of classification-based analysis is particularly helpful for practical trading applications, given that most decisions depend on the potential direction of the stock rather than necessarily on the exact size of the movement. Powerful metrics of classification can be if used in a financial context in which often the pre-eminent goal is not exact return values but just to correctly anticipate directional movements. The results regarding these classification metrics are shown in table 3.1 below:

*Table 3.1. Classification based metrics for each of the machine learning models*

| <b>Metric</b> | <b>RF</b> | <b>GBR</b> | <b>NN1</b> | <b>NN2</b> | <b>SR</b> | <b>VR</b> |
|---------------|-----------|------------|------------|------------|-----------|-----------|
| Accuracy      | 0.64      | 0.64       | 0.69       | 0.69       | 0.67      | 0.69      |
| Precision     | 0.65      | 0.65       | 0.70       | 0.70       | 0.70      | 0.68      |
| Recall        | 0.70      | 0.71       | 0.71       | 0.73       | 0.69      | 0.76      |
| F1 score      | 0.67      | 0.68       | 0.70       | 0.71       | 0.69      | 0.71      |

To further enhance the analysis of classification metrics, I constructed confusion matrices for each of the predictive models, showed in figure 3.1. The confusion matrix is a valuable tool descriptor in a classification task as it points to the number of correct and incorrect predictions across classes. It plots true positives, true negatives, false positives, and false negatives as output, giving a clear picture of areas in which the model performs accordingly, and areas in which it needs improvements.

Classification metrics and confusion matrices indicate that neural networks and ensemble techniques, in general, tend to slightly outperform tree-based models on most of the relevant performance metrics. In general, given an accuracy rate for all models of between 65% and 70%, it is safe to say that adopting similar models offers visible benefits in terms of correct predictions, compared to a random situation of 50%/50%. In this regard, however, it is also necessary to consider that, on average, stocks tend to increase in value, so taking 50/50 as a benchmark is not entirely correct. However, especially in the short term, a percentage of correct predictions close to 70% is certainly a good starting point. NN2 and Voting Regressor have the highest recall and F1 score, which means these have a better balance between precision and recall. This is especially so for NN2, which seems to have done a pretty good job of reducing the number of false negatives. The Voting Regressor is slightly better in recall, which means an even better ability to identify positive upward movement of a stock, which could be very helpful in the stock return prediction when capturing upward movements is more important. Also the other models, seems to have a good ability of predicting positive returns. This is also confirmed by good percentages of precision. It can also be inferred, from the confusion matrices, that the models are less adept at anticipating downward movements of the stocks under analysis, with respect to the already discussed upward movements.

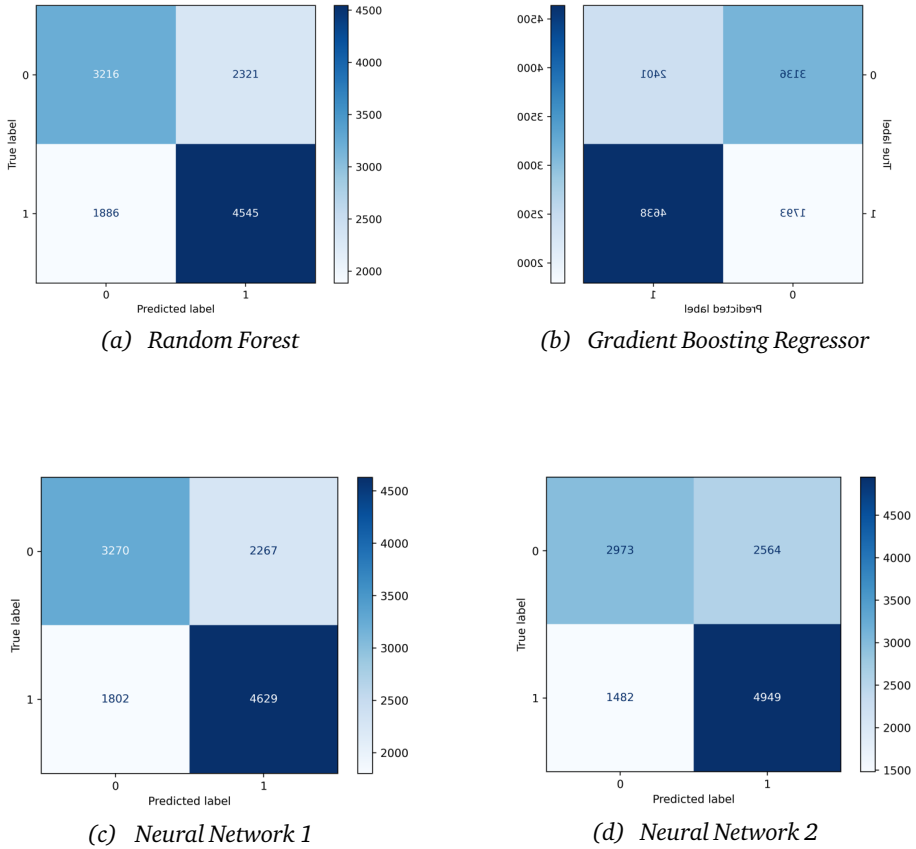


Figure 3.1: Confusion matrices. The matrices are constructed by comparing the target monthly return and the predictions of each of the machine learning models in order to show the exact number of correct and incorrect predictions

Secondly i deploy statistical analysis to quantify the models' predictive accuracy, i.e., how well predicted returns reflect realized returns. i will use different metrics, explained below.

In addition to the common  $R^2$  Out of sample, that measures the proportion of the variance in the actual returns that is explained by our predictors, as Fieberg (2021) and Cakici et. al. (2023) i will deploy a common metric to measure the capability of a model to predict the magnitude of the returns, the out-of-sample (pseudo)  $R^2$ , which is defined as:

$$R_{\text{Pseudo OOS}}^2 = 100 \cdot \left( 1 - \frac{\sum_t \sum_i (\hat{r}_{i,t} - r_{i,t})^2}{\sum_t \sum_i r_{i,t}^2} \right)$$

This approach of assessing predictive accuracy compares predicting errors between different models to a benchmark prediction of zero. This strategy is consistent with the methodology employed by GKX. The use of zero-prediction as a benchmark rather than the historical mean commonly utilized in Out of Sample R-squared calculations is justified by zero-prediction's greater success at forecasting stock returns. The  $R^2_{OOS}$  metric spans from  $-\infty$  to 100, with negative values indicating a model that performs worse than the zero-prediction benchmark. A number of zero indicates that the two forecasts have identical precision, however positive values show that the model surpasses the zero prediction, with a maximum score of 100 reflecting a perfect forecast.

Other two statistical metrics used to complete the analysis are Mean Squared Error (MSE) and Mean Absolute Error (MAE):

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2$$

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i|$$

The use of both measures allows us to examine the impact that outlier predictions have on model performance. MSE squares the errors before averaging, making it more sensitive to big deviations, which can severely penalize models with a few substantial errors. In contrast, MAE computes the average of the absolute disparities between actual and predicted values, assigning equal weight to all errors regardless of magnitude. As a result, MAE is a more robust measure of average prediction accuracy, especially in the presence of outliers. The results are stored in table 3.2.

From the point of view of forecast errors, the analysis of MAE and MSE shows that all models behave similarly. Surprisingly, ensemble methodologies, Stacking Regressor and Voting Regressor, are the ones that exhibit higher forecast errors, albeit slightly, under the point of view of both metrics. It is also crucial to emphasize that for all models the MSE is lower than the MAE: this indicates that the magnitude of most errors is relatively small, which is most likely due to the fact that there are no significant outliers or large deviations in the predictions. This can in part be explained, for Random Forest and Gradient Boosting Regressor, by the fact that the predictions generated by the models exhibit relatively low standard deviation, as shown by the table A.4 available in APPENDIX A.

