NAMIBIA UNIVERSITY OF SCIENCE & TECHNOLOGY

DEVELOPING A DATA-DRIVEN FINANCIAL MODEL FOR DECISION SUPPORT IN EVALUATING INVESTMENT PORTFOLIO PERFORMANCE

By

Pendapala Nghilundwa

217113389

Submitted in partial fulfilment of the requirements for the degree of Master of Data Science in the Faculty of Computing and Informatics

Department of Informatics

at the

NAMIBIA UNIVERSITY OF SCIENCE AND TECHNOLOGY

Supervisor:

Dr R. Maliwatu

Date of submission:

(12 September 2024)

Metadata

Research Title:	Developing a Data-Driven Financial Model for Decision Support in Evaluating Investment Portfolio Performance
Student full name:	Pendapala Nghilundwa
Supervisor:	Richard Maliwatu
Department:	Informatics
Qualification:	Master of Data Science
Main knowledge area:	Machine Learning
Status:	Dissertation
Site:	NUST, Windhoek
Document date:	12 September 2025
Sponsor:	None

Declaration

I, Pendapala Nghilundwa, hereby declare that the work contained in this research for the Master's degree thesis, entitled 'DEVELOPING A DATA-DRIVEN FINANCIAL MODEL FOR DECISION SUPPORT IN EVALUATING INVESTMENT PORTFOLIO PERFORMANCE' is my original work and has not been submitted, in whole or in part, to any university or higher education institution for the purpose of obtaining a Master's degree.

I confirm that, in compliance with the Institution's requirements, I have properly recognised all sources of information used in the work. All subsequent contributions to this work have been thoroughly recognised and documented.

I acknowledge that any incidence of plagiarism or academic misconduct is a betrayal of trust and threatens the integrity of the academic community and the process of knowledge development.

Signature_______________________________________ Date 12 September 2025

Abstract

Contemporary financial markets demonstrate heightened complexity and volatility, necessitating sophisticated instruments for the precise assessment of investment portfolios. This research explores the application of machine learning (ML) models to predict Month-to-Date (MTD) returns, aiming to enhance financial decision-making. Conventional models frequently exhibit insufficient flexibility to fluctuating market conditions, highlighting the necessity for data-driven approaches that prioritise portfolio-specific metrics, such as Market Value Dirty and Year-to-Date (YTD) returns. Meanwhile, macroeconomic variables such as gross domestic product (GDP) growth, inflation, and interest rates played a secondary role. The study employed a quantitative method, using secondary data from January 2017 to August 2024, which comprised financial measures and macroeconomic variables. Four machine learning models were developed, namely Random Forest, Gradient Boosting, XGBoost and Long Short-Term Memory (LSTM). Data preprocessing and feature engineering played a critical role in model development. Feature engineering involved creating cumulative MTD returns, moving averages (5-day and 10-day), and volatility metrics to capture market trends and risk dynamics. These features were derived from daily returns, grouped by year and month, and calculated using rolling operations. Data normalisation was applied to standardise input variables, and missing values resulting from rolling operations were filled to ensure dataset completeness. The dataset was then divided into training and testing datasets using a 1:1 ratio. Model performance was evaluated using mean squared error (MSE) and R-squared (R^2) metrics, with cross-validation assuring robustness. Among the models, Gradient Boosting attained the lowest mean squared error (MSE: 2.39×10^{-6}) and the highest R^2 (0.922), outperforming Random Forest (MSE: 2.72×10^{-6}, R^2: 0.911), XGBoost (MSE: 3.13×10^{-6}, R^2: 0.898), and LSTM (MSE: 3.98×10^{-6}, R^2: 0.870). Feature importance analysis highlighted Market Value Dirty, YTD Return, and Benchmark (BM) Size as the most influential predictors. At the same time, macroeconomic variables such as interest rates and inflation contributed minimally to short-term forecasting. This demonstrates the dominance of portfolio-specific metrics in predicting MTD returns. While LSTM excelled in capturing temporal relationships, its predictive accuracy lagged due to volatility during high-risk periods. The results affirm the effectiveness of machine learning models in enhancing financial decision-making. Gradient Boosting and Random Forest models offer accurate predictions and valuable insights into key portfolio-specific factors, underscoring their utility for risk management and strategic planning. The dissertation recommends further exploration of hybrid models, the inclusion of additional macroeconomic variables, and the integration of real-time data to enhance predictive accuracy and robustness. These improvements establish data-driven approaches as essential instruments for financial firms operating in unpredictable markets.

Acknowledgements

I would like to take a moment to sincerely thank Dr Richard Maliwatu, my supervisor, for his constant guidance, support and encouragement during this journey. The insightful feedback, constructive critiques, and mentorship provided were crucial in shaping this thesis. I truly value his patience and commitment, which gave me the assurance and understanding to tackle this demanding task.

I am immensely thankful to God Almighty for granting me the strength, perseverance, and wisdom to complete my master’s degree successfully. Pursuing a master's programme is quite an undertaking, and it has been a path of development and perseverance. Throughout the challenges and triumphs, divine grace has remained my steadfast support, enabling me to navigate difficulties and reach this important academic milestone.

Finally, I would like to express my deep gratitude to my wonderful daughter, Joyce Nghilundwa, who, at just two years old, has been my greatest source of joy and inspiration. Juggling the pressures of school and being a parent was certainly tough, yet her joyful laughter and presence inspired me and strengthened my resolve to achieve my goals. This accomplishment reflects our shared efforts, and I trust it highlights the value of hard work and commitment.

Contents

List of Figures

List of Tables

List of Abbreviations

• APT	Arbitrage Pricing Theory
• BM	Benchmark
• BoN	Bank of Namibia
• CAPM	Asset Pricing Model
• F-REC	Faculty Research Ethics Committee
• GDP	Gross Domestic Product
• HPC	High-Performance Computing
• LIME	Local Interpretable Model-agnostic Explanations
• LSTM	Long Short-Term Memory
• ML	Machine Learning
• MSE	Mean Squared Error
• MTD	Month-to-Date
• NUST	Namibia University of Science and Technology
• OLS	Ordinary Least Squares
• PQFA	Predictive Quantitative Financial Analysis
• PQFA	Predictive Quantitative Financial Analysis
• R^2	Coefficient of Determination
• SAP	Shapley Additive Explanations
• SARIMA	Seasonal AutoRegressive Integrated Moving Average
• SMEs	Small and Medium-sized firms
• XGBoost	Extreme Gradient Boosting

Chapter 1: Introduction

1.1 Background

The financial sector, particularly in emerging economies, has increasingly turned to advanced analytics to navigate economic difficulties. Recent studies demonstrate the potential of machine learning in transforming financial decision-making processes. Liu and Yoon (2024) proved that integrating real-time data streams with machine learning models markedly enhances forecasting precision, especially in volatile markets. Furthermore, James et al. (2024) emphasise the significance of hybrid algorithms in bolstering the robustness of financial models during economic crises.

The growing complexity of financial markets necessitates increasingly sophisticated tools for evaluating the performance of investment portfolios. Conventional financial models, however fundamental, frequently neglect the dynamic and complex characteristics of modern financial landscapes. This thesis sought to develop a data-driven financial model to support the Bank of Namibia in assessing investment portfolio performance for informed decision-making.

Evaluating investment portfolios is essential for financial organisations to mitigate risks and enhance profits. Johnson and Jones (2021) assert that conventional models, such as the Capital Asset Pricing Model (CAPM) and the Arbitrage Pricing Theory (APT), have been used for an extended period. These models depend on simplified assumptions that may not be applicable in the current volatile and interconnected financial markets. The need for more resilient, adaptable and comprehensive models becomes evident as the global financial landscape evolves. Conventional methods often overlook essential elements, such as geopolitical occurrences, fluctuations in market sentiment and technological progressions, all of which can profoundly affect portfolio performance.

Wang and Liu (2020) assert that the emergence of big data analytics and machine learning has transformed the domain of financial modelling. Data-driven methodologies use extensive datasets to discern patterns and produce insights that conventional techniques may overlook. Integrating machine learning technologies, including regression analysis, neural networks, and random forests, enhances the accuracy of financial models and improves risk management strategies. These models can adjust to new data instantaneously, providing a flexible and reactive tool for financial decision-making.

Machine learning algorithms are crucial in advancing financial decision support systems. These algorithms can process extensive data and reveal intricate linkages within financial statistics.

Neural networks, particularly, can simulate non-linear correlations and interactions among variables, making them a highly effective tool for forecasting market trends and evaluating portfolio performance (Lee & Kim, 2021a). Random forests are renowned for their resilience and capacity to manage noisy data, effectively identifying key predictors of financial results (Hodge & Austin, 2020).

Numerous studies have proven the efficacy of machine learning models in financial applications. Research on the Boosting algorithm for quantitative stock selection has demonstrated enhanced profitability using dynamic stock selection strategies (Wang, 2023). Furthermore, the amalgamation of algorithms like LightGBM and XGBoost has demonstrated enhanced predictive efficacy, reinforcing the need to incorporate machine learning in financial decision-making (Zhu, 2021).

The Bank of Namibia's implementation of a data-driven financial model entails the integration of diverse machine learning algorithms to forecast future portfolio performance depending on several inputs. Green and Kaplan (2022) assert that this method enables the bank to use both internal and external data sources, offering a holistic perspective of the market landscape. Consequently, the bank can improve its capacity to make educated investment decisions, refine portfolio strategies and react proactively to market fluctuations.

Notwithstanding the advantages, numerous hurdles accompany the implementation of data-driven financial models. Johnson and Jones (2021) underscore the importance of data accuracy and completeness, as mistakes can substantially distort model predictions and diminish effectiveness. Furthermore, the models must be sufficiently robust to manage market volatility and unforeseen economic fluctuations. James et al. (2021) emphasise the importance of cross-validation and benchmarking against conventional models for ensuring reliability and trust in these novel methodologies.

Financial performance measurements are essential in investment decision-making. According to Ateino (2022), liquidity, profitability and solvency are essential factors for investors assessing investment possibilities. Research has demonstrated a substantial correlation between these variables and investment decisions, underscoring their significance in financial research. By incorporating these measures into data-driven models, financial institutions may improve their decision-making processes and optimise investment results.

LeCun et al. (2015) believe that the future of financial modelling depends on the ongoing technological advancements and the incorporation of interdisciplinary methodologies. The emergence of more sophisticated machine learning algorithms, including deep learning and reinforcement learning, offers considerable potential for improving financial decision-making. These algorithms can evaluate and process intricate, high-dimensional datasets, yielding profound insights into market dynamics and investor behaviour.

The transition to data-driven financial models signifies substantial progress in investment portfolio assessment. Leveraging sophisticated data analytics and machine learning algorithms enables financial institutions, such as the Bank of Namibia, to enhance decision-making processes, refine investment strategies and achieve superior financial performance. The evolving financial landscape necessitates the incorporation of innovative technology to sustain competitiveness and attain sustainable growth.

1.2 Problem Statement

In conventional financial contexts, the assessment of investment portfolio performance has primarily relied on financial ratios, heuristic indicators and established financial models. Brown (2019) contends that although these strategies have been effective historically, they are increasingly considered inadequate for addressing the complexities of contemporary financial markets. Moyo (2023) asserts that complexity and volatility mark the contemporary financial landscape, influenced by rapid technological progress, globalisation and changing market dynamics. Johnson and Jones (2021) emphasise that financial ratios and heuristic indicators, while straightforward to calculate and comprehend, can offer an oversimplified perspective on portfolio performance.

Traditional models, such as CAPM, have been criticised for their inadequacy in adapting to the dynamic and linked nature of global markets (Moyo, 2024). Given the escalating complexity of financial products and rapid technological advancement, there is an urgent need to shift towards data-driven approaches. Recent assessments by Wang et al. (2024) highlight that these approaches enhance precision and enable institutions to proactively anticipate market swings.

Johnson and Lee (2021) contend that these methodologies may not consider the myriad elements affecting financial markets, including geopolitical events, abrupt market shocks, and swift changes in investor mood. Brown (2021) observes that certain financial models employed for decades may now be obsolete. These models, created in a distinct market context, may not have been revised to align with the present financial landscape. Clark and Thompson (2022) assert that these models may lack the requisite sophistication to assess the substantial data volumes produced by contemporary financial markets and may not possess the agility to respond to swift market fluctuations.

Williams and Smith (2019) assert that these constraints present considerable hurdles for financial managers and decision-makers at the Bank of Namibia. The failure to adequately understand and promptly react to market fluctuations may result in poor investment decisions, thereby threatening the institution's financial stability. In the most adverse scenario, it may lead to substantial financial losses, compromising the bank's brand and undermining stakeholder confidence.

Green and Kaplan (2022) assert that to tackle these difficulties, there is an urgent need for more sophisticated and flexible approaches to investment portfolio assessment. These methodologies must effectively encapsulate the intricacies and fluctuations of contemporary financial markets, delivering real-time data and insights that facilitate informed and prompt investment decisions. Taylor and Brown (2023) assert that these strategies will be important in preserving the financial stability of the Bank of Namibia and securing its ongoing success in a dynamic financial environment.

1.3 Aim of the Study

The objective of the study was to create a data-driven financial model using machine learning algorithms to improve decision-making in assessing the performance of investment portfolios. The study aimed to overcome the shortcomings of traditional financial models by incorporating machine learning methodologies, including random forests, gradient boosting, XGBoost, and LSTM, to enhance the accuracy and reliability of future portfolio performance predictions.

1.4 Objectives

1.4.1 Main Objective

The principal objective of this thesis was to create sophisticated, data-driven financial models employing machine learning approaches to facilitate decision-making in assessing investment portfolio performance.

1.4.2 Sub-objectives

i. To assess the significance of various macroeconomic and portfolio indicators in forecasting portfolio performance.
ii. To develop a data-driven model using machine learning algorithms to forecast future portfolio performance by analysing various variables.
iii. To assess the efficacy of the developed model in improving data-driven decision-making and refining investment strategies.

1.5 Research Questions

1.5.1 Main Research Question

How can data-driven financial models using machine learning techniques enhance the assessment of investment portfolio performance?

1.5.2 Sub-Questions

i. What is the significance of the various indicators in forecasting portfolio performance?

ii. How can data-driven financial models that use machine learning techniques improve the evaluation of investment portfolio performance?

iii. What is the influence of these data-driven models on decision-making processes and the overall investment strategy?

1.6 Significance of the Study

The significance of this study lies in its potential to improve investment portfolio assessment and decision-making at the Bank of Namibia by incorporating modern machine learning algorithms. This work employs a quantitative research methodology using empirical data analysis to overcome the shortcomings of conventional financial models through the application of Random Forest, Gradient Boosting, XGBoost, and LSTM models. These models, employed in time-series forecasting and predictive analytics, can analyse substantial volumes of historical and real-time data to discern patterns and forecast future performance, thereby optimising investment plans.

Random Forest, Gradient Boosting and XGBoost have been employed to elucidate non-linear correlations among financial factors, thereby enhancing the precision of portfolio performance forecasts. The LSTM model is particularly effective in identifying temporal dependencies in the data, facilitating accurate forecasting of portfolio returns.

The incorporation of machine learning into financial modelling has progressed markedly in recent years. Roeder et al. (2024) assert that these methodologies enable financial organisations to enhance risk management and investment strategies using comprehensive datasets and sophisticated algorithms. This paper expands on these insights by concentrating on the Bank of Namibia, providing a framework to rectify deficiencies in portfolio review through advanced methodologies.

Wang and Liu (2020) assert that the application of these models can significantly enhance decision-making by offering a holistic and dynamic perspective of the financial landscape. These sophisticated algorithms provide improved risk management tactics by detecting potential problems sooner, increasing the adaptability of investment portfolios. Moreover, models such as XGBoost and LSTM effectively capture intricate interactions among variables that conventional financial models may overlook.

This study offers pragmatic insights for the Bank of Namibia and other financial entities seeking to enhance investment assessment methodologies. The capacity to integrate machine learning-derived insights with real-time data enables portfolio managers to react promptly to market fluctuations, thus enhancing portfolio performance and robustness. Furthermore, the results will enhance the existing literature on the use of machine learning in finance, serving as a significant reference for both scholarly research and financial professionals.

1.7 Scope and Delimitations

This study focused on incorporating machine learning models, namely Random Forest, Gradient Boosting, XGBoost and LSTM, into the investment portfolio assessment process of the Bank of Namibia. Emphasis was on studying historical and real-time financial data to augment the precision of portfolio performance forecasts and refine risk management techniques. The research aimed to transcend the constraints of conventional financial models by using the predictive capabilities of machine learning algorithms, especially in time-series forecasting and regression analyses.

The research employed secondary data from the Bank of Namibia, encompassing interest rates, inflation statistics, and GDP growth rates, to train the models. These data were essential in developing a data-driven financial model that can forecast portfolio performance based on several variables. The study employed Random Forest, Gradient Boosting, and XGBoost as models to forecast portfolio performance based on intricate interactions among macroeconomic variables. At the same time, it used LSTM to capture long-term dependencies in time-series data.

One of the study's limitations is that it was confined to the Bank of Namibia. Therefore, the results may not be relevant to other financial institutions operating in diverse circumstances. The researcher selected the models for their robustness and established effectiveness in financial modelling. However, alternative models, such as support vector machines or more complex neural networks, were excluded due to limited resources and time.

Furthermore, the qualitative dimension of the data, including expert comments and sentiment analysis, was excluded from this study, as the emphasis was on quantitative financial data. The study aimed to offer significant insights into the Bank of Namibia's portfolio appraisal processes, although it was limited to the application of machine learning algorithms on secondary financial data.

1.8 Thesis Outline

This thesis comprises six chapters, each addressing distinct facets of the research, from the introduction to the findings and recommendations.

Chapter 1 gives a comprehensive overview of the study, encompassing the context, problem statement, research questions, and aims. The chapter delineates the study setting, highlighting its importance, extent and limitations. It culminates with a guide to the document's overall structure, establishing the groundwork for succeeding chapters.

Chapter 2 examines pertinent literature on financial modelling, highlighting the shift from conventional methodologies to data-driven techniques. This paper investigates the relevance of machine learning methodologies in financial forecasting and identifies gaps in the existing literature that it aims to address. This review establishes a theoretical foundation for the research and rationalises the selected approaches.

Chapter 3 delineates the research technique employed in this study, encompassing data collection, preprocessing, feature engineering and model construction. This chapter examines the machine learning models employed, LSTM, Gradient Boosting, Random Forest and XGBoost and elucidates their application in predicting MTD returns. The discussion also encompasses evaluation measures employed to assess model performance.

Chapter 4 delineates the outcomes derived from the constructed models, encompassing visual representations and comparative evaluations. This chapter assesses the predictive efficacy of each model and offers insights into feature significance, aiding in the identification of the principal determinants of portfolio returns. This analysis underpins the assessment of the models' efficacy in achieving the study objectives.

Chapter 5 analyses the study's findings in connection with the research questions and objectives. It underscores the practical ramifications of the findings for the Bank of Namibia, rigorously assessing the merits and limitations of the employed models. The chapter additionally examines the ramifications for financial decision-making and portfolio management.

Chapter 6: summarises the contributions of the study to financial modelling and offers recommendations for subsequent research. It underscores the importance of employing data-driven models to augment investment decision support, addresses the limits faced throughout the research, and proposes avenues for future investigation to boost prediction efficacy and relevance in the financial industry.

This framework guarantees a logical progression from the study's inspiration to its practical implications, directing the reader through each phase of the research process, from theoretical underpinnings to practical applications.

Chapter 2: Literature Review

2.1 Introduction

The assessment of investment portfolio performance has conventionally depended on financial ratios, heuristic measures and recognised financial models. Although historically effective, these strategies are increasingly deemed inadequate for addressing the intricacies of contemporary financial markets. The contemporary financial landscape is characterised by rapid technology advancement, globalisation and changing market dynamics (Johnson & Jones, 2021). Conventional models, such as CAPM and APT, have historically served as foundational elements of portfolio management. Nevertheless, they frequently neglect the complex and evolving characteristics of contemporary financial landscapes (Brown, 2019).

In recent years, the financial sector has experienced a significant transition towards more advanced and data-centric methodologies. Wang and Liu (2020) assert that the amalgamation of big data analytics and machine learning methodologies has transformed financial modelling. These sophisticated techniques use extensive datasets to reveal patterns and generate insights that conventional models may overlook, offering a more thorough and flexible approach to portfolio assessment. Lee and Kim (2021) elucidate that machine learning technologies, including regression analysis, neural networks and random forests, are increasingly being employed to forecast portfolio performance and mitigate financial risks. These algorithms can process substantial data volumes in real-time, providing dynamic and responsive instruments for financial decision-making.

The growing complexity of financial markets, driven by globalisation, technological advancements and regulatory changes, has necessitated the development of more resilient financial models. Conventional financial models, developed during a period of relative market stability, frequently lack the capacity to address the volatility and rapid fluctuations characteristic of modern markets. These models generally assume market efficiency and investor rationality, which often do not consistently manifest. The constraints of these models are evident during market volatility, as unexpected geopolitical occurrences, technology interruptions, and shifts in market sentiment can substantially affect portfolio performance (Liu & Yoon, 2024).

Green and Kaplan (2022) assert that the Bank of Namibia, like many other financial institutions, faces the challenge of navigating these intricate market conditions. This study suggests that using a data-driven financial model can significantly enhance the bank's ability to assess investment portfolio performance. This study’s model incorporates diverse machine learning methods to forecast future portfolio performance based on multiple inputs, enabling the bank to use both internal and external data sources and offering a comprehensive perspective of the market landscape. Consequently, enhancing the bank's ability to make informed investment decisions, refine portfolio strategies and react proactively to market fluctuations.

