

ADDIS ABABA UNIVERSITY

COLLEGE OF BUSINESS AND ECONOMICS

SCHOOL OF COMMERCE

Predictive Analytics in ATM Cash Demand Forecasting Using Machine Learning Approach: A Case Study of the Commercial Bank of Ethiopia

A Thesis Submitted to the School of Commerce Addis Ababa University in Partial Fulfillment of the Requirements for the Degree of Masters of Business Administration and Information Systems

By: Merdiya Taha Musa

Advisor: Minale Ashagrie (PhD)

June 2026

ADDIS ABABA UNIVERSITY

COLLEGE OF BUSINESS AND ECONOMICS

SCHOOL OF COMMERCE

Predictive Analytics in ATM Cash Demand Forecasting Using Machine Learning Approach: A Case Study of the Commercial Bank of Ethiopia

By: Merdiya Taha Musa

Approved by Board of Examiners

Signature: _____ Date _____

Minale Ashagrie (PhD), Advisor

Signature: _____ Date: _____

Kuribachew Gizaw (PhD), Examiner

Signature: _____ Date: _____

Yosef (PhD), Examiner

Declaration

I hear by declare that *Predictive Analytics in ATM Cash Demand Forecasting Using Machine Learning Approach: A Case Study of The Commercial Bank of Ethiopia* project is my original work, and the study has not been submitted for any degree program in this or other institution**.**

Table Contents

List of Tables VII
List of Figures VII
Acknowledgement VIII
List of Abbreviations IX
Abstract XI
CHAPTER ONE 1
INTRODUCTION 1
1.1 Background of the Study 1
1.2 Background of the Company 2
1.3 Statement of the Problem 3
1.4 Research Questions 4
1.5 Objective of the Study 4
1.5.1 General Objective 4
1.5.2 Specific Objective 4
1.6 Significance of the Study 4
1.7 Scope (Delimitation) of the Study 5
1.8 Limitation of the Study 5
1.9 Definition of Key Terms 5
1.10 Organization of the Report 6
CHAPTER TWO 7
REVIEW OF THE RELATED LITERATURE 7
2.1 Theoretical Review 7
2.1.1 Forecasting 7
2.1.2 Demand Forecasting 8
2.1.3 Predictive Analytics 8
2.1.4 Machine Learning 9
2.1.5 Time Series Analysis 9
2.1.6 Linear Regression 10

2.1.7 Random Forest 10
2.1.8 Extreme Gradient Boosting XGBoost 11
2.1.9 Model Evaluation and Generalization 11
2.1.10 ATM Cash Management 12
2.1.11 Decision Support System 12
2.1.12 Big Data Analytics in Banking 13
2.1.13 Ethiopian Banking and Digital Transformation Context 13
2.1.14 Data Governance and Organizational Readiness for Analytics in Banking 13
2.1.15 Business IT Alignment 14
2.2 Empirical Review 14
RESEARCH METHODOLOGY 19
3.1 Research Approach 19
3.2 Research Design 19
3.3 Population 20
3.4 Sampling Size 20
3.5 Sampling Technique 20
3.6 Data Source and Type 21
3.7 Data collection procedure 21
3.8 Data Analysis 21
3.9 Evaluation Metrics 22
3.10 Tools and Software 23
3.11 Validity 24
3.12 Reliability 24
3.13 Ethical Consideration 25
CHAPTER FOUR 26
DATA ANALYSIS AND RESULT DISCUSSION 26
4.1 Dataset Description 26
4.2 Data cleaning 27
4.2.1 Data Type Correction and Standardization 27

4.2.2 Impossible Values 27
4.2.3 Outlier Detection and Treatment 27
4.3 Exploratory Data Analysis 28
4.3.1 Demand by Day of Week 28
4.3.2 Demand by Month 29
4.3.3 Effect of Key Events 29
4.3.4 Additional Data Visualization 30
4.4 Feature Engineering 33
4.4.1 Calendar Features 34
4.4.2 Salary Cycle Feature 34
4.4.3 Holiday Features 35
4.4.4 Lag features 35
4.4.5 Feature Correlation 35
4.5 Data Preprocessing 36
4.5.1 Removing Incomplete Records 36
4.5.2 Train and Test Split 37
4.5.3 Feature Scaling 37
4.6 Model Development and Hyperparameter Selection 37
4.7 Model Result and comparison 38
4.7.1 Linear Regression Result 38
4.7.2 Random Forest Result 38
4.7.3 XGBoost Result 40
4.7.4 Model Comparisons 41
CHAPTER FIVE 43
SUMMARY, CONCLUSION AND RECOMMENDATION 43
5.1 Summary of Findings 43
5.2 Conclusion 44
5.3 Recommendation 44
5.4 Limitation and Future area of study 45

