UNIVERSITÀ POLITECNICA DELLE MARCHE

FACOLTÀ DI ECONOMIA "GIORGIO FUÀ"

Bachelor's Degree Course in

Digital Economics and Business (DEB)

Artificial Intelligence versus Traditional Data Analytics in Financial Decision-Making: A Comparative Analysis of Methods and Applications

Supervisor:
Prof. Francesco James Mazzocchini

Final Report of:
Fasika Kejella Kumta

Academic Year 2025/2026

INDEX

Introduction…….. 4
Chapter 1: Financial Decision-Making and the Role of Data Analytics 5
1.1 Definition and Foundations of Financial Decision-Making 5
1.2 Importance of Financial Decision-Making in Modern Business 5
1.3 Major Areas of Financial Decision-Making 6
1.3.1 Investment and Capital Budgeting 6
1.3.2 Credit and Risk Assessment 7
1.3.3 Investment Decisions and Portfolio Management 8
1.3.4 Financial Forecasting and Risk Management 8
1.4 Evolution of Data in Finance 9
1.5 Role of Data Analytics in Financial Decision-Making 10
Chapter 2: Traditional Data Analytics in Financial Decision-Making 11
2.1 Traditional Data Analytics: Definition and Characteristics 11
2.2 Main Traditional Data Analysis Methods in Finance 11
2.2.1 Descriptive Statistics and Measures of Dispersion 11
2.2.2 Linear and Logistic Regression Analysis 12
2.2.3 Time-Series Analysis and Forecasting 13
2.2.4 Hypothesis Testing and Correlation Analysis 14
2.3 Applications in Financial Decision-Making 14
2.3.1 Credit Risk Assessment and Scoring 14
2.3.2 Investment Decisions and Portfolio Management 16
2.3.3 Capital Budgeting and Cash Flow Forecasting 16
2.4 Advantages and Limitations of Traditional Analytics 17
2.4.1 Advantages of Traditional Data Analytics 17
2.4.2 Limitations of Traditional Data Analytics 17
Chapter 3: AI-Based Data Analytics in Financial Decision-Making 19
3.1 Introduction to AI-Based Data Analytics 19

3.1.1 Definition and Scope of AI in Data Analytics 19
3.1.2 Evolution from Traditional Analytics to AI-Based Approaches 19
3.2 AI-Based Data Analysis Methods 20
3.2.1 Machine Learning 20
3.2.2 Deep Learning and Neural Networks 21
3.2.3 Natural Language Processing and Sentiment Analysis 21
3.2.4 Unsupervised and Reinforcement Learning 22
3.3 Applications in Financial Decision-Making 23
3.3.1 Credit Scoring and Risk Assessment 23
3.3.2 Investment and Portfolio Management 24
3.3.3 Financial Forecasting 24
3.3.4 AI in Fraud Detection and Anomaly Detection 25
3.4 Advantages and Limitations of AI-Based Analytics 25
3.4.1 Advantages of AI-Based Analytics 25
3.4.2 Limitations of AI-Based Analytics 26
Chapter 4: Comparative Analysis of Traditional and AI-Based Analytics 27
4.1 Research Design and Methodology 27
4.1.1 Research Design 27
4.1.2 Dataset and Data Preprocessing 28
4.1.3 Model Selection and Evaluation Criteria 28
4.2 Empirical Comparison: Credit Default Prediction 29
4.2.1 Logistic Regression 29
4.2.2 Random Forest and XGBoost 30
4.2.3 Model Performance and Comparative Analysis 32
4.3 Broader Comparison of Traditional and AI-Based Analytics 33
4.4 Regulatory, Governance and Ethical Considerations 34
4.5 Hybrid Approaches 35
Conclusions 36
References 37

INTRODUCTION

Financial decision-making has always depended on the availability and interpretation of relevant information. Decisions concerning investment, credit, portfolio management and financial risk require decision-makers to analyze available data, evaluate alternatives and make informed financial decisions. Traditionally, this process has relied mainly on statistical and econometric methods applied to structured and historical financial data. However, the increasing volume and complexity of financial information have expanded the range of analytical methods available. In particular, the development of artificial intelligence and machine learning has introduced new approaches that can identify complex patterns, process larger and more diverse datasets, and support prediction in financial applications.

The purpose of this thesis is to compare traditional and AI-based data analytics in financial decision-making and to understand the strengths and limitations of each approach. The comparison focuses not only on their analytical methods, but also on how they are applied to financial problems and the differences that arise in terms of predictive performance, flexibility and interpretability. Rather than assuming that AI-based methods are necessarily superior to traditional approaches, the thesis considers the circumstances in which each can provide useful support for financial decisions. In this way, the study provides a balanced comparison and considers whether AI should be viewed as a replacement for traditional analytics or as an extension of the existing analytical tools used in finance.

To complement the theoretical comparison, the thesis includes an empirical analysis based on credit-default prediction. The analysis uses the *Default of Credit Card Clients* dataset from the UCI Machine Learning Repository and compares three classification models applied to the same data. Logistic Regression is used as the traditional statistical benchmark, while Random Forest and XGBoost represent the machine-learning approaches. The models are evaluated using accuracy, precision, recall, F1-score and ROC-AUC, allowing the comparison to consider different aspects of predictive performance rather than relying only on overall accuracy. Credit-default prediction provides a practical setting for this comparison because identifying clients with different levels of default risk is directly relevant to financial decision-making.

The thesis is organized into four main chapters, followed by the conclusions. Chapter 1 introduces the foundations of financial decision-making and discusses the role of data analytics in major financial areas, including investment, credit assessment, portfolio management and financial forecasting. Chapter 2 examines traditional data analytics, focusing on the main statistical methods, their financial applications, and their advantages and limitations. Chapter 3 considers AI-based analytics, introducing the main techniques and examining their applications in financial decision-making together with their strengths and limitations. Chapter 4 brings the two approaches together through the empirical credit-default analysis and provides a broader comparison of traditional and AI-based analytics, including considerations related to governance, ethics and hybrid approaches. Finally, the Conclusions summarize the main findings, limitations and possible future developments.

Chapter 1: Financial Decision-Making and the Role of Data Analytics

1.1 Definition and Foundations of Financial Decision-Making

Finance can broadly be understood as the study and management of money, credit, investment, and the allocation of financial resources. Its scope is broad and encompasses several interconnected activities related to how financial resources are obtained, allocated, and managed. In the context of organizations, financial decision-making concerns the choices through which firms acquire assets, obtain the necessary financing, manage their financial resources, and seek to create long-term value (Brigham and Houston, 2019).

Financial activities are generally organized into three closely related areas: financial management, capital markets, and investments. Financial management, also referred to as corporate finance, focuses on decisions concerning the acquisition and management of assets, the financing required to support those assets, and the overall management of the firm with the objective of maximizing its value. Capital markets concern the mechanisms through which capital is supplied and allocated, including the determination of interest rates and the prices of financial securities. Investments, in turn, focus on the analysis and management of financial assets, including the valuation of securities, portfolio construction, and the analysis of financial markets (Brigham and Houston, 2019). Although these areas can be distinguished conceptually, they are strongly interconnected in practice.

Financial decision-making therefore involves evaluating alternative courses of action under conditions of uncertainty and selecting those that best support an organization's financial objectives. Such decisions may concern investment and capital budgeting, financing, risk management, credit assessment, portfolio management, and financial planning. The quality of these decisions depends substantially on the availability and interpretation of relevant financial information.

As financial markets and organizations have become increasingly data-intensive, the role of information in financial decision-making has expanded. Decision-makers are no longer dependent solely on traditional financial statements and historical accounting information; they can increasingly access large volumes of structured and unstructured data from financial markets, transactions, digital platforms, and other sources. Consequently, analytical methods have become an increasingly important component of financial decision-making. The development and application of these methods provide the foundation for examining the differences between traditional data analytics and AI-driven approaches in finance.

1.2 Importance of Financial Decision-Making in Modern Business

Financial decision-making plays a fundamental role in modern business because nearly every organizational activity has financial implications. Decisions concerning production, marketing, human resources, investment, and business expansion ultimately affect the firm's costs, revenues, cash flows, financing requirements, and overall value. Finance therefore provides a common framework through which managers can evaluate the economic consequences of decisions made across different areas of an organization.

A central objective of financial management is the creation of long-term firm value. Managers must therefore consider not only the expected profitability of a decision, but also the timing and uncertainty of future cash flows and the risks associated with them. Since financial resources are limited, firms must compare competing alternatives and allocate capital to activities that are expected to contribute effectively to their strategic objectives. Financial information and performance measures support this process by helping managers assess alternatives, monitor performance and evaluate how efficiently resources are being used.

Financial decisions also take place within broader governance and ethical responsibilities. The separation between ownership and management may create conflicts of interest when managerial objectives differ from those of shareholders, making appropriate governance mechanisms important for aligning decisions with the long-term interests of the firm. At the same time, financial decisions must respect legal and ethical considerations, since practices such as misleading reporting or inappropriate risk-taking can damage both stakeholder interests and firm value. Long-term value creation therefore involves not only financial performance but also responsible decision-making within appropriate governance and ethical boundaries.

The importance of effective financial decision-making becomes particularly evident under uncertain and changing market conditions. Changes in interest rates, asset prices, economic conditions and consumer demand can alter the expected outcomes of financial decisions. Managers consequently require reliable information to compare alternatives, assess potential risks and form expectations about future financial outcomes.

In this context, data analysis provides an important source of support for financial decision-making. Analytical methods allow financial information to be transformed into evidence that can be used to evaluate performance, estimate risk and generate forecasts. As the amount and complexity of available financial data have increased, the techniques used to analyse this information have also developed, creating a progression from established statistical methods toward more advanced AI-based approaches.

1.3 Major Areas of Financial Decision-Making

1.3.1 Investment and Capital Budgeting

Investment decisions represent one of the fundamental areas of financial decision-making because they involve committing resources in the present with the expectation of obtaining economic benefits in the future. In The Capital Budgeting Decision, Bierman and Smidt (2007) describe investment as the commitment of resources in anticipation of future benefits. In a corporate context, capital budgeting therefore involves identifying investment opportunities, assessing their potential consequences, and determining whether the expected benefits justify the resources committed to them.