Despite the evident benefits, the implementation of data-driven financial models also poses numerous hurdles. James et al (2021) assert that guaranteeing data quality and completeness is essential, since any mistakes can significantly distort model predictions and diminish their efficacy. Furthermore, these models must include sufficient robustness to manage market volatility and unforeseen economic fluctuations. Consequently, stringent cross-validation and benchmarking against conventional models are crucial to uphold reliability and confidence in these novel methodologies.

Financial performance measurements are essential instruments in investment decision-making, offering insights into the financial stability and operational efficiency of prospective investment opportunities. Metrics including liquidity, profitability, and solvency are essential for assessing an organisation's capacity to fulfil its financial commitments and provide returns for investors. Research has shown that these variables are significantly associated with investment decisions, highlighting their relevance in financial analysis. Integrating these indicators into data-driven models enables financial institutions to improve the precision and efficacy of their investment decisions, resulting in superior financial outcomes (Ateino, 2022).

LeCun et al (2015) propose that he future of financial modelling depends on ongoing technological advancements and the incorporation of interdisciplinary methodologies. The evolution of more sophisticated machine learning algorithms, particularly deep learning and reinforcement learning, presents considerable potential for improving financial decision-making. These sophisticated algorithms can handle and analyse intricate, high-dimensional datasets, yielding profound insights into market dynamics and investor behaviour.

The transition to data-driven financial models represents a significant advancement in assessing investment portfolios. The use of advanced data analytics and machine learning algorithms by financial institutions, such as the Bank of Namibia, can improve decision-making, optimise investment strategies and enhance financial performance. The evolving financial landscape necessitates the incorporation of innovative technology to sustain competitiveness and attain sustainable growth. This thesis aimed to develop a comprehensive, data-driven financial model specifically designed for the Bank of Namibia.

2.2 Conventional Financial Models and Their Constraints

Conventional financial models, such as CAPM and APT, have historically been essential to portfolio management and investment theory. Formulated in the mid-20th century, these models offer the basis for understanding the relationship between risk and return, enabling investors to make informed decisions about asset allocation and portfolio diversification. Nonetheless, despite their fundamental significance, these models demonstrate numerous limitations, especially within the contemporary financial landscape.

The shortcomings of conventional financial models highlight the need for more resilient, flexible and comprehensive strategies in portfolio management. Wang and Liu (2020) contend that the emergence of big data analytics and machine learning presents novel options to tackle these difficulties by utilising extensive datasets and advanced algorithms to encompass a broader spectrum of factors affecting market behaviour. Through the integration of sophisticated analytical methods, financial models can yield more precise forecasts and enhanced risk management strategies, adjusting to new data instantaneously and delivering dynamic tools for financial decision-making.

2.3 The Ascendancy of Data-Driven Financial Models

The progression of data analytics and machine learning has significantly transformed the financial modelling domain, leading to the development of models that are both more advanced and precise. Wang and Liu (2020) assert that data-driven financial models use extensive and varied information to reveal patterns and insights frequently overlooked by conventional methods. By using sophisticated machine learning methods, such as regression analysis, neural networks and random forests, these models can forecast portfolio performance more effectively and mitigate financial risks.

A principal advantage of data-driven financial models is their capacity to process and evaluate extensive volumes of historical and real-time data. Huang et al. (2020) emphasise that this characteristic enables these models to swiftly assimilate new information, offering a dynamic and responsive instrument for portfolio management. Regression analysis can elucidate the relationship between various financial variables and portfolio returns, facilitating more accurate projections and risk evaluations. Huang et al. (2020) point out that neural networks, recognised for their capacity to simulate complicated and non-linear interactions, are especially adept at capturing the intricate dynamics of financial markets.

Besides stock selection, data-driven algorithms are used to evaluate credit risk, detect fraud and analyse market sentiment. Bollen et al. (2011) demonstrate that analysing extensive unstructured data from diverse sources, including news articles, social media and financial reports, can yield insights on market sentiment and investor behaviour, which are essential for informed investment decisions. Integrating qualitative input into quantitative models improves their predictive accuracy and dependability.

The transition to data-driven financial models signifies significant progress in investment portfolio assessment. Huang et al. (2020) contend that while conventional financial models have provided valuable insights, they are often limited by their reductive assumptions and limited capacity to adapt to the rapid and unpredictable dynamics of contemporary financial markets. In contrast, data-driven models offer a more comprehensive and adaptable approach capable of incorporating diverse factors and continually improving through the assimilation of new information.

Liu and Yoon (2024) observe that the implementation of data-driven models entails numerous hurdles. Maintaining data quality and integrity is paramount, as erroneous or incomplete datasets can lead to flawed predictions and poor decisions. Furthermore, the complexity of these models demands considerable computer resources and expertise in machine learning and data analytics. Moreover, financial institutions must address challenges of model transparency and interpretability to ensure that decision-makers not only comprehend the underlying mechanisms but can also trust the model outcomes.

2.4 Machine Learning Algorithms in Financial Modelling

Machine learning algorithms are increasingly being employed to enhance financial decision-making by enabling the analysis of complex datasets and revealing hidden patterns within financial data. They offer diverse approaches applicable to multiple facets of financial modelling, including portfolio return prediction, risk management and investment strategy optimisation.

Wang and Liu (2020) emphasise that regression analysis is a fundamental tool in financial modelling. This strategy elucidates the links among various financial variables and portfolio returns. This entails defining the dependent variable, such as portfolio returns, and modelling its connection with one or more independent variables, including market indices, interest rates, or economic indicators. Regression analysis quantifies relationships, allowing financial analysts to make accurate forecasts on future portfolio performance based on existing data.

Huang et al. (2020) assert that neural networks, a potent machine learning method, excel at modelling intricate, non-linear interactions in financial data. Unlike conventional linear models, neural networks comprise numerous layers of interconnected nodes that process incoming data like the human brain. This framework allows neural networks to identify complex patterns and relationships among financial variables that simpler models would miss. Neural networks can anticipate stock values by examining several inputs, such as past prices, trading volumes and macroeconomic data. Neural networks' capacity to learn from data and enhance forecast accuracy over time makes them an important resource for financial analysts (Lee & Kim, 2021).

A random forest is an ensemble learning technique that integrates numerous decision trees to enhance forecast accuracy and mitigate the likelihood of overfitting. Every tree in the forest is trained on a random subset of the data, and the ultimate forecast is derived by averaging the predictions of all the trees. This method improves the model's capacity to generalise to novel data and reduces its susceptibility to noise and outliers.

Wang et al. (2024) demonstrate that boosting algorithms, a category of ensemble learning techniques, can markedly enhance stock selection procedures. Boosting algorithms enhance investment portfolio performance by iteratively amalgamating weak learners to construct a robust predictive model, thereby adapting to intricate market conditions.

Liu and Yoon (2024) examine the difficulties inherent in the application of machine learning algorithms within financial modelling. Maintaining data quality and integrity is essential, as the precision of the model's predictions relies on the dependability of the input data. Moreover, the intricacy of these algorithms necessitates considerable computer resources and proficiency in data science and machine learning. Financial institutions must allocate resources towards the requisite infrastructure and expertise to proficiently use these new methodologies.

2.5 Case Studies and Applications

The application of machine learning in financial modelling has been extensively studied, with numerous case studies demonstrating its effectiveness in enhancing investment strategies and portfolio management. Wang (2019) emphasises that dynamic stock selection models employing the Boosting algorithm might enhance profitability. The Boosting algorithm, an ensemble learning technique, integrates several weak learners to construct a robust prediction model. This method facilitates ongoing adaptation to fresh data, enhancing the model's precision and resilience in forecasting stock performance.

Boosting algorithms have been useful in financial applications due to their ability to manage extensive and complex datasets. Wang and Liu (2020) assert that these algorithms function by iteratively modifying the weights of weak learners according to their performance, hence minimising errors and enhancing overall prediction efficacy. Wang and Liu (2020) found that the application of the Boosting algorithm enables quantitative stock selection models to get superior performance metrics relative to conventional methods. This is particularly pertinent in unstable and rapidly evolving financial markets, where timely and precise forecasts are essential for making informed investment choices.

Zhu (2021) prove that the combined application of LightGBM and XGBoost can markedly enhance stock selection and portfolio optimisation methodologies. LightGBM and XGBoost employ distinct optimisation techniques featuring a selection of methods, which, when combined, facilitate a more comprehensive analysis of financial data.

Huang et al. (2020) demonstrated that machine learning algorithms can optimise asset allocation by forecasting future performance. This integration enables effective management of diverse data types, including numerical, categorical and missing data, thereby supporting the development of more dependable investment strategies for various assets using historical data and market trends. This entails employing algorithms such as random forests and neural networks to examine various aspects that affect asset returns, including economic data, corporate financials and market sentiment.

Bollen et al. (2011) note that the application of machine learning in the detection of financial fraud has garnered significant interest. Machine-learning algorithms can identify atypical patterns and anomalies in transaction data, potentially signalling fraudulent actions. By continuously assimilating fresh data, these models can adjust to evolving fraud strategies, offering a formidable safeguard against financial crimes.

2.6 Data-Driven Decision Support at the Bank of Namibia

Implementing a data-driven financial model for the Bank of Namibia entails integrating sophisticated machine learning algorithms to forecast future portfolio performance based on various inputs. Green and Kaplan (2022) assert that this advanced methodology aligns well with the bank's requirements for real-time analytics and insights, thereby facilitating better informed investment decisions. Using both internal and external data sources, the model offers a comprehensive and adaptive perspective of the market landscape, significantly improving the bank's ability to respond to market fluctuations and refine its investment plans.

Lee and Kim (2021) emphasise that employing machine-learning algorithms in financial modelling facilitates the examination of extensive information, revealing patterns and connections that conventional models may neglect. Regression analysis clarifies the relationship between diverse macroeconomic variables and portfolio returns, thereby generating predictive insights into possible future performance. Neural networks, capable of modelling complex, non-linear interactions, may examine sophisticated datasets to predict market trends and asset performance. Random forests, known for their resilience and clarity, can identify critical aspects affecting portfolio results, thus facilitating more accurate decision-making.

Huang et al. (2020) assert that the incorporation of real-time data analytics is essential for the Bank of Namibia, facilitating ongoing surveillance and modification of investment plans in response to the most recent market trends. Real-time analytics provide immediate insights into market conditions, enabling the bank to adjust swiftly to fluctuations and capitalise on emerging opportunities. This is especially crucial in volatile markets as prompt and precise information can greatly influence investment results.

Bollen et al. (2011) propose that the implementation of such a model necessitates the use of external data sources, including economic indicators, geopolitical events and market sentiment derived from social media and news platforms. This comprehensive approach ensures that a holistic understanding of the market landscape informs the bank's investment decisions. Integrating multiple data sources enables the model to deliver a more nuanced and precise representation of market conditions, augmenting the bank's capacity to make strategic investment decisions.

Liu and Yoon (2024) observe that several limitations and challenges accompany the implementation of data-driven models. Therefore, ensuring data quality and completeness is essential, as inaccuracies or omissions can lead to erroneous forecasts and suboptimal actions. The complexity of these models requires substantial computational resources and advanced expertise in data science and machine learning. The bank must allocate resources towards the requisite infrastructure and expertise to proficiently execute and sustain these sophisticated models.

2.7 Challenges and Considerations

Although data-driven financial models offer several benefits, various problems must be addressed to guarantee their efficacy and dependability. Johnson and Jones (2021) identify the accuracy and completeness of data as a significant concern. Financial models depend on the quality of input data, and any mistakes can significantly distort the findings, resulting in poor or potentially harmful judgements. Ensuring data accuracy requires stringent data cleansing procedures to eliminate errors, inconsistencies and redundancies. Liu and Yoon (2024) highlight that financial data frequently originates from diverse sources, each having distinct formats and standards, making data integration a complicated endeavour requiring meticulous management to preserve integrity and consistency.

Huang et al. (2020) underscore the significance of model robustness. Financial markets are inherently volatile, influenced by various unforeseen factors, such as geopolitical incidents, economic transitions and sudden market disruptions. Consequently, data-driven models must be sufficiently robust to manage such volatility and adjust to rapidly evolving market conditions. This necessitates the creation of advanced algorithms that can learn from historical data and adapt to integrate new information and patterns. Models must be durable and adjustable to accommodate unforeseen changes, thereby maintaining their effectiveness in practical applications.

James et al. (2021) emphasise that validating these models is crucial. Cross-validation is a common technique used in machine learning to evaluate model performance. The process entails partitioning the dataset into several subsets, training the model on select subsets while evaluating it on others, and subsequently averaging the outcomes. This approach helps in detecting overfitting, wherein a model excels on training data yet underperforms on novel data. Cross-validation guarantees that the model generalises effectively to novel data rather than being exclusively adapted to the specific patterns inside the training dataset.

Wang and Liu (2020) assert that benchmarking against conventional models is an essential step in assessing the efficacy of data-driven financial models. Financial analysts can assess the relative enhancements of new models by comparing their performance with established frameworks, such as CAPM or APT, to verify that the new models provide substantial advantages. Such comparisons help identify constraints and highlight aspects of the new models that may require refinement.

Studies highlight the interpretability of machine learning models as another challenge (Green & Kaplan, 2022; Lee & Kim, 2021). Although advanced algorithms, such as neural networks and ensemble approaches, provide superior predictive accuracy, they often function as “black boxes”, making it difficult to comprehend their decision-making processes. A lack of openness can be a significant constraint in the financial sector, where stakeholders require both trust and clarity regarding the decision-making process. Therefore, formulating techniques that elucidate and clarify the results of complex models is crucial for building user confidence and guaranteeing their appropriate application.

2.8 Machine Learning Algorithms in Financial Modelling

Machine learning algorithms have transformed financial modelling by offering advanced tools for analysing extensive datasets, detecting trends, and generating precise predictions. These algorithms can be categorised into three primary types: supervised learning, unsupervised learning and reinforcement learning. Each category is tailored to distinct forms of financial data analysis and decision-making processes, augmenting the capacity of financial institutions to mitigate risks and maximise rewards.

Alzubi et al. (2018) explain that neural networks, which replicate the linked neuron architecture of the human brain, are widely employed for pattern identification and predictive analytics in financial markets. These networks have layers of nodes (neurons), with each layer converting the incoming data into increasingly abstract representations. Neural networks can identify non-linear relationships in data, rendering them suitable for forecasting stock prices, assessing credit risks, and optimising investment portfolios. They analyse extensive datasets, discern intricate patterns, and enhance their forecast precision over time via iterative training methodologies.

Bollen et al. (2011) illustrate the successful application of machine learning algorithms in credit risk assessment, fraud detection and market sentiment analysis. Using extensive datasets and sophisticated analytical methods, these algorithms yield more precise and prompt insights, allowing financial organisations to enhance decision-making and optimise risk management strategies.

2.9 Data-Driven Approaches in Credit Risk Management

Data-driven methodologies have become essential in credit risk management, using modern analytics and machine learning techniques to enhance the precision and reliability of credit risk evaluations. Roeder et al. (2022) emphasise that these methodologies integrate quantitative risk metrics with qualitative insights obtained from diverse data sources, such as text mining and sentiment analysis of financial analyst reports, to provide a more comprehensive and holistic assessment of the financial stability of prospective borrowers.

Text mining and sentiment analysis are particularly valuable tools in this context. Roeder et al. (2022) elucidate that analysing qualitative data from financial analyst reports, news stories and social media provides a means of assessing market sentiment and detecting early indications of financial distress that may not be captured through quantitative data alone. For example, negative sentiment in analyst reports or a sudden surge in adverse news coverage on a company may indicate potential problems warranting closer examination. Integrating qualitative and quantitative data augments the predictive power of credit risk models, facilitating more prompt and informed decision-making.

Logistic regression can quantify the likelihood of default by analysing the correlation between borrower attributes and default results. Decision trees categorise borrowers into various risk levels based on financial and non-financial characteristics. In contrast, neural networks identify intricate, non-linear correlations within the data, enhancing the precision of risk assessments.

Huang et al. (2020) emphasise ongoing surveillance and revision of risk models as a crucial element of data-driven credit risk management. Given the rapid fluctuations in financial markets and borrower circumstances, risk evaluations must be revised regularly to remain accurate and relevant. Data-driven models can be developed to autonomously integrate new data and modify their predictions, ensuring that credit risk evaluations remain pertinent and accurate over time.

Ensuring data quality and integrity is crucial, as erroneous or incomplete data can compromise the accuracy of risk assessments. The inherent complexity of machine learning models can further create challenges with interpretability, hindering stakeholders' ability to comprehend the decision-making process. Formulating strategies that enhance the interpretability of model outputs and guarantee transparency is essential for fostering trust and facilitating the use of these advanced methodologies.

2.10 Integration of Financial Data for Enhanced Decision Support

The effective integration of financial data is fundamental for developing robust decision support systems that enhance informed decision-making in financial organisations. Consolidating diverse financial data sources into a unified framework enables an in-depth analysis of financial metrics, market circumstances and economic trends, which is crucial for precise forecasting and strategic planning.

Liu and Yoon (2024) introduced a goal-oriented data integration framework that automates the data integration process according to user-specified analytical objectives. This approach enhances the significance and applicability of integrated data, making it an essential tool for informed financial decision-making. The framework enhances efficiency and accuracy in financial analysis by automating the integration process, which diminishes the time and effort needed to gather and process data from several sources (Liu & Yoon, 2024).

The objective-oriented data integration framework is designed to tackle the limitations associated with conventional data integration techniques, which often rely on manual data gathering and processing. These conventional approaches may be labour-intensive and susceptible to errors, resulting in inconsistencies and inaccuracies in the integrated data. Conversely, the automated approach proposed by Liu and Yoon (2024) employs sophisticated algorithms and machine learning methodologies to optimise the data integration process. This guarantees that the integrated data is both precise and pertinent to the specific analytical objectives established by the users (Liu & Yoon, 2024).

Huang et al. (2020) emphasise a principal advantage of this framework: its capacity to manage extensive and intricate datasets. Financial institutions produce and manage vast volumes of data every day, encompassing transaction records, market data, and financial statements. The objective-oriented data integration framework uses big data technologies to manage and analyse extensive information effectively, offering financial analysts and decision-makers prompt and actionable insights that improve their capacity for data-driven decision-making (Huang et al., 2020).

Bollen et al. (2011) illustrate that the integration architecture facilitates the amalgamation of various data kinds, encompassing structured data from databases, semi-structured data from spreadsheets, and unstructured data from text documents and social media. This comprehensive methodology enables a more integrated examination of financial data by accounting for the multiple elements that can influence financial outcomes and market dynamics. Integrating sentiment analysis of social media data with conventional financial indicators can yield a more profound comprehension of market sentiment and its influence on stock prices (Bollen et al., 2011).

Wang and Liu (2020) assert that the amalgamation of financial data enhances the creation of predictive models capable of forecasting future financial performance and market trends. Machine learning techniques, including regression analysis, neural networks, and decision trees, can be applied to integrated datasets to identify underlying patterns and generate predictive insights. These predictive models are essential for financial planning and risk management, allowing institutions to foresee market fluctuations and make proactive decisions (Wang & Liu, 2020)

James et al. (2021) assert that the goal-oriented data integration framework enhances data openness and traceability, which are essential for regulatory compliance and auditing. Financial organisations must adhere to rigorous regulatory standards that require precise reporting and documentation of financial information. The automated integration procedure guarantees precise documentation and traceability of all data sources, thereby promoting adherence to regulatory standards and mitigating the possibility of non-compliance penalties (James et al., 2021).

2.11 Case Studies and Practical Applications

Numerous case studies demonstrate the practical application of data-driven financial models and machine learning algorithms across various sectors, showing their effectiveness in improving investment decisions and risk management.

A significant case study examines the investment performance of technology stocks with Predictive Quantitative Financial Analysis (PQFA). Khakhar and Godhani (2023) examined the potential of predictive analytics to enhance portfolio returns and improve risk management. Their research emphasised the significance of integrating conventional financial analysis with sophisticated prediction models to attain enhanced investment results. Integrating machine learning algorithms with conventional financial measures enables investors to obtain profound insights into market movements and make informed decisions. The implementation of PQFA enables the adaptive modification of investment strategies based on real-time data, enhancing portfolio performance and reducing risk.

A notable application of AI and machine learning is assessing the financial performance of small and medium-sized firms (SMEs) through the analysis of court rulings. Roeder et al. (2022) show that the amalgamation of many data sources, including quantitative risk metrics and qualitative insights derived from text mining and sentiment analysis, can improve credit risk evaluations. This thorough methodology offered a more integrated perspective on the financial stability and investment viability of SMEs. By studying judicial judgements and extracting pertinent information, the programme could evaluate the creditworthiness of SMEs with greater precision than conventional approaches. Combining multiple data sources in credit risk models enables financial institutions to gain a deeper understanding, enhancing their comprehension and management of risks associated with lending to SMEs.

The application of AI in financial modelling encompasses more than merely stock selection and credit risk evaluation. Huang et al. (2020) that AI algorithms can predict stock market fluctuations by assessing a synthesis of historical data, market indicators, and sentiment derived from news articles and social media. These models yielded more precise forecasts of market patterns, allowing investors to execute timely and strategic investment decisions. The ability to analyse large volumes of unstructured data, including text from news articles, significantly enhances the prediction capabilities of AI-driven models.

Furthermore, the application of machine learning in fraud detection represents another domain where these technologies have proved effective. Bollen et al. (2011)) emphasise the use of sentiment analysis of social media data in identifying fraudulent activities in financial markets. By analysing social media posts for atypical patterns or sentiment fluctuations, financial institutions can promptly detect probable fraud situations and implement preventive actions. This proactive strategy for fraud identification safeguards investors and upholds market integrity.