In line with the results achieved by Marsì (2021) from the analysis of  $R^2_{OOS}$  metrics is it possible to infer that ML algorithms perform quite bad in explaining monthly returns of the stocks contained in the major European index. Among all,

NN1 and NN2 have the highest  $R^2 OOS$  and  $pseudo - R^2 OOS$  values, with NN1 recording the highest values of  $R^2$ , respectively of 0.26% and 0.29%, suggesting that it is able to explain some level of variance in the out-of-sample predictions and that again, neural based models are able to outperform tree-based models. Even for these models, however, the relatively low positive  $R^2$  values obtained imply poor generalization performance, if we consider the financial metrics under observation. All the other models exhibit negative  $R^2$ , meaning they perform worse than naive predictions of zero return. Although ensembles models might be useful to decrease error via model averaging, in this case they are not capturing the general, fundamental relationships across the data well, exhibiting the lowest  $R^2$ . The findings of the current study go against the ones obtained by Kelly and Xiu (2020) and Drobetz and Otto (2020). The researchers discovered modestly positive  $R^2$  values for most deployed models, including RF, GBR, and NN. They employed more predictors, which is most likely why the disparities exist. Nevertheless, the values of  $R^2$  for neural network models is relatively similar.

These findings highlight the need of including complex predictor interactions, which are inherent in neural network models but tree-model based methodologies. The results also reveal that in the monthly return situation, the upsides of "deep" learning are constrained. The two-layer model fail to improve over the single layer one.

Table 3.2. Statistical analysis metrics for each of the machine learning model

| Metric               | RF    | GBR   | NN1  | NN2  | SR    | VR    |
|----------------------|-------|-------|------|------|-------|-------|
| MSE (%)              | 1.16  | 1.15  | 1.10 | 1.11 | 1.14  | 1.18  |
| MAE (%)              | 6.90  | 6.86  | 6.72 | 6.74 | 6.85  | 7.05  |
| $R^2$ OOS (%)        | -5.25 | -4.29 | 0.26 | 0.14 | -3.38 | -6.40 |
| Pseudo $R^2$ OOS (%) | -5.23 | -4.26 | 0.29 | 0.17 | -3.35 | -6.38 |

However, comparing these indicators separately does not allow us to determine whether one model is considerably superior to another. Following Fieberg (2023) and Otto (2021), i use the popular Diebold and Mariano (DM) test (1995) to compare the prediction accuracy of two models. The null hypothesis for equal prediction errors is  $H_0 : E[dt] = 0$ , where  $dt$  represents the loss differential between two forecasts at time  $t$ . The test statistic is defined as following:

$$DM = \frac{\bar{d}_{1,2}}{\hat{\sigma}_{\bar{d}}} \sim \mathcal{N}(0,1), \quad \text{with}$$

$$\bar{d}_{1,2} = \frac{1}{T} \sum_{t=1}^T \hat{d}_{1,2,t}$$

$$\hat{d}_{1,2,t} = \frac{1}{N} \sum_{i=1}^N (e_{1,i,t}^2 - e_{2,i,t}^2)$$

$$\hat{\sigma}_{\bar{d}} = \sqrt{\frac{1}{T} \sum_{t=1}^T (\hat{d}_t - \bar{d})^2}$$

Table 3.3 presents the test statistic for the DM test on the full sample of stocks. \*, \*\*, \*\*\* represents respectively statistical significance at 10%, 5% and 1% level. The models listed in the rows are compared against the models listed in the columns. A positive value suggests that the model in the column performs better than the model in the row, on average, according to the DM test.

Table 3.3. *T-test of pairwise Diebold-Mariano (DM) tests. This table reports t-test of pairwise Diebold-Mariano (DM) tests comparing stock-level return forecasts of various models against the null hypothesis of equal predictive accuracy. Positive numbers show that the column model outperforms the row model. \*\*\* measures the levels of statistical significance, 10%, 5% and 1%.*

|     | RF | GB     | NN1     | NN2     | SR       | VR       |
|-----|----|--------|---------|---------|----------|----------|
| RF  | –  | -1.81* | 4.22*** | 4.17*** | 0.95     | -2.47**  |
| GB  |    | –      | 4.02*** | 3.94*** | 1.43     | -0.47    |
| NN1 |    |        | –       | -1.2    | -4.15*** | -4.50*** |
| NN2 |    |        |         | –       | -40.3*** | -4.53*** |
| SR  |    |        |         |         | –        | -0.8     |
| VR  |    |        |         |         |          | –        |

The DM test validates the empirical conclusions generated from the  $R^2$  values shown in Table 3.2. The first conclusion to draw is that neural network models, according to the DM test, can outperform tree-based and ensemble models at 1% statistical significance, with NN1 remaining the highest performing model. In the realm of tree-based models, RF appears to be the most valuable, outperforming GRB at the 10% statistical significance level and the stacking regressor model at the 5% level. For other pairwise comparisons, the null hypothesis of identical prediction accuracy cannot be rejected.

### 3.2 Variable Importance

Subsequently, although the various forecast models exhibit similar levels of complexity (at least within each model family), it is important to explore whether different techniques highlight distinct predictors as most significant for predicting future returns. To assess the importance of specific variables, I adopt the methodology introduced by Kelly et al. (2019), which entails computing the variable importance matrix for each model in two stages. Initially, the absolute variable importance is determined by measuring the decrease in  $R^2$  when a

particular predictor's values are all set to zero in the training dataset. Subsequently, I standardize the absolute variable importance values to ensure they sum up to one, indicating the relative contribution of each variable to the model. Figure 3.2 depicts the relative variable importance measures for each forecast model separately.

On average, in line with Cakici (2023) and Drobetz and Otto (2021) I find that the various forecast models identify similar variables as the most valuable (the momentum features,  $dp$ ,  $dy$ ,  $eps\_5y$ , etc.). However, some models concentrate on a small number of factors. In aversion to Drobetz and Otto (2021) I find that neural network models give more importance to a smaller number of predictors, opposed to tree-based models. Furthermore, the same sort of models has a very close pattern of varying relevance, in terms of which predictors and the behavior of their concentration. For what concerns momentum features, my findings resemble the ones of Cakici (2023) on the European stock market and the ones of Gu et al. (2020) on the US market. For example, short-term-reversal ( $mom\_1m$ ) scored second and first respectively in the above-mentioned researches, while in the present studies it ranks third, as it is exhibited in Figure 3.3. Apart from this, many other financial indicators do not seem to have much relevance for the models considered. This may be surprising as, for example, related to ROE, financial evaluation models of companies have been developed based mainly on this measure and how it changes over time. A similar argument can be made for the operating profit margin or the gross profit margin, indicators that are often taken into account to calculate the true value of a company and consequently its stock price. The reason why they are not particularly important may be the type of data used (monthly) or the time frame considered. Also worth mentioning is the importance of the dividend payout ratio and dividend yield, which have been shown (by means of linear regression studies) to be important for investors and thus able to explain stock returns, as mentioned above in the Literature review section.