Capital-budgeting decisions can have both strategic and tactical implications for a firm and typically involve a significant time dimension, since resources are committed today while the resulting benefits may occur over several years. This creates uncertainty because managers must make decisions without knowing the exact future consequences of their investments. Bierman and

Smidt (2007) therefore emphasize the importance of making informed capital-budgeting decisions while recognizing the uncertainty surrounding future outcomes.

According to The Capital Budgeting Decision, a sound capital-budgeting process should take into consideration four fundamental factors: the time value of money, risk considerations, alternative investments, and future opportunities (Bierman and Smidt, 2007). The time value of money recognizes that financial benefits occurring at different points in time cannot be treated as equivalent. Risk considerations arise because future outcomes are uncertain, while alternative investments must be considered because committing resources to one opportunity may prevent their use elsewhere. Future opportunities are also relevant because an investment made today may create additional possibilities for the firm later.

The interaction between time and risk is particularly important in capital budgeting because investment decisions involve committing resources before their future outcomes are fully known. As a result, decision-makers must consider not only expected financial benefits but also their timing, associated risks and contribution to the long-term value of the firm. Bierman and Smidt (2007) also emphasize that quantitative analysis cannot fully eliminate the uncertainty surrounding major investment decisions, particularly when some future consequences cannot be measured reliably in advance. Capital-budgeting decisions therefore combine quantitative evaluation with managerial judgement when assessing how current investments may contribute to future firm value.

1.3.2 Credit and Risk Assessment

Credit risk assessment is an important area of financial decision-making because lending institutions must evaluate the likelihood that borrowers will meet their financial obligations before granting credit. In consumer lending, this process involves not only deciding whether credit should be granted to a new applicant but also determining how existing borrowers should be managed over time. Credit scoring supports the initial lending decision, whereas behavioral scoring assists lenders in decisions concerning existing customers, such as changes in credit limits and responses to repayment difficulties (Thomas, 2000).

Credit assessment depends heavily on the availability and analysis of information about borrowers. According to Thomas (2000), lenders may use information provided in credit applications, data obtained from credit-reference agencies, and historical information concerning previous applicants and their subsequent repayment performance. The objective is essentially a classification problem: using observable borrower characteristics to distinguish applicants who are likely to meet their obligations from those who present a greater risk of default.

The methods used to support credit decisions have evolved considerably over time. Traditionally, lending decisions relied heavily on the judgement of credit analysts, who evaluated factors such as the applicant's character, capital, collateral, repayment capacity and prevailing market conditions. As the volume of credit applications increased, particularly with the expansion of consumer credit, financial institutions increasingly adopted quantitative credit-scoring systems. Statistical techniques such as discriminant analysis, linear regression, logistic regression and classification methods became important tools for predicting default risk and supporting more

systematic lending decisions (Thomas, 2000). This progression also illustrates how credit assessment has become increasingly data-driven, creating an important area of application for both traditional statistical approaches and AI-based methods.

1.3.3 Investment Decisions and Portfolio Management

Investment decisions concern the allocation of financial resources among different assets with the objective of achieving an appropriate balance between expected return and risk. Rather than evaluating investments only based on their individual expected returns, portfolio management considers how different assets interact when they are combined. This principle forms the basis of modern portfolio theory, introduced by Markowitz (1952), who argued that investment decisions should consider both the expected return and the variance of returns. Under this framework, investors select portfolios by considering the trade-off between the return they expect to obtain and the level of risk they are willing to accept.

A central principle of portfolio management is diversification. Holding several securities can reduce portfolio risk when their returns are not perfectly positively correlated. Consequently, the risk of a portfolio depends not only on the risk of each individual security but also on the relationships among their returns. Traditional portfolio analysis therefore relies on statistical measures such as expected returns, variance, covariance and correlation to evaluate alternative combinations of assets. These measures are generally estimated using historical financial data and provide the quantitative basis for determining portfolio weights and constructing efficient portfolios. The efficient frontier represents portfolios that offer the highest expected return for a given level of risk, from which investors can select a portfolio according to their risk-return preferences.

The increasing availability of financial data and advances in computational methods have expanded the analytical approaches available for portfolio decisions. More recent data-driven methods can support activities such as return prediction, risk estimation and portfolio construction, while also allowing more complex relationships within financial data to be considered. However, their application also raises challenges related to data quality, model complexity and interpretability (Bartram et al., 2021).

1.3.4 Financial Forecasting and Risk Management

Financial forecasting and risk management are closely connected components of financial decision-making. Forecasting involves using historical and current information to form expectations about future financial conditions, while risk management focuses on identifying and evaluating potential adverse outcomes. In financial markets, forecasts of asset returns, volatility, and return distributions are particularly important because they influence investment decisions, portfolio allocation, and the assessment of financial risk. Andersen et al. (2003) emphasize that the distributional characteristics of asset returns are relevant to financial instrument pricing, portfolio allocation, performance evaluation, and managerial decision-making. Changes in volatility are particularly important because they provide information about the uncertainty surrounding future financial outcomes.

Traditionally, financial forecasting has relied extensively on statistical and econometric time-series methods that use historical observations to identify patterns and estimate future financial outcomes. Established approaches include autoregressive and ARIMA-type models for time-series forecasting, as well as models specifically designed to estimate and forecast financial volatility. Andersen et al. (2003), for example, identify GARCH as a widely used approach in academic applications and Risk Metrics as an established procedure among financial practitioners. Volatility forecasts can also contribute to risk measures such as Value-at-Risk (VaR), which are used to assess potential portfolio losses and support financial risk management.

Advances in data availability and computational methods have broadened the range of approaches used for financial forecasting and risk assessment. More recent research has increasingly explored machine-learning and deep-learning techniques alongside established statistical methods, particularly where financial data are large or relationships are complex (Sutiene et al., 2024). These developments provide additional possibilities for financial prediction, although the reliability and interpretability of analytical results remain important considerations when forecasts are used to support financial decisions.

1.4 Evolution of Data in Finance

The nature and availability of financial data have changed substantially with developments in information technology and the digitalization of financial activities. Traditionally, financial analysis relied mainly on structured and periodically reported information, such as financial statements, accounting records, market prices and other historical indicators. These data provided the basis for established statistical and econometric analysis. The expansion of digital financial systems, however, has significantly increased the volume, variety and frequency of information available to financial decision-makers.

The development of electronic markets, digital transactions and high-frequency data has also changed the speed and level of detail at which financial information can be observed. Andersen et al. (2003), for example, demonstrate the value of high-frequency intraday data for measuring and forecasting market volatility. Their findings illustrate how more granular financial information can improve the measurement of changing market conditions and provide useful inputs for forecasting and risk assessment.

At the same time, financial analysis has expanded beyond conventional numerical datasets. Modern financial environments generate substantial amounts of unstructured and alternative information, including corporate reports, financial news, conference calls and social-media content. Bartram et al. (2021) describe how natural language processing and related machine-learning techniques can transform textual information into signals relevant to financial prediction and investment analysis. As financial data have become larger, more diverse and more complex, the ability to extract useful information from them has become increasingly important for financial decision-making.

1.5 Role of Data Analytics in Financial Decision-Making

Data analytics provides the link between the availability of financial information and its practical use in decision-making. Financial data alone do not necessarily improve decisions; they must first be processed, analysed and interpreted in ways that provide meaningful information. Analytical methods enable financial decision-makers to identify patterns and relationships in historical data, evaluate financial performance, estimate risk and generate forecasts concerning future outcomes.

As illustrated by the financial areas discussed in the previous sections, analytical methods can support different types of decisions, from assessing the creditworthiness of borrowers to evaluating investments, constructing portfolios and forecasting financial risks. The increasing volume and complexity of financial data have also broadened the range of techniques available for these purposes. Traditional statistical and econometric approaches continue to provide structured and interpretable frameworks for financial analysis, while more recent AI-based methods offer additional possibilities for analyzing larger and more complex datasets and identifying relationships that may be difficult to capture through conventional models.

The development of AI-based analytics does not necessarily imply that traditional approaches have become obsolete. Instead, the two approaches differ in their analytical capabilities, data requirements, interpretability and predictive performance, making their suitability dependent on the particular financial problem being addressed. Understanding these differences provides the basis for the comparison developed in the following chapters, which examine traditional and AI-based data analytics and their application to financial decision-making.

Chapter 2: Traditional Data Analytics in Financial Decision-Making

2.1 Traditional Data Analytics: Definition and Characteristics

Traditional data analytics refers to the use of established statistical and quantitative methods to collect, organize, summarize and interpret data in ways that support decision-making. Statistical analysis generally includes both descriptive and inferential approaches: descriptive statistics summarize the main characteristics of observed data, while inferential methods use sample information to draw conclusions or make predictions beyond the available observations (Mishra et al., 2019). In financial decision-making, these approaches are commonly applied to historical financial data to identify relationships, classify observations, estimate risk and generate forecasts.

Traditional financial analytics generally follows a structured modelling process in which variables, model specifications and analytical objectives are defined in advance. Financial information such as accounting measures, financial ratios, borrower characteristics and historical market data can be incorporated into statistical models to identify relationships or estimate outcomes relevant to financial decisions. Its long-standing application can be observed in areas such as corporate failure prediction, where methods including multivariate discriminant analysis and logistic regression have been used to evaluate firms' financial characteristics and distinguish between failing and non-failing firms (Balcaen and Ooghe, 2006). Similarly, logistic regression can relate borrower characteristics to the probability of default, while traditional portfolio analysis uses measures such as expected returns, variances and covariances to evaluate portfolio risk and return (Thomas, 2000).

An important characteristic of traditional analytical methods is their reliance on predefined statistical structures and assumptions. This structured approach can support interpretation because estimated relationships can often be connected to economically meaningful variables. However, the reliability of the results depends on the specification of the model and the characteristics of the available data. Different statistical techniques may require assumptions concerning distributions, relationships among variables and other properties of the data. In financial applications, issues such as multicollinearity, outliers, non-stationarity, data instability and sampling selectivity can affect model performance (Balcaen and Ooghe, 2006; Mishra et al., 2019). Traditional analytics therefore provides a structured and relatively interpretable framework for financial analysis, although its effectiveness depends on the characteristics of the data and the stability of the relationships being examined.