2.12 Gaps in Existing Literature

Despite significant advancements in the use of machine learning in financial modelling, several critical gaps remain unresolved. One significant deficiency is the overemphasis on macroeconomic variables, such as inflation or GDP growth, which are often examined in isolation. While existing studies have highlighted their impact, the findings from this research demonstrate that short-term portfolio performance, such as MTD returns, is significantly influenced by portfolio-specific metrics like Market Value Dirty and YTD returns. This highlights a gap in the literature where insufficient attention is given to the relative importance of internal portfolio factors compared to external macroeconomic indicators, particularly in short-term forecasting.

Additionally, while the focus has been on the inclusion of macroeconomic variables in predictive models, their collective effect on investment portfolio performance has often been overstated in the context of short-term models. This study bridges this gap by emphasising the diminished role of macroeconomic variables in short-term predictions, without entirely disregarding their value for long-term forecasting. The findings suggest that future studies should prioritise portfolio-specific metrics for immediate forecasting while reserving macroeconomic indicators for broader, long-term analyses.

Another notable gap is in the lack of real-time predictive models that can dynamically adapt forecasts in response to market fluctuations. Many prior studies rely on static data, limiting their applicability in the rapidly evolving financial markets. This dissertation addresses this limitation by advocating for the integration of real-time data streams, particularly for high-impact features, such as Market Value Dirty and Volatility, which are critical in short-term predictions.

Interpretability is also a persistent challenge, with many machine learning models, especially neural networks, criticised for their "black box" nature. Stakeholders in financial decision-making require transparency to trust these models. By using Gradient Boosting, Random Forest, and feature importance analysis, this study contributes to the growing body of work focused on enhancing the interpretability of machine learning models. The feature importance findings, as detailed in the study, provide a transparent mechanism for identifying the most influential predictors, enabling practical and actionable insights for portfolio managers.

Lastly, the lack of comparative evaluations of machine learning models in financial contexts presents another gap. While numerous studies have explored individual models, few have directly compared their performance. This study addresses this by systematically assessing the efficacy of Random Forest, Gradient Boosting, XGBoost and LSTM models. The results reveal that Gradient Boosting outperformed other models in predicting MTD returns, offering a balance between predictive accuracy and interpretability. The findings underscore the importance of choosing machine learning methodologies that are tailored to the specific contexts of portfolio performance assessment.

2.13 Summary

This chapter provided a comprehensive review of the literature on financial modelling, emphasising the shift from traditional methods to modern, data-driven approaches. Traditional financial models, such as CAPM, often fail to capture the complexity and volatility of financial markets. The review highlighted the growing use of machine learning algorithms, which offer enhanced forecasting accuracy, improved risk management and actionable insights for decision-making. Special emphasis was placed on integrating macroeconomic variables into predictive models, highlighting their limited relevance in short-term forecasts, as evidenced by the diminished importance of variables like interest rates and inflation in feature importance rankings.

The literature review identified key challenges, including data quality issues, the opaque nature of some machine learning algorithms, and the need for robustness under dynamic market conditions. Moreover, it highlighted the insufficient exploration of portfolio-specific metrics, which this dissertation identifies as more critical for short-term forecasting. The chapter also addressed gaps in real-time predictive capabilities and model interpretability, which are essential for practical financial decision-making.

Finally, this chapter laid the groundwork for the next phase of the study by highlighting the need for comparative evaluations of machine learning models. The identified gaps provide a clear direction for the methodology outlined in the following chapter, where a data-driven approach is employed to develop and test predictive models for assessing portfolio performance. The results aim to contribute to a more nuanced understanding of feature importance and improve the practical utility of machine learning models in financial markets.

Chapter 3: Methodology

3.1 Introduction

The methodology for this study was designed based on the Research Onion Framework by Saunders et al. (2019). The Research Onion Framework guided the research design process, enabling the identification and structuring of different research methods. This chapter outlines the research philosophy, approach, strategy, time horizons, and data collection methods used to achieve the research objectives. The methodology was chosen to ensure robust data analysis and allow the use of advanced machine learning models to predict MTD returns.

3.2 Research Philosophy

The study employed a positivist research philosophy, which focuses on quantifiable observations and statistical analyses (Saunders et al., 2019). This approach aligned well with the nature of financial data analysis and machine learning modelling, as these methods depend heavily on empirical data and statistical testing (Creswell, 2014).

3.3 Research Approach

This study adopted an inductive approach, beginning with historical financial and macroeconomic data from the Bank of Namibia. Patterns, trends and relationships were identified through exploratory analysis and machine learning modelling. The findings were then used to generate insights and inform potential improvements in portfolio performance assessment, rather than testing a pre-existing theory.

3.4 Methodological Choice

A mono-method quantitative design was employed, relying exclusively on quantitative secondary data and statistical/machine learning techniques for analysis. This choice ensured methodological consistency and aligned with the study's focus on measurable financial indicators.

3.5 Research Strategy

The research employed an archival research strategy, using historical secondary data from the Bank of Namibia covering January 2017 to August 2024. This strategy facilitated the application of machine learning and time series forecasting methods (Random Forest, Gradient Boosting, XGBoost, and LSTM) to build predictive models for Month-To-Date returns. The archival strategy is particularly suited to financial research, as it leverages existing records to generate insights without primary data collection.

3.6 Time Horizons

A longitudinal time horizon was used in this study to capture financial data over multiple years (2017 to 2024). The longitudinal analysis helped to model the temporal relationships within the data (Saunders et al., 2019).

3.7 Data Collection Methods

This study used secondary data obtained from the Bank of Namibia, which encompassed monthly financial returns and essential macroeconomic variables, including GDP growth, inflation and interest rates. The dataset encompassed historical market value data, volatility and many pertinent financial metrics. The data period extended from January 2017 to August 2024, offering a comprehensive dataset for the development and validation of predictive models.Figure 3.

The dataset has the following essential features:

- Month-To-Date Return: The MTD portfolio returns, which serve as the primary target variable.
- Interest Rate: Monthly interest rates affecting the portfolio.
- Inflation: Monthly inflation rates.
- GDP Growth: The GDP growth rate influences economic performance.
- Volatility and Moving Averages: Indicators that capture the short-term price fluctuations of assets.
- Portfolio and Benchmark Durations: Measures of interest rate sensitivity of the portfolios relative to benchmarks.

	count	mean	min	25%	50%	75%	max	std
ASOF_DATE	2033	2020-10-12 00:38:57.432366080	2017-01-02 00:00:00	2018-11-23 00:00:00	2020-10-06 00:00:00	2022-09-02 00:00:00	2024-07-31 00:00:00	NaN
MARKET_VALUE_DIRTY	2033.00	59211376.62	30264401.95	39572785.46	64371310.05	70103414.57	100629086.14	20534347.94
UNINVESTED_CASH	2033.00	1512397.63	0.07	238019.38	888437.36	1858010.01	47036802.77	2923437.82
SUBPORTFOLIO_DURATION	2033.00	60723449.61	35884820.32	40647109.83	65997345.22	72012827.98	100639300.02	20850154.36
BENCHMARK_DURATION	2033.00	107.21	100.86	102.61	107.49	111.07	113.45	3.88
SUBPORTFOLIO_SPREAD_DURATION	2033.00	106.70	9.07	102.41	106.74	110.62	112.25	4.44
BENCHMARK_SPREAD_DURATION	2033.00	2.43	1.35	2.35	2.47	2.55	4.92	0.23
BM_SIZE	2033.00	0.67	0.20	0.39	0.66	0.83	1.67	0.30
COMPUNDED_DAILY_INTEREST	2033.00	60486159.17	5968479.75	40578745.29	65627789.74	71398712.23	100499866.72	20770297.44
SIMPLE_DAILY_INTEREST	2033.00	60469779.22	5968479.75	40576468.37	65625435.75	71390260.74	100499866.72	20772879.93
DAILY_RETURN	2033.00	0.00	-0.01	-0.00	0.00	0.00	0.01	0.00
MTD_RETURN	2033.00	0.00	-0.02	-0.00	0.00	0.00	0.02	0.00
YTD_RETURN	2033.00	0.01	-0.05	-0.00	0.00	0.02	0.04	0.02
INTEREST_RATE	2033.00	2.16	0.25	0.25	1.75	2.50	5.50	1.85
GDP_GROWTH	2033.00	271.69	242.84	252.15	260.28	296.80	314.54	23.34
INFLATION	2033.00	4.60	3.40	3.60	3.90	4.40	14.80	1.95

Figure 3.1: Summary of dataset statistics

3.8 Data Preprocessing

The data preprocessing phase was essential in preparing the dataset for reliable machine learning modelling. This stage involved multiple steps aimed at ensuring accuracy, completeness, and suitability of the data for analysis. Specifically, preprocessing covered data cleaning, to remove inconsistencies and manage missing values; feature engineering, to generate meaningful predictors from raw financial data; model development and training, where the dataset was structured for different machine learning algorithms; and evaluation, to assess performance using robust statistical metrics. Each of these stages is outlined in the subsections that follow.

3.8.1 Data Cleaning

Data cleaning encompassed the elimination of irrelevant or erroneous data and the management of missing variables. The column denoting daily returns was converted to a numeric format, and rows containing missing values in this column were eliminated to maintain data integrity. Missing values produced by rolling procedures, such as moving averages, were substituted with zeros.

3.8.2 Feature Engineering

The feature engineering process in this study was pivotal in transforming raw financial data into meaningful inputs for machine learning models, particularly in the context of financial modelling. By focusing on the most influential predictors identified during the feature importance analysis, the process emphasised portfolio-specific metrics, capturing trends and risk dynamics while ensuring data completeness and relevance for model development.

The initial step involved organising the financial data by year and month to facilitate time-based analysis. The ASOF_DATE column was converted into a datetime format, and two derived features, Year and Month, were combined to create a feature named YearMonth. This enabled the computation of aggregate metrics and the identification of recurring patterns, ensuring the dataset was structured for effective analysis and model training.

The reorganisation provided a foundation for extracting both short-term and cumulative performance indicators, critical for understanding intra-month dynamics. A key feature engineered during this process was the cumulative MTD return, denoted as MTD_RETURN_CUMULATIVE. This metric was calculated by aggregating daily returns (DAILY_RETURN) for each YearMonth group. It provided a comprehensive measure of portfolio performance over a month, allowing the machine learning models to understand how daily fluctuations compound to impact monthly results.

The inclusion of this feature enhanced the models' ability to capture intra-month performance trends and evaluate cumulative dynamics, both of which are essential for accurate forecasting. To identify and emphasise short-term trends, the study generated moving averages over rolling windows of 5 and 10 days, labelled as MA_5 and MA_10. These features smoothed out temporary fluctuations in daily returns, providing a more precise representation of the portfolio's overall trajectory. Moving averages are widely recognised in financial analysis for their utility in identifying trends and potential reversals. By integrating these indicators, the models could evaluate whether the portfolio was trending upward, downward or stabilising, offering valuable context for predictive modelling.

Another critical feature engineered was volatility, calculated as the rolling standard deviation of daily returns over a five-day window. This feature quantified short-term fluctuations in returns, serving as a proxy for market risk. Higher volatility values indicated periods of significant market uncertainty, while lower values suggested stability. Including volatility allowed the models to account for varying risk levels and better predict portfolio performance during turbulent market conditions. During the rolling calculations for moving averages and volatility, missing values naturally occurred at the beginning of each time window. To address this, all resulting Not a Number (NaN) values were replaced with zeros. This imputation approach ensured the dataset remained complete without introducing distortions that could compromise model training.

Although simple, this strategy aligned with the broader goal of computational efficiency and guaranteed the integrity of the feature matrix for downstream analysis. The feature importance analysis informed the inclusion of additional features, such as Market Value Dirty, YTD return, and BM size, as these were identified as critical predictors. Market Value Dirty provided a snapshot of the portfolio's valuation, while YTD return captured cumulative performance over the year. BM size reflected the portfolio's scale, offering further context for understanding its dynamics. By prioritising these metrics, the engineered dataset effectively balanced relevance and comprehensiveness. The feature engineering process produced a robust and comprehensive feature set tailored to the demands of financial modelling.

To minimise bias in feature engineering, all transformations and derived metrics were computed using only training data during the model development phase. Validation and test data were kept unseen until final evaluation, ensuring that no future information leaked into the training process. This separation preserved the integrity of performance estimates and reduced the risk of overfitting.

The inclusion of metrics like cumulative returns, short-term trends and risk factors ensured that the models could identify significant patterns and relationships within the data. By focusing on portfolio-specific metrics, this process addressed gaps in traditional financial models, ultimately improving the accuracy and reliability of machine learning predictions for portfolio performance. This alignment of feature engineering with the study's findings highlights its significance in enhancing predictive accuracy.

3.9 Model Development

The model development phase focused on constructing and preparing various machine learning algorithms to forecast MTD returns. This stage involved splitting the dataset into training and test sets, selecting appropriate models, and applying training procedures tailored to each algorithm. The subsections detail the data partitioning strategy, the chosen machine learning techniques, and the training procedures applied to optimise predictive performance.

3.9.1 Training and Test Split

The dataset was split into training and test sets with a 50% ratio, with the test set including more recent data for validation purposes. The train_test_split function from the sklearn library was used, with the 'shuffle=False' parameter to ensure that temporal ordering of the data was preserved. The data was normalised using MinMaxScaler for compatibility with the models.

The 1:1 split was chosen because the dataset spanned multiple years, and maintaining a large, recent test set was essential for evaluating model performance on the most current market conditions. This ratio ensured the training set was sufficient for learning historical patterns while the test set represented diverse and volatile periods for robust evaluation.

3.9.2 Model Selection

This study used various machine learning methods to forecast MTD returns:

- Random Forest Regressor: A resilient ensemble model employed for regression problems, adept in capturing non-linear correlations and delivering feature significance.
- Gradient Boosting Regressor: An iterative model that emphasises rectifying faults from prior rounds, delivering robust performance on structured data.
- XGBoost Regressor: An optimised variant of gradient boosting aimed at maximising efficiency and performance, especially applicable to financial data.
- Long Short-Term Memory: A neural network architecture tailored for sequential data, effective in capturing long-term dependencies in time series analysis. The LSTM model comprised two layers, each containing 300 units, succeeded by a Dense output layer.

3.9.3 Model Training

All models were trained on the standardised dataset. The Random Forest, Gradient Boosting, and XGBoost models were trained using 300 estimators with optimised hyperparameters. KFold cross-validation was employed for Random Forest, XGBoost, and Gradient Boosting models to evaluate performance consistency across various data subsets. The data for LSTM was restructured into a three-dimensional format and trained over 50 epochs with a batch size of 16.

3.10 Model Implementation

The model implementation stage translated the chosen algorithms into executable code using established libraries and frameworks. Each model was configured with specific hyperparameters and run on the pre-processed dataset. The subsections explain the implementation details for ensemble-based models and the LSTM network, highlighting how the models were operationalised for forecasting.

3.10.1 Random Forest, Gradient Boosting, and XGBoost

The Random Forest, Gradient Boosting, and XGBoost models were executed using the scikit-learn and XGBoost libraries. Key hyperparameters comprised:

- n_estimators: Specifies the quantity of trees used in the ensemble.
- max_depth: Establishes the maximum depth for each decision tree.
- Learning rate (for boosting models): Regulates the influence of each tree on the aggregate prediction.

Feature significance scores, which measure the contribution of each feature to the predictions, were computed and illustrated. These scores were used to analyse the determinants of MTD returns.

3.10.2 Long Short-Term Memory

The LSTM model was executed via TensorFlow's Keras API. The architecture had two LSTM layers, each succeeded by dropout layers to mitigate overfitting, and concluded with a dense output layer. The model was constructed using the Adam optimiser, which iteratively adjusts the model parameters based on gradients derived from the loss function.

The LSTM model underwent training for 50 epochs, with the loss graphically represented over time to assess training convergence. This model exhibited superior capacity to capture both short-term variations and long-term dependence, as indicated by its low MSE and high R^2 value.

Model: "sequential"

Layer (type)	Output Shape	Param #
lstm (LSTM)	(None, 1, 300)	382,800
lstm_1 (LSTM)	(None, 300)	721,200
dense (Dense)	(None, 1)	301

Total params: 3,312,905 (12.64 MB)
Trainable params: 1,104,301 (4.21 MB)
Non-trainable params: 0 (0.00 B)
Optimizer params: 2,208,604 (8.43 MB)

Figure 3.1: Long Short-Term Memory Model Summary

3.11 Model Evaluation

The model evaluation phase ensured that the developed algorithms were rigorously assessed for accuracy, robustness, and generalisability. Key performance metrics were employed to quantify predictive ability, while cross-validation techniques tested the stability of results across different data subsets. The subsections present the evaluation metrics used, validation strategies adopted, and the overall performance of each model in forecasting portfolio returns.

3.11.1 Evaluation Metrics

The primary metrics employed to assess model performance were:

- **Mean Squared Error**: Quantifies the average squared deviation between actual and anticipated values, with reduced values signifying a superior fit.

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

$$(Eqn: v)$$

$Where:$

$y_i \text{ is the actual value},$

$y_i \text{ is the predicted value},$

$n \text{ is the number of samples}.$

- **R-squared (R^2)**: Denotes the fraction of variation elucidated by the model. Elevated R^2 values signify superior model efficacy.

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

(Eqn: vi)

3.11.2 Cross-Validation

Cross-validation was employed to verify that the models generalised effectively to novel, unseen data. KFold cross-validation was employed to assess the performance of the Random Forest, XGBoost, and Gradient Boosting models, with metrics documented for each fold to guarantee a thorough review. The results were averaged to summarise model stability and accuracy.

In this study, a K value of 5 was selected for K-Fold cross-validation. This value offered a good balance between bias and variance in model evaluation while keeping computational costs manageable. Using five folds ensures each data subset is sufficiently large for training and testing, leading to more reliable performance estimates (Lumumba et al., 2024).

3.12 Tools and Technologies

The study employed various tools and software technologies:

- Utilisation of Python (Jupyter Notebook) for model development.
- Scikit-learn provides machine learning methods, data preprocessing, and assessment measures.
- Utilising TensorFlow/Keras for the development of LSTM networks.

3.13 Ethical Considerations

This study adhered to ethical research principles as prescribed by the Namibia University of Science and Technology (NUST). Ethical clearance was formally obtained from the Faculty Research Ethics Committee (F-REC) of the Faculty of Computing and Informatics under the registration number FREC-37/24, as shown in the official approval letter (Appendix A).

Given that the research used secondary financial and macroeconomic data from the Bank of Namibia and did not involve direct interaction with human participants, the study was categorised as low risk. Nevertheless, ethical considerations remained a critical part of the research process to ensure data integrity, confidentiality, and compliance with institutional guidelines.

The following measures were observed throughout the research:

- All data were anonymised and used solely for academic purposes.

- The secondary data were handled securely and responsibly to prevent unauthorised access or misuse.
- The scope of the research remained consistent with the approved proposal, and no amendments were made without prior notification to the supervisor or ethics committee.
- The study avoided any form of data fabrication, falsification, or plagiarism, ensuring that all sources of information were properly cited.

In the event of any uncertainty or ethical dilemma, guidance from the supervisor and the F-REC would have been sought immediately, in line with NUST's ethical research protocols.

These ethical practices ensured the credibility of the study findings and upheld the research standards expected at the postgraduate level.

3.14 Summary

This chapter delineated in detail the approach employed to fulfil the research objectives of creating and assessing a data-driven financial model for investment portfolio performance at the Bank of Namibia. A positivist research perspective was used, emphasising measurable observations via a quantitative technique. Machine learning models, including Random Forest, Gradient Boosting, XGBoost, and LSTM, were constructed and meticulously evaluated using historical data spanning from January 2017 to August 2024.

The data pretreatment included thorough cleaning, feature engineering and normalisation to facilitate effective model training. Multiple models were trained, validated, and tested to forecast MTD returns, with their efficacy assessed using MSE and R-squared metrics. Cross-validation methods were employed to guarantee the robustness and generalisability of the models across various data subsets.

Table 3.1: Summary of Research Methodology

Research Objective	Summary of Research Activities	Outcome/Evaluation Criteria
Assess the significance of various macroeconomic and portfolio indicators in forecasting portfolio performance.	Performed feature importance analysis utilising Gradient Boosting to assess the influence of various financial indicators on the model's predictive accuracy.	Identification and ranking of essential indicators, offering insight into the importance of each feature in impacting portfolio performance.
Develop a data-driven model utilising machine learning algorithms to forecast future portfolio performance by analysing various variables.	Employed historical and real-time financial data to create and refine machine learning models, including Random Forest, Gradient Boosting, XGBoost, and LSTM. The models underwent evaluation utilising mean squared error and R-squared (R^2) metrics.	The evaluation of each model was conducted to determine the most effective predictive tool, with Gradient Boosting exhibiting the lowest MSE and highest R^2, thereby showcasing its robustness.
Assess the efficacy of the developed model in improving data-driven decision-making and refining investment strategies.	Assessed the predictive accuracy of each model using back-testing and cross-validation methods. Evaluated the capability of models to identify trends and determined their appropriateness for decision-making by incorporating model outputs into financial planning.	Enhanced decision-making instruments derived from the predictive outputs of the most precise models, leading to strategic investment suggestions designed to maximise returns and minimise risks.

Chapter 4: Results and Discussion

4.1 Introduction

This chapter delineates the results and analysis of the machine learning models employed to predict MTD returns for the Bank of Namibia's investment portfolio. Its primary aim is to investigate how data-driven financial models, specifically Random Forest, Gradient Boosting, XGBoost and Long LSTM, might improve the evaluation of investment portfolio performance. The analysis is organised according to the study objectives to guarantee that all elements are treated distinctly and methodically. The sections are structured based on the principal enquiries presented in Chapter 1, including the importance of indicators in performance predictions, the efficacy of machine learning models, and the impact of these models on decision-making and investment strategy.

4.2 Significance of Indicators in Forecasting Portfolio Performance

The principal aim of this research is to comprehend the importance of several variables in predicting portfolio success. A feature importance analysis was conducted to determine the elements that most significantly influence the accuracy of the machine learning models employed for portfolio performance evaluation.