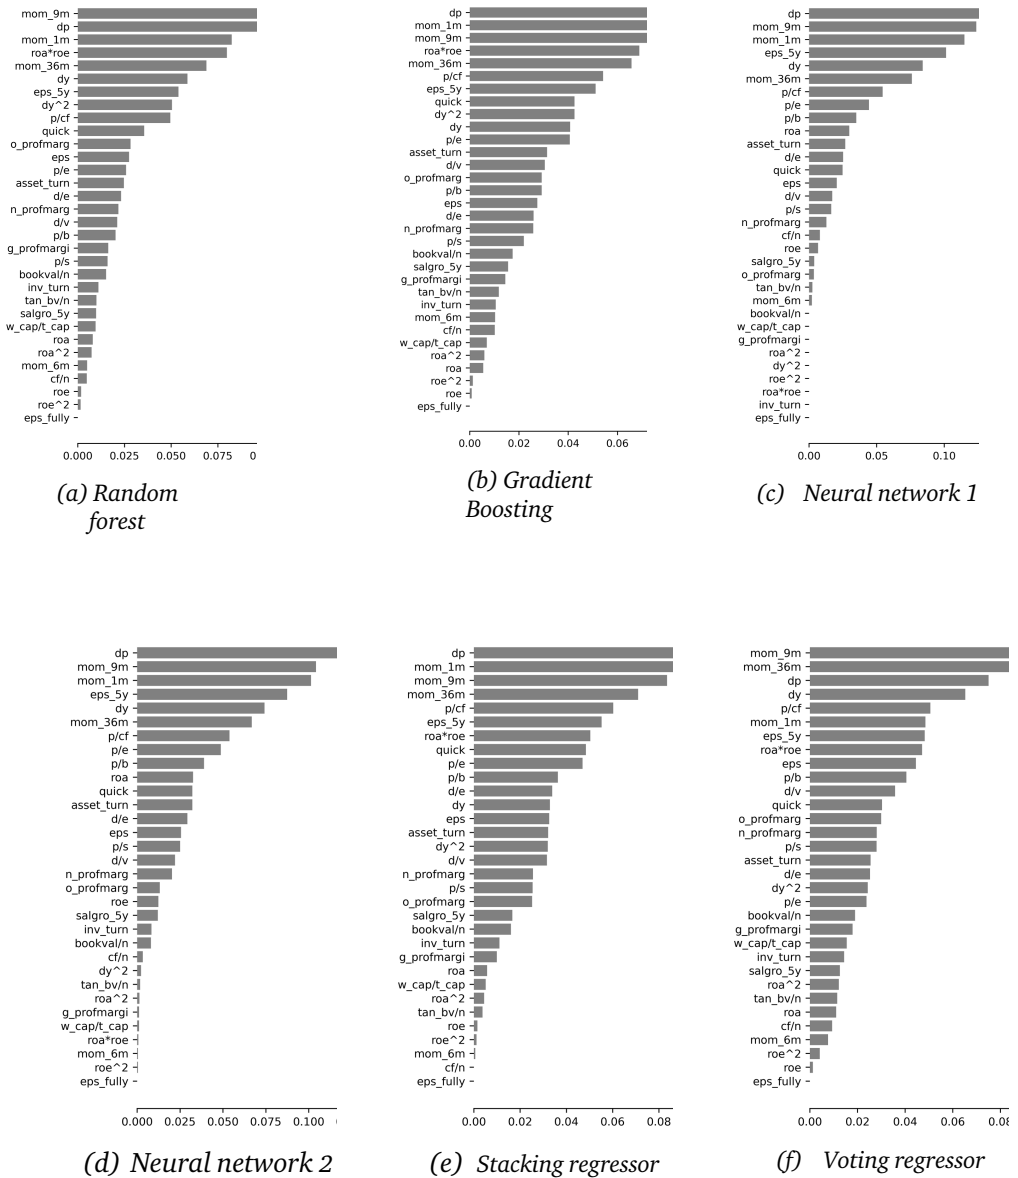


Figure 3.2: Model-specific relative variable importance. This figure depicts the time series average of relative variable importance indicators of each machine learning model

To improve inter-model comparability, i rank the average relative contribution of each variable within a single model and add the ranks across all models to get an overall rank (greater variable importance = higher rank). Figure 3.3 depicts a

heat map showing relative variable relevance rankings. The rows are arranged in descending order according to overall rank. As a result, the higher a variable is ranked, the more significant it is overall. Darker cell colors indicate that the associated variable is more important to the specific model.




Figure 3.3: Variable importance heatmap. This figure depicts a heat map of the average relative variable importance ranks of each machine learning model described in the "Forecast models" section. The rankings are obtained by ranking the relative contribution of each variable within a specific model and summing the ranks across all models to obtain an overall rank. The rows are arranged in descending order according to overall rank. Darker cell colors indicate that the associated variable is more important to the model.

### 3.3 Portfolio Performances

Statistical performance is of secondary relevance to an investor. According to Leitch and Tanner (1991), statistical measurements are only loosely related to the economic potential gain of a forecast. As a result, it is useful to determine if differences in statistical predictive performance transfer into differences in predictive capacity from an economic standpoint for a realistic trading strategy.

In order to do so, i assess the economic performance of our models by portfolio sorts, designing a new set of portfolios to directly, exploit machine learning forecasts. For each stock in our dataset, machine learning models generate predictions of future returns on a monthly basis. Using these predictions, i sort stocks into deciles. At the end of each month, i sort stocks into deciles based on their predicted returns from the machine learning models. The stocks in decile 1 (Low) are those with the lowest predicted returns, while the stocks in decile 10 (High) are those with the highest predicted returns. Each decile represents a portfolio of stocks ranked by their expected performance for the following month, as forecasted by the models. Once the stocks are sorted into deciles, i calculate average return, realized return, standard deviation and Sharpe ratio, on a monthly basis. These metrics help to evaluate how well the machine learning models' predictions align with the actual stock returns. In addition to evaluating the individual deciles, a High-Low (H-L) strategy is constructed. This strategy involves creating a long-short portfolio by going long on the stocks in decile 10 (those predicted to have the highest returns) and short on the stocks in decile 1 (those predicted to have the lowest returns). The results are display in table 3.4.

Table 3.4. Decile Portfolios based on predicted expected returns from each model. This table presents the time series averages of predicted and realized returns of decile portfolios based on the expected returns obtained from each machine learning model

|           | NN2   |       |       |       | VR    |      |       |       | SR    |       |       |       |
|-----------|-------|-------|-------|-------|-------|------|-------|-------|-------|-------|-------|-------|
|           | Pred  | Avg   | SD    | SR    | Pred  | Avg  | SD    | SR    | Pred  | Avg   | SD    | SR    |
| 1 (low)   | -2.11 | -0.86 | 6.49  | -0.15 | -6.31 | 0.17 | 6.85  | 0.01  | -3.91 | -0.36 | 7.24  | -0.06 |
| 2         | -0.85 | -0.54 | 6.11  | -0.11 | -3.22 | 0.28 | 6.69  | 0.03  | -1.63 | 0.23  | 6.43  | 0.02  |
| 3         | -0.29 | -0.38 | 6.31  | -0.08 | -1.99 | 0.36 | 6.28  | 0.04  | -0.87 | 0.26  | 6.42  | 0.03  |
| 4         | 0.07  | 0.32  | 6.17  | 0.04  | -1.03 | 0.08 | 6.16  | -0.00 | -0.33 | 0.32  | 6.45  | 0.03  |
| 5         | 0.36  | 0.53  | 5.76  | 0.07  | -0.20 | 0.09 | 6.23  | -0.00 | 0.09  | 0.40  | 6.08  | 0.05  |
| 6         | 0.62  | 0.68  | 6.08  | 0.10  | 0.58  | 0.57 | 6.31  | 0.07  | 0.46  | 0.31  | 6.02  | 0.04  |
| 7         | 0.89  | 0.93  | 6.25  | 0.13  | 1.31  | 0.43 | 6.03  | 0.05  | 0.82  | 0.50  | 5.73  | 0.07  |
| 8         | 1.22  | 0.83  | 6.40  | 0.11  | 2.17  | 0.63 | 6.33  | 0.08  | 1.18  | 0.47  | 5.90  | 0.06  |
| 9         | 1.66  | 1.24  | 6.43  | 0.18  | 3.35  | 0.70 | 6.25  | 0.10  | 1.61  | 0.97  | 6.44  | 0.13  |
| 10 (high) | 2.68  | 1.68  | 6.56  | 0.24  | 6.16  | 1.09 | 6.00  | 0.16  | 2.58  | 1.30  | 6.29  | 0.19  |
| H-L       | 4.80  | 2.55  | 13.05 | 0.39  | 12.47 | 0.92 | 12.86 | 0.15  | 6.49  | 1.66  | 13.52 | 0.25  |