References 46
Appendices 49

List of Tables

Table 1 Table 4.1 Statistical summary 17
Table 2Table 4.2: New Features Description 24
Table 3Table 4.3 Model Performance Comparison on Test Set 32

List of Figures

Figure 12.4 Conceptual Framework 12
Figure 2Figure 4.1: Average ATM Cash Demand by Day of Week 19
Figure 3Figure 4.2: Average ATM Cash Demand by Month 20
Figure 4Figure 4.4 Average Daily ATM Cash Demand Over Time 22
Figure 5Figure 4.5 Time-series Decomposition of Daly Demand 23
Figure 6Figure 4.6 Distribution of ATM Cash Demand by Day of Week 23
Figure 7Figure 4.7 Average ATM Cash Demand by Day of the Month 24
Figure 8Figure 4.8 Correlation Heatmap of Features 27
Figure 9Figure 4.9 Actual vs Predicted Demand for Linear Regression 29
Figure 10Figure 4.10 Actual vs Predicted Demand for Random Forest 30
Figure 11Figure 4.11 Random Forest feature Importance 31
Figure 12Figure 4.12 Actual vs Predicted Demand for XGBoost 32
Figure 13Figure 4.13 Model performance Comparison of MAE and R2 for The Three models 33

Acknowledgement

First, I would like to express my deepest gratitude to almighty Allah for granting me the power, courage and wisdom to finish this research paper.

I would like to express my deepest gratitude to my advisor Dr. Minale Ashagrie for supervising me and providing me with valuable comments and constructive suggestions without which this work would not have been brought to light.

I would also like to express my deepest gratitude and indebtedness to my father, my mother, my sisters and my brother Mohammedamin Taha who have always encouraged me for higher success.

Finally, I would like to forward my special thanks the Commercial Bank of Ethiopia employees. Many thanks go particularly to Research and Development, Card Banking and MIS team.

List of Abbreviations

ANN	Artificial Neural Network
ARIMA	Autoregressive Integrated Moving Average
ATM	Automated Teller Machine
CBE	Commercial Bank of Ethiopia
CDNN	Clustered Deep Neural Network
CNN	Convolutional Neural Network
EDA	Exploratory Data Analysis
ETB	Ethiopian Birr
DSS	Decision Support System
DQN	Double Deep Q-Network
GMDH	Group Method of Data Handling
GRNN	General Regression Neural Network
IT	Information Technology
KNN	K-Nearest Neighbors
L1	Lasso (L1) Regularization
L2.	Ridge (l2) Regularization
LSTM	Long Short-Term Memory (Neural Network)
MAE	Mean Absolute Erro
MAPE	Mean Absolute Percentage Error
ML	Machine Learning
MLP	Multilayer Perceptron
OLS	Ordinary Least Squares
RNN	Recurrent Neural Network
RF	Random Forest
RMSE	Root Mean Square Error
ROA	Return on Asset
ROE	Return on Equity
R^2	Coefficient of Determination
SARIMA	Seasonal Autoregressive Integrated Moving Average
SNN	Shallow neural network

SMAPE Symmetric Mean Absolute Percentage Error

SVM Support Vector Machine

SVR Support Vector Regression

XGBoost Extreme Gradient Boosting

Abstract

Reliable cash availability at Automated Teller Machine (ATM) is important for customer service and for controlling the cost of holding and moving cash. The purpose of this study was to examine if predictive analytics and machine learning models can forecast daily ATM cash demand at CBE and support better replenishment planning. We used quantitative predictive approach to achieve the objective of the study. Secondary data was collected from the bank. Daily ATM transaction records from 15 ATMs covering two years period were used. Based on the raw data from the bank, data were cleaned preprocessed and new features were created. Three machine learning models were trained and compared namely, Linear Regression as baseline, Random Forest and Extreme Gradient Boosting (XGBoost). Model performance was measured using MAE, RMSE, R2 and MAPE with using MAE and R2 as primary metrics. XGBoost was the best model with MAE of 268,821 and R2 of o.6682 and Random Forest was a close second. The analysis showed that recent demand history is the strongest driver of future demand, while salary days, pre-holiday days holiday events also shaped withdrawal. The study concludes that CBE already holds that data needed to move from experience based to data driven cash management and recommends adopting the model as a decision support tool supported by stronger data governance and staff capacity.

***Key words**