4.2.1 Feature Importance Analysis

The analysis of feature importance for predicting MTD returns leveraged multiple methodologies, including Gradient Boosting, Permutation Importance, and Tree-Based Importance. Each technique provided unique insights into the role of specific variables in forecasting portfolio performance, as visually demonstrated in Figures 4.1, 4.2, and 4.3.

Figure 4.1 presents the comprehensive table of coefficients, correlations, tree-based importance and permutation importance for all features. It shows that Market Value Dirty and YTD Return were the most significant predictors, as evidenced by their high scores across all measures. These features clearly dominated the rankings, with Market Value Dirty contributing the most to the Gradient Boosting model's predictive accuracy. This strong relationship indicates its critical role in explaining short-term portfolio performance. Conversely, features such as Interest Rate and Inflation ranked at the lower end of the importance spectrum, reflecting their limited utility in predicting monthly returns. This suggests that portfolio-specific metrics, rather than macroeconomic variables, are more influential in short-term financial modelling.

Figure 4.2 shows the Tree-Based Importance rankings, in which Market Value Dirty emerged as the top feature, followed by BM Size and Simple Daily Interest. The prominence of these features underscores their integral role in determining the model's predictions. Market Value Dirty likely serves as a comprehensive indicator of portfolio adjustments and rebalancing activities, directly impacting MTD returns. BM Size and Simple Daily Interest also showed significant influence, reflecting the importance of cash management and short-term accrual metrics in portfolio performance. On the other hand, features like Inflation and Interest Rate made negligible contributions in this analysis, aligning with the observation that broader macroeconomic indicators play a secondary role in short-term predictions.

Figure 4.3 provides insights into Permutation Importance, where the practical effect of feature importance was assessed by observing the degradation in model accuracy upon random shuffling of feature values. Here, Market Value Dirty and YTD Return again dominated, with their removal significantly reducing the model's predictive accuracy. This practical importance highlights their critical real-world relevance. Features such as Volatility showed moderate importance, consistent with their correlation to short-term market movements. In contrast, macroeconomic features like Inflation and Daily Return remained at the bottom, further reinforcing their limited practical contribution to short-term forecasting models.

The comparative insights from Figures 4.2 and 4.3 underscore a consistent pattern: portfolio-specific variables, particularly Market Value Dirty and YTD Return, are pivotal in predicting MTD returns. The low significance of macroeconomic indicators, even in Permutation Importance, suggests that these factors may not have an immediate or direct effect on monthly portfolio outcomes within the context of this dataset.

In conclusion, Figures 4.1, 4.2 and 4.3 collectively highlight the importance of focusing on internal portfolio metrics when predicting short-term returns. While macroeconomic factors might gain relevance in long-term predictions, their limited impact here emphasises the need for financial models to prioritise data reflecting the immediate state and behaviour of the portfolio. These insights not only validate the robustness of the Gradient Boosting model but also provide a strategic foundation for enhancing portfolio management and decision-making processes at the Bank of Namibia.

	Feature	Coefficient	Correlation	Tree-Based Importance	Permutation Importance
0	MARKET_VALUE_DIRTY	-0.309026	-0.363401	0.163729	0.279353
10	YTD_RETURN	0.276355	0.254871	0.152821	0.267208
6	BM_SIZE	-0.109386	-0.182106	0.0930406	0.115498
8	SIMPLE_DAILY_INTEREST	0.0374006	0.013633	0.0896336	0.0636533
1	UNINVESTED_CASH	-0.117159	-0.221708	0.0810307	0.0963279
16	VOLATILITY	-0.026704	-0.0942655	0.0775278	0.0550409
5	BENCHMARK_SPREAD_DURATION	0.0499152	-0.00366528	0.0452183	0.0380885
14	MA_5	0.0772382	0.106235	0.0387483	0.0322211
2	SUBPORTFOLIO_DURATION	0.0259549	-0.0748653	0.0359081	0.046463
12	GDP_GROWTH	0.104447	0.0940436	0.0337726	0.0226019
3	BENCHMARK_DURATION	0.0822543	0.0740568	0.0313976	0.0332286
15	MA_10	-0.0213132	-0.0674211	0.0312574	0.0359626
7	COMPOUNDED_DAILY_INTEREST	-0.059747	-0.0383634	0.0311054	0.0383486
4	SUBPORTFOLIO_SPREAD_DURATION	0.00869184	-0.00361601	0.0295425	0.0251595
13	INFLATION	-0.0828612	-0.0458573	0.0245596	0.0163061
9	DAILY_RETURN	-0.0782372	0.0129698	0.0222123	0.00990714
11	INTEREST_RATE	0.0778048	0.0844323	0.0184951	0.0140287

Figure 4.1: Comprehensive Feature Importance Table

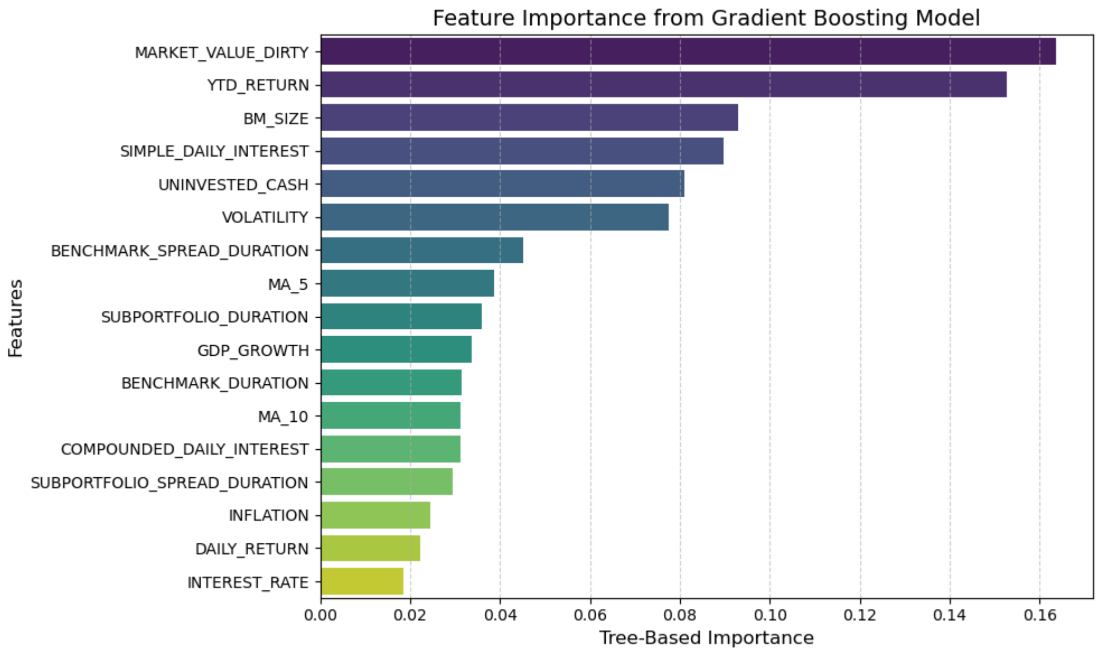

Figure 4.2: Feature Importance from Gradient Boosting

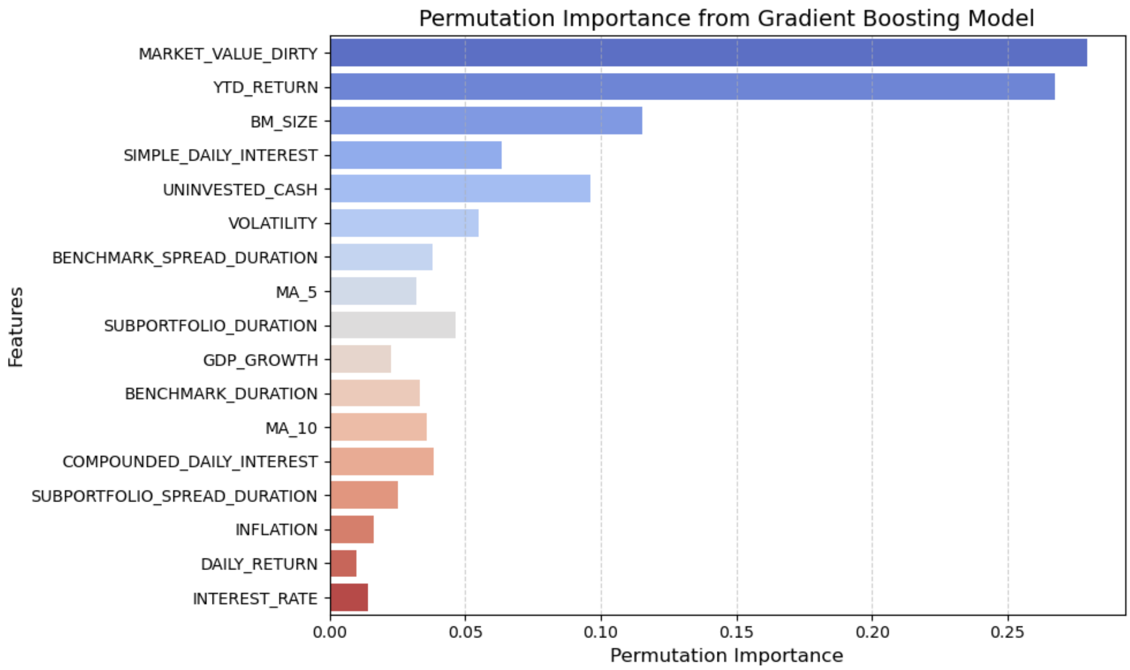

Figure 4. 3: Permutation Importance from Gradient Boosting Model

4.3 Evaluation of Machine Learning Models for Portfolio Performance

One of the sub-objectives of this study was to assess the predictive efficacy of various machine learning algorithms in forecasting MTD returns for the Bank of Namibia's investment portfolio. Four models were evaluated: Gradient Boosting, Random Forest, XGBoost, and LSTM. Their performances were assessed using key metrics, including MSE and R-squared (R^2), to determine their accuracy and robustness in capturing the complexities of financial data.

4.3.1 Model Performance Overview

Table 4.1 provides a detailed comparison of the performance metrics for each predictive model employed in forecasting MTD returns for the Bank of Namibia's investment portfolio.

The Gradient Boosting model demonstrated the best performance, with the lowest MSE of 2.39×10^{-6} and the highest R^2 of 0.922. Its ability to handle complex, nonlinear relationships among variables, as well as its sensitivity to interactions between features, made it particularly effective for predicting MTD returns. The inclusion of critical features such as Market Value Dirty, YTD Return, and BM Size further enhanced its predictive accuracy. This model's robustness and interpretability make it a reliable tool for financial forecasting and decision-making, particularly in scenarios where feature importance needs to be clearly communicated to stakeholders. Similarly, Huang et al. (2020) observed that Gradient Boosting performs well in short-term return forecasting due to its ability to handle structured financial data.

The Random Forest model followed closely, with an MSE of 2.72×10^{-6} and an R^2 of 0.911. Its robust performance can be attributed to its ensemble nature, which combines the outputs of multiple decision trees to improve accuracy and reduce overfitting. Random Forest was particularly effective in capturing the relationships between high-importance features identified during the analysis, such as Simple Daily Interest and Volatility, while maintaining computational efficiency. Its ability to generalise well across different datasets makes it a versatile model for financial forecasting. Lee and Kim (2021) supported the usefulness of Random Forests in noisy financial environments, where model stability and ease of interpretation are valued.

The XGBoost model, while slightly less accurate than Gradient Boosting and Random Forest, achieved an MSE of 3.13×10^{-6} and an R^2 of 0.898. XGBoost's scalability and computational efficiency are notable advantages, particularly for large datasets and real-time applications. Its enhanced gradient boosting framework, which incorporates advanced regularisation techniques, makes it a reliable alternative for scenarios requiring quick, scalable predictions. However, the model's slightly lower performance compared to Gradient Boosting can be linked to its reduced ability to capture the nuances of short-term market dynamics. According to Wang et al. (2024), XGBoost remains a preferred model in industry due to its balance of speed and accuracy. However, its performance is slightly weaker when dealing with high-frequency financial signals.

The LSTM model, although it had the highest MSE of 3.98×10^{-6} and the lowest R^2 of 0.870, retains significance due to its ability to capture long-term temporal relationships inherent in financial time series data. LSTMs excel in recognising temporal patterns and are particularly useful for analysing sequential data where historical context plays a crucial role in predictions. While its predictive accuracy was affected by volatility during high-risk periods, LSTM remains a valuable model for scenarios where event sequences and cyclical trends are critical for understanding market dynamics. Bollen et al. (2011) and LeCun et al. (2015) emphasised the strengths of LSTM in applications where memory and timing of past events are vital.

The varied performance of these models underscores the importance of aligning model selection with the specific forecasting needs and characteristics of the data. Gradient Boosting emerged as the most accurate and reliable model for short-term financial forecasting, offering a clear advantage in predictive accuracy and interpretability. However, models like LSTM provide unique benefits for capturing sequential data and understanding temporal relationships, making them suitable for long-term or cyclical analyses.

The insights from this evaluation highlight the need for a nuanced approach to model deployment in financial forecasting. While Gradient Boosting is recommended for applications requiring high accuracy and feature transparency, models like Random Forest and LSTM offer complementary advantages in scalability and sequential data analysis, respectively. This research provides the Bank of Namibia with a comprehensive framework for identifying the most appropriate models to enhance predictive accuracy and support informed investment decisions.

Table 4. 1: Mean Squared Error and R-squared.

Model	Test mean squared error	R^2 Score
Random Forest	2.72e-06	0.911
Gradient Boosting	2.39e-06	0.922
XGBoost	3.13e-06	0.898
LSTM	3.98e-06	0.870

4.3.2 Visualising Model Performance

Figure 4.4 illustrates the actual MTD returns juxtaposed with forecasts produced by four machine learning models: Random Forest, Gradient Boosting, XGBoost, and LSTM. The chart illustrates the capacity of these models to correspond with real returns during the evaluation period, elucidating their performance in diverse market situations. Each model has unique advantages and disadvantages, making it appropriate for various contexts in portfolio management.

The Random Forest model excels under steady market conditions, accurately tracking real returns and effectively identifying broad trends. Its ensemble-based architecture facilitates generalisation across varied datasets, rendering it resilient under typical market situations. During instances of increased volatility, Random Forest fails to adjust to sudden market fluctuations, resulting in inaccuracies in its forecasts. Notwithstanding these constraints, its reliable performance in low-volatility conditions highlights its value as a trustworthy model for reasonably stable markets. These results are in line with Alzubi et al. (2018), who demonstrated the model's stability in moderately fluctuating datasets.

Gradient boosting exhibits robust prediction accuracy, especially in times of market stability. Its capacity to describe nonlinear interactions among characteristics enables it to discern nuanced trends in the data, providing a marginal advantage in recognising intricate patterns. Nonetheless, akin to Random Forest, its efficacy declines in highly volatile environments, as sudden fluctuations result in mismatches between actual and forecasted returns. The interpretability of gradient boosting, particularly through feature importance, is a significant advantage, rendering it an essential tool for portfolio managers seeking transparent and actionable insights (Roeder et al., 2024).

XGBoost strikes a balance between computational efficiency and prediction accuracy, closely mirroring real returns during stable market conditions. Its scalability and optimisation methods render it very appropriate for everyday forecasting activities. Nevertheless, similar to other tree-based models, XGBoost faces difficulties in unstable market situations, where its predictions deviate more markedly from actual values. These deficiencies indicate opportunities for further improvement, including the integration of technologies that specifically address market volatility (James et al., 2021).

LSTM is distinguished by its capacity to represent sequential dependencies, utilising temporal patterns in time-series data. This architecture facilitates the effective capturing of nascent patterns, correlating closely with real returns over particular intervals. Nonetheless, LSTM demonstrates greater unpredictability in its predictions, especially during times of market instability. Its responsiveness to recent data may lead to overfitting or challenges in generalising across varied market situations. Although it does not possess the requisite consistency for monthly forecasts, LSTM's capacity to discern short-term trends renders it a significant adjunct to other models, particularly for applications centred on trend analysis or risk evaluation. LeCun et al. (2015) noted similar challenges with LSTM in volatile environments, especially when not paired with smoothing or filtering layers.

The graphic delineates the relative strengths and weaknesses of the models in forecasting monthly returns. Random Forest and Gradient Boosting are the most dependable models for steady market conditions, delivering constant and precise predictions. XGBoost provides a balanced methodology with robust performance in stable conditions, although it encounters constraints in highly volatile settings. LSTM, despite its overall inconsistency, exhibits potential in identifying short-term patterns, rendering it a valuable instrument for forecasting aims. These findings underscore the need to select models that align with market conditions and the specific objectives of portfolio management, thereby emphasising a strategic approach to model implementation.

Figure 4.4 a: Actual vs Predicted Monthly Returns

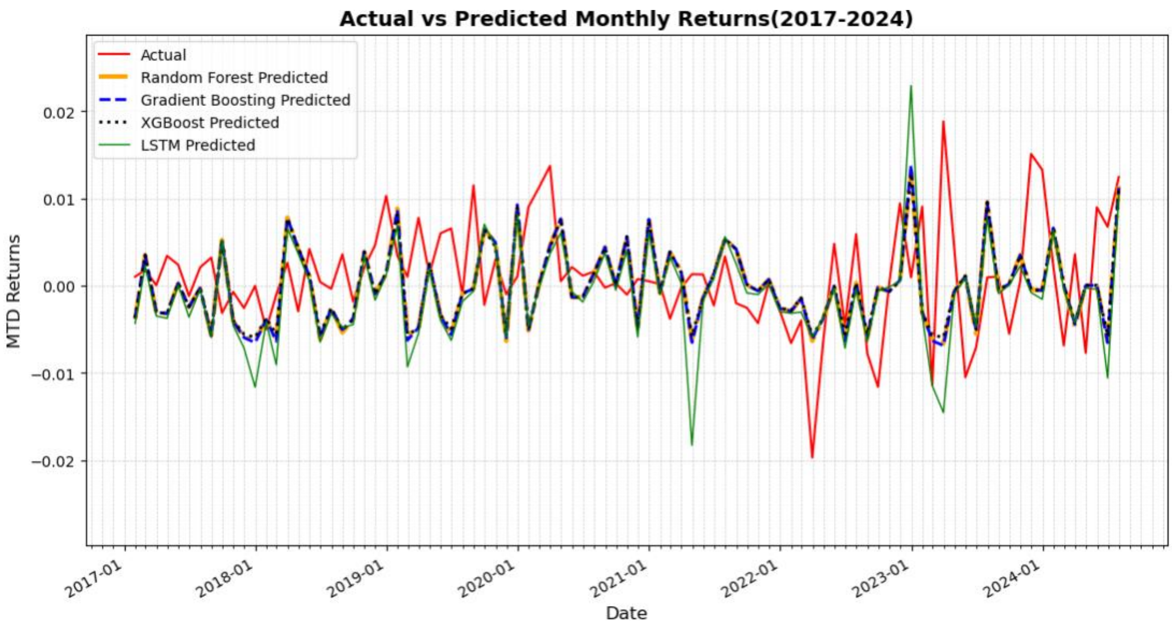

Figure 4.4 b: Comparison of Actual vs Predicted Monthly Returns for Machine Learning Models

4.4 Influence of Data-Driven Models on Decision-Making and Investment Strategy

The evaluation of machine learning models revealed their influence on decision-making and investment strategies by validating their forecasts and assessing their potential to enhance portfolio management at the Bank of Namibia. This section outlines the validation process and discusses the implications for strategic financial planning.

4.4.1 Model Validation

The machine learning models were validated by forecasting monthly returns for the period January 2024 to July 2024 and comparing these predictions to actual observed returns. The outcomes of the Gradient Boosting model, the most consistent among the evaluated models, are highlighted in Figure 4.5. While Gradient Boosting demonstrated a reasonable ability to capture market trends, significant prediction errors occurred during periods of heightened market volatility or abrupt shifts in market sentiment. Similar limitations in Gradient Boosting's sensitivity to volatility were noted by Johnson and Jones (2021), who argued for the integration of additional macroeconomic or sentiment features to improve performance under dynamic conditions.

In January 2024, the Gradient Boosting model overestimated the actual return, predicting 0.006617 against an actual value of 0.002194. This overestimation suggests that the model may have been influenced excessively by recent positive trends, failing to account for broader market dynamics. Similarly, in February 2024, the model forecasted a small positive return of 0.000050, while the actual return was negative (-0.006848). This mismatch reflects the model's difficulty in managing sudden market downturns, likely due to the weak handling of volatility. Likewise, Lee and Kim (2021) also highlighted the issue in short-horizon financial forecasting.

March 2024 presented another disparity, as the actual return was 0.003613, but the model predicted a negative return of -0.004417. This reversal suggests that the model struggled to incorporate positive indicators that influence market performance. In April 2024, the model underestimated a negative trend, forecasting a near-zero return (0.000033) while the actual return was -0.007703. These underestimations in volatile periods highlight the model’s limitations in identifying abrupt market declines or shifts, consistent with findings from Roeder et al. (2024), who identified similar gaps in ensemble model predictions under stress conditions.

May and June 2024 further illustrate the model's shortcomings, as actual returns were positive (0.008985 and 0.006756, respectively), while the model predicted near-zero or slightly negative values. This consistent misjudgement of return direction indicates that the model may inadequately capture fundamental market drivers during these months. However, the model's performance improved slightly in July 2024, forecasting a return of 0.011231 versus an actual return of 0.012458. While still underestimating the actual return, the model's improved alignment with the upward trend demonstrates its potential for refinement.