2.2 Main Traditional Data Analysis Methods in Finance

2.2.1 Descriptive Statistics and Measures of Dispersion

Descriptive statistics provide a fundamental starting point for data analysis by summarizing the main characteristics of a dataset in a form that can be more easily interpreted. Rather than drawing conclusions beyond the available observations, descriptive analysis focuses on organizing and presenting the information contained in the data. Mishra et al. (2019) classify descriptive statistics

into measures of frequency, central tendency and dispersion, which together provide an initial understanding of the distribution and variability of observed data.

Measures of central tendency describe the typical or central position of observations within a dataset and commonly include the mean, median and mode. The mean represents the arithmetic average, while the median identifies the middle value and the mode the most frequently occurring value. The choice between these measures can depend on the characteristics of the data, since the mean is more sensitive to extreme observations, whereas the median can provide a more representative measure when distributions are skewed (Mishra et al., 2019).

Measures of dispersion describe how widely observations vary around their central values. Common measures include variance, standard deviation, range and interquartile range. Variance and standard deviation are particularly useful for assessing the degree of variability within a dataset, while measures such as skewness and kurtosis provide additional information about its distributional shape. Examining these characteristics can help identify unusual observations and determine whether the data are suitable for subsequent statistical analysis (Mishra et al., 2019).

In financial analysis, descriptive statistics provide an initial framework for examining historical information before more complex models are applied. Variables such as asset returns, revenues, financial ratios and borrower characteristics can be summarized in terms of their central tendency, variability and distribution. For example, average historical returns can provide information about investment performance, while measures of dispersion indicate the variability associated with those returns. Descriptive analysis can therefore help identify patterns, extreme observations and other characteristics that may influence subsequent techniques such as regression and time-series analysis.

2.2.2 Linear and Logistic Regression Analysis

Regression analysis is an established statistical method used to examine relationships between variables and determine how changes in one or more explanatory variables are associated with an outcome of interest. In linear regression, a dependent variable is expressed as a function of one or more independent variables together with an error term. Multiple linear regression extends this approach by considering several explanatory variables simultaneously, making it possible to assess their individual relationships with the dependent variable while holding other variables constant. This is particularly useful in financial analysis, where outcomes such as firm performance, returns or other financial measures may be associated with several factors. Buşe, Ganea and Cîrciumaru illustrate this application by using multiple linear regression to examine the relationship between firm-level financial indicators and economic performance.

The reliability of linear regression depends partly on the specification of the model and the assumptions underlying the statistical analysis. These may concern the form of the relationship between variables and the behaviour of the model errors. In the analysis by Buşe, Ganea and Cîrciumaru, for example, assumptions concerning linearity, constant error variance and independence of errors were evaluated when assessing the estimated model. This illustrates an important characteristic of traditional regression analysis: its results must be interpreted in relation to both the selected variables and the assumptions of the model.

While linear regression is appropriate for continuous outcomes, many financial decisions involve categorical outcomes. Logistic regression is particularly useful when the dependent variable is binary, such as whether a borrower defaults or repays a loan. Rather than directly predicting a continuous outcome, logistic regression estimates the probability of an event, allowing borrower or firm characteristics to be related to the likelihood of a particular financial outcome. Thomas (2000) identifies logistic regression as an important conventional method in credit scoring, where borrower information can be used to estimate credit risk.

Logistic regression continues to provide an important benchmark for financial classification problems, particularly in credit-risk analysis. Gunnarsson et al. (2021) describe it as a conventional statistical approach for binary credit-scoring outcomes, in which estimated default probabilities can be used to classify borrowers according to a selected threshold. Its structured and relatively interpretable form makes it useful for examining how explanatory variables relate to financial outcomes. At the same time, dependence on model specification and predefined relationships may limit its ability to capture more complex patterns in financial data. These characteristics make logistic regression particularly relevant to the later empirical comparison with machine-learning models.

2.2.3 Time-Series Analysis and Forecasting

Time-series analysis is a statistical approach used to examine observations recorded sequentially over time. Unlike cross-sectional analysis, time-series methods consider the temporal ordering of data and the possibility that current observations are related to their previous values. This is particularly relevant in finance and economics, where variables such as asset prices, market indices, exchange rates, interest rates and volatility are observed over time. By identifying patterns and relationships in historical observations, time-series models can be used to support inference and forecasting (Anderson, 1994).

Traditional time-series forecasting includes methods such as autoregressive (AR), moving-average (MA), exponential smoothing and Autoregressive Integrated Moving Average (ARIMA) models. Among these, ARIMA has been widely applied to economic and financial series and provides a structured framework for modelling temporal dependence (Siami-Namini and Siami Namin, 2018). An ARIMA model is commonly represented as ARIMA (p,d,q), where the autoregressive component (p) captures relationships with previous observations, the integrated component (d) represents differencing used to address non-stationarity, and the moving-average component (q) incorporates information from previous forecast errors.

Stationarity is an important consideration in time-series modelling because the statistical characteristics of financial and economic series may change over time. When trends or other forms of non-stationarity are present, historical relationships may not provide a stable basis for forecasting, and transformations such as differencing can be used before modelling the series. Simpler approaches such as exponential smoothing instead update forecasts as new observations become available, generally giving greater importance to recent information. Andersen et al. (2003) note the practical use of smoothing approaches in financial forecasting because of their simplicity and ease of implementation.

Despite their established role in financial forecasting, traditional time-series models also have limitations. ARIMA primarily represents linear temporal relationships and may therefore have difficulty capturing more complex nonlinear patterns. Forecasting performance can also be affected by changes in economic conditions, instability in financial markets and the statistical characteristics of the observed series (Siami-Namini and Siami Namin, 2018). Nevertheless, traditional time-series techniques remain an important part of financial analytics and provide useful benchmarks for comparison with more flexible AI-based forecasting approaches.

2.2.4 Hypothesis Testing and Correlation Analysis

Hypothesis testing is a traditional statistical method used to determine whether observed patterns in sample data provide sufficient evidence to support a particular conclusion. It generally involves formulating a null hypothesis (H_0), representing the absence of an effect or difference, and an alternative hypothesis (H_1), representing the presence of an effect or difference. The evidence is evaluated using a test statistic and its associated p-value, which is compared with a predetermined significance level to determine whether the null hypothesis should be rejected (Sabarieswar and Rajaseharan, 2026). In financial analysis, hypothesis testing can be used to evaluate differences in financial outcomes, investment strategies, pricing models and assumptions concerning financial risk.

Correlation analysis, in contrast, examines the degree of association between variables. In finance, it is particularly relevant when analyzing the relationship between the returns of different assets. Correlation provides a standardized measure of the relationship between two variables and is closely related to covariance, allowing relationships to be compared across different assets (Fabozzi et al., 2008). This is particularly important in portfolio analysis, where the degree to which asset returns move together affects diversification and portfolio risk. Correlation can also be examined across time through autocorrelation, which measures the relationship between observations of the same variable at different time periods (Anderson, 1994).

Although hypothesis testing and correlation analysis provide useful information for financial analysis, their results require careful interpretation. Statistical significance does not necessarily imply that an observed relationship is economically important, while statistical tests can also be affected by factors such as sampling bias and multiple testing. Statistical evidence should therefore be interpreted together with the characteristics of the data and the financial context in which the analysis is conducted (Sabarieswar and Rajaseharan, 2026).

2.3 Applications in Financial Decision-Making

2.3.1 Credit Risk Assessment and Scoring

Credit risk assessment represents an important application of traditional data analytics in financial decision-making. Lending institutions must evaluate information about potential borrowers in order to estimate creditworthiness and support decisions concerning whether credit should be granted. Credit scoring provides a quantitative approach to this problem by using information about applicants and the historical performance of previous borrowers to support lending

decisions. Thomas (2000) explains that the information available for credit-scoring decisions can include details provided in credit applications, information obtained from credit-reference agencies, and historical data concerning previous applicants and their subsequent repayment behavior.

From an analytical perspective, credit scoring can be treated as a classification problem in which borrower characteristics are used to distinguish applicants according to their expected credit performance. Historical observations of previous borrowers provide information about the characteristics associated with repayment or default, allowing statistical models to establish relationships between applicant information and credit outcomes. The resulting model can then be applied to new applicants to support the lending decision. In this way, historical credit data are transformed into quantitative information that assists financial institutions in assessing default risk more systematically (Thomas, 2000).

The use of statistical credit-scoring models also represents an important development from earlier judgement-based lending practices. Traditionally, credit analysts evaluated applicants using factors such as character, capital, collateral, repayment capacity and market conditions. However, the increasing scale of consumer lending created a greater need for systematic and automated methods of evaluating applications. Thomas (2000) describes how statistical and operational-research techniques consequently became important in the construction of credit scorecards, with methods such as discriminant analysis and logistic regression being used to classify borrowers according to credit risk.

Among these methods, logistic regression has become particularly important in traditional credit scoring. The method estimates the relationship between borrower characteristics and the probability of a particular credit outcome, making it suitable for situations in which applicants must be classified according to whether they are expected to repay or default. Thomas (2000) identifies logistic regression as a major approach to classification in credit scoring, while Gunnarsson et al. (2021) similarly classify logistic regression as a traditional statistical credit-scoring model.

Traditional credit-scoring models therefore illustrate how statistical analysis can directly support financial decision-making by allowing institutions to use historical borrower information to evaluate applications more systematically. Gunnarsson et al. (2021) note that empirical credit-scoring models can support faster and more consistent assessment, while improvements in predictive accuracy may also have financial consequences for institutions managing large loan portfolios. Advances in computing and data availability have subsequently introduced alternative approaches, including random forests, XGBoost and neural networks, creating the basis for comparing traditional statistical credit-scoring models with AI-based methods.

2.3.2 Investment Decisions and Portfolio Management

Traditional data analytics has long played an important role in investment and portfolio management by providing quantitative measures for evaluating the relationship between expected return and risk. Portfolio decisions require investors to determine how financial resources should be distributed among different assets while considering both their return objectives and tolerance for risk. Within traditional portfolio analysis, historical financial information and statistical measures are therefore used to transform past asset performance into inputs that can support asset-allocation decisions.