|           | NN1   |       |       |       | RF    |       |       |       | GBR   |       |       |       |
|-----------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|
|           | Pred  | Avg   | SD    | SR    | Pred  | Avg   | SD    | SR    | Pred  | Avg   | SD    | SR    |
| 1 (low)   | -2.00 | -0.82 | 6.36  | -0.14 | -3.65 | -0.29 | 7.15  | -0.05 | -3.16 | -0.26 | 7.13  | -0.05 |
| 2         | -0.70 | -0.56 | 6.16  | -0.11 | -1.71 | 0.15  | 6.61  | 0.01  | -1.23 | -0.07 | 6.36  | -0.03 |
| 3         | -0.13 | -0.03 | 6.23  | -0.02 | -1.03 | 0.30  | 6.38  | 0.03  | -0.65 | 0.50  | 6.17  | 0.07  |
| 4         | 0.24  | 0.07  | 6.18  | -0.01 | -0.54 | 0.23  | 6.36  | 0.02  | -0.26 | 0.33  | 6.46  | 0.04  |
| 5         | 0.53  | 0.41  | 6.14  | 0.05  | -0.14 | 0.32  | 6.10  | 0.04  | 0.07  | 0.25  | 6.30  | 0.02  |
| 6         | 0.78  | 0.43  | 5.85  | 0.06  | 0.21  | 0.65  | 6.07  | 0.09  | 0.34  | 0.21  | 6.30  | 0.02  |
| 7         | 1.02  | 0.94  | 6.26  | 0.13  | 0.54  | 0.35  | 5.73  | 0.04  | 0.59  | 0.42  | 5.80  | 0.06  |
| 8         | 1.26  | 1.03  | 6.34  | 0.15  | 0.89  | 0.44  | 6.00  | 0.06  | 0.83  | 0.70  | 6.08  | 0.10  |
| 9         | 1.57  | 1.23  | 5.93  | 0.19  | 1.31  | 0.80  | 6.12  | 0.11  | 1.12  | 1.02  | 5.97  | 0.15  |
| 10 (high) | 2.35  | 1.75  | 6.74  | 0.24  | 2.23  | 1.43  | 6.47  | 0.21  | 1.84  | 1.31  | 6.49  | 0.19  |
| H-L       | 4.34  | 2.56  | 13.10 | 0.39  | 5.89  | 1.72  | 13.62 | 0.26  | 5.00  | 1.57  | 13.62 | 0.24  |


Out-of-sample portfolio performance is remarkably similar to the previously disclosed results on machine learning forecast accuracy. Machine learning forecasts from all methods grows in proportion with realized returns monotonically with exception of decile 8 of NN2, decile 4, 5,7 of VR, decile 6 and 8 of SR, from 3 to 6 of GRB and 4, 7 and 8 of RF. Neural network methods once again outperform tree-based techniques and ensemble techniques due to the similarity of decile predictions with realized returns. In addition, for all except the most extreme deciles, the quantitative fit between anticipated returns and average realized returns employing neural networks is quite close, in particular for NN1. The greatest High – Low strategy is from NN1, which returns on average 2.56% every month. Its monthly volatility is 13.10% resulting in a monthly out-of-sample Sharpe ratio of 0.39%. Another remarkable high – low strategy is the one from NN2, with a monthly average return of 2.55%, a monthly standard deviation of 13.05% and a Sharpe ratio of 0.39%. The discussed techniques capture cross-sectional diversity in realized excess returns, although their ability to provide effective trading signals varies significantly.

I will deploy four different ways to assess the economic significance of the machine learning models used. First, will be create two 'long only' portfolios. The originated portfolios take long positions in equities that the model predicts will have the greatest returns, ignoring low deciles. Two portfolios will be created using solely long positions and an equal weight criterion: the top 10% portfolio, which selects the top 10% of the best performing stocks, and the top 30% portfolio, which selects the top 30% of the best performing stocks. This twofold strategy allows to see whether there are economic gains even when evaluating a bigger universe of stocks to invest in (i.e. the top 30%). If the universe of equity in which to invest is limited, high-capital trading techniques may face liquidity issues and huge bid-ask spreads, rendering the strategy unprofitable. Table 3.5 summarizes the performance metrics for the two long methods.


*Table 3.5. Risk and return characteristics of top 10% and top 30% long strategy*

| Metrics                   | Top 10% |      |      |      |      |      | Top 30% |      |      |      |      |      | MKT  |
|---------------------------|---------|------|------|------|------|------|---------|------|------|------|------|------|------|
|                           | RF      | GBR  | NN1  | NN2  | VR   | SR   | RF      | GBR  | NN1  | NN2  | VR   | SR   | MKT  |
| Cumulative Return         | 1.94    | 1.89 | 2.51 | 2.33 | 1.73 | 1.82 | 1.63    | 1.81 | 2.15 | 2.03 | 1.62 | 1.66 | 1.23 |
| Annual Return             | 0.14    | 0.13 | 0.18 | 0.17 | 0.11 | 1.12 | 0.10    | 0.12 | 0.15 | 0.14 | 0.10 | 0.10 | 0.05 |
| Annual Standard Deviation | 0.19    | 0.19 | 0.18 | 0.18 | 0.18 | 0.18 | 0.17    | 0.16 | 0.17 | 0.17 | 0.16 | 0.17 | 0.17 |
| Annual Sharpe Ratio       | 0.73    | 0.70 | 1.02 | 0.91 | 0.63 | 0.67 | 0.60    | 0.73 | 0.89 | 0.82 | 0.60 | 0.62 | 0.27 |

Figure 3.4 shows the results of table 3.5. It plots the cumulative performance of expected return-sorted portfolios with monthly restructuring of the market portfolio and of each of the portfolios constructed using the machine learning models described in the section Empirical Analysis, for both the top 30% and top 10% long only strategy.



(a) Top 10%



(b) top 30%

Figure 3.4. Cumulative performance of expected return-sorted portfolios with monthly restructuring – Long only strategy. This figure shows the performances of the top 10% and top 30% strategy going long on the best performing stocks. Each month the equal weights portfolios are rebalanced to match the best performing deciles.


In line with Leitch and Tanner (1993), it is observed that the poor statistical performance achieved in the preceding sections are not always a sign of poor economic performance of the machine learning models under consideration. It can be seen that all models in both strategies outperform the market portfolio.