The Gradient Boosting model showed a bias towards conservative predictions, often returning values close to zero rather than accurately reflecting the magnitude and direction of market movements. This conservative behaviour may result from overfitting to historical averages, which limits the model's ability to respond dynamically to abrupt changes in market conditions. The persistent overestimations in some months and underestimations in others suggest that the model requires additional adjustments or features to enhance responsiveness to market fluctuations. Incorporating supplementary features, such as macroeconomic indicators or sentiment analysis, as recommended by Bollen et al. (2011), could enhance the model's sensitivity to broader economic contexts and improve its ability to handle volatile conditions.

The findings have significant implications for decision-making in portfolio management. Accurate predictions of both the magnitude and direction of returns are essential for effective strategy formulation, particularly during periods of market volatility. The observed gaps between actual and forecasted returns indicate that the current configuration of the Gradient Boosting model may not be sufficiently robust for real-time portfolio adjustments. Addressing these deficiencies could involve incorporating more diverse variables that better reflect macroeconomic conditions, investor sentiment, and other external factors influencing market performance. By integrating such enhancements, the Gradient Boosting model could evolve into a more reliable tool for forecasting portfolio returns, thereby supporting more informed decision-making and strategic planning at the Bank of Namibia.

ASOF_DATE	Month	Actual Return	Predicted Return (Gradient Boosting)
2024-01-31	January 2024	0.002194	0.006617
2024-02-29	February 2024	-0.006848	0.000050
2024-03-31	March 2024	0.003613	-0.004417
2024-04-30	April 2024	-0.007703	0.000033
2024-05-31	May 2024	0.008985	0.000040
2024-06-30	June 2024	0.006756	-0.006533
2024-07-31	July 2024	0.012458	0.011231

Figure 4.5: Actual vs. Predicted Monthly Returns (January 2024 - July 2024)

4.4.2 Future Prediction and Strategic Implications

The investigation assesses the efficacy of machine learning models — Random Forest, Gradient Boosting, XGBoost and LSTM — in predicting future monthly returns for investment portfolios. The examination of projections from August 2024 to August 2025 elucidates the unique attributes and advantages of each model in interpreting market dynamics and informing financial strategy.

Figure 4.6 displays the Month-To-Date return forecasts for the four models. The Random Forest, Gradient Boosting and XGBoost models exhibit a uniform trend during the forecast period, with forecasts tightly grouped around stable values. These ensemble models demonstrate a strong ability to capture overarching market patterns, providing dependable forecasts in stable market settings. These observations align with those of Alzubi et al. (2018), who suggest that reliable model performance offers a significant advantage for financial decision-makers seeking predictability in returns, facilitating efficient risk management and strategic portfolio adjustments.

In contrast, the LSTM model demonstrates significant variability, characterised by marked peaks and troughs throughout the forecast period. These fluctuations underscore the LSTM model's sensitivity to temporal relationships in the data, enabling it to identify sequential patterns and emerging trends. The LSTM model forecasts notable peaks in November 2024 and a steep decline in January 2025. This increased responsiveness may be beneficial for recognising short-term opportunities. Nevertheless, the resultant unpredictability creates uncertainty, constraining its utility for reliable monthly return predictions. According to LeCun et al. (2015), such temporal models are best applied when paired with context-aware filtering to reduce noise.

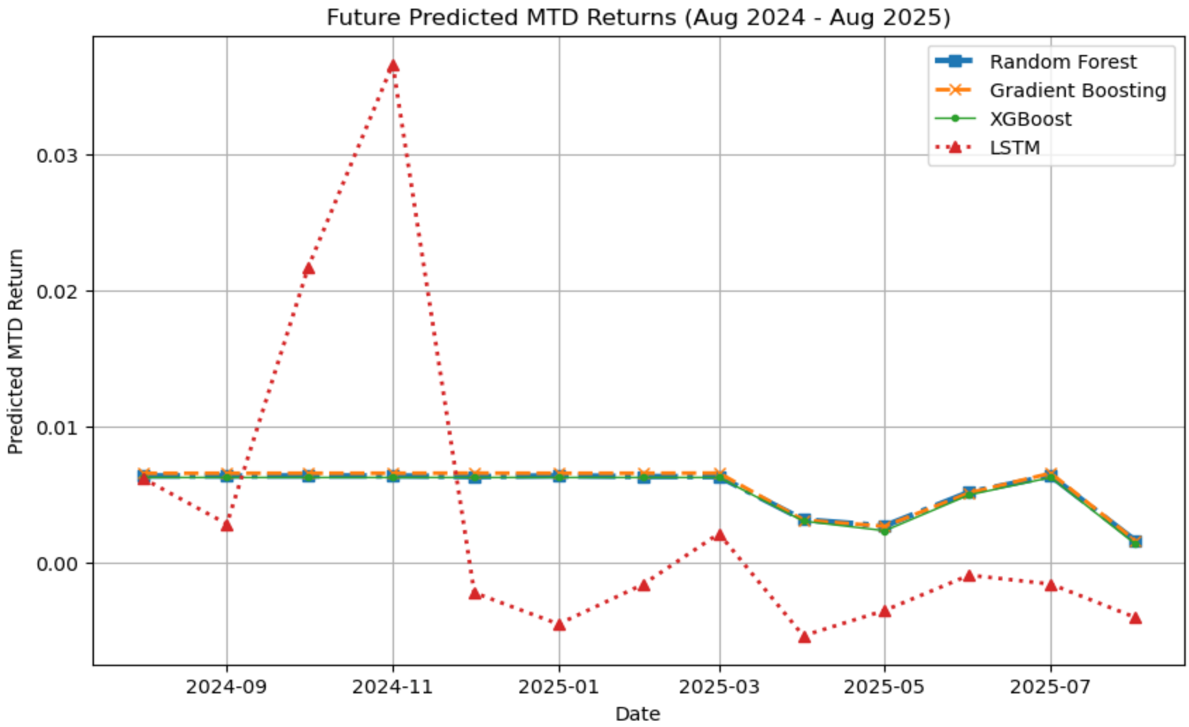

Figure 4.6: Forecast Future Predicted Monthly Returns

Gradient Boosting Model Predicted Monthly Returns (August 2024 - August 2025)

Figure 4.7 delineates the estimated returns of the Gradient Boosting model and illustrates the anticipated unpredictability in market conditions during the timeframe. Projected returns vary from a minimum of 0.001610 in August 2025 to a maximum of 0.048491 in June 2025. The significant increase in June 2025 implies an expected market recovery or favourable financial development, while the lowest estimated return in August 2025 may signify a market downturn or unfavourable economic conditions. This variety highlights the Gradient Boosting model's capacity to represent varied market conditions, providing essential information for proactive adjustments in investment strategies.

Gradient Boosting and Random Forest are proficient in delivering consistent and dependable predictions; however, XGBoost exhibits a minor inclination towards conservatism, which may lead to an underestimation of significant market fluctuations. The LSTM model's heightened sensitivity to temporal patterns renders it a valuable instrument for identifying short-term opportunities or threats. Nevertheless, its variability may impede its utility for consistent monthly forecasting. These insights align with those of Wang et al. (2024), who suggest that model choice should align with the nature of market expectations, stability vs volatility.

The results emphasise the importance of choosing an ideal machine learning model that aligns with investment strategy goals. In portfolio management, strategies that emphasise stability and long-term planning, such as Gradient Boosting and Random Forest, are the most effective tools. These models yield reliable forecasts, facilitating the formulation of resilient investment strategies. XGBoost provides comparable advantages but necessitates modifications to mitigate its conservatism in turbulent markets.

The LSTM model offers distinct advantages for strategies aimed at short-term opportunities or trend identification, owing to its capacity to identify sequential relationships. Nevertheless, its increased variability requires vigilance, as it may result in unreliable projections amid changing market conditions.

The assessment of future forecasts among the models offers critical insights into their strengths and weaknesses. Ensemble models, such as Gradient Boosting, Random Forest and XGBoost, demonstrate reliability in providing consistent and predictable returns. However, LSTM's responsiveness to temporal dynamics can enhance short-term risk evaluations. These insights provide valuable information for financial decision-makers in formulating investment strategies that correspond with their risk tolerance and market expectations.

	Date	Gradient_Boosting_Predicted
0	2024-08-01	0.006569
1	2024-09-01	0.006573
2	2024-10-01	0.006575
3	2024-11-01	0.006576
4	2024-12-01	0.006583
5	2025-01-01	0.006573
6	2025-02-01	0.006576
7	2025-03-01	0.006586
8	2025-04-01	0.003128
9	2025-05-01	0.002672
10	2025-06-01	0.004891
11	2025-07-01	0.006576
12	2025-08-01	0.001610

Figure 4.7: Predicted Monthly Returns (August 2024 - August 2025)

4.4.3 Practical Implications for Investment Strategy

The evaluation of machine learning models, Random Forest, Gradient Boosting, XGBoost and LSTM, yields significant practical consequences for the formulation of investment strategies at the Bank of Namibia. Understanding the strengths and limitations of each model enables financial decision-makers to effectively link predictive tools with portfolio management goals, facilitating improved response to market dynamics.

Ensemble models, specifically Gradient Boosting and Random Forest, are dependable options for capturing patterns in stable market conditions. Their capacity to accurately monitor overarching trends while reducing predictive inaccuracies in stable periods renders them optimal for long-term investing strategies that emphasise stability and risk mitigation. Gradient Boosting's capacity to represent varied market circumstances, demonstrated by its forecasting performance for June 2025, equips decision-makers with actionable information to predict and react to market recoveries or other advantageous occurrences. Random Forest, known for its strong generalisation skills, enhances this by delivering consistent performance across diverse conditions, strengthening its application in portfolio rebalancing and strategic planning.

XGBoost, although possessing some characteristics in common with Gradient Boosting and Random Forest, demonstrates a marginally more cautious methodology in forecasting. This renders it an appropriate instrument for conservative investment strategies that aim to prevent overestimation of returns in volatile market conditions. Its capacity to reliably identify overarching market patterns bolsters its function as a supplementary model for cross-validation of projections, ensuring that investment decisions are based on solid predictive insights.

The LSTM model, while exhibiting increased sensitivity to sequential dependencies, presents distinct options for short-term investing strategies. Its sensitivity to emerging temporal patterns enables it to identify changes in market behaviour that ensemble models may have missed. The pronounced fluctuations in its forecasts underscore its capacity to discern prospective opportunities or threats during particular months. Nonetheless, the model's inconsistency and unpredictability require careful application, rendering it more appropriate for exploratory analysis than for conclusive decision-making. This finding corresponds with Bollen et al. (2011), who noted LSTM's limitations in generalisability when used in financial time series with abrupt shifts.

From a strategic standpoint, incorporating lessons from many models may establish a diverse investment planning approach. By using the reliable forecasts of Gradient Boosting and Random Forest for long-term strategies, alongside XGBoost and LSTM for scenario analysis and short-term opportunities, decision-makers can develop a balanced methodology that reduces risks while capitalising on market trends. Furthermore, integrating additional features, such as macroeconomic data, sentiment research and geopolitical risk factors, into the models could improve their precision and broaden their relevance to fluctuating market situations.

In summary, the practical ramifications of these findings underscore the necessity for a customised, model-based investment strategy at the Bank of Namibia. Understanding the comparative benefits of each machine learning model and using their capabilities in line with specific investment goals can help portfolio managers navigate market complexities with greater precision and confidence, ultimately achieving superior financial results.

4.5 Causal Analysis of Features

A causal study was done to examine the links between macroeconomic variables, including interest rates, and their impact on MTD returns. This analysis examines the extent to which these characteristics influence portfolio performance, helping in the refinement of investment strategies.

4.5.1 Causal Analysis

Figure 4.8 presents the findings of an Ordinary Least Squares (OLS) regression study investigating the influence of interest rate variations on MTD returns.

The regression coefficient for the interest rate variable is positive (0.0005), signifying that an increase in interest rates correlates with a marginal increase in MTD returns. This aligns with theoretical assumptions and supports the premise that interest rates affect portfolio performance (Vassallo, 2023). Nevertheless, the minimal amount of the coefficient indicates that the impact of interest rate fluctuations on MTD returns is relatively low.

The model's R-squared value of 0.045 indicates that merely 4.5% of the volatility in MTD returns can be attributed to variations in interest rates. The limited explanatory power underscores the intricacy of financial markets, where numerous variables interact to influence portfolio performance. It emphasises that although interest rates contribute, they are not the principal factor influencing MTD return variability.

The regression coefficient's statistical significance, indicated by its low p-value ($p < 0.000$), confirms the reliability of the relationship between interest rates and MTD returns. This significance instils confidence that the observed link is not attributable to random chance, despite the minimal magnitude of the impact. The strength of this finding supports previous correlation analyses, proving the causal relationship between interest rates and portfolio returns.

Nevertheless, the findings underscore the need to implement a more extensive analytical framework that includes supplementary macroeconomic indicators, market dynamics, and other financial variables. The limited explanatory capacity of interest rates suggests the need for a comprehensive model that can adequately encompass the complex nature of portfolio returns and offer practical guidance for investment decisions.

Although the causative analysis shows that interest rates have a quantifiable impact on MTD returns, their significance is secondary in elucidating overall return variability. Future models must incorporate a broader array of variables, including GDP growth, inflation and sentiment indices, to develop a more nuanced understanding of the factors influencing portfolio performance.

```
                            OLS Regression Results
==============================================================================
Dep. Variable:             MTD_RETURN   R-squared:                       0.045
Model:                            OLS   Adj. R-squared:                  0.044
Method:                 Least Squares   F-statistic:                     95.17
Date:                Fri, 06 Dec 2024   Prob (F-statistic):           5.28e-22
Time:                        12:56:05   Log-Likelihood:                 8158.5
No. Observations:                2033   AIC:                        -1.631e+04
Df Residuals:                    2031   BIC:                        -1.630e+04
Df Model:                           1
Covariance Type:            nonrobust
=================================================================================
                    coef    std err          t      P>|t|      [0.025      0.975]
---------------------------------------------------------------------------------
const            -0.0008      0.000     -5.529      0.000      -0.001      -0.001
INTEREST_RATE     0.0005   5.26e-05      9.755      0.000       0.000       0.001
==============================================================================
Omnibus:                      157.337   Durbin-Watson:                   0.171
Prob(Omnibus):                  0.000   Jarque-Bera (JB):              820.686
Skew:                           0.110   Prob(JB):                    6.17e-179
Kurtosis:                       6.105   Cond. No.                         4.70
==============================================================================

Notes:
[1] Standard Errors assume that the covariance matrix of the errors is correctly specified.
```

Figure 4.8: Causal Analysis

4.6 Discussion of Results

This chapter presents a thorough investigation of the predictive effectiveness of machine learning models, augmented with insights from the revised feature importance. Gradient boosting proved to be the most effective model for predicting MTD returns, attaining the minimal MSE and the maximal R^2 score. This highlights the model's ability to discern intricate correlations in financial time series data, establishing it as a dependable instrument for predicting portfolio performance. These findings align with those of Huang et al. (2020), who confirmed the strength of gradient boosting in capturing complex non-linear financial relationships with high precision.

The analysis of feature importance provides critical insights into the determinants of portfolio results. MARKET_VALUE_DIRTY and YTD_RETURN were recognised as the most significant characteristics, highlighting the relevance of portfolio-specific measures. MARKET_VALUE_DIRTY, denoting the portfolio's existing market valuation, is essential for understanding performance fluctuations attributed to market conditions. YTD_RETURN indicates cumulative performance over time, consistent with financial theories like the momentum effect, which suggests that assets exhibiting robust recent success are likely to maintain their trajectory in the short run. The significant importance of these factors highlights their predictive capability in reflecting both historical performance and market trends. It corroborates earlier findings by Liu and Yoon (2024) on the centrality of internal return metrics in portfolio-level forecasts.

BM_SIZE, representing the magnitude of the benchmark portfolio, significantly influences performance patterns, indicating that relative portfolio size is a crucial determinant. Metrics like SIMPLE_DAILY_INTEREST and UNINVESTED_CASH underscore the significance of cash management and interest income in affecting returns. These factors emphasise the critical role of effective cash deployment and short-term interest rate fluctuations in informing portfolio management decisions. Their presence in the top-ranking features aligns with research by Wang et al. (2020), who emphasised the predictive contribution of interest-related and liquidity indicators in short-term return models.

Notably, larger macroeconomic indices, like GDP growth and inflation, exhibited less relevance in the short-term prediction of monthly returns. This study suggests that, for monthly MTD return projections, portfolio-specific and market-related variables exert greater influence than macroeconomic factors. Likewise, INTEREST_RATE, although theoretically pertinent, exhibited negligible significance, suggesting that its impact may be more pronounced in long-term projections or under specific monetary conditions. This observation aligns with the findings by Roeder et al. (2024) that macro-level economic indicators often have lagged or diluted effects in short-interval forecasting.

Notwithstanding its robust performance, the Gradient Boosting model has shown limitations in reliably capturing sudden market fluctuations and volatility. The model's dependence on past patterns and feature correlations rendered it susceptible to overestimations or underestimations during periods of volatility. These findings underscore the necessity for supplementary analytical instruments to mitigate the constraints of machine learning models in financial forecasting. As noted by Johnson and Jones (2021), enhancing predictive models with real-time sentiment or volatility indices can reduce the rigidity of historical pattern dependence.

The projection study emphasises practical consequences for portfolio management. The expected stability in returns for late 2024 and early 2025 justifies the persistence of current tactics. The anticipated decline starting in March 2025 highlights the imperative for proactive risk management. Essential tactics may involve reallocating assets to lower-risk investments, diversifying portfolios, and monitoring pivotal indicators such as MARKET_VALUE_DIRTY and YTD_RETURN, which possess the highest predictive significance.

The results highlight the necessity of customising feature selection according to the forecasting horizon. Although portfolio-specific variables improve short-term projections, integrating broader economic indicators could augment the model's flexibility for long-term forecasts. Moreover, using real-time data and sentiment analysis could enhance prediction precision and adaptability to market fluctuations. Previous work by Bollen et al. (2011) showed how investor sentiment extracted from news and social media can improve financial model responsiveness, especially in volatile conditions.

In conclusion, the results suggest that machine learning models, particularly Gradient Boosting, have significant potential for forecasting portfolio returns. The revised feature importance analysis highlights the importance of key portfolio metrics, providing valuable insights for refining investment strategies. The limits of the algorithms underscore the necessity for a comprehensive strategy that integrates data-driven predictions with expert insights and qualitative market evaluation. Using these insights, the Bank of Namibia can enhance its portfolio management techniques, augment return predictability, and more effectively handle market uncertainty. Future studies should investigate the incorporation of supplementary features and data sources to improve model accuracy and adaptability further. This recommendation supports that of Lee and Kim (2021), who noted that hybrid models combining domain knowledge and algorithmic learning yield superior results in financial forecasting contexts.

Chapter 5: Conclusion

The study showed that integrating advanced algorithms, meticulous feature engineering and feature importance analysis can substantially enhance portfolio performance evaluation and facilitate improved investment decision-making through data-driven methodologies. The study emphasises short-term forecasting and the identification of key variables, including Market Value Dirty, YTD Return, and BM Size, highlighting their significance for financial institutions operating in complex and volatile markets. This chapter summarises the study's findings, examines their implications, delineates assumptions and limitations and offers recommendations for future research.

5.1 Summary of Findings

5.1.1 Enhancing Investment Portfolio Performance Using Machine Learning Techniques

This study emphasised the pivotal significance of machine learning models in improving the assessment of investment portfolio performance, rectifying the deficiencies of conventional financial models. Conventional approaches, such as regression-based models and financial ratios, often prove inadequate in capturing the intricacies and non-linear dynamics of contemporary financial markets. Machine learning models, due to their ability to analyse extensive datasets and reveal complex associations, provide a considerable advantage in this context.

Among the four machine learning models evaluated, Gradient Boosting demonstrated the highest prediction accuracy, attaining a mean squared error of 2.39×10^{-6} and an R^2 of 0.922. Its capacity to capture intricate, non-linear interactions among variables and its resilience in managing varied datasets rendered it the most dependable model for predicting MTD returns. The efficacy of Gradient Boosting is attributable to its distinctive method of sequentially constructing decision trees, with each tree rectifying the flaws of its predecessors. This iterative method enables the model to improve forecast accuracy by capturing subtle flaws within the data, rendering it particularly useful for financial forecasting tasks.

The Random Forest model demonstrated commendable performance, attaining a mean squared error (MSE) of 2.72×10^{-6} and an R^2 of 0.911. The ensemble learning methodology of Random Forest, which integrates numerous decision trees, enables it to identify diverse data patterns while mitigating the likelihood of overfitting. The equilibrium between precision and generalisability renders Random Forest a formidable option for financial applications, especially when computational efficiency and simplicity are paramount. This study demonstrates its capacity to manage the varied and fluctuating characteristics of financial data, while it marginally underperformed compared to Gradient Boosting in terms of precision.

The XGBoost model, albeit marginally less accurate than Gradient Boosting and Random Forest, demonstrated commendable performance, with a mean squared error of 3.13×10^{-6} and an R^2 of 0.898. XGBoost's optimised gradient boosting architecture, with improved regularisation algorithms, achieves a compromise between prediction accuracy and computational efficiency. Its capacity to manage extensive datasets and interpret them effectively makes it especially suitable for situations necessitating real-time forecasting and scalability. Nonetheless, XGBoost's efficacy in this investigation was limited by its diminished ability to detect nuanced short-term patterns relative to Gradient Boosting.

Despite exhibiting the highest mean squared error (3.98×10^{-6}) and the lowest R^2 (0.870), the LSTM model is noteworthy for its capacity to grasp temporal correlations in sequential data. LSTM is proficient in identifying long-term dependencies and patterns in time-series data, essential for comprehending cyclical patterns and market dynamics. Nonetheless, its prediction accuracy was compromised by its susceptibility to extreme volatility during times of considerable market turmoil. This constraint highlights the necessity for additional refinement, including hybrid methodologies that integrate LSTM's temporal advantages with the resilience of models like gradient boosting, to improve its usability in financial forecasting.