A fundamental framework for this process is the mean-variance approach developed by Markowitz. Rather than evaluating securities only according to their individual characteristics, the approach considers how each asset contributes to the risk and expected return of the portfolio. Expected portfolio return is determined by the expected returns of the individual assets and their respective portfolio weights, while risk is assessed using measures such as variance and standard deviation. Importantly, portfolio risk also depends on the covariance or correlation between asset returns, meaning that the relationships among investments influence the overall level of risk (Fabozzi, Markowitz and Gupta, 2008).

These statistical relationships provide the basis for diversification. When asset returns are not perfectly positively correlated, combining them can reduce the overall risk of a portfolio without necessarily producing a corresponding reduction in expected return. Traditional portfolio analysis therefore uses estimates of expected returns, variances and covariances or correlations to examine different asset combinations. Through mean-variance analysis, these inputs can be used to construct an efficient frontier representing portfolios that provide the highest expected return for a given level of risk. Investors can subsequently select among these efficient portfolios according to their individual risk and return preferences (Fabozzi, Markowitz and Gupta, 2008).

Traditional portfolio analytics, however, depends strongly on the quality of the statistical inputs used in the analysis. As the number of securities increases, estimating expected returns and variance-covariance relationships becomes increasingly demanding. Bartram et al. (2021) identify difficulties in estimating these inputs as important shortcomings of classical Markowitz portfolio optimization.

Overall, traditional portfolio analytics provides a systematic framework for evaluating expected return, risk and diversification. However, its dependence on statistical assumptions and estimated inputs can create practical limitations, providing an important basis for comparison with AI-based portfolio approaches.

2.3.3 Capital Budgeting and Cash Flow Forecasting

Traditional data analysis plays an important role in capital budgeting, where firms evaluate whether long-term investment opportunities are expected to generate sufficient future benefits. Since these decisions depend on outcomes occurring over several periods, historical financial

information can be used to estimate future cash flows and support investment evaluation. These forecasts can subsequently be incorporated into conventional capital-budgeting measures, such as net present value and internal rate of return, to compare alternative investments.

The relationship between financial information, forecasting and investment decisions is also illustrated by Gilchrist and Himmelberg (1995), who examine the role of cash flow in explaining firms' investment behaviour. Their analysis uses traditional econometric methods to incorporate historical firm-level information into estimates of future investment opportunities and shows that cash flow can provide information about both expected investment opportunities and the availability of internal finance. The study therefore illustrates how structured financial data and statistical forecasting can support corporate investment decisions, while also showing that financial constraints may influence how firms respond to investment opportunities.

2.4 Advantages and Limitations of Traditional Analytics

2.4.1 Advantages of Traditional Data Analytics

Traditional data analytics provides an established and systematic framework for organizing financial information, examining relationships and supporting quantitative decision-making. Methods such as descriptive statistics, regression, correlation and time-series analysis allow financial data to be summarized, relationships between variables to be evaluated and forecasts to be generated. Their structured statistical foundations also allow analysts to assess how variables and model specifications contribute to the resulting conclusions (Mishra et al., 2019).

Another important advantage is the relative interpretability of many traditional statistical methods. Their predefined structures and clearly specified variables can make analytical results easier to understand and communicate, which is particularly valuable in financial decision-making where the reasoning behind an estimate or prediction may be important. Traditional analytics therefore remains useful across areas such as credit-risk assessment, portfolio analysis and financial forecasting by providing structured and interpretable quantitative evidence to support financial decisions.

2.4.2 Limitations of Traditional Data Analytics

Despite their established role in financial decision-making, traditional analytical methods have limitations that can affect their reliability and predictive performance. One important limitation is their dependence on predefined model structures and statistical assumptions concerning factors such as data distributions, relationships among variables and stability of the underlying process. When these assumptions are not adequately satisfied, the resulting estimates and conclusions may become less reliable (Mishra et al., 2019).

A further limitation concerns the complexity and changing nature of financial data. Relationships identified from historical observations may not remain stable as firms, markets and economic conditions change. Balcaen and Ooghe (2006) identify non-stationarity, data instability and sampling choices as important problems in classical financial models. Traditional approaches may

also have difficulty representing complex relationships when restrictive functional forms are imposed. These challenges become increasingly relevant as financial datasets contain larger numbers of potential predictors and interactions, which can make some conventional statistical approaches more difficult to apply (Athey and Imbens, 2019).

These limitations do not make traditional analytics ineffective, but they can become particularly important when financial data are high-dimensional, relationships are nonlinear or unstable, or conventional statistical assumptions are difficult to maintain. Such challenges have contributed to increasing interest in more flexible computational approaches, particularly machine learning and other AI-based methods, which are examined in the following chapter.

Chapter 3: AI-Based Data Analytics in Financial Decision-Making

3.1 Introduction to AI-Based Data Analytics

3.1.1 Definition and Scope of AI in Data Analytics

Artificial intelligence has become increasingly integrated with data analytics as organizations seek to extract useful information from growing volumes and varieties of data. In this context, AI-based data analytics can be understood as the application of artificial intelligence techniques to the processing and analysis of data in order to identify patterns, generate predictions and produce information that can support decision-making. Zong and Guan (2025) describe the integration of AI with data analytics as *intelligent data analytics*, in which large datasets are analysed to identify patterns, correlations and potential opportunities that may be difficult to detect through conventional analytical approaches. AI can therefore extend the analytical process beyond the summarization of historical information by supporting more automated and adaptive forms of analysis.

The scope of artificial intelligence in data analytics is broader than machine learning alone. Cao, Yang and Yu (2021) describe AI and data science as encompassing a range of computational approaches, including machine learning, deep learning, natural language processing, pattern recognition and intelligent decision-support systems. Machine learning therefore represents an important component of AI rather than the entirety of the field. This distinction is relevant to financial analytics because different AI-based techniques can be applied depending on the type of data and the financial problem being addressed.

Within finance, AI-based analytics can be applied across activities including banking, lending, trading, wealth management and risk management (Cao et al., 2021). More specifically, AI-related methods can support tasks such as financial forecasting, portfolio optimization, risk scoring, fraud detection and sentiment analysis. This broad range of applications illustrates how AI-based analytics can support different forms of financial decision-making using both conventional financial information and more complex sources of data.

3.1.2 Evolution from Traditional Analytics to AI-Based Approaches

Financial data analysis has traditionally relied on statistical and econometric methods to examine relationships between variables and support financial decision-making. Techniques such as regression analysis and time-series models provided structured frameworks for explaining economic relationships and forecasting financial outcomes. Hoang and Wiegratz (2023) note that traditional econometric analysis is primarily concerned with estimating the direction and strength of relationships between economic variables, with linear regression being one of its dominant methods because of its flexibility and interpretability.

The increasing availability and complexity of financial data, however, has contributed to the growing use of machine-learning approaches. Financial analysis increasingly involves datasets containing large numbers of observations and variables as well as unconventional information such as text, images and other forms of unstructured data. Hoang and Wiegratz (2023) explain that machine-learning methods can provide particular benefits when dealing with high-dimensional data, where numerous variables, nonlinearities and interaction effects may be relevant. They can also extract economically relevant information from unconventional data that are difficult to incorporate directly into conventional analytical models.

This development has broadened the analytical approaches available in finance. While traditional econometric methods often emphasize explaining relationships between variables, machine learning is frequently oriented toward prediction and the identification of patterns within data. Supervised learning can use observed data to predict outcomes for previously unseen observations, while unsupervised learning can identify underlying structures without requiring a predefined outcome variable. Machine learning should therefore not necessarily be viewed as replacing traditional statistical and econometric methods, but rather as extending the financial analytical toolkit with approaches suited to prediction, high-dimensional information and more complex forms of data (Hoang and Wiegratz, 2023).

3.2 AI-Based Data Analysis Methods

3.2.1 Machine Learning

Machine learning represents one of the main components of AI-based data analytics and includes methods that learn patterns from observed data to perform tasks such as prediction and classification. In financial analysis, supervised machine learning is particularly relevant when historical observations contain a known outcome that can be used to train a predictive model. Hoang and Wiegratz (2023) explain that supervised learning uses labelled data, consisting of predictor variables and a target variable, to construct models that can generate predictions for previously unseen observations. Compared with traditional econometric analysis, supervised machine learning generally places greater emphasis on out-of-sample predictive performance.

Decision trees are a widely used supervised machine-learning technique that generate predictions by repeatedly dividing observations according to the values of predictor variables. Their structure allows the variables and thresholds involved in the prediction process to be observed, making individual trees relatively interpretable. Decision trees can also capture interactions among predictor variables without requiring these relationships to be specified in advance (Hoang and Wiegratz, 2023). However, more advanced tree-based methods generally combine multiple trees to improve predictive performance.

Random Forest is an ensemble method that constructs multiple decision trees using different samples of the available data and combines their predictions. Boosting follows a different approach by constructing trees sequentially, with later trees focusing more strongly on observations that previous trees predicted poorly. These ensemble approaches can improve predictive performance compared with individual decision trees, although combining many trees also reduces model

transparency. Hoang and Wiegratz (2023) therefore highlight an important trade-off between model complexity, predictive performance and interpretability.

A prominent boosting method is Extreme Gradient Boosting (XGBoost), which has been applied to financial classification problems such as credit scoring. Gunnarsson et al. (2021) compare several credit-scoring techniques and report strong predictive performance for XGBoost across the datasets examined, while Random Forest also represents an important benchmark among machine-learning classifiers. However, both models are generally less transparent than conventional approaches such as logistic regression. Random Forest and XGBoost therefore provide useful examples of how machine-learning methods can improve predictive flexibility while introducing greater model complexity, a trade-off examined empirically in Chapter 4.

3.2.2 Deep Learning and Neural Networks

Artificial neural networks (ANNs) are machine-learning models composed of interconnected processing units, or neurons, typically organized into input, hidden and output layers. Connections between neurons are associated with weights that are adjusted during learning, allowing the network to identify relationships between input data and the desired output. With activation functions, neural networks can represent nonlinear relationships between variables, making them suitable for problems involving complex patterns in data (Artificial Intelligence for Business; Gunnarsson et al., 2021).