Again, neural models are proved superior to tree-based models and ensemble techniques, as also anticipated by the decile construction. In particular, given the greater proximity to realized returns, the neural models are able to more accurately predict which stocks to select each month as top performers, both in the case of the top 10% and the top 30% strategy. NN1 has the highest cumulative return of 2.51, followed by NN2 with 2.33, which is considerably larger than the market's cumulative return of 1.23. Furthermore, the Sharpe ratios of these models beat that of the market: 1.02 for NN1 and 0.91 for NN2, compared to the market Sharpe ratio of only 0.27. All models had an annual volatility of roughly 0.18, which is in line with the market volatility, indicating that these models do not introduce undue risk while providing higher returns. The outperformance shown here highlights the power of machine learning techniques and, in particular, neural networks in picking up more sophisticated patterns in the data than traditional models or market strategies. It is worth mentioning that even tree-based models, had a rather good performance, too, with cumulative returns of 1.94 and 1.89, respectively. Even in the context of economic performances analysis, the ensemble techniques seem to perform the worst with cumulative returns of and 1.73 and 1.83, respectively, but still above the market return. When the algorithm selects the top 3 deciles, with the op 30% strategy, returns drop, but still remain above the market portfolio. Again, the neural models prove superior, with NN1 achieving a cumulative return of 2.15 and a Sharpe ratio of 0.82. Overall, the standard deviation levels remain the same, resulting in lower Sharpe ratio levels for all the models considering the lower returns achieved. It is interesting to note that when increasing the investment pool from 10% to 30%, machine learning models seems to converge in performance. In this context, the cumulative returns of Random Forest, Gradient Boosting, and Neural Networks are much closer to each other. This suggests that with more stocks in the investment pool, the predictiveness of the machine learning models becomes somewhat watered down.


In addition to the two long only strategies, there are two mirror strategies that allow investors to take short positions. The main reasons for distinguishing between these two types of strategies are as follows:

- It is possible that some investors are not in a position to take short positions on equities: therefore, the aim of this analysis is also to test whether the use of these models is valid for an average investor with limited access to economic resources and trading strategies.
- It's crucial to test the quality of the predictions made by the models in the event of a bear market. If the cumulative return on the use of short positions is lower than the corresponding strategy when the only option is to go long, this means that the model is not capable of correctly

predicting downward stock movements, but is more likely to correctly predict a return when the latter is positive.



(a) Top-worst 10%



(b) Top-worst 30%

Figure 3.5. Cumulative performance of expected return-sorted portfolios with monthly restructuring – Long short strategy. This figure shows the performances of the top-worst 10% and top-worst 30% strategy going long on the best performing stocks and short on the worst performing stocks. Each month the equal weights portfolios are rebalanced to match the best and the worst performing deciles.

Two distinct portfolios will be constructed for the long-short strategy. The first will be a top-worst 10% portfolio, comprising the 10% of the best-performing stocks and the 10% of the worst-performing stocks. The second will be a top-worst 30% portfolio, comprising the 30% of the best-performing stocks and the 30% of the worst-performing stocks. It is evident that the strategy will entail going long in the top percentage and short in the worst. The results for each model are summarized in Table 3.6, with particular focus on the contribution that the short position, labelled 'Annual short return', makes to the total cumulative return of the portfolio. Figure 3.5 shows the results of table 3.6. It plots the cumulative performance of expected return-sorted portfolios with monthly restructuring of the market portfolio and of each of the portfolios constructed using the machine learning models described in the section Empirical Analysis, for both the top-worst 30% and top-worst 10% long and short strategy.

Table 3.6. Risk and return characteristics of top-worst 10% and top-worst 30% long short strategy

| Metrics                   | Top-worst 10% |      |       |       |       |      | Top-worst 30% |      |       |       |      |      | MKT  |
|---------------------------|---------------|------|-------|-------|-------|------|---------------|------|-------|-------|------|------|------|
|                           | RF            | GBR  | NN1   | NN2   | SR    | VR   | RF            | GBR  | NN1   | NN2   | SR   | VR   | MKT  |
| Cumulative Return         | 2.16          | 1.97 | 3.63  | 3.62  | 2.01  | 1.43 | 1.47          | 1.62 | 2.67  | 2.73  | 1.49 | 1.22 | 1.23 |
| Annual Return             | 0.14          | 0.13 | 0.24  | 0.24  | 0.13  | 0.07 | 0.07          | 0.09 | 0.18  | 0.18  | 0.08 | 0.08 | 0.05 |
| Annual Standard Deviation | 0.11          | 0.12 | 0.14  | 0.12  | 0.12  | 0.14 | 0.09          | 0.09 | 0.11  | 0.10  | 0.09 | 0.08 | 0.18 |
| Annual Sharpe Ratio       | 1.27          | 1.08 | 1.77  | 1.96  | 1.07  | 0.54 | 0.85          | 0.72 | 1.63  | 1.92  | 0.83 | 0.46 | 0.27 |
| Annual Long Return        | 0.14          | 0.13 | 0.18  | 0.17  | 0.12  | 0.11 | 0.10          | 0.12 | 0.15  | 0.14  | 0.11 | 0.10 | -    |
| Annual Short Return       | -0.01         | 0.00 | -0.06 | -0.07 | -0.01 | 0.04 | 0.03          | 0.03 | -0.03 | -0.04 | 0.03 | 0.06 | -    |

In the top-worst 10% strategy shown in figure 3.5, neural networks models are again performing significantly better than the market and other machine learning models. The cumulative return for NN1 peaked 3.63, 3.62 for NN2. These returns are significantly higher than the market benchmark of 1.23, indicating that the capabilities to select top performers and correctly short underperformers. The annual short returns, (that are the contribution to the cumulative performance of going shorts on the selected stocks) for NN1 and NN2 are -0.06 and -0.07, respectively, signaling that the models are able to identify the worst stock performers and, hence, generate profits from shorting positions. This clearly outperforms the RF and GBR models, which have near-zero annual short returns of -0.01 and 0.00, respectively. The latter models were more conservative in their shorting strategy by opening positions with limited downside but also limiting the potential gain from short positions. This due to their lack of ability to predict returns accurately in line with the realized ones, as it was shown in table 3.4. The Sharpe ratios demonstrate the improvements with respect to the long only strategy. NN1 has a Sharpe ratio of 1.77, which is significantly greater than the 1.02 value achieved in the long-only strategy and far above the market's Sharpe ratio of 0.27. This increase in the Sharpe ratio supports the concept that short positions boost the risk-adjusted return of neural networks through higher returns as well as risk diversification. Ensemble models,

such as the Stacking and Voting ensemble models, perform noticeably worse, with cumulative returns of 2.01 and 1.43, respectively, and lower, but still high Sharpe ratios. This underperformance might partly be because of their generally weaker shorting ability; the Voting Ensemble has a positive short return of 0.04, indicating incorrect selection of the 10% worst performing stocks. Figure 3.5 shows that expanding the investment pool to the top-worst 30%, reduces the cumulative returns for most of the models, though neural networks still ranked first among all models. The annual short returns remain overall negative, with the models still apparently effective in their choice of stocks to short against. NN1 and NN2 report short returns of -0.03 and -0.04, respectively, which again shows that even within a wider universe of stocks, the neural networks are able to select those stocks that would ultimately will depicts negative returns. On the other hand, the short returns for both Random Forest and Gradient Boosting are slightly positive at 0.03, which suggests they are incorrectly shorting the stocks that instead have an increase in stock price. This weakness on the short side explains their relatively lower cumulative returns of 1.62 and 1.67, respectively, compared to the neural networks. The worst-performing model in this case is the Voting Regressor, which, due to an annual short return contribution of 0.06, underperforms the market portfolio, but with a lower standard deviation of 0.1%.

### Transaction costs

Transaction expenses reduce the profitability of any investment. Investing strategies with high turnover rates may not outperform in terms of returns. Realistic transaction costs can be difficult to obtain, time-consuming, and expensive (Lesmond et al. (1999)). Collins and Fabozzi (1991) argue that genuine transaction costs are intrinsically unobservable. Therefore, estimating transaction costs is likely to be erroneous. The present study considers the average monthly turnover of the top 10% and top 30% of the investment strategies, both in the long-only and long-short strategy. The turnover is calculated as follows:

$$\text{Turnover} = \frac{1}{T} \sum_{t=1}^T \left( \sum_i \left| w_{i,t+1} - \frac{w_{i,t}(1 + r_{i,t+1})}{\sum_j w_{j,t}(1 + r_{j,t+1})} \right| \right),$$

where  $w_{i,t}$  is the weight of stock  $i$  in the portfolio at time  $t$ .