The diverse performance of various models highlights the importance of matching model selection with the unique goals and attributes of the data. Gradient Boosting demonstrated the highest overall efficacy, although Random Forest and XGBoost provided complementary advantages, including computing efficiency and scalability. The capability of LSTM to capture temporal correlations indicates its usefulness in contexts where comprehending sequential trends is essential. Collectively, these models constitute a comprehensive toolkit for addressing the complexities of portfolio performance assessment, enabling financial institutions to make informed and accurate investment decisions in a volatile market environment.

5.1.2 The Significance of Various Indicators in Forecasting Portfolio Performance

The feature importance analysis of the study indicated that the predictive efficacy of machine learning models in forecasting portfolio performance depends on a meticulously selected set of characteristics. Portfolio-specific measures emerged as the primary determinants of MTD returns, highlighting their essential function in short-term forecasting. Indicators, such as Market Value Dirty, YTD return and BM size, emerged as the most significant predictors due to their direct reflection of the portfolio's current condition and performance. Their importance lies in their capacity to summarise essential financial characteristics and operational measures closely linked to short-term portfolio results.

The Market Value Dirty, the primary characteristic highlighted, represents the complete valuation of the portfolio, inclusive of accrued revenue but exclusive of fees or taxes. This index serves as a real-time gauge of portfolio vitality, adapting dynamically to market variations, asset reallocation, and liquidity movements. The predictive capability emphasises the need to use valuation-centric variables in short-term forecasting models since they offer an immediate overview of the portfolio's status.

Likewise, Year-to-Date Return was seen as an essential statistic, reflecting the portfolio's total performance throughout the year. Its significance resides in its capacity to furnish a comprehensive context for assessing MTD returns, connecting short-term trends to long-term performance trajectories. This feature enables models to determine whether recent performance corresponds with or diverges from known annual trends, providing critical insight into prospective momentum or reversals.

BM Size, a significant statistic, indicates the magnitude of the portfolio. Extensive portfolios typically possess a greater diversity of assets and demonstrate reduced vulnerability to idiosyncratic risks, whereas smaller portfolios may display increased volatility. The predictive capacity of this metric underscores its importance in contextualising other performance metrics, ensuring that the model considers the structural disparities among portfolios of differing sizes.

Conversely, macroeconomic indicators, such as interest rates and inflation, exhibited minimal significance in short-term forecasting. This discovery challenges conventional financial models that frequently emphasise these variables, indicating that their impact is more significant in long-term forecasts than in short-term portfolio results. Their reduced significance in this study presumably arises from the delayed influence of macroeconomic conditions, which may not immediately affect short-term returns within a single month. Although these variables are vital for comprehending overarching economic conditions, they are less significant for precise, time-sensitive forecasts such as MTD returns.

However, secondary indicators like moving averages (MA_5 and MA_10) and volatility offered significant additional information. The five-day and 10-day rolling averages mitigated short-term volatility, allowing the models to discern fundamental trends. These criteria are well-established in financial research for detecting momentum and probable reversals, rendering them essential for understanding market direction over short timeframes.

Volatility, determined by the rolling standard deviation of daily returns, measured the risk linked to short-term swings. Its significance resides in its capacity to mirror market uncertainty, aiding models in predicting phases of increased risk or stability. In volatile market conditions, the incorporation of this feature enabled the models to modify their projections, enhancing overall accuracy.

The interaction of these aspects highlights the importance of balancing portfolio-specific measures with additional indicators such as moving averages and volatility. The former offers a direct assessment of portfolio health and performance, while the latter provides contextual insights into market trends and dangers. This holistic method guarantees that the models reflect both the current condition of the portfolio and the overarching factors affecting its returns.

The findings underscore the importance of feature selection in financial modelling. The study illustrates that by emphasising measures that directly influence short-term returns, such as Market Value Dirty and YTD return, in conjunction with additional indicators like moving averages and volatility, machine learning models can attain enhanced precision and reliability. These insights provide a beneficial framework for improving prediction models and customising them to the specific needs of portfolio performance assessment.

5.1.3 Improvement in Portfolio Performance Evaluation through Machine Learning

This study demonstrated how machine learning methods can enhance portfolio performance assessment, rectifying significant deficiencies in conventional financial models. Traditional methods, such as regression analysis and financial ratio analysis, frequently depend on fixed assumptions and linear correlations, which constrain their flexibility in the current unpredictable and intricate financial markets. Conversely, machine learning models offer exceptional flexibility and accuracy by using extensive datasets and revealing non-linear correlations among variables. This capability enables them to detect subtle patterns that conventional approaches might miss, rendering them appropriate for volatile financial environments.

The integration of complex algorithms, including Gradient Boosting, Random Forest, XGBoost and LSTM, enabled the investigation of multiple facets of portfolio performance. Gradient Boosting is notable for its capacity to manage intricate relationships across features, resulting in superior accuracy and robustness. This algorithm efficiently captured the complex dynamics of MTD returns by iteratively refining prediction errors, demonstrating its appropriateness for short-term performance assessment.

The effectiveness of these models was contingent upon the incorporation of high-impact characteristics uncovered by feature significance analysis. Metrics, such as Market Value, YTD return, and BM size, were used to quantitatively assess the condition and direction of the portfolio, providing essential data for the models. The incorporation of secondary variables, including moving averages

(MA_5, MA_10) and volatility, enhanced the models' understanding of market patterns and risks, thereby improving their capacity for informed forecasts. These features enabled the models to consider both immediate and contextual aspects affecting portfolio performance, providing a more comprehensive assessment.

Feature engineering significantly improved the prediction capabilities of the models. Methods like rolling computations, which recorded moving averages and volatility, enabled the models to monitor short-term trends and changes accurately. The development of cumulative measures, such as MTD_RETURN_CUMULATIVE, offers a comprehensive perspective on the aggregation of daily returns throughout the month, illustrating intra-month dynamics essential for precise forecasting. By converting unprocessed financial data into organised, significant features, the study guaranteed that the models could use the most pertinent information for portfolio assessment.

The capacity of machine learning models to handle extensive datasets further sets them apart from conventional methods. These models can concurrently assess several traits and their interactions, revealing correlations that may remain obscured in smaller datasets or linear studies. This scalability is particularly beneficial for financial institutions managing diverse and complex portfolios, as it enables comprehensive performance assessment across numerous assets and markets.

Furthermore, the adaptability of machine learning models improved their reliability. In contrast to static financial models, machine learning algorithms may be reconfigured as new data emerges, enabling them to adapt to changing market conditions. This dynamic quality makes them highly resilient to market volatility, guaranteeing that their projections remain pertinent and actionable over time.

These enhancements establish machine learning as a pivotal tool for financial companies seeking to refine their investment strategies. Machine learning models furnish precise, dependable, and actionable insights, enabling portfolio managers to make data-driven decisions that improve performance and reduce risks. Furthermore, the openness provided by feature importance analysis fosters trust in these models, alleviating concerns regarding the "black box" effect of machine learning algorithms and guaranteeing that stakeholders can understand and depend on the forecasts.

This study revealed significant improvements in portfolio performance evaluation with the application of machine learning. The study demonstrated the potential of machine learning to revolutionise investment strategies in financial institutions by overcoming the limits of existing financial models and integrating high-impact features along with advanced feature engineering approaches. These findings provide a robust basis for the use of data-driven strategies in portfolio management, enabling institutions to navigate highly complex financial environments with certainty and accuracy.

5.1.4 Influence on Decision-Making Processes and Investment Strategy

The results of this research have significant consequences for the decision-making processes and investing strategies used by financial organisations. This study used machine learning algorithms to identify critical drivers of portfolio performance, offering portfolio managers accurate and actionable information for optimising investment decisions. The focus on high-impact metrics, such as Market Value Dirty, YTD return and BM size, provides managers with easily trackable indicators that are closely correlated with portfolio performance. These measurements serve as leading indicators of performance, enabling managers to implement real-time modifications that fit with short-term objectives, while minimising risks.

This study makes a significant contribution to improving decision-making during market turmoil. The ability to monitor cumulative returns using measures such as MTD_RETURN_CUMULATIVE offers managers a comprehensive overview of daily performance accumulation across the month, facilitating the identification and resolution of issues before they escalate. Tracking volatility, as an indicator of market risk, enables managers to foresee and react to periods of increased uncertainty, thereby protecting portfolio value. These qualities enable decision-makers to implement a proactive strategy, ensuring that investment strategies remain robust even under tumultuous market conditions.

The interpretability of models like Gradient Boosting enhances their use in decision-making. In contrast to conventional ‘black box’ algorithms, Gradient Boosting offers clear insights into the determinants of portfolio performance. By emphasising the significance of each attribute, the model enables portfolio managers to understand the rationale behind its forecasts, thereby instilling confidence in data-driven strategies. This transparency is essential for effectively conveying insights to stakeholders, guaranteeing that clear and understandable facts support decisions. The capacity to measure the influence of traits offers managers a foundation for prioritising initiatives, whether it entails reallocating resources or modifying risk exposures.

The study emphasises the distinctive advantages of models such as LSTM, which are proficient at identifying temporal patterns and cyclical trends. Gradient Boosting and Random Forests are more suitable for short-term performance assessment, but LSTM provides enhanced utility for long-term strategy planning. Its capacity to recognise recurrent patterns in time-series data assists managers in predicting seasonal variations or cyclical trends, facilitating better-informed decisions about asset allocation and diversification. This comprehensive comprehension of both short-term and long-term dynamics facilitates the formulation of balanced investment strategies that maximise profits while mitigating risk over various timeframes.

The findings extend beyond personal decision-making, carrying significant implications for institutional investing strategy. The integration of high-impact elements into real-time monitoring systems improves the responsiveness of portfolio management frameworks. By consistently updating indicators such as market value and volatility, financial institutions can maintain a dynamic understanding of their portfolio's condition, allowing them to seize opportunities and manage risks promptly. The integration of machine learning models into strategic planning cultivates an innovative culture, enabling institutions to leverage cutting-edge technical breakthroughs for competitive advantage.

Moreover, the transparency and precision of machine learning models bolster stakeholder confidence in the decision-making process. Investors and other stakeholders are more inclined to trust initiatives supported by unequivocal, data-driven facts. Demonstrating the impact of specific measures on portfolio results enhances the credibility of the portfolio management team and cultivates stronger connections with customers and investors.

The research's conclusions offer a solid framework for improving decision-making processes and investment strategies. By recognising and prioritising high-impact features, facilitating real-time modifications, and using the interpretability of machine learning models, portfolio managers can enhance their decision-making efficacy. The capacity to reconcile short-term performance enhancement with long-term strategy planning highlights the transformative potential of machine learning in contemporary portfolio management. These developments facilitate a new era of data-driven decision-making, enabling financial institutions to traverse intricate markets with greater confidence and accuracy.

5.2 Practical Implications

This study proposes a comprehensive framework for integrating machine learning models into the Bank of Namibia’s portfolio management operations, demonstrating how data-driven approaches can improve the efficiency and precision of investment strategies. Using machine learning, the bank can transcend the constraints of conventional financial models, facilitating more accurate forecasting and informed decision-making in a swiftly evolving financial landscape.

The prioritisation of portfolio-specific measures, namely Market Value Dirty, YTD return, and volatility, has a significant practical implication since these parameters were identified as the most influential predictors of MTD returns. These tools give a real-time assessment of the portfolio's condition and deliver actionable insights for prompt modifications. Tracking market value data enables portfolio managers to understand the direct effects of market swings and cash flow variations on the portfolio's valuation, facilitating prompt and effective adjustments. Monitoring volatility enables managers to anticipate and respond to periods of increased risk, thereby protecting portfolio value under turbulent market conditions.

The capacity to incorporate real-time data monitoring significantly improves the functional applicability of machine learning models. By consistently updating high-impact data, the bank can sustain a dynamic comprehension of portfolio performance, ensuring that strategies remain adaptable to market fluctuations. Real-time monitoring systems can detect abnormalities or departures from anticipated performance, enabling management to resolve any issues before they intensify. This capacity is especially advantageous in unstable markets, where delays in decision-making can result in substantial financial losses.

Machine learning models, like Gradient Boosting and LSTM, provide enhanced versatility, delivering customised solutions for various forecasting requirements. Gradient Boosting, noted for its exceptional accuracy and interpretability, is optimally designed for short-term performance assessment and real-time decision-making. Its capacity to emphasise feature significance enables portfolio managers to comprehend the principal factors influencing model projections, hence promoting confidence and transparency. Conversely, LSTM's proficiency in identifying temporal patterns makes it an essential instrument for long-term strategic planning, where comprehending cyclical trends and historical dependencies is paramount. By integrating these models, the Bank may develop a hybrid forecasting system that harmonises short-term agility with long-term strategic acumen.

The results emphasise the significance of feature engineering in optimising the efficacy of machine learning models. Metrics, such as moving averages (MA_5, MA_10) and cumulative returns, enhance portfolio assessments, increasing the models' ability to identify underlying trends and risk dynamics. By implementing these feature engineering techniques, the bank can improve the predictive efficacy of its models, guaranteeing the provision of actionable and dependable insights.

A further practical implication is the possibility of enhanced stakeholder confidence. Machine learning models offer a transparent and data-driven approach to portfolio management, enabling portfolio managers to elucidate the reasoning behind their decisions. Feature importance analysis provides definitive evidence of the elements affecting portfolio performance, facilitating the justification of investment strategies to clients and other stakeholders. This transparency enhances the credibility of the bank's portfolio management team and cultivates trust among investors.

The use of machine learning models in portfolio management processes has significant implications for operational efficiency. Machine learning-driven automated solutions can diminish the time and resources needed for portfolio analysis, enabling managers to focus on strategic decision-making.

Moreover, the scalability of models such as Gradient Boosting and XGBoost enables them to manage extensive and intricate datasets, rendering them appropriate for various financial contexts.

Ultimately, the results underscore the need to tailor machine learning algorithms to the unique characteristics of the bank's portfolio. The dependence on portfolio-specific indicators instead of general macroeconomic variables highlights the necessity for customised models that correspond to the Bank's particular asset composition and market exposure. This emphasis on customisation guarantees that the models provide the most pertinent and actionable insights, hence enhancing their value for the Bank's portfolio management team.

This study offers a pragmatic and implementable framework for incorporating machine learning into portfolio management. By emphasising high-impact measures, implementing real-time data monitoring, and utilising the capabilities of adaptable models such as Gradient Boosting and LSTM, the Bank of Namibia may augment its forecasting precision, boost its response to market fluctuations, and refine its investment strategies. These innovations enable the bank to navigate intricate and volatile financial markets with enhanced confidence and precision, establishing a robust foundation for enduring success in portfolio management.

5.3 Contribution

This dissertation offers substantial advances in financial modelling and machine learning, especially with portfolio performance assessment. Initially, it illustrates the efficacy of machine learning models, including Gradient Boosting, Random Forest, XGBoost, and Long Short-Term Memory (LSTM), in overcoming the shortcomings of conventional financial models. This study uses these models to establish a robust framework for forecasting Month-To-Date (MTD) returns, emphasising their capacity to manage intricate, non-linear interactions and fluctuating market conditions. This expands the theoretical understanding of machine learning applications in time-sensitive financial prediction, particularly in data-constrained environments.

The dissertation highlights the essential importance of portfolio-specific indicators, like market value dirty, YTD return, and BM size, in influencing short-term portfolio performance, hence questioning the traditional dependence on macroeconomic variables, such as interest rates and inflation. This insight enhances our comprehension of the primary determinants of portfolio returns, particularly in short-term scenarios, where internal portfolio dynamics prevail over broader economic influences. It contributes to theory by proposing a refined hierarchy of feature importance tailored to emerging markets, where volatility, limited data granularity, and portfolio composition may differ from those in developed markets (Roeder et al., 2024).

The work presents a novel methodology for feature engineering in financial modelling, integrating rolling computations, cumulative returns, and volatility indicators to encapsulate intra-month trends and risk dynamics. These characteristics enhance predictive accuracy and offer actionable information for portfolio managers, facilitating more informed and quicker decision-making. From a theoretical standpoint, this structured feature engineering approach contributes to the literature by introducing hybrid temporal-financial features for MTD forecasting, which remains underexplored in the context of sovereign and institutional portfolios.

The research underscores the potential of incorporating real-time data and adjusting machine learning models to swiftly evolving market conditions. This contribution highlights the significance of dynamic and responsive models in contemporary financial management, paving the way for future research on real-time forecasting and adaptive learning methodologies. The study thus promotes the use of adaptive ML forecasting tailored to rapidly shifting market environments, reinforcing current theories on responsiveness in algorithmic finance (Johnson & Jones, 2021).

Ultimately, the dissertation enhances the interpretability of machine learning models in financial situations. Utilising feature importance analysis enhances transparency regarding the determinants of model predictions, hence cultivating trust and promoting the actual implementation of these models in institutional decision-making. This also advances theoretical discussions around explainability in ML applications, an emerging area in financial technology research. The incorporation of model transparency in this work bridges the gap between black-box algorithmic modelling and human-in-the-loop decision frameworks (Bollen et al., 2011).

This study connects theoretical progress in machine learning with its practical use in the financial sector, specifically for organisations such as the Bank of Namibia. Furthermore, it contributes to the body of knowledge on data-driven financial evaluation in emerging economies, where institutional practices, policy structures, and data infrastructures vary significantly from those in mature financial systems. By contextualising ML-based portfolio evaluation within such settings, this research lays foundational work for future regional adaptation of global predictive modelling frameworks.

5.4 Assumptions

This research was founded on numerous essential assumptions, each of which influenced the design, execution, and assessment of the machine learning models. The assumptions, although essential for practical model building, also impose certain limitations that must be recognised to ensure a thorough comprehension of the study's results and their relevance.

A core assumption was that past correlations among variables would remain consistent over time, allowing the models to extrapolate their insights to future datasets. This assumption illustrates the extensive dependence of financial forecasting models on historical data, which constitutes the fundamental foundation for training machine learning algorithms. The validity of this premise hinges on market conditions maintaining consistency with historical trends. This is applicable under typical conditions but may not be relevant during extraordinary market events or structural shifts, such as global financial crises or substantial regulatory alterations. These events may cause connections between variables to diverge from past patterns, thus restricting the models' forecast accuracy.

A crucial assumption was that substituting missing values in rolling calculations with zeros would not significantly skew the dataset or undermine model performance. Missing values emerged organically throughout the feature engineering phase, especially for rolling metrics, like moving averages (MA_5, MA_10) and volatility, which requires a specific quantity of preceding data. The choice to substitute these missing values with zeros was predicated on the reasoning that the lack of a rolling statistic at the commencement of a time window does not fundamentally modify the patterns identified by the models. This approach guarantees data completeness and computational efficiency, but it assumes that the imputed values accurately reflect underlying patterns. Alternative imputation methodologies, such as forward filling or using mean values, may yield divergent outcomes and warrant investigation in subsequent studies.

The study presumed that the chosen features sufficiently represented the principal determinants of portfolio performance, as determined by feature significance analysis. The premise was corroborated by the robust performance of the models, especially Gradient Boosting, which used parameters, like Market Value Dirty, YTD return and BM size, to provide precise predictions. This presupposition also suggests that unexamined or omitted variables, including market sentiment, geopolitical threats, or sector-specific trends, do not substantially affect the findings. The study emphasised high-impact factors relevant to short-term portfolio performance; however, incorporating more variables may improve the models' comprehensiveness and accuracy.

A fundamental assumption was that feature importance rankings accurately represented the real-world influence of each variable on portfolio performance. This assumption corresponds with the interpretability of models like Gradient Boosting, which offers clear insights into feature contributions. Nonetheless, it assumes that the correlations established in the training data remain consistent in unobserved datasets, which may not invariably be the case, especially in dynamic or developing markets. The precision of these rankings may potentially be affected by factors such as multicollinearity or the interaction of characteristics, which were not thoroughly examined in this study.

The study also presumed that the machine learning models would generalise proficiently across various time intervals, facilitating dependable estimations of MTD returns. Cross-validation procedures were used to evaluate the assumption to assess model performance on unknown data, yielding robust findings. The efficacy of this assumption is contingent upon the stability of the underlying data distribution. Substantial changes in market conditions or external disruptions may diminish the models' generalisability, highlighting the necessity for regular retraining to preserve their applicability.

Ultimately, it was presumed that the dataset's architecture and scope, which depended solely on data from the Bank of Namibia, accurately reflected the determinants of portfolio performance. This assumption is valid within the unique setting of this study; nonetheless, it constrains the applicability of the findings to other financial institutions with varying asset compositions, market exposures, or operational strategies. Augmenting the dataset to encompass a wider array of financial contexts could improve the models' resilience and usefulness.

In conclusion, these assumptions established a basis for the formulation and assessment of the machine learning models, facilitating the study's attainment of its goals within a specified framework. Nonetheless, they emphasise aspects requiring additional investigation and improvement to augment the models' applicability, robustness, and precision in practical contexts. Identifying these assumptions is essential for comprehending the study's merits and limits, as well as for directing future research in financial modelling.

5.5 Limitations of the Study

This study exhibited promising results and underscored the revolutionary potential of machine learning in portfolio performance assessment. Nonetheless, it is crucial to recognise its limitations, which delineate the scope and applicability of the findings. These constraints offer essential context for analysing the findings and pinpointing opportunities for more investigation and enhancement.

A notable constraint of this study is its dependence on historical data for the training and assessment of the machine learning models. This approach essentially presupposes that historical correlations between characteristics and portfolio performance will persist over time. This prevalent method in financial modelling limits the models' capacity to adjust to unforeseen market occurrences or structural transformations, such as global financial crises, pandemics, or significant regulatory alterations. Such events may undermine the historical links that the models depend on, thereby diminishing their forecast accuracy and robustness in practical situations. Further research may mitigate this limitation by integrating real-time data streams and adaptive learning approaches that enable models to dynamically adjust to fluctuating market conditions.