Deep learning extends this approach by using neural networks with multiple hidden layers, allowing models to learn increasingly complex representations from data. Different architectures can also be designed for particular types of information. Recurrent neural networks (RNNs), for example, process sequential observations by allowing information from previous observations to influence subsequent processing. Long Short-Term Memory (LSTM) networks are a development of recurrent networks designed for sequential data and have consequently been applied to financial problems involving information observed over time (Artificial Intelligence for Business).

In finance, deep-learning methods have been applied to areas including credit scoring, financial forecasting and portfolio analysis. Gunnarsson et al. (2021), for example, discuss LSTM applications to sequential credit-scoring data, while Sutiene et al. (2024) identify LSTM and other deep-learning architectures among methods used in financial forecasting. However, greater model complexity does not necessarily produce better predictive performance. Gunnarsson et al. (2021) found that deeper neural networks did not generally outperform shallower alternatives in their credit-scoring comparison, illustrating that the usefulness of deep learning depends on the characteristics of the data and the financial problem being addressed.

3.2.3 Natural Language Processing and Sentiment Analysis

Natural Language Processing (NLP) is an area of artificial intelligence concerned with making human language accessible to computational processing and analysis. NLP provides techniques for transforming unstructured textual information into a form from which useful information can be extracted. This capability is particularly relevant in finance, where potentially valuable

information is contained not only in numerical datasets but also in financial reports, company communications, news and social media (Artificial Intelligence for Business).

One important application of NLP is sentiment analysis, which uses textual information to identify opinions or attitudes that may be classified, for example, as positive, negative or neutral. In finance, sentiment analysis can transform qualitative information about companies, securities or financial markets into indicators that can contribute to financial analysis (Sutiene et al., 2024). Bartram et al. (2021) discuss the use of NLP to extract sentiment and other predictive information from sources such as annual reports, analyst conference calls, news and social-media content, which can then contribute to return prediction and investment analysis.

However, extracting reliable information from financial language remains challenging because meaning can depend strongly on context and financial terminology may be highly specialized. Sutiene et al. (2024) also identify limited availability of labelled data as a challenge for financial sentiment analysis. Moreover, textual information from sources such as social media can contain substantial noise. Bartram et al. (2021) report mixed evidence regarding the ability of social-media sentiment to predict subsequent returns, illustrating that the usefulness of NLP depends on the quality and context of the textual data as well as the methods used to analyse it.

3.2.4 Unsupervised and Reinforcement Learning

Unlike supervised learning, unsupervised learning operates without a predefined target variable. Hoang and Wiegratz (2023) explain that unsupervised learning uses unlabeled data to identify underlying structures and characteristics within a dataset. One of its principal techniques is clustering, which groups observations according to their similarities. In finance, clustering can reveal similarities among financial assets and reduce the complexity of financial data, making it useful for areas such as portfolio analysis and optimization (Sutiene et al., 2024).

Unsupervised learning can also support anomaly or outlier detection, where the objective is to identify observations that differ substantially from the general structure of the data. Hoang and Wiegratz (2023) identify outlier detection as an unsupervised-learning approach and highlight its usefulness in high-dimensional settings. In financial analytics, identifying unusual patterns can contribute to applications such as abnormal trading and fraud detection (Cao, Yang and Yu, 2021).

Reinforcement learning (RL), in contrast, focuses on sequential decision-making. An agent interacts with an environment, performs actions and receives feedback in the form of rewards or penalties, gradually learning which actions are most effective for achieving a particular objective (*Artificial Intelligence for Business*). This differs from supervised learning because the correct action is not provided through labelled examples; instead, the agent learns through interaction and experience.

This characteristic makes reinforcement learning relevant to financial decisions that involve repeated actions over time. Cao et al. (2021) identify applications including portfolio and trading-strategy optimization, while Bartram et al. (2021) discuss its use in areas such as portfolio management and optimal trade execution. Overall, unsupervised learning provides tools for

discovering structures and unusual observations in financial data, while reinforcement learning provides an approach for learning sequential financial decisions.

3.3 Applications in Financial Decision-Making

3.3.1 Credit Scoring and Risk Assessment

Credit scoring represents an important application of artificial intelligence in financial decision-making because lending institutions must evaluate the probability that borrowers will meet their repayment obligations. As discussed in Chapter 1, credit-scoring models use information about applicants and historical loans to support decisions concerning whether credit should be granted. From an analytical perspective, the problem can be treated as a classification task in which models distinguish applicants who are likely to repay from those who are more likely to default. Gunnarsson et al. (2021) explain that a credit-scoring model can be constructed from historical loans containing borrower and loan characteristics together with their observed outcomes, allowing the estimated probability of default to be compared with a decision threshold when assessing creditworthiness.

The development of machine learning has expanded the range of models available for this task. Credit scoring has historically relied heavily on conventional statistical classification techniques, particularly logistic regression, while decision trees, neural networks and, more recently, ensemble methods have provided alternative approaches (Gunnarsson et al., 2021). Machine-learning methods therefore allow lenders to construct predictive models from historical borrower information and use the identified patterns to estimate the credit risk associated with new applications. The increasing use of these models is particularly relevant in consumer lending, where large numbers of applications make automated and consistent evaluation important. Thomas (2000) similarly describes how the expansion of consumer credit contributed to the automation of lending decisions and the development of quantitative credit-scoring techniques.

Evidence from credit-scoring research illustrates the potential of machine-learning approaches but also shows that more complex models are not necessarily more effective. Gunnarsson et al. (2021) compare conventional, ensemble and deep-learning classifiers across multiple credit-scoring datasets and performance measures. In their analysis, XGBoost achieved the strongest overall classification performance among the models examined, while deep neural networks did not outperform their shallower counterparts and required substantially greater computational resources. The authors also refer to previous large-scale evidence suggesting that Random Forest can provide a useful benchmark for evaluating credit-scoring classifiers. These findings demonstrate that the value of AI in credit assessment does not depend simply on increasing model complexity; the suitability of the model depends on the characteristics of the credit data and the objective of the analysis.

At the same time, predictive performance is not the only consideration in lending decisions. Ensemble models such as XGBoost and Random Forest can be difficult to interpret because their predictions result from the combination of many underlying models. Gunnarsson et al. (2021) consequently note that conventional approaches such as logistic regression may remain preferable when interpretability is a primary requirement, whereas XGBoost may be attractive when classification performance is the principal objective. Credit scoring therefore illustrates an

important trade-off in AI-based financial analytics: machine-learning models can provide strong predictive capabilities, but their practical use must also consider transparency, computational requirements and the ability of financial institutions to understand and justify lending decisions.

3.3.2 Investment and Portfolio Management

AI-based analytics has become increasingly relevant to investment and portfolio management by supporting decisions concerning asset selection, allocation and portfolio adjustment. Machine-learning techniques can contribute to different stages of the investment process, including signal generation, portfolio construction and trade execution (Bartram et al., 2021). By analysing financial and market data, these methods can provide additional information for selecting assets and allocating capital under changing market conditions.

Different AI techniques can contribute to portfolio decisions in different ways. Machine learning can support asset-price prediction, while dimensionality-reduction methods can identify patterns across large sets of asset prices and contribute to portfolio diversification. Deep-learning techniques can also be applied to portfolio optimization, while reinforcement learning is particularly relevant to sequential decisions in which portfolio allocations are adjusted as market conditions change (Sutiene et al., 2024). In addition, NLP and sentiment analysis can incorporate information from financial news and other textual sources, extending portfolio analysis beyond conventional numerical market data.

However, the use of AI does not guarantee superior investment performance. Bartram et al. (2021) conclude that machine-learning techniques show considerable potential for active portfolio management, particularly because they can capture nonlinear patterns and support different stages of the investment process. Nevertheless, their examination of active exchange-traded funds using machine learning finds mixed investment performance, demonstrating that promising analytical capabilities do not necessarily translate into consistently superior portfolio returns. AI-based portfolio management can also create challenges concerning transparency and explainability, since investment decisions generated by complex models may be difficult for portfolio managers and investors to interpret (Sutiene et al., 2024).

3.3.3 Financial Forecasting

Financial forecasting is an important application of AI-based analytics because many financial decisions depend on expectations about future market and economic conditions. Machine-learning methods have been applied to variables such as asset returns, prices, volatility and macroeconomic indicators. A key advantage is their ability to incorporate large sets of potential predictors and model complex relationships within high-dimensional data. Masini, Medeiros and Mendes (2021) review the use of techniques including regularization, nonlinear models and ensemble methods in economic and financial forecasting, while Sutiene et al. (2024) discuss machine-learning applications to predictor selection and volatility forecasting.

Deep learning has further expanded the methods available for financial time-series forecasting, particularly through recurrent architectures such as LSTM and GRU that can process sequential observations. Sutiene et al. (2024) report applications in which LSTM achieved stronger

forecasting results than ARIMA. However, such findings should not be interpreted as evidence that deep learning is universally superior, since forecasting performance depends on the dataset, forecasting objective and model specification. Masini et al. (2021) similarly emphasize the importance of comparing alternative methods based on their predictive performance rather than assuming that a particular approach will consistently perform better.

AI-based forecasting can therefore support investment, portfolio and risk-management decisions by expanding the range of data and relationships that can be incorporated into financial forecasts. Rather than simply replacing established time-series models, AI provides additional approaches for modelling complex financial information. Its usefulness ultimately depends on the reliability of its out-of-sample predictions and appropriate model evaluation.

3.3.4 AI in Fraud Detection and Anomaly Detection

Financial fraud detection is an important application of AI-based analytics because fraudulent activities can involve complex and changing patterns that may be difficult to identify through predefined rules alone. Yang, Shukur and Sahran (2026), in their systematic review of 122 studies, identify machine learning, ensemble learning, deep learning and other AI-based approaches across areas including credit-card fraud, insurance fraud, loan fraud and anti-money laundering. Methods such as Random Forest and XGBoost can use historical financial data to distinguish fraudulent from legitimate observations, while more complex approaches can capture nonlinear, temporal and relational patterns within financial data.

AI-based anomaly detection is also useful for identifying observations that differ substantially from normal financial behaviour. This is particularly relevant when fraud patterns evolve over time or previously unseen forms of fraudulent behaviour differ from historical cases. However, fraud-detection models face challenges such as changing fraud patterns, limited labelled data and highly imbalanced datasets, since fraudulent observations typically represent only a small proportion of total transactions (Yang, Shukur and Sahran, 2026).