Furthermore, in accordance with the methodology proposed by Cakici et al. (2023), i also calculated the breakeven trading costs as the mean portfolio return divided by the mean turnover. Breakeven trading costs are those that refer to the amount beyond which a trading strategy remains profitable. This is extremely valuable in strategies with heavy rebalancing or high turnover; wherein modest transaction costs can shrink net returns considerably. Comparing breakeven trading costs with the actual or estimated trading costs in the market allows an

investor or the portfolio manager to see how robust and practically applicable the trading strategy may be beyond its theoretically computed returns.

The results are showed in table 3.7 for both the long only and the long-short strategy.

Table 3.7. Turnover and Break-even trading cost for the long only strategies and long short strategies

| Metrics                  | L-Top 10% |      |      |      |      |      | L-Top 30% |      |      |      |      |      |
|--------------------------|-----------|------|------|------|------|------|-----------|------|------|------|------|------|
|                          | RF        | GBR  | NN1  | NN2  | VR   | SR   | RF        | GBR  | NN1  | NN2  | VR   | SR   |
| Avg monthly Turnover (%) | 76        | 74   | 59   | 60   | 61   | 75   | 67        | 66   | 56   | 57   | 58   | 66   |
| BE Transaction Costs (%) | 1.49      | 1.47 | 2.52 | 2.33 | 1.55 | 1.33 | 1.26      | 1.51 | 2.21 | 2.05 | 1.41 | 1.33 |

| Metrics                  | LS-Top-worst 10% |      |      |      |      |      | LS-Top-worst 30% |      |      |      |      |      |
|--------------------------|------------------|------|------|------|------|------|------------------|------|------|------|------|------|
|                          | RF               | GBR  | NN1  | NN2  | VR   | SR   | RF               | GBR  | NN1  | NN2  | VR   | SR   |
| Avg monthly Turnover (%) | 144              | 142  | 118  | 121  | 124  | 143  | 132              | 130  | 112  | 115  | 116  | 130  |
| BE transaction Costs (%) | 0.82             | 0.74 | 1.69 | 1.62 | 0.49 | 0.76 | 0.45             | 0.58 | 1.34 | 1.33 | 0.28 | 0.48 |

In the L-Top 10% approach, neural networks have a considerably lower turnover of about 59% for NN1 and 60% for NN2, compared to others such as RF and GBR, which have turnovers of around 74-76%. Lower levels of such turnover would indicate that the conducted neural networks had more stable trading activity than the other models, supporting prior assessments that found greater performance and greater cumulative returns. These models also have higher break-even transaction costs (2.52% for NN1 and 2.33% for NN2), allowing them to absorb more transaction fees without losing profitability. This is because to their higher returns, which allow them to better absorb the impact of trading expenses than models with significantly lower returns, such as RF and GBR. The top-worst 10% long-short strategy is more complex because it manages both long and short positions; hence, all turnover increases. For neural network models, turnover rises to 118-121%, and for Random Forest and Gradient Boosting Regressor, it equals 144-142%, respectively. In opposition to Cakici (2023) and in line with Fieberg et al. (2023), neural network models show the lowest scores, while tree-based models the highest ones. Despite the high turnovers, the break-even transaction costs remain high for neural networks: 1.69% for NN1 and 1.62% for NN2, confirming their dominance in overall performance. Models that exhibited difficulties in handling short positions, such as the Voting Ensemble, show very low break-even transaction costs of 0.49%, underlining its vulnerability to transaction costs in high-turnover environments. The top-worst 30% strategies show similar trends: the overall turnover is lower, but still higher for the long-short version with respect to the long-only. Neural networks continue to highly outperform due to stronger returns, while tree-based models and ensemble techniques remain more cost-sensitive.

## CONCLUSION

This research has demonstrated that machine learning models, especially neural networks, hold significant advantages in predicting stock returns within the European equity market. It is presented that suitably tuned ML models outperform basic market-wide heuristics portfolios on long only and long-short strategies, even if a limited number of predictors is considered. The results point to the possibility that machine learning can successfully uncover complex nonlinear relationships in financial data, which would ensure enhanced predictive efficiency and better performance in portfolios creation. The study hereby provides evidence that machine learning models can scan achieve notable risk-adjusted outperformance in European markets, with a relatively low statistical significance from the out-of-sample tests. Given the relatively small sample of firms, the constrained time span, and the limited set of predictors in number and nature, this research contributes to the existing body of literature related to asset pricing using machine learning in European markets. In all the aspects of the analysis, neural network methods outperformed consistently compared with Ensemble techniques and tree-based methods. The scope of this research could be extended considering the following concepts:

- implementation of transaction costs for portfolios construction;
- consideration of stock liquidity and size in portfolios construction;
- rolling or expanding window data split, while considering a limited sample size and number of predictors;

## APPENDIX A

**Additional Tables and Figures****Empirical Analysis***Table A.1. NACE Codes and Descriptions*

| <b>NACE Code</b> | <b>Description</b>   |
|------------------|--|
| A                | Agriculture, Forestry and Fishing  |
| B                | Mining and Quarrying   |
| C                | Manufacturing  |
| D                | Electricity, Gas, Steam and Air Conditioning Supply                        |
| E                | Water Supply; Sewerage, Waste Management and Remediation Activities        |
| F                | Construction   |
| G                | Wholesale and Retail Trade; Repair of Motor Vehicles and Motorcycles       |
| H                | Transportation and Storage   |
| I                | Accommodation and Food Service Activities                                  |
| J                | Information and Communication  |
| K                | Financial and Insurance Activities   |
| L                | Real Estate Activities   |
| M                | Professional, Scientific and Technical Activities                          |
| N                | Administrative and Support Service Activities                              |
| O                | Public Administration and Defense; Compulsory Social Security              |
| P                | Education  |
| Q                | Human Health and Social Work Activities                                    |
| R                | Arts, Entertainment and Recreation   |
| S                | Other Service Activities   |
| T                | Activities of Households as Employers; Undifferentiated Goods and Services |
| U                | Activities of Extraterritorial Organizations and Bodies                    |




Figure A.1. Predictors correlation matrix

Table A.2. Predictors time-series correlations with returns

| Predictor   | Corr /w return |
|-------------|----------------|
| roa         | 0.035049       |
| o_profmargi | 0.033561       |
| eps_basic   | 0.028434       |
| eps_fully   | 0.028329       |
| p/s         | 0.024785       |
| eps_5y      | 0.024633       |
| asset_turn  | 0.019406       |
| roe         | 0.017956       |
| quick       | 0.017420       |
| roic        | 0.016616       |
| p/e         | 0.015837       |
| current     | 0.015732       |
| g_profmargi | 0.014576       |
| salgro_5y   | 0.007344       |
| w_cap/t_cap | 0.006884       |
| inv_turn    | 0.006158       |
| p/cf        | 0.004585       |
| p/b         | 0.004499       |
| cf/n        | 0.002114       |
| d/v         | -0.001920      |
| n_profmargi | -0.005147      |
| d/e         | -0.007664      |
| tan_bv/n    | -0.008129      |
| bookval/n   | -0.011999      |
| eps         | -0.013412      |
| dp          | -0.025744      |
| dy          | -0.167749      |