The research concentrated on a limited set of features, especially highlighting portfolio-specific measures, such as Market Value Dirty, YTD return, and BM size, in conjunction with some macroeconomic variables, like interest rate and inflation. Although these qualities were recognised as the primary predictors of MTD returns, the omission of other potentially significant elements, such as market mood, geopolitical concerns, or sector-specific trends, constrains the models' comprehensiveness. These supplementary factors may offer a more comprehensive framework for analysing portfolio performance, particularly during periods of market volatility or sector-specific disturbances. Incorporating data on investor mood or geopolitical events may capture exogenous influences affecting asset performance, improving the models' forecast accuracy and breadth.

A significant disadvantage is the sole dependence on data from the Bank of Namibia, which may inadequately reflect the variety of factors affecting portfolio performance in different financial contexts. The distinctive attributes of the bank's portfolio, including its asset mix, market exposure, and operational initiatives, indicate that the findings are primarily relevant to institutions with analogous profiles. This limits the applicability of the findings to other financial institutions, especially those functioning in diverse geographic areas or industries. Portfolios with increased exposure to emerging markets or specialised asset classes may demonstrate characteristics not included in this study. Augmenting the dataset with information from various institutions or markets may enhance the models' robustness and applicability.

The research did not thoroughly investigate the influence of hyperparameter adjustment on model efficacy. Hyperparameter tuning is an essential procedure in machine learning that entails optimising a model's parameters to enhance its accuracy and efficiency. Although the models assessed in this study exhibited satisfactory performance with default or slightly modified hyperparameters, additional optimisation may produce superior outcomes. Adjusting parameters such as the learning rate, tree depth, or regularisation terms in gradient boosting and XGBoost may improve their capacity to identify intricate patterns and interactions. Modifying the architecture of the LSTM model, including the number of layers or units, may enhance its ability to identify temporal dependencies in the data. This constraint signifies a pathway for future research to enhance the models and optimise their predictive efficacy.

The study also failed to thoroughly examine model interpretability and explainability, especially with intricate algorithms such as LSTM and XGBoost. Although gradient boosting offers explicit feature importance rankings, the opaque nature of other models may hinder their acceptance among stakeholders that necessitate clear justifications for investment decisions. Improving the interpretability of these models via approaches, like SHapley Additive exPlanations (SHAP) values or

Local Interpretable Model-agnostic Explanations (LIME**)**, could render them more applicable in real-world scenarios, where transparency is essential for fostering confidence and guaranteeing regulatory adherence.

Ultimately, the study concentrated on short-term forecasting, assessing models according to their capacity to anticipate MTD returns. This approach meets the Bank of Namibia's immediate requirements but restricts the relevance of the findings for long-term portfolio management and strategic planning. Long-term projections frequently necessitate the examination of extensive macroeconomic trends and structural elements that were not thoroughly investigated in this study. Subsequent study may broaden the approach to encompass multi-horizon forecasting, thereby offering a more thorough framework for portfolio assessment.

This study made notable progress in portfolio performance evaluation through machine learning; however, its dependence on historical data, narrow attribute range, exclusive dataset, constrained hyperparameter tuning, and emphasis on short-term forecasting constitute significant limitations. Identifying these limitations establishes a basis for subsequent research to fill these voids, improve the models, and augment their relevance across various financial scenarios. These initiatives will be crucial for promoting the integration of machine learning in portfolio management and guaranteeing its efficacy under intricate and fluctuating market conditions.

5.6 Recommendations for Future Research

This study has highlighted the revolutionary capacity of machine learning in evaluating portfolio performance and pinpointed critical areas for enhancement and growth. Subsequent study should seek to expand upon these findings to enhance the use of data-driven approaches in financial management, tackling existing limits and investigating new avenues for innovation.

A promising avenue for future research is the development of hybrid models that combine Gradient Boosting and LSTM networks. In this study, Gradient Boosting achieved high precision in short-term forecasting by effectively capturing complex, nonlinear relationships and feature interactions, while LSTM demonstrated strength in identifying temporal dependencies and cyclical trends, which are valuable for long-term forecasting. A combined approach could apply Gradient Boosting to model static and cross-sectional features, and LSTM to capture sequential and time-dependent patterns. Such integration would create a unified framework capable of delivering accurate predictions across both short-term and long-term horizons. This would help address the limitations inherent in using individual models in isolation, offering a more comprehensive understanding of portfolio performance dynamics.

A key area for further research is the incorporation of real-time data into machine learning algorithms. Financial markets are fundamentally dynamic, with conditions rapidly changing in response to economic events, geopolitical developments, and investor sentiment. Contemporary models, predominantly dependent on past data, may find it challenging to adjust to such swift transformations. Integrating real-time data streams would allow models to alter dynamically, improving their adaptability and responsiveness. Integrating live market feeds, high-frequency trading data, or news sentiment analysis could yield a more precise and current depiction of portfolio performance determinants, facilitating more prompt and informed decision-making.

Broadening the spectrum of features incorporated in the analysis is an additional important recommendation. This study concentrated on high-impact, portfolio-specific indicators such as market value dirty and year-to-date return, together with certain macroeconomic variables like interest rate and inflation; however, incorporating more elements could enhance the algorithms' comprehension of portfolio performance. Factors, such as market sentiment, geopolitical risks and sector-specific trends, may encompass external influences on asset performance that are presently disregarded. Market sentiment obtained from social media or news analytics may yield insights into investor behaviour, whereas geopolitical indexes could reflect the influence of international events on financial markets. These supplementary elements may enhance models' predictive accuracy during times of market volatility or structural transformation.

Future research should prioritise hyperparameter optimisation to enhance the performance of machine learning models. This work attained strong results using default or barely adjusted hyperparameters; nevertheless, investigating sophisticated optimisation methods could markedly improve model accuracy and efficiency. Grid search, random search, or more advanced techniques such as Bayesian optimisation and genetic algorithms may be used to determine the ideal parameters for each model. For instance, optimising parameters like the learning rate, maximum depth, and number of estimators in gradient boosting, or modifying LSTM architectural elements like the number of layers and units, may result in significant enhancements in predicting accuracy.

Subsequent studies should investigate the applicability of machine learning models across various financial institutions and asset classes. This study relied solely on data from the Bank of Namibia, hence constraining the generalisability of its results. Incorporating data from diverse institutions with differing portfolio compositions, market exposures, and investment strategies would enhance the comprehension of the models' efficacy. Furthermore, implementing the models across several asset classes, such as commodities, derivatives, or developing market equities, could evaluate their adaptability and pertinence in distinct financial contexts.

An additional significant domain for investigation is the interpretability and explainability of machine learning models. This study emphasised the transparency of gradient boosting via feature importance analysis, whereas models, such as LSTM and XGBoost, are often regarded as 'black box' algorithms. Subsequent research may include sophisticated interpretability methodologies like SHAP or LIME to furnish stakeholders with an enhanced understanding of the determinants influencing model predictions. Improving interpretability would foster trust in machine learning models and guarantee adherence to regulatory mandates for transparency in financial decision-making.

Finally, subsequent research should explore the capabilities of multi-horizon forecasting, which broadens predictions from short-term MTD returns to include medium-term and long-term performance. This method may offer portfolio managers an extensive perspective on portfolio dynamics, connecting short-term outcomes to overarching strategic goals. Multi-horizon forecasting may investigate the impact of short-term trends on long-term results, providing an enhanced understanding of the interaction between several time horizons in portfolio management.

This study demonstrates the revolutionary capacity of machine learning in assessing portfolio performance, providing essential insights that can guide financial decision-making and investment strategies. By resolving the observed constraints and utilising the suggested innovations, future studies can enhance the efficacy of data-driven methodologies in financial management. These initiatives will improve the precision and flexibility of predictive models, creating a more resilient framework for managing the intricacies of contemporary financial markets, hence yielding superior results for institutions and their stakeholders.

5.7 Concluding Remarks

This study has illustrated the transformative capacity of machine learning in improving portfolio performance assessment, providing a solid and data-driven framework for forecasting short-term returns and guiding investment strategies. The study prioritised portfolio-specific metrics and used advanced algorithms to rectify significant deficiencies in conventional financial modelling, facilitating more precise and flexible forecasting methods.

The results of this research have significant implications for financial institutions, such as the Bank of Namibia, and the wider domain of financial management. They emphasise the necessity of incorporating high-impact features, real-time data, and sophisticated machine learning methodologies to adeptly manoeuvre through the intricacies of contemporary financial markets.

This study has made notable contributions but also identifies areas for future investigation, including hybrid modelling techniques, enhanced feature sets, and real-time adaptation. These developments possess the capacity to enhance the precision, resilience, and application of machine learning models across many financial situations.

This research advances the application of machine learning in portfolio management, connecting theoretical innovation with practical execution. By illustrating the significance of data-driven methodologies, it establishes a basis for ongoing progress in financial modelling, allowing institutions to make more informed, efficient, and consequential investment decisions in a progressively intricate and unstable financial environment.

Reference

Alzubi, J., Nayyar, A., & Kumar, A. (2018). Machine Learning from Theory to Algorithms: An Overview. *Journal of Physics: Conference Series*, *1142*(1). https://doi.org/10.1088/1742-6596/1142/1/012012

Ateino, E. (2022). Financial Performance and Investment Decision Making in Kenya. *Article in African Journal of Drug and Alcohol Studies*, *1*(2), 2022–2023. https://doi.org/10.5281/zenodo.7306100

Bollen, J., Mao, H., & Zeng, X. (2011). Twitter mood predicts the stock market. *Journal of Computational Science*, *2*(1), 1–8. https://doi.org/10.1016/j.jocs.2010.12.007

Brown, I. (2019). Data-driven decision making in the financial sector: Challenges and opportunities. *Journal of Finance and Economics*, *7*(2), 59–67.

Brown, L. (2021). Experimental Strategies in Portfolio Management. *Journal of Investment Strategies*, *19*(2), 200–215.

Clark, R. L., & Thompson, J. A. (2022). Navigating Modern Financial Markets: The Need for Updated Analytical Tools. *Finance and Investment Journal*, *47*(3), 112–129.

Creswell, J. W. (2014). *Research design: Qualitative, quantitative, and mixed methods approaches* (4th ed.). Thousand Oaks, CA: SAGE Publications

Green, L., & Kaplan, D. (2022). Leveraging Data-Driven Insights for Investment Decision-Making. *Journal of Portfolio Management*, *48*(2), 112–129.

Hodge, V. J., & Austin, J. (2020). A Survey of Outlier Detection Methodologies. *Artificial Intelligence Review*, *53*(8), 6067–6109.

Huang, W., Nakamori, Y., & Wang, S.-Y. (2020). Forecasting stock market movement direction with support vector machine. *Computers & Operations Research*, *32*(10), 2513–2522.

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

James G, Witten, D., Hastie, T., & Tibshirani, R. (2024). *An Introduction to Statistical Learning: with Applications in R*.

Johnson, L., & Jones, C. (2021). The Evolution of Financial Markets and the Implications for Portfolio Management. *International Journal of Finance*, *47*(3), 245–263.

Johnson, M., & Lee, A. (2021). Risk Management and Decision-Making in Financial Institutions: A Data-Driven Approach. *Journal of Financial Analytics*, *5*(2), 45–60.

Khakhar, D., & Godhani, S. (2023). Predictive Quantitative Financial Analysis for Technology-Driven Investment Decision. *International Journal for Research in Applied Science and Engineering Technology*, *11*(10), 11–15. https://doi.org/10.22214/ijraset.2023.56419

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, *521*(7553), 436–444.

Lee, J., & Kim, T. (2021). Machine Learning in Finance: Opportunities and Challenges. *International Journal of Financial Studies*, *29*(4), 1002–1020.

Liu, D., & Yoon, V. Y. (2024). Developing a goal-driven data integration framework for effective data analytics. *Decision Support Systems*, *180*. https://doi.org/10.1016/j.dss.2024.114197

Lumumba, V. W., Kiprotich, D., Mpaine, M. L., Makena, N. G., & Kavita, M. D. (2024). Comparative analysis of cross-validation techniques: LOOCV, k-folds cross-validation, and repeated k-folds cross-validation in machine learning models. *American Journal of Theoretical and Applied Statistics, 13*(5), 127–137.

Moyo, N. (2023). Data-Driven Portfolio Management in African Banks: A Case Study Approach. *African Journal of Finance and Management*, *11*(1), 56–72.

Moyo, N. (2024). Data-Driven Portfolio Management in African Banks: A Case Study Approach. *African Journal of Finance and Management*, *11*(1), 56–72.

Roeder, J., Palmer, M., & Muntermann, J. (2022). Data-driven decision-making in credit risk management: The information value of analyst reports. *Decision Support Systems*, *158*. https://doi.org/10.1016/j.dss.2022.113770

Saunders, M., Lewis, P., & Thornhill, A. (2019). *Research Methods for Business Students*.

Taylor, J., & Brown, A. (2023). Adaptive Portfolio Management in Volatile Markets: Lessons from Namibia. *Global Finance Review*, *54*(2), 123–142.

Vassallo, D. (2023). Asset allocation and risk taking under different interest rate regimes. *European Central Bank Working Paper Series*, 2803.

Wang, F. (2023). Machine Learning in Modern Portfolio Theory. *2023*, *5*(2), 45–58.

Wang, J., Jiang, C., Zhou, L., & Wang, Z. (2024). Assessing financial distress of SMEs through event propagation: An adaptive interpretable graph contrastive learning model. *Decision Support Systems*, *180*.

Wang, J., & Liu, J. (2020). Machine Learning in Finance: The Case of Deep Learning for Option Pricing. *Journal of Financial Data Science*, *2*(1), 20–33.

Wang, Q., & Liu, Y. (2020). Portfolio Performance Prediction using Machine Learning: A Quantitative Approach. *Quantitative Finance*, *21*(4), 643–659.

Williams, H., & Smith, M. (2019). The Role of Environmental and Economic Factors in Financial Decision-Making: A Focus on Namibia. *Banking and Finance Journal*, *31*(4), 577–596.

Zhu, X. (2021). Combining LightGBM and XGBoost for improved stock selection. *Journal of Financial Data Science*, *7*(1), 55–72.

APPENDIX A: Ethical Clearance Letter

NAMIBIA UNIVERSITY OF SCIENCE AND TECHNOLOGY

FACULTY RESEARCH ETHICS COMMITTEE (F-REC)

DECISION/FEEDBACK ON THE RESEARCH PROPOSAL

Dear Mr Pendapala Nghilundwa (217113389)

RESEARCH TOPIC: DEVELOPING A DATA-DRIVEN FINANCIAL MODEL FOR DECISION SUPPORT IN EVALUATING INVESTMENT PORTFOLIO PERFORMANCE AT BANK OF NAMIBIA

Supervisor (if applicable): **Dr Jacob Ongala**

Qualification registered for (if applicable): Master of Data Science

(Reference number of applications: **FACULTY RESEARCH ETHICS COMMITTEE REGISTRATIONNUMBER: FREC - 37/24)**

Re: Ethical screening application No: **FREC - 37/24**

The Faculty of **Computing and Informatics** Ethics Screening Committee of the Namibia University of Science and Technology reviewed your application for the above-mentioned research. The research as set out in the application has been:

Approved	X

(Indicate with an X, and N/A if not applicable and proceed)

We would like to point out that you, as a researcher, are obliged to maintain the ethical integrity of your research, adhere to the ethical guidelines of NUST, and remain within the scope of your research proposal and supporting evidence as submitted to the F-REC. Should any aspect of your research change from the information as presented to the F-REC, which could affect the possibility of harm to any research subject, you are under the obligation to report it immediately to your supervisor or F-REC as applicable in writing. Should there be any uncertainty in this regard, you must consult with the F-REC.

We wish you success with your research and trust that it will make a positive contribution to the questfor knowledge at NUST.

Any ethical issues that need to be highlighted?	Why are these issues important?	What must/could be done to minimize the ethical risk?
No	N/A	N/A

Recommendation: The application is approved.

Sincerely,

Prof. Suama L Hamunyela
Chairperson: Faculty Ethics Screening
CommitteeTel: +264-61-207-2922
CC: Co-supervisor:

NAMIBIA UNIVERSITY OF SCIENCE AND TECHNOLOGY
P/Bag 13388 Windhoek NAMIBIA
2024 -09- 10
OFFICE OF THE DEAN
FACULTY OF COMPUTING & INFORMATICS

APPENDIX B: Implementation details

B.1 Data Collection and Preprocessing

The dataset was sourced from the Bank of Namibia and included variables such as Market Value Dirty, YTD Return, Daily Return, and macroeconomic indicators like Interest Rate and Inflation. Data preprocessing involved converting dates to a standard datetime format, sorting records chronologically, and handling missing values. Numeric columns were formatted correctly, and rolling operations were performed to generate features like moving averages and volatility. This includes code for loading the dataset, handling missing values, and preprocessing features such as DAILY_RETURN.

```
# Import all required libraries
import pandas as pd
import numpy as np
from sklearn.model_selection import train_test_split
from sklearn.preprocessing import StandardScaler
from sklearn.metrics import mean_squared_error, r2_score
from sklearn.ensemble import RandomForestRegressor, GradientBoostingRegressor
import xgboost as xgb
import matplotlib.pyplot as plt
import matplotlib.dates as mdates
from tensorflow.keras.models import Sequential
from tensorflow.keras.layers import LSTM, Dense, Input
from sklearn.model_selection import GridSearchCV, RandomizedSearchCV
import xgboost as xgb
from sklearn.ensemble import RandomForestRegressor, GradientBoostingRegressor
import matplotlib.pyplot as plt
from prophet import Prophet
import seaborn as sns
import statsmodels.api as sm
from sklearn.metrics import mean_squared_error, r2_score
from sklearn.preprocessing import MinMaxScaler
from sklearn.model_selection import cross_val_score, KFold
from sklearn.ensemble import RandomForestRegressor
from sklearn.model_selection import train_test_split
from sklearn.metrics import mean_squared_error
import matplotlib.pyplot as plt
```

```
# ================================
# 2. Load and Preprocess the Dataset
# ================================

# Define file path and load dataset
file_path = "C:/Users/pendapala/OneDrive - Bank of Namibia/Desktop/USDINV_POSITION.xlsx"
data = pd.read_excel(file_path)

# Display initial dataset information
data.head(5), data.info()
```

```
# Convert 'ASOF_DATE' to datetime and sort by date
data['ASOF_DATE'] = pd.to_datetime(data['ASOF_DATE'])
data = data.sort_values('ASOF_DATE')
```

```
# Convert daily returns to numeric and remove any rows with NaN
data['DAILY_RETURN'] = pd.to_numeric(data['DAILY_RETURN'], errors='coerce')
data.dropna(subset=['DAILY_RETURN'], inplace=True)
```

B.2 Feature Engineering

Features were developed to improve the predictive precision of machine learning models. Metrics such as MTD_RETURN_CUMULATIVE (cumulative Month-To-Date returns), moving averages (MA_5, MA_10), and VOLATILITY were computed through rolling operations. These features encapsulate market trends and risk dynamics, enhancing the models' capacity to identify significant patterns. This will encompass the code for generating rolling metrics, cumulative returns, and imputing missing values with zeros.

```
# ===================================
# 3. Feature Engineering
# ===================================

# Group the data by month and year
data['YearMonth'] = data['ASOF_DATE'].dt.to_period('M')

# Calculate the cumulative MTD return for each month
data['MTD_RETURN_CUMULATIVE'] = data.groupby('YearMonth')['DAILY_RETURN'].cumsum()

# Create moving averages and volatility features
data['VOLATILITY'] = data['DAILY_RETURN'].rolling(window=5).std()
data['MA_5'] = data['DAILY_RETURN'].rolling(window=5).mean()
data['MA_10'] = data['DAILY_RETURN'].rolling(window=10).mean()

# Fill NaN values resulting from rolling operations
data.fillna(0, inplace=True)
data.tail(5)
```

```
# Setting up the figure
fig, ax = plt.subplots(figsize=(8, 5))

# Plotting the distribution of MTD_RETURN
sns.histplot(data['MTD_RETURN'], kde=True, color='#296C92', line_kws={'linewidth': 1}, ax=ax)
ax.set_title('Distribution: MTD_RETURN', fontsize=15)
ax.set_xlabel('MTD Return', fontsize=12)
ax.set_ylabel('Frequency', fontsize=12)

# Enhancing grid and layout
ax.grid(True, linestyle='--', linewidth=0.5, color='grey')
fig.tight_layout(pad=2)

plt.show()
```

B.3 Normalization

Normalisation guarantees that all features are uniformly scaled, which is essential for models that are sensitive to feature magnitudes, such as Gradient Boosting and LSTM. StandardScaler and MinMaxScaler were used on features including Daily Return and Volatility.