Class imbalance also makes model evaluation particularly important. Yang, Shukur and Sahran (2026) emphasize that high overall accuracy can provide a misleading indication of fraud-detection performance when most observations belong to the legitimate class. Measures such as precision, recall and F1-score can therefore provide more informative assessments of the model's ability to identify fraudulent observations. AI can consequently strengthen fraud and anomaly detection, but its practical usefulness depends on appropriate model evaluation, data quality and the ability to respond to changing patterns of financial fraud.

3.4 Advantages and Limitations of AI-Based Analytics

3.4.1 Advantages of AI-Based Analytics

AI-based analytics offers several advantages in financial analysis, particularly when problems involve complex relationships, large numbers of variables or diverse forms of data. One important strength is its predictive orientation. Machine-learning methods can be designed to improve

performance on previously unseen observations and can capture nonlinear relationships and interactions that may be difficult to represent through simpler statistical specifications. These characteristics make AI-based methods particularly useful for financial problems involving prediction and classification, although their suitability ultimately depends on the analytical objective and the characteristics of the data (Hoang and Wiegratz, 2023).

AI-based approaches can also process high-dimensional and unconventional information. Large numbers of potential predictors can be incorporated into financial models, while techniques such as NLP allow information from textual sources to contribute to financial analysis. Hoang and Wiegratz (2023) therefore describe machine learning as extending the financial analytical toolkit by allowing more complex forms of information and relationships to be considered. AI-based analytics should consequently be viewed as expanding the range of analytical possibilities available for financial decision-making rather than providing a universally superior alternative to traditional methods.

3.4.2 Limitations of AI-Based Analytics

Despite their analytical capabilities, AI-based methods also present important limitations in financial decision-making. One of the most significant concerns is interpretability. As model complexity increases, understanding how individual predictions are generated can become more difficult. Hoang and Wiegratz (2023) identify a trade-off between predictive performance and interpretability, with more complex machine-learning methods generally being less transparent than simpler alternatives. This is particularly relevant in financial contexts where decisions may require explanation and accountability.

A further limitation concerns data quality and generalization. Since AI models learn from historical observations, their reliability depends on the quality and representativeness of the training data, while relationships identified historically may change under new financial conditions. Yang, Shukur and Sahran (2026), for example, identify challenges including concept drift, class imbalance, inconsistent labelling and distributional changes in financial fraud detection. Strong historical performance therefore does not necessarily guarantee equally reliable predictions when the underlying data change.

AI-based models may also require greater computational and implementation resources, particularly as model complexity increases. More sophisticated models can require additional computing resources, training and parameter tuning without necessarily producing corresponding improvements in performance. The suitability of AI-based analytics therefore depends on the balance between predictive performance, interpretability, data quality and practical implementation. These trade-offs provide the basis for the empirical comparison developed in Chapter 4.

Chapter 4: Comparative Analysis of Traditional and AI-Based Analytics

4.1 Research Design and Methodology

This chapter provides an empirical comparison between traditional and AI-based analytical methods in the context of credit default prediction. The analysis focuses on whether differences between the two approaches can be observed when they are applied to the same financial dataset and evaluated under the same conditions. Logistic Regression is used to represent the traditional statistical approach, while Random Forest and XGBoost are used as machine learning alternatives.

Credit default prediction was selected because it provides a practical setting in which data analysis directly supports financial decision-making. Financial institutions need to distinguish between borrowers with different levels of credit risk, and the quality of such predictions can influence lending and risk-management decisions. The empirical analysis therefore focuses on a binary classification problem: predicting whether a credit-card client will default on the following month's payment.

Rather than attempting to identify a universally superior model, the analysis compares the predictive behaviour of the three methods using a common set of evaluation measures. This makes it possible to examine not only overall predictive accuracy but also differences in the models' ability to identify clients who actually default.

4.1.1 Research Design

The empirical analysis follows a comparative quantitative approach. Three classification models are estimated using the same credit-card client data: Logistic Regression, Random Forest and XGBoost. Logistic Regression provides the traditional benchmark against which the two machine learning methods are compared.

To ensure comparability, the dataset was divided into a training set and a test set. Eighty percent of the observations were used for model training and the remaining 20% for testing. This resulted in 24,000 observations in the training set and 6,000 observations in the test set. A stratified split was used so that the proportion of default and non-default cases remained similar in both samples. A fixed random state was also applied to make the split reproducible.

The same training and test observations were used across the three models. This is important for the comparison because differences in performance can then be attributed more directly to the modelling approaches rather than to different test samples.

The target variable is *default payment next month*, where a value of 1 represents default and 0 represents non-default. The explanatory variables contain information on the clients' credit limits,

demographic characteristics, repayment history, bill statement amounts and previous payment amounts.

4.1.2 Dataset and Data Preprocessing

The empirical analysis uses the Default of Credit Card Clients dataset available through the UCI Machine Learning Repository. The dataset was originally developed by Yeh and Lien in the context of predicting the probability of default among credit-card clients in Taiwan. The version used in this analysis contains 30,000 observations and 23 explanatory variables, in addition to an identification variable and the binary default outcome.

The explanatory variables cover several aspects of the clients' financial and demographic profiles. These include the amount of credit granted, sex, education, marital status and age. The dataset also contains six months of repayment-status information, six bill statement amounts and six previous payment amounts. Together, these variables provide information on both client characteristics and past payment behaviour.

Before estimating the models, the dataset was checked for missing and duplicate observations. No missing values were identified. No completely duplicated rows were found when the identification variable was included. When the ID variable was excluded from the duplicate check, 35 observations shared identical values across the remaining variables. These observations were retained because identical recorded characteristics do not necessarily imply that they represent the same client. The ID variable itself was excluded from the predictors because it serves only as an observation identifier and does not provide meaningful information for predicting default.

The categorical variables also required some preprocessing. In the original documentation, education is represented by four main categories, while the released dataset contains additional codes of 0, 5 and 6. These undocumented categories accounted for a relatively small part of the sample and were grouped into category 4, representing the broader "other" category. After recoding, the education variable contained 10,585 observations in category 1, 14,030 in category 2, 4,917 in category 3 and 468 in category 4. Similarly, the undocumented value 0 in marital status was combined with category 3 ("other"), resulting in 13,659 married, 15,964 single and 377 other observations.

The dependent variable showed a moderate class imbalance. Of the 30,000 clients, 23,364 (77.88%) did not default, while 6,636 (22.12%) defaulted. This distribution is particularly relevant when evaluating the models because a high overall accuracy can partly result from correctly predicting the much larger non-default group. For this reason, model performance is assessed using several measures rather than accuracy alone.

4.1.3 Model Selection and Evaluation Criteria

Three classification models were selected for the empirical comparison: Logistic Regression, Random Forest and XGBoost. Logistic Regression was chosen as the traditional statistical benchmark, while Random Forest and XGBoost represent machine learning approaches. Using

these three models allows the analysis to compare a relatively simple and interpretable statistical method with more flexible tree-based methods.

Logistic Regression is widely used for binary classification problems such as credit scoring because it estimates the probability of an event occurring based on a set of explanatory variables. Yeh and Lien (2009) describe one of its main advantages as the ability to produce a relatively simple probabilistic classification model, while also noting its limitations in capturing nonlinear relationships and interactions among explanatory variables. In the present analysis, Logistic Regression therefore serves as the traditional benchmark for predicting whether a client defaults.

Random Forest and XGBoost were selected as the AI-based alternatives. Both are tree-based ensemble methods capable of modelling more complex relationships among predictors without requiring the same linear structure as Logistic Regression. Random Forest combines predictions from multiple decision trees, while XGBoost builds trees sequentially, with later trees attempting to improve the errors made by earlier ones. The purpose here is not to identify the most sophisticated possible model, but to examine whether these machine learning approaches provide an improvement over the traditional benchmark when applied to the same data.

Model performance was evaluated on the same test sample using accuracy, precision, recall, F1-score and ROC-AUC. Accuracy measures the overall proportion of correct classifications, while precision indicates how frequently predicted defaults are actually defaults. Recall measures the proportion of actual defaulters correctly identified, and the F1-score combines precision and recall into a single measure. ROC-AUC evaluates the model's ability to distinguish between default and non-default cases across different classification thresholds.

Using several evaluation measures is particularly important for this dataset because only 22.12% of the observations belong to the default class. A model could therefore achieve relatively high accuracy mainly by predicting the majority non-default class correctly. Yeh and Lien (2009) encountered the same issue in their credit-default analysis and argued that simple error rates were not sufficiently sensitive when the majority of clients were non-risky. For this reason, the comparison in this study does not rely on accuracy alone.

4.2 Empirical Comparison: Credit Default Prediction

Following the preprocessing and train-test split described in the previous section, the three selected models were trained using the same training sample and evaluated on the same test sample. The test set contained 6,000 observations, of which 4,673 were non-default cases and 1,327 were default cases. Using the same observations for evaluation provides a consistent basis for comparing the predictive performance of Logistic Regression, Random Forest and XGBoost.

4.2.1 Logistic Regression

Logistic Regression was used as the traditional benchmark model. This choice is consistent with its long-standing application in credit scoring and default prediction. Yeh and Lien (2009), for example, include Logistic Regression among the statistical methods used for credit-risk prediction and highlight its ability to provide a relatively simple probabilistic classification model. At the

same time, they point out that its linear structure can make it less suitable for capturing nonlinear relationships and interactions among explanatory variables.

When applied to the test sample, the Logistic Regression model achieved an accuracy of 80.90% and a precision of 69.30%. However, its recall was considerably lower at 24.49%, meaning that the model identified only around one quarter of the clients who actually defaulted. The resulting F1-score was 36.19%, while the ROC-AUC was 0.710, indicating a moderate ability to distinguish between defaulting and non-defaulting clients.

The confusion matrix provides a clearer view of this result. Of the 4,673 clients who did not default, 4,529 were correctly classified, while 144 were incorrectly predicted as defaulters. Among the 1,327 clients who actually defaulted, however, only 325 were correctly identified, while 1,002 were classified as non-defaulters.