Table A.3. Variable importance factors. This table is constructed by the studying the correlation with returns of the factors that are still in the dataset after the elimination based on the correlation matrix. Since there are no values higher than 5 these are the predictors that will be deployed in the analysis

| Predictor       | VIF      |
|-----------------|----------|
| bookval/n       | 4.964133 |
| roa             | 3.861362 |
| tan_bv/n        | 2.702279 |
| cf/n            | 2.512495 |
| roe             | 2.334957 |
| eps_fully       | 2.278092 |
| p/b             | 1.804407 |
| asset_turn      | 1.577246 |
| p/s             | 1.459381 |
| g_profmargi     | 1.394601 |
| w_cap/t_cap     | 1.268088 |
| dy              | 1.259769 |
| o_profmargi     | 1.247353 |
| p/cf            | 1.231761 |
| dp              | 1.206332 |
| d/e             | 1.168680 |
| eps_5y          | 1.162011 |
| d/v             | 1.106803 |
| eps             | 1.068407 |
| salgro_5y       | 1.049787 |
| n_profmargi     | 1.044080 |
| p/e             | 1.043180 |
| monthly_returns | 1.033329 |
| inv_turn        | 1.002257 |

Table A.4. Summary statistics for Market, Random Forest, Gradient Boosting, and Neural Networks

| Metric                 | Actual (%) | RF (%) | GBR (%) | NN 1 (%) | NN 2 (%) |
|------------------------|------------|--------|---------|----------|----------|
| Mean Return (Monthly)  | 0.44       | 0.20   | 0.16    | 0.77     | 0.43     |
| Std Dev (Monthly)      | 8.34       | 3.43   | 3.22    | 5.92     | 5.94     |
| Sharpe Ratio (Monthly) | 4.06       | 2.84   | 1.99    | 11.37    | 5.47     |
| Mean Return (Annual)   | 5.27       | 2.37   | 1.97    | 9.29     | 5.10     |
| Std Dev (Annual)       | 28.90      | 11.87  | 11.15   | 20.52    | 20.57    |
| Sharpe Ratio (Annual)  | 14.07      | 9.85   | 6.89    | 39.40    | 18.96    |

## REFERENCES

- [1] E. F. Fama. “Efficient capital markets”. In: *Journal of finance* 25.2 (1970), pp. 383–417.
- [2] P. A. Samuelson. “Proof that properly discounted present values of assets vibrate randomly”. In: *The Bell Journal of Economics and Management Science* (1973), pp. 369–374.
- [3] M. G. Kendall. “The Analysis of Economic Time Series—Part 1: Prices”. In: *Journal of the Royal Statistical Society. Series A (General)* 116.1 (1953), pp. 11–34.
- [4] A. Cowles. “A Revision of Previous Conclusions Regarding Stock Price Behavior”. In: *Econometrica* 28.4 (1960), pp. 909–915.
- [5] M. F. Osborne. “Brownian motion in the stock market”. In: *Operations research* 7.2 (1959), pp. 145–173.
- [6] E. F. Fama. “The behavior of stock-market prices”. In: *The journal of Business* 38.1 (1965), pp. 34–105.
- [7] E. F. Fama. “Efficient capital markets: II”. In: *The journal of finance* 46.5 (1991), pp. 1575–1617.
- [8] E. F. Fama and K. R. French. “Common risk factors in the returns on stocks and bonds”. In: *Journal of financial economics* 33.1 (1993), pp. 3–56.
- [9] F. Black, M. C. Jensen, and M. Scholes. “The Capital Asset Pricing Model: Some Empirical Tests”. In: *Studies in the Theory of Capital Markets* (1972). Ed. by M. C. Jensen, pp. 79–121.
- [10] M. S. Rozeff. “Dividend yields are equity risk premiums”. In: *Journal of Portfolio management* (1984), pp. 68–75.
- [11] J. Y. Campbell and R. J. Shiller. “The dividend-price ratio and expectations of future dividends and discount factors”. In: *The review of financial studies* 1.3 (1988), pp. 195–228.
- [12] J. H. Cochrane. “Production-Based Asset Pricing and the Link Between Stock Returns and Economic Fluctuations”. In: *Journal of Finance* 46.1 (1991), pp. 209–237.
- [13] W. N. Goetzmann and R. G. Ibbotson. “The Performance of Mutual Funds in the Period 1965–1984”. In: *Journal of Business* 66.2 (1993), pp. 133–156.
- [14] R. J. Hodrick. “Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement”. In: *Review of Financial Studies* 5.3 (1992), pp. 357–386.
- [15] J. Lewellen. “Predicting Returns with Financial Ratios”. In: *Journal of Financial Economics* 74.2 (2004), pp. 209–235.
- [16] S. Basu. “Investment Performance of Common Stocks in Relation to Their Price-Earnings Ratios: A Test of the Efficient Market Hypothesis”. In: *Journal of Finance* 32.3 (1977), pp. 663–682.
- [17] A. Chattopadhyay, M. R. Lyle, and C. C. Wang. “Accounting data, market values, and the cross section of expected returns worldwide”. In: *Harvard Business School Accounting & Management Unit Working Paper 15-092* (2016).

- [18] S. Kheradyar, I. Ibrahim, and F. M. Nor. "Stock return predictability with financial ratios". In: *International Journal of Trade, Economics and Finance* 2.5 (2011), p. 391.
- [19] X. Chen, K. A. Kim, T. Yao, and T. Yu. "On the predictability of Chinese stock returns". In: *Pacific-Basin Finance Journal* 18.4 (2010), pp. 403–425.
- [20] C. Venkates, M. Tyagi, and L. Ganesh. "Fundamental analysis and stock returns: An Indian evidence". In: *Global Advanced Research Journal of Economics, Accounting and Finance* 1.2 (2012), pp. 033–039.
- [21] I. N. Agustin. "The integration of fundamental and technical analysis in predicting the stock price". In: *Jurnal Manajemen Maranatha* 18.2 (2019), pp. 93–102.
- [22] M. B. Khan, S. Gul, S. U. Rehman, N. Razzaq, and A. Kamran. "Financial ratios and stock return predictability (Evidence from Pakistan)". In: *Research Journal of Finance and Accounting* 3.10 (2012), pp. 1–6.
- [23] B. Lev and S. R. Thiagarajan. "Fundamental information analysis". In: *Journal of Accounting Research* 31.2 (1993), pp. 190–215.
- [24] U. S. Akbar and N. A. Bhutto. "Predictability of Stock Returns using Financial Ratios: Evidence from Pakistan Stock Exchange". In: *Pakistan Research Journal of Social Sciences* 2.3 (2023).
- [25] M. Dempsey. "The book-to-market equity ratio as a proxy for risk: evidence from Australian markets". In: *Australian Journal of Management* 35.1 (2010), pp. 7–21.
- [26] A. J. Hatta and B. S. Dwiyanto. "The company fundamental factors and systematic risk in increasing stock price". In: *Journal of Economics, Business, and Accountancy Ventura* 15.2 (2012), pp. 245–256.
- [27] D. Martani and R. Khairurizka. "The effect of financial ratios, firm size, and cash flow from operating activities in the interim report to the stock return". In: *Chinese Business Review* 8.6 (2009), p. 44.
- [28] I. Welch and A. Goyal. "A comprehensive look at the empirical performance of equity premium prediction". In: *The Review of Financial Studies* 21.4 (2008), pp. 1455–1508.
- [29] J. Gu, Z. Wang, J. Kuen, L. Ma, A. Shahroudy, B. Shuai, T. Liu, X. Wang, G. Wang, J. Cai, et al. "Recent advances in convolutional neural networks". In: *Pattern Recognition* 77 (2018), pp. 354–377.
- [30] S. Gu, B. Kelly, and D. Xiu. "Empirical asset pricing via machine learning". In: *The Review of Financial Studies* 33.5 (2020), pp. 2223–2273.
- [31] E. F. Fama and K. R. French. "Dissecting anomalies". In: *The Journal of Finance* 63.4 (2008), pp. 1653–1678.
- [32] W. Drobetz and T. Otto. "Empirical Asset Pricing via Machine Learning: Evidence from the European Stock Market". In: *Journal of Asset Management* 22.7 (2021), pp. 507–538.
- [33] O. Tobek and M. Hronec. "Does it pay to follow anomalies research? Machine learning approach with international evidence". In: *Journal of Financial Markets* 56 (2021), p. 100588.