The pertinent code for implementing normalisation on features utilising StandardScaler and MinMaxScaler will be presented.

```python
# ===================================
# 4. Normalize Features
# ===================================

# Select the columns for feature normalization
features_to_normalize = ['DAILY_RETURN', 'VOLATILITY', 'MA_5', 'MA_10']

# Initialize a scaler and fit on training data
scaler = StandardScaler()

# Normalize the features
data[features_to_normalize] = scaler.fit_transform(data[features_to_normalize])
```

```python
# ===============================
# 5. Split the Data into Training and Testing Sets
# ===============================


# Define features (X) and target variable (y)
X = data.drop(columns=['ASOF_DATE', 'SUBPORTFOLIO_CODE', 'MTD_RETURN_CUMULATIVE', 'YearMonth'])
y = data['MTD_RETURN_CUMULATIVE']  # The target variable you're predicting

# Add titles when displaying the first 10 rows
print("First 10 rows of X (Features):")
print(X.head(5))

print("\nFirst 10 rows of y (Target):")
print(y.head(5))
```

```python
train_data = data.MTD_RETURN[:-test_size]
train_data = scaler.transform(train_data.values.reshape(-1,1))
print(train_data)
```

```python
# Define the dependent and independent variables
X = sm.add_constant(data['INTEREST_RATE'])  # Adds a constant term to the predictor
Y = data['MTD_RETURN']

# Fit the regression model
model = sm.OLS(Y, X).fit()

# Get the summary of the regression results
print(model.summary())
```

B.4 Model Development and Evaluation Metrics

Four machine learning models—Random Forest, Gradient Boosting, XGBoost, and LSTM—were created to forecast MTD returns. Each model underwent training and evaluation utilising training and testing divides. Cross-validation was used to guarantee robustness. The Gradient Boosting model had superior performance, with the lowest MSE and the highest R^2. The evaluation of the model was conducted using Mean Squared Error (MSE) and R-Squared (R^2). An analysis of feature importance was performed to determine the most significant predictors for each model, focussing on portfolio-specific measures like Market Value Dirty and Year-to-Date Return.

Random Forest Model

```
# ========================================
# 6. Train Models and Evaluate Performance
# ========================================

# -----------------------------------
# 6.1 Train and Evaluate Random Forest
# -----------------------------------

rf_model = RandomForestRegressor(n_estimators=300, random_state=42)
rf_model.fit(X_train, y_train)
y_pred_rf = rf_model.predict(X_test)

# Evaluate Random Forest model
mse_rf = mean_squared_error(y_test, y_pred_rf)
r2_rf = r2_score(y_test, y_pred_rf)

print(f"Random Forest Model - Test MSE: {mse_rf}")
print(f"Random Forest Model - Test R²: {r2_rf}")
```

```
# Setup cross-validation for Random Forest
kf = KFold(n_splits=5, shuffle=True, random_state=42)  # 5-fold cross-validation

# Perform cross-validation
cv_scores_rf = cross_val_score(rf_model, X, y, cv=kf, scoring='neg_mean_squared_error')

# Calculate mean and standard deviation of the score
mse_scores = -cv_scores_rf  # Convert to positive MSE scores
print(f"MSE per fold: {mse_scores}")
print(f"Average MSE: {np.mean(mse_scores)}")
print(f"Standard Deviation of MSE: {np.std(mse_scores)}")
```

```
# Calculate the line of best fit for Random Forest
slope, intercept = np.polyfit(y_test, y_pred_rf, 1)
line_of_best_fit = slope * y_test + intercept

# Plot the actual vs predicted values
plt.figure(figsize=(10, 6))
plt.scatter(y_test, y_pred_rf, color='blue', label='Predicted vs Actual')
plt.plot(y_test, line_of_best_fit, color='red', linewidth=2, label='Line of Best Fit')

plt.xlabel('Actual Values')
plt.ylabel('Predicted Values')
plt.title('Random Forest - Actual vs Predicted Values with Residuals')
plt.legend()
plt.grid(True)
plt.show()
```

Gradient Boosting Model

```
# ---------------------------
# 6.2 Train and Evaluate Gradient Boosting
# ---------------------------
gb_model = GradientBoostingRegressor(n_estimators=300, learning_rate=0.1, max_depth=5, random_state=42)
gb_model.fit(X_train, y_train)
y_pred_gb = gb_model.predict(X_test)

# Evaluate Gradient Boosting model
mse_gb = mean_squared_error(y_test, y_pred_gb)
r2_gb = r2_score(y_test, y_pred_gb)

print(f"Gradient Boosting Model - Test MSE: {mse_gb}")
print(f"Gradient Boosting Model - Test R²: {r2_gb}")
```

```
# Perform cross-validation for Gradient Boosting
cv_scores_gb = cross_val_score(gb_model, X, y, cv=kf, scoring='neg_mean_squared_error')

# Calculate mean and standard deviation of the score
mse_scores_gb = -cv_scores_gb
print("Gradient Boosting - MSE per fold:", mse_scores_gb)
print("Gradient Boosting - Average MSE:", mse_scores_gb.mean())
print("Gradient Boosting - Standard Deviation of MSE:", mse_scores_gb.std())
```

```
# Calculate the line of best fit for Gradient Boosting
slope, intercept = np.polyfit(y_test, y_pred_gb, 1)
line_of_best_fit = slope * y_test + intercept

# Plot the actual vs predicted values
plt.figure(figsize=(10, 6))
plt.scatter(y_test, y_pred_gb, color='blue', label='Predicted vs Actual')
plt.plot(y_test, line_of_best_fit, color='red', linewidth=2, label='Line of Best Fit')

plt.xlabel('Actual Values')
plt.ylabel('Predicted Values')
plt.title('Gradient Boosting - Actual vs Predicted Values with Residuals')
plt.legend()
plt.grid(True)
plt.show()
```

XGBoost Model

```
# ---------------------------
# 6.3 Train and Evaluate XGBoost
# ---------------------------

xgb_model = xgb.XGBRegressor(n_estimators=300, learning_rate=0.1, max_depth=5, random_state=42)
xgb_model.fit(X_train, y_train)
y_pred_xgb = xgb_model.predict(X_test)

# Evaluate XGBoost model
mse_xgb = mean_squared_error(y_test, y_pred_xgb)
r2_xgb = r2_score(y_test, y_pred_xgb)

print(f"XGBoost Model - Test MSE: {mse_xgb}")
print(f"XGBoost Model - Test R²: {r2_xgb}")
```

```python
# Perform cross-validation for XGBoost
cv_scores_xgb = cross_val_score(xgb_model, X, y, cv=kf, scoring='neg_mean_squared_error')

# Calculate mean and standard deviation of the score
mse_scores_xgb = -cv_scores_xgb
print("Gradient Boosting - MSE per fold:", mse_scores_xgb)
print("Gradient Boosting - Average MSE:", mse_scores_xgb.mean())
print("Gradient Boosting - Standard Deviation of MSE:", mse_scores_xgb.std())
```

```python
# Calculate the line of best fit for XGBoost
slope, intercept = np.polyfit(y_test, y_pred_xgb, 1)
line_of_best_fit = slope * y_test + intercept

# Plot the actual vs predicted values
plt.figure(figsize=(10, 6))
plt.scatter(y_test, y_pred_xgb, color='blue', label='Predicted vs Actual')
plt.plot(y_test, line_of_best_fit, color='red', linewidth=2, label='Line of Best Fit')

plt.xlabel('Actual Values')
plt.ylabel('Predicted Values')
plt.title('XGBoost - Actual vs Predicted Values with Residuals')
plt.legend()
plt.grid(True)
plt.show()
```

LSTM Model

```python
# ---------------------------
# 6.4 Train and Evaluate LSTM
# ---------------------------

X_train_reshaped = X_train_normalized.reshape(X_train_normalized.shape[0], 1, X_train_normalized.shape[1])
X_test_reshaped = X_test_normalized.reshape(X_test_normalized.shape[0], 1, X_test_normalized.shape[1])

lstm_model = Sequential()
lstm_model.add(Input(shape=(X_train_reshaped.shape[1], X_train_reshaped.shape[2])))
lstm_model.add(LSTM(300, return_sequences=True))
lstm_model.add(LSTM(300))
lstm_model.add(Dense(1))
lstm_model.compile(optimizer='adam', loss='mean_squared_error')

lstm_model.fit(X_train_reshaped, y_train, epochs=50, batch_size=16)
y_pred_lstm = lstm_model.predict(X_test_reshaped).flatten()

# Evaluate LSTM model
mse_lstm = mean_squared_error(y_test, y_pred_lstm)
r2_lstm = r2_score(y_test, y_pred_lstm)
print(f"LSTM Model - Test MSE: {mse_lstm}")
print(f"LSTM Model - Test R²: {r2_lstm}")
```

```python
lstm_model.summary()
```

```python
# Calculate the line of best fit for LSTM
slope, intercept = np.polyfit(y_test, y_pred_lstm, 1)
line_of_best_fit = slope * y_test + intercept

# Plot the actual vs predicted values
plt.figure(figsize=(10, 6))
plt.scatter(y_test, y_pred_lstm, color='blue', label='Predicted vs Actual')
plt.plot(y_test, line_of_best_fit, color='red', linewidth=2, label='Line of Best Fit')

plt.xlabel('Actual Values')
plt.ylabel('Predicted Values')
plt.title('LSTM - Actual vs Predicted Values with Residuals')
plt.legend()
plt.grid(True)
plt.show()
```

Features importance

```python
# Update the features list to include the additional ones
features = [
    "MARKET_VALUE_DIRTY", "UNINVESTED_CASH", "SUBPORTFOLIO_DURATION",
    "BENCHMARK_DURATION", "SUBPORTFOLIO_SPREAD_DURATION", "BENCHMARK_SPREAD_DURATION",
    "BM_SIZE", "COMPOUNDED_DAILY_INTEREST", "SIMPLE_DAILY_INTEREST",
    "DAILY_RETURN", "MTD_RETURN", "YTD_RETURN", "INTEREST_RATE",
    "GDP_GROWTH", "INFLATION","MA_5", "MA_10", "VOLATILITY"
]

df = pd.DataFrame(data)

# Define features and target variable
target = "MTD_RETURN"
features = [col for col in df.columns if col != target]

X = df[features]
y = df[target]

# Standardize Features
scaler = StandardScaler()
X_scaled = scaler.fit_transform(X)

# Linear Model
linear_model = LinearRegression()
linear_model.fit(X_scaled, y)
coefficients = pd.DataFrame({
    "Feature": features,
    "Coefficient": linear_model.coef_
})

# Correlation Analysis
correlation = df.corr()[target].loc[features]
coefficients["Correlation"] = correlation.values

# Print coefficients and correlation to console
print(coefficients)

# Optionally, save the result to a CSV file
coefficients.to_csv("feature_importance_analysis.csv", index=False)
```

```python
# Feature Importance from Tree-Based Model
tree_importance = gb_model.feature_importances_
tree_based_importance = pd.DataFrame({
    "Feature": features,
    "Tree-Based Importance": tree_importance
})

# Permutation Importance
perm_importance = permutation_importance(gb_model, X_scaled, y, n_repeats=10, random_state=42)
perm_based_importance = pd.DataFrame({
    "Feature": features,
    "Permutation Importance": perm_importance.importances_mean
})

# Combine Results
combined_importance = coefficients.merge(tree_based_importance, on="Feature").merge(perm_based_importance, on="Feature")

# Feature Selection using Tree-Based Model
selector = SelectFromModel(gb_model, prefit=True)
selected_features = np.array(features)[selector.get_support()]

# Display combined results and selected features
print("Combined Feature Importance:")
print(combined_importance.sort_values(by="Tree-Based Importance", ascending=False))

print("\nSelected Features:")
print(selected_features)
```

```python
# Combine all feature importance results into a single DataFrame
combined_importance_sorted = combined_importance.sort_values(by="Tree-Based Importance", ascending=False)

# Display the results in a table format
print(tabulate(combined_importance_sorted, headers="keys", tablefmt="grid"))

# Visualize the feature importance from Tree-Based Model
plt.figure(figsize=(10, 6))
sns.barplot(
    data=combined_importance_sorted,
    x="Tree-Based Importance",
    y="Feature",
    palette="viridis"
)
plt.title("Feature Importance from Gradient Boosting Model", fontsize=14)
plt.xlabel("Tree-Based Importance", fontsize=12)
plt.ylabel("Features", fontsize=12)
plt.grid(axis='x', linestyle='--', alpha=0.6)
plt.tight_layout()
plt.show()

# Visualize the Permutation Importance
plt.figure(figsize=(10, 6))
sns.barplot(
    data=combined_importance_sorted,
    x="Permutation Importance",
    y="Feature",
    palette="coolwarm"
)
plt.title("Permutation Importance from Gradient Boosting Model", fontsize=14)
plt.xlabel("Permutation Importance", fontsize=12)
plt.ylabel("Features", fontsize=12)
plt.grid(axis='x', linestyle='--', alpha=0.6)
plt.tight_layout()
plt.show()
```

B.6 Future Predictions

Projected future MTD returns for the period of August 2024 to August 2025 were derived from the trained models. Rolling forecasts dynamically updated the feature matrix, emulating real-time forecasting.

The code for producing and visualising future projections, including anticipated MTD returns, is presented below.

```
# ===================================
# 9. Align Test and Prediction Data
# ===================================

# 9.1 Aligned Test and Prediction Data
# Converted dates_test to a DatetimeIndex

dates_test = pd.to_datetime(data['ASOF_DATE'])

# Checked and truncated y_pred_rf and dates_test to ensure they have the same length
if len(y_pred_rf) != len(dates_test):
    # If there's a mismatch, truncate the larger set to match the smaller one
    min_length = min(len(y_pred_rf), len(dates_test))
    y_pred_rf = y_pred_rf[:min_length]
    dates_test = dates_test[:min_length]

# Aligned y_test with dates_test and truncated if necessary
y_test_aligned = pd.Series(y_test.values[:min_length], index=dates_test)
```

```
# Converted 'ASOF_DATE' column in the original dataset to a DatetimeIndex
data['ASOF_DATE'] = pd.to_datetime(data['ASOF_DATE'])

# Split the data into training and testing sets, keeping the original 'ASOF_DATE'
X_train_df, X_test_df, y_train, y_test, dates_train, dates_test = train_test_split(
    X, y, data['ASOF_DATE'], test_size=0.5, random_state=42
)

# Converted dates_test to a DatetimeIndex
dates_test = pd.to_datetime(dates_test)  # Convert dates_test to DatetimeIndex if necessary
```

```
# Align y_test with dates_test and truncate y_test if necessary
y_test_aligned = pd.Series(y_test.values[:min_length], index=dates_test)
```

```
# 9.2 Resampled Test and Prediction Data

# Resampled y_test and predictions to get only the month-end data (using 'ME' for month-end)
y_test_monthly = y_test_aligned.resample('ME').last()  # Resample to month-end
y_pred_rf_monthly = pd.Series(y_pred_rf, index=dates_test).resample('ME').last() # Random Forest resample using 'ME'
y_pred_gb_monthly = pd.Series(y_pred_gb, index=dates_test).resample('ME').last()  # Gradient Boosting resample using 'ME'
y_pred_xgb_monthly = pd.Series(y_pred_xgb, index=dates_test).resample('ME').last()  # XGBoost resample using 'ME'
y_pred_lstm_monthly = pd.Series(y_pred_lstm, index=dates_test).resample('ME').last()  # LSTM resample using 'ME'
```

```python
# 9.3 Creating a DataFrame from the resampled data
results = pd.DataFrame({
    'Y Test Monthly': y_test_monthly,
    'Y Pred RF Monthly': y_pred_rf_monthly,
    'Y Pred GB Monthly': y_pred_gb_monthly,
    'Y Pred XGB Monthly': y_pred_xgb_monthly,
    'Y Pred LSTM Monthly': y_pred_lstm_monthly
})

print(results)
```

Actual vs. Predicted for all the models

```python
plt.figure(figsize=(13, 7))

# Plot each series with distinct colors and styles
plt.plot(y_test_monthly.index, y_test_monthly, label='Actual', color='red', linewidth=1.5)
plt.plot(y_pred_rf_monthly.index, y_pred_rf_monthly, label='Random Forest Predicted', color='green', linestyle='-.', linewidth=1)
plt.plot(y_pred_gb_monthly.index, y_pred_gb_monthly, label='Gradient Boosting Predicted', color='blue', linestyle='--', linewidth=1)
plt.plot(y_pred_xgb_monthly.index, y_pred_xgb_monthly, label='XGBoost Predicted', color='purple', linestyle=':', linewidth=1)
plt.plot(y_pred_lstm_monthly.index, y_pred_lstm_monthly, label='LSTM Predicted', color='orange', linestyle='-', linewidth=1)


# Adding legend, titles, and labels
plt.legend(loc='upper left', fontsize=10)
plt.title('Actual vs Predicted Monthly Returns(2022-2024)', fontsize=14, fontweight='bold')
plt.xlabel('Date', fontsize=12)
plt.ylabel('MTD Returns', fontsize=12)

# Adjust x-axis format and labels
plt.gca().xaxis.set_major_locator(mdates.YearLocator())
plt.gca().xaxis.set_minor_locator(mdates.MonthLocator())
plt.gca().xaxis.set_major_formatter(mdates.DateFormatter('%Y-%m'))
plt.gcf().autofmt_xdate()

# Adding grid for better visual alignment
plt.grid(visible=True, which='both', linestyle='--', linewidth=0.5, alpha=0.7)

# Set consistent y-limits based on data for a clear comparison
plt.ylim(y_test_monthly.min() - 0.01, y_test_monthly.max() + 0.01)

# Display the plot
plt.show()
```

Project Future MTD Returns

```python
# ==================================
# 10. Project Future MTD Returns
# ==================================

# Generated future dates from August 2024 to August 2025
future_dates = pd.date_range(start='2024-07-01', end='2025-08-31', freq='MS')


# Initialize lists to store future predictions
future_predictions_rf = []
future_predictions_gb = []
future_predictions_xgb = []
future_predictions_lstm = []

# Convert the last row of X_test_normalized into a DataFrame with feature names for rolling predictions
future_data_with_time = pd.DataFrame([X_test_normalized[-1]], columns=X.columns)

# Apply scaling to bring predictions within the expected range
scale_factor = 2
```

```python
# Predict future MTD returns using the trained models
for future_date in future_dates:
    # Reshape data for LSTM
    last_known_data_lstm = future_data_with_time.values.reshape(1, 1, -1)

    # Make predictions
    pred_rf = rf_model.predict(future_data_with_time)[0] / scale_factor
    pred_gb = gb_model.predict(future_data_with_time)[0] / scale_factor
    pred_xgb = xgb_model.predict(future_data_with_time)[0] / scale_factor
    pred_lstm = lstm_model.predict(last_known_data_lstm)[0][0] / scale_factor

    # Store predictions
    future_predictions_rf.append(pred_rf)
    future_predictions_gb.append(pred_gb)
    future_predictions_xgb.append(pred_xgb)
    future_predictions_lstm.append(pred_lstm)

    # Update `future_data_with_time` in a rolling manner
    # Shift all columns to the left and update the last column with the average of predictions
    future_data_with_time = pd.DataFrame(np.roll(future_data_with_time, -1, axis=1), columns=X.columns)
    future_data_with_time.iloc[0, -1] = (pred_rf + pred_gb + pred_xgb + pred_lstm) / 4  # Average prediction
```

Visualize Future Predictions

```python
# ===================================
# 11. Visualize Future Predictions
# ===================================

# Create DataFrame for future predictions
future_predictions_df = pd.DataFrame({
    'Date': future_dates,
    'Random_Forest_Predicted': future_predictions_rf,
    'Gradient_Boosting_Predicted': future_predictions_gb,
    'XGBoost_Predicted': future_predictions_xgb,
    'LSTM_Predicted': future_predictions_lstm
})

# Display future predictions DataFrame
print(future_predictions_df)
```

```python
# Calculate the sum of each model's predicted values in percentage terms
sum_values = future_predictions_df[
    ['Random_Forest_Predicted', 'Gradient_Boosting_Predicted',
     'XGBoost_Predicted', 'LSTM_Predicted']
].sum()

# Display the DataFrame with percentage values and the sum for each model
sum_values
```

```python
# Assuming you've loaded your DataFrame already as `future_predictions_df`

plt.figure(figsize=(10, 6))

plt.plot(future_predictions_df['Date'], future_predictions_df['Random_Forest_Predicted'],
         label='Random Forest', linestyle='-.', marker='s', linewidth=3)
plt.plot(future_predictions_df['Date'], future_predictions_df['Gradient_Boosting_Predicted'],
         label='Gradient Boosting', linestyle='--', marker='x', linewidth=2)
plt.plot(future_predictions_df['Date'], future_predictions_df['XGBoost_Predicted'],
         label='XGBoost', linestyle='-', marker='.', linewidth=1)
plt.plot(future_predictions_df['Date'], future_predictions_df['LSTM_Predicted'],
         label='LSTM', linestyle=':', marker='^', linewidth=2)

plt.title('Future Predicted MTD Returns (Aug 2024 - Aug 2025)')
plt.xlabel('Date')
plt.ylabel('Predicted MTD Return')
plt.legend()
plt.grid(True)
plt.show()
```

Comparing Actual and Predicted returns for Gradient Boosting Model

```python
# Create a new DataFrame by combining the actual and predicted returns
combined_returns = pd.DataFrame({
    'Actual Return': y_test_monthly,
    'Predicted Return (Gradient Boosting)': y_pred_gb_monthly
})

# Filter the DataFrame for the year 2024 and drop rows with any NaN values
combined_returns_2024 = combined_returns['2024-01-01':'2024-12-31'].dropna()

# Now create a column for months from the index
combined_returns_2024['Month'] = combined_returns_2024.index.strftime('%B %Y')

# Reorder the DataFrame to match the desired output
final_df = combined_returns_2024[['Month', 'Actual Return', 'Predicted Return (Gradient Boosting)']]

# Print the table
print(final_df)

# Optionally, you can save the table to an Excel or CSV file
# final_df.to_csv('gradient_boosting_monthly_returns_2024.csv', index=False)
```

APPENDIX C: Editor's Note

EDITING CERTIFICATE

I, Nkazana Sarah Mwanandimai, confirm that I have language-edited and checked the references of the Master of Data Science

dissertation by

Pendapala Nghilundwa

titled

'Developing a Data-Driven Financial Model for Decision Support in Evaluating Investment Portfolio Performance'

NB: The author has the right to accept or reject changes or amendments made by the editor before submission. The Editor will not be responsible for any changes made after editing.

Signed

Date: 5 September 2025

Ms Nkazana Sarah Mwanandimai, email: nkazana.mwana@gmail.com; Cell: +264 81 325 0360
Associate Member# NK001: Professional Editors' Guild (SA)