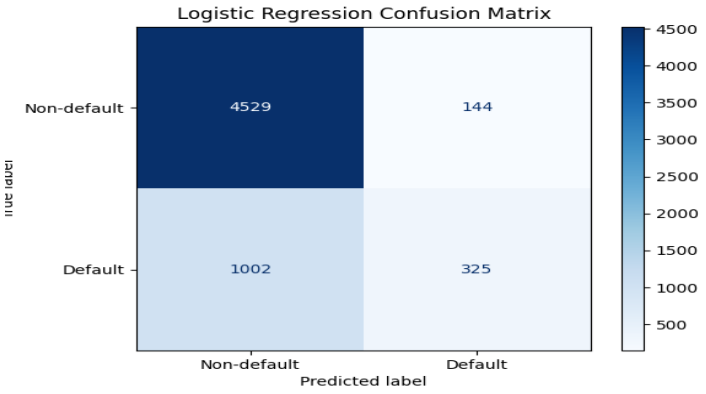


Figure 1. Logistic Regression Confusion Matrix

This result shows why accuracy alone can give an incomplete picture of model performance. Although the model correctly classified more than 80% of the test observations, much of this performance came from correctly identifying the larger non-default group. Its relatively low recall shows that detecting actual default cases was considerably more difficult. This is particularly relevant in credit-risk analysis, where failing to identify a potential defaulter may carry greater consequences than simply achieving a high overall classification rate. Yeh and Lien (2009) similarly warn that overall error rates can be insensitive when the dataset contains a much larger proportion of non-risky clients.

4.2.2 Random Forest and XGBoost

Random Forest and XGBoost were applied as the machine-learning alternatives to the Logistic Regression benchmark. Both models were trained and evaluated using the same training and test samples used for Logistic Regression, allowing their predictive performance to be compared under consistent conditions.

The Random Forest model achieved an accuracy of 81.33%, slightly higher than the 80.90% obtained by Logistic Regression. Its precision was 63.67%, while recall increased to 36.32% and

the F1-score reached 46.26%. The model also produced a ROC-AUC of 0.7584, indicating a stronger ability to distinguish between default and non-default cases than Logistic Regression.

The Random Forest confusion matrix shows that 4,398 of the 4,673 non-defaulting clients were correctly classified, while 275 were incorrectly classified as defaulters. Among the 1,327 clients who actually defaulted, 482 were correctly identified, whereas 845 were incorrectly classified as non-defaulters. Compared with Logistic Regression, Random Forest therefore identified substantially more actual default cases, although this improvement was accompanied by a larger number of false-positive default predictions.

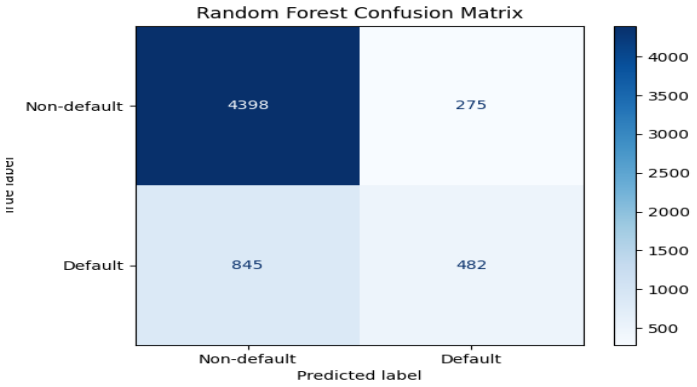


Figure 2. Random Forest Confusion Matrix

The XGBoost model produced the highest overall accuracy among the three models, at 81.87%. Precision was 66.71%, recall was 35.95%, and the F1-score was 46.72%. XGBoost also obtained a ROC-AUC of 0.7785, the highest value observed in the empirical analysis.

Its confusion matrix shows that XGBoost correctly classified 4,435 non-defaulting clients and incorrectly classified 238 as defaulters. Among actual defaulters, 477 were correctly identified, while 850 were classified as non-defaulters.

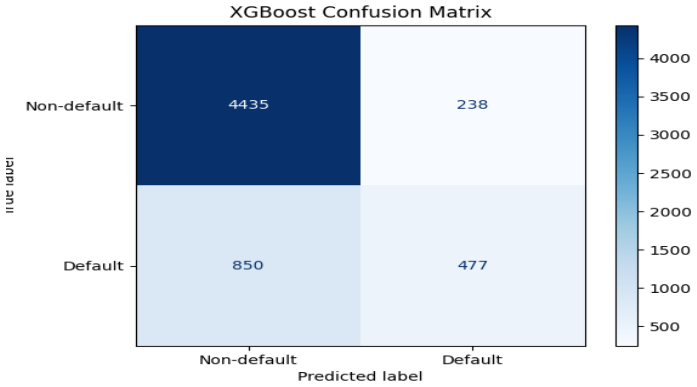


Figure 3. XGBoost Confusion Matrix

Overall, the two machine-learning models showed similar classification patterns. Random Forest achieved a marginally higher recall than XGBoost (36.32% compared with 35.95%), while XGBoost performed better in terms of accuracy, precision, F1-score and ROC-AUC. Importantly, however, neither model identified the majority of actual default cases at the standard classification threshold. This suggests that the machine-learning models improved default detection relative to Logistic Regression, but the classification problem remained challenging.

4.2.3 Model Performance and Comparative Analysis

Table 1 summarizes the performance of the three models on the test dataset. Although the differences in overall accuracy were relatively small, larger differences appeared in the other evaluation measures, particularly recall, F1-score and ROC-AUC.

Model	Accuracy	Precision	Recall	F1-score	ROC-AUC
Logistic Regression	80.90%	**69.30%**	24.49%	36.19%	0.7100
Random Forest	81.33%	63.67%	**36.32%**	46.26%	0.7584
XGBoost	**81.87%**	66.71%	35.95%	**46.72%**	**0.7785**

Table 1. Performance comparison of Logistic Regression, Random Forest and XGBoost

In terms of accuracy, all three models performed relatively similarly. Logistic Regression correctly classified 80.90% of the test observations, compared with 81.33% for Random Forest and 81.87% for XGBoost. If accuracy were considered alone, the improvement obtained from the machine-learning models would therefore appear quite limited. However, given the unequal distribution between default and non-default cases in the dataset, the additional evaluation measures provide a more informative comparison.

The clearest difference concerns the identification of actual defaulters. Logistic Regression achieved a recall of only 24.49%, whereas Random Forest and XGBoost reached 36.32% and 35.95%, respectively. This can also be seen from the confusion matrices: Logistic Regression correctly identified 325 of the 1,327 actual defaulters in the test sample, compared with 482 for Random Forest and 477 for XGBoost. The machine-learning models therefore detected considerably more default cases than the traditional benchmark.

This improvement, however, involved a trade-off. Logistic Regression obtained the highest precision, at 69.30%, compared with 63.67% for Random Forest and 66.71% for XGBoost. In other words, when Logistic Regression classified a client as a defaulter, that prediction was more frequently correct. Its higher precision was accompanied by much lower recall, showing that the model made fewer positive default predictions but missed a larger proportion of clients who defaulted.

Considering the balance between precision and recall, the machine-learning models performed more strongly. Their F1-scores were 46.26% for Random Forest and 46.72% for XGBoost, compared with 36.19% for Logistic Regression. A similar pattern appears in ROC-AUC, which increased from 0.7100 for Logistic Regression to 0.7584 for Random Forest and 0.7785 for XGBoost. Among the models examined, XGBoost therefore provided the strongest overall discrimination between default and non-default clients, while Random Forest achieved the highest recall.

These findings suggest that the AI-based models provided an improvement over the traditional model in this particular credit-default prediction task, especially in their ability to identify default risk and distinguish between the two classes. Nevertheless, the results do not indicate that the traditional approach was inferior in every respect. Logistic Regression maintained the highest precision and produced an overall accuracy close to those of the two machine-learning models. The empirical comparison therefore points to a trade-off rather than an absolute superiority of one analytical approach: more flexible machine-learning models improved several predictive measures, while the traditional model remained competitive on others.

This interpretation is also consistent with the broader credit-scoring literature used earlier in the thesis. Gunnarsson et al. (2021) found XGBoost to be a strong-performing classifier in their empirical comparisons, but their results also show that increasing model complexity does not automatically lead to superior performance across all models and datasets. This reinforces the need to evaluate financial prediction models according to the specific objective and performance measure rather than assuming that a more complex method will always be preferable.

4.3 Broader Comparison of Traditional and AI-Based Analytics

The empirical analysis of credit default prediction illustrates some of the broader differences between traditional and AI-based analytical approaches in finance. Although all three models produced relatively similar levels of overall accuracy, their performance differed more clearly when other evaluation measures were considered. In particular, XGBoost achieved the highest ROC-AUC and F1-score, while Logistic Regression produced the highest precision. Random Forest also improved recall considerably compared with Logistic Regression. These findings suggest that the distinction between traditional and AI-based analytics cannot be reduced to a simple question of which approach is more accurate. Rather, the choice depends on the purpose of the analysis and on the balance between predictive performance, flexibility and interpretability.

One important distinction concerns the type and complexity of relationships that the models are designed to capture. Traditional econometric methods are generally developed around specified relationships between variables and remain particularly valuable when the objective is to understand and interpret these relationships. Machine-learning methods place greater emphasis on prediction and can exploit nonlinearities, interactions and large numbers of variables without requiring the researcher to specify all of these relationships in advance. Hoang and Wiegratz (2023) identify this as an important advantage of machine learning in financial applications, particularly when the available data are high-dimensional. They also point out that machine-learning techniques can extract information from unconventional sources such as text and images, extending the range of information that can be incorporated into financial analysis.

This additional flexibility, however, does not imply that AI-based methods should automatically replace traditional models. The results of Gunnarsson et al. (2021) provide a useful example from credit scoring. Across the classifiers examined in their study, XGBoost achieved the strongest overall predictive performance, while deeper neural networks did not consistently outperform their simpler counterparts and required considerably greater computational effort. The authors also emphasize that ensemble methods such as XGBoost are less straightforward to interpret than conventional approaches such as Logistic Regression. Consequently, the appropriate model depends partly on whether predictive performance or interpretability is given greater priority.