- [34] M. Leippold, Q. Wang, and W. Zhou. “Machine learning in the Chinese stock market”. In: *Journal of Financial Economics* 145.2 (2022), pp. 64–82.
- [35] R. Chopra and G. D. Sharma. “Application of artificial intelligence in stock market forecasting: a critique, review, and research agenda”. In: *Journal of Risk and Financial Management* 14.11 (2021), p. 526.
- [36] A. Rubesam. “Machine learning portfolios with equal risk contributions: Evidence from the Brazilian market”. In: *Emerging Markets Review* 51 (2022), p. 100891.
- [37] S. Lalwani and V. Meshram. “Stock Return Prediction in India Using Traditional and Machine Learning Models”. In: *Indian Journal of Finance* 16.5 (2022), pp. 35–52.
- [38] M. P. E. Martens and O. Penninga. “Predicting Bond Returns”. In: *Financial Analysts Journal* 77.3 (2021), pp. 133–155.
- [39] S. R. Das, M. Kalimipalli, and S. Nayak. “Did CDS Trading Improve the Market for Corporate Bonds?” In: *Journal of Financial Economics* 111.2 (2013), pp. 495–525.
- [40] G. A. Karolyi. “Home Bias, an Academic Puzzle”. In: *Review of Finance* 20.6 (2016), pp. 2049–2078.
- [41] A. Goyal and S. Wahal. “Is Momentum an Echo?” In: *Journal of Financial and Quantitative Analysis* 50.6 (2015), pp. 1237–1267.
- [42] A. G. Woodhouse, R. W. Sias, and H. J. Turtle. “What Drives the Cross-Section of Stock Returns in Global Markets?” In: *Journal of Financial Economics* 126.3 (2017), pp. 620–643.
- [43] H. Jacobs and S. Müller. “Anomalies Across the Globe: Once Public, No Longer Existent?” In: *Journal of Financial Economics* 135.1 (2020), pp. 213–230.
- [44] C. R. Harvey. “The Scientific Outlook in Financial Economics”. In: *Journal of Finance* 72.4 (2017), pp. 1393–1442.
- [45] K. Hou, C. Xue, and L. Zhang. “Replicating Anomalies”. In: *Review of Financial Studies* 31.1 (2018), pp. 129–178.
- [46] C. R. Harvey, Y. Liu, and H. Zhu. “... and the Cross-Section of Expected Returns”. In: *Review of Financial Studies* 29.1 (2016), pp. 5–68.
- [47] C. Fieberg, D. Metko, T. Poddig, and T. Loy. “Machine Learning Techniques for Cross-Sectional Equity Returns’ Prediction”. In: *OR Spectrum* 45.1 (2023), pp. 289–323.
- [48] A. Marsi. “Stock Return Predictability Using Machine Learning: Evidence from the EURO STOXX 50 Index”. In: *Journal of Financial Data Science* 3.2 (2021), pp. 45–62.
- [49] T. Masters. *Practical Neural Network Recipes in C++*. Burlington, MA, US: Morgan Kaufmann Publishers, 1993.
- [50] M. J. Van der Laan, E. C. Polley, and A. E. Hubbard. “Super Learner”. In: *Statistical Applications in Genetics and Molecular Biology* 6.1 (2007), pp. 1–21.
- [51] S. Lessmann, B. Baesens, H. V. Seow, and L. C. Thomas. “Benchmarking State-of-the-Art Classification Algorithms for Credit Scoring: An Update of

- Research”. In: *European Journal of Operational Research* 247.1 (2015), pp. 124–136.
- [52] L. I. Kuncheva. *Combining Pattern Classifiers: Methods and Algorithms*. Hoboken, NJ, US: Wiley-Interscience, 2004.
- [53] L. Rokach. “Ensemble-based classifiers”. In: *Artificial Intelligence Review* 33.1–2 (2010), pp. 1–39.
- [54] J. Patel, S. Shah, P. Thakkar, and K. Kotecha. “Predicting Stock and Stock Price Index Movement Using Trend Deterministic Data Preparation and Machine Learning Techniques”. In: *Expert Systems with Applications* 42.1 (2015), pp. 259–268.
- [55] Y. Dong and W. Bao. “Explaining Black-Box Machine Learning Predictions in Finance Using Voting Ensembles”. In: *Journal of Forecasting* 39.6 (2020), pp. 878–892.
- [56] N. Jegadeesh and S. Titman. “Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency”. In: *Journal of Finance* 48.1 (1993), pp. 65–91.
- [57] M. M. Carhart. “On Persistence in Mutual Fund Performance”. In: *Journal of Finance* 52.1 (1997), pp. 57–82.
- [58] B. Han, C. Lee, and M. Worah. “Momentum and Machine Learning: A Framework for Financial Prediction”. In: *Journal of Financial Data Science* 1.3 (2018), pp. 112–130.
- [59] J. Nordvig and N. Firoozye. “Momentum Trading and Machine Learning Models in Financial Markets”. In: *Journal of Financial Markets* 45 (2019), pp. 200–219.
- [60] Y. Xia, X. Hu, and T. Sargent. “Incorporating Momentum in Machine Learning Models for Financial Predictions”. In: *Journal of Machine Learning in Finance* 4.2 (2020), pp. 154–176.
- [61] M. Falk, A. Marx, and T. Gneiting. “Spline Interpolation for Financial Time Series: A Case Study on Missing Value Imputation”. In: *Journal of Applied Statistics* 46.9 (2019), pp. 1567–1583.
- [62] S. Beretta and A. Santaniello. “Nearest Neighbor Imputation Algorithms: A Comprehensive Review and Application to Financial Data”. In: *Journal of Big Data* 3.1 (2016), p. 9.
- [63] J. Lewellen. “The Cross-Section of Expected Stock Returns”. In: *Critical Finance Review* 4.1 (2015), pp. 1–44.
- [64] W. Drobetz and T. Otto. “Empirical Asset Pricing via Machine Learning: Evidence from the European Stock Market”. In: *Journal of Asset Management* 22.7 (2021), pp. 507–538.
- [65] B. Kelly, S. Pruitt, and Y. Su. “Characteristics are Covariances: A Unified Model of Risk and Return”. In: *Journal of Financial Economics* 134.3 (2019), pp. 501–524.
- [66] N. Cakici, C. Fieberg, D. Metko, and A. Zaremba. “Machine Learning Goes Global: Cross-Sectional Return Predictability in International Stock Markets”. In: *Journal of Economic Dynamics and Control* 155 (2023), p. 104725.

- [67] F. Diebold and R. Mariano. “Comparing Predictive Accuracy”. In: *Journal of Business & Economic Statistics* 13.3 (1995), pp. 253–263.
- [68] G. Leitch and J. E. Tanner. “Economic Forecast Evaluation: Profits Versus the Conventional Error Measures”. In: *American Economic Review* 81.3 (1991), pp. 580–590.
- [69] D. A. Lesmond, J. P. Ogden, and C. A. Trzcinka. “A New Estimate of Transaction Costs”. In: *The Review of Financial Studies* 12.5 (1999), pp. 1113–1141.
- [70] B. M. Collins and F. J. Fabozzi. “A Methodology for Measuring Transaction Costs”. In: *Financial Analysts Journal* 47.2 (1991), pp. 27–36.