Interpretability is particularly relevant in financial decision-making because model outputs may ultimately influence decisions involving credit, investment or risk. More complex models can identify patterns that simpler specifications may overlook, but understanding how a prediction was reached can become more difficult. Similar concerns appear in portfolio management, where Sutiene et al. (2024) note that AI techniques can process large amounts of information and address complex portfolio problems, while their black-box nature creates additional difficulties in explaining decisions to investors, managers and other stakeholders.

Overall, the comparison suggests that traditional and AI-based analytics offer different strengths rather than representing two completely competing approaches. Traditional methods remain useful when transparency, simplicity and interpretation are central to the decision, whereas machine-learning methods become particularly attractive when prediction involves complex relationships or high-dimensional information. This interpretation is also consistent with the empirical results obtained in this study: the AI-based models provided better identification of default cases and stronger overall discrimination, but Logistic Regression remained competitive and offered a simpler analytical framework. The evidence therefore does not support the view that greater model complexity is always preferable. Instead, model selection in financial decision-making should reflect the specific objective, characteristics of the available data and the relative importance placed on prediction and interpretability.

4.4 Regulatory, Governance and Ethical Considerations

The increasing use of AI in financial decision-making also creates challenges related to transparency, fairness and accountability. Complex machine-learning models may produce useful predictions while providing limited insight into how a particular decision was reached. This is especially relevant in finance, where automated outputs can influence decisions affecting customers and investors. Sutiene et al. (2024) emphasize that the use of AI in financial applications is accompanied by requirements for transparency, fairness and explainability, making the ability to understand and justify model outputs an important consideration alongside predictive performance.

These concerns make governance an important part of implementing AI-based financial analytics. Financial institutions need to consider not only whether a model performs well, but also whether its decisions can be appropriately monitored and explained. Explainable AI can help address some of these difficulties by providing greater insight into the predictions generated by complex models. Nevertheless, the growing use of AI does not eliminate the need for human oversight, particularly where automated predictions contribute to decisions with financial consequences.

4.5 Hybrid Approaches

The comparison between traditional and AI-based analytics also suggests that the two approaches do not necessarily need to be treated as substitutes. Machine learning can complement established statistical methods by improving the predictive components of an analysis or by identifying complex patterns that may be difficult to capture with simpler specifications. Hoang and Wiegratz (2023), for example, identify the extension of the existing econometric toolkit as one of the important applications of machine learning in finance. In this sense, AI can strengthen existing analytical processes rather than completely replace traditional methods.

A similar argument appears in credit scoring. Gunnarsson et al. (2021) note that even when interpretability makes conventional methods attractive, high-performing black-box models can still be useful for identifying nonlinear relationships and interactions in the data. These insights can subsequently support the development of more interpretable models. A hybrid approach can therefore combine the predictive flexibility of machine learning with the transparency and interpretability associated with traditional analysis. This is also consistent with the findings of the present study: XGBoost and Random Forest improved several predictive measures, while Logistic Regression remained competitive and easier to interpret. Rather than identifying one approach as universally superior, the comparison supports selecting or combining analytical methods according to the requirements of the financial decision.

Conclusions

This thesis examined traditional and AI-based data analytics in financial decision-making, with the aim of understanding their respective strengths, limitations and practical relevance. The theoretical analysis showed that traditional statistical methods continue to provide structured and relatively interpretable approaches to financial analysis, particularly when understanding and explaining relationships between financial variables is important. At the same time, the growing volume and complexity of financial data have increased the relevance of AI-based methods, which offer greater flexibility in identifying nonlinear relationships, interactions and patterns within data.

The empirical analysis provided a direct comparison of these approaches through credit-card default prediction. Logistic Regression was used as the traditional benchmark and compared with two machine-learning models, Random Forest and XGBoost. Although differences in overall accuracy were relatively small, the other performance measures revealed more noticeable differences. Logistic Regression achieved an accuracy of 80.90% and the highest precision, at 69.30%, but its recall was only 24.49%. Random Forest increased recall to 36.32%, while XGBoost obtained the highest overall accuracy (81.87%), F1-score (46.72%) and ROC-AUC (0.7785). These results show that evaluating the models only through accuracy would have hidden important differences in their ability to identify clients who actually defaulted.

Overall, the findings suggest that the relationship between traditional and AI-based analytics should not be understood simply in terms of one replacing the other. Their usefulness depends on the financial problem being addressed and on the relative importance placed on predictive performance, interpretability or a balance between the two. AI-based models can provide greater predictive flexibility, while traditional methods remain valuable when transparency and interpretation are particularly important. This also supports the use of hybrid approaches that combine the strengths of both methods. As AI becomes increasingly integrated into financial decision-making, explainability, governance and human oversight will also remain important considerations.

The empirical analysis is limited to one credit-default dataset and three classification models, and the results should therefore not be generalized to every financial application. Future research could extend the comparison to different datasets, financial problems and analytical techniques. Overall, the evidence considered in this thesis suggests that the development of AI does not make traditional data analytics obsolete. Instead, AI expands the analytical tools available for financial decision-making. The choice of method should ultimately reflect the characteristics of the financial problem, the available data and the purpose of the analysis rather than the assumption that greater model complexity necessarily leads to better financial decisions.

REFERENCES

Akerkar R. (2019), *Artificial Intelligence for Business*, Springer International Publishing, Cham.

Andersen T.G., Bollerslev T., Diebold F.X., Labys P. (2003), "Modeling and Forecasting Realized Volatility", *Econometrica*, vol. 71, no. 2, pp. 579–625.

Anderson T.W. (1994), *The Statistical Analysis of Time Series*, Wiley-Interscience, New York.

Athey S., Imbens G.W. (2019), "Machine Learning Methods That Economists Should Know About", *Annual Review of Economics*, vol. 11, pp. 685–725.

Balcaen S., Ooghe H. (2006), "35 Years of Studies on Business Failure: An Overview of the Classic Statistical Methodologies and Their Related Problems", *The British Accounting Review*, vol. 38, no. 1, pp. 63–93.

Bartram S.M., Branke J., De Rossi G., Motahari M. (2021), "Machine Learning for Active Portfolio Management", *The Journal of Financial Data Science*, vol. 3, no. 3, pp. 9–30.

Bierman H., Smidt S. (2007), *The Capital Budgeting Decision: Economic Analysis of Investment Projects*, 9th ed., Routledge.

Brigham E.F., Houston J.F. (2019), *Fundamentals of Financial Management*, Cengage Learning.

Buşe L., Ganea M., Cîrciumaru D. (2010), "Using Linear Regression in the Analysis of Financial-Economic Performances", *Annals of University of Craiova – Economic Sciences Series*, vol. 2, no. 38, pp. 32–43.

Cao L., Yang Q., Yu P.S. (2021), "Data Science and AI in FinTech: An Overview", *International Journal of Data Science and Analytics*, vol. 12, no. 2, pp. 81–99.

Fabozzi F.J., Markowitz H.M., Gupta F. (2008), "Portfolio Selection", in Fabozzi F.J. (ed.), *Handbook of Finance*, John Wiley & Sons.

Gilchrist S., Himmelberg C.P. (1995), "Evidence on the Role of Cash Flow for Investment", *Journal of Monetary Economics*, vol. 36, no. 3, pp. 541–572.

Gunnarsson B.R., vanden Broucke S., Baesens B., Óskarsdóttir M., Lemahieu W. (2021), "Deep Learning for Credit Scoring: Do or Don't?", *European Journal of Operational Research*, vol. 295, no. 1, pp. 292–305.

Hoang D., Wiegratz K. (2023), "Machine Learning Methods in Finance: Recent Applications and Prospects", *European Financial Management*, vol. 29, no. 5, pp. 1657–1701.

Markowitz H.M. (1952), "Portfolio Selection", *The Journal of Finance*, vol. 7, no. 1, pp. 77–91.

Masini R.P., Medeiros M.C., Mendes E.F. (2021), “Machine Learning Advances for Time Series Forecasting”, *Journal of Economic Surveys*, vol. 37, no. 1, pp. 76–111.

Mishra P., Pandey C.M., Singh U., Gupta A., Sahu C., Keshri A. (2019), “Descriptive Statistics and Normality Tests for Statistical Data”, *Annals of Cardiac Anaesthesia*, vol. 22, no. 1, pp. 67–72.

Sabarieswar V., Rajaseharan V. (2026), “Hypothesis Testing in Business Analytics: A Statistical Framework for Data-Driven Decision Making”, *International Journal of Engineering Research & Technology*, vol. 15, no. 5.

Siami-Namini S., Tavakoli N., Siami Namin A. (2018), “A Comparison of ARIMA and LSTM in Forecasting Time Series”, in *2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pp. 1394–1401.

Sutiene K., Schwendner P., Sipos C., Lorenzo L., Mirchev M., Lameski P., Kabasinskas A., Tidjani C., Ozturkkal B., Cerneviciene J. (2024), “Enhancing Portfolio Management Using Artificial Intelligence: Literature Review”, *Frontiers in Artificial Intelligence*, vol. 7, article 1371502.

Thomas L.C. (2000), “A Survey of Credit and Behavioural Scoring: Forecasting Financial Risk of Lending to Consumers”, *International Journal of Forecasting*, vol. 16, no. 2, pp. 149–172.

Yang H., Shukur Z., Sahran S. (2026), “A Review of Artificial Intelligence for Financial Fraud Detection”, *Applied Sciences*, vol. 16, no. 4, article 1931.

Yeh I.-C., Lien C.-h. (2009), “The Comparisons of Data Mining Techniques for the Predictive Accuracy of Probability of Default of Credit Card Clients”, *Expert Systems with Applications*, vol. 36, no. 2, pp. 2473–2480.

Zong Z., Guan Y. (2025), “AI-Driven Intelligent Data Analytics and Predictive Analysis in Industry 4.0: Transforming Knowledge, Innovation, and Efficiency”, *Journal of the Knowledge Economy*, vol. 16, no. 1, pp. 864–903.

DATASET / WEBSITE

UCI Machine Learning Repository, *Default of Credit Card Clients*, dataset contributed by I-Cheng Yeh, DOI: 10.24432/C55S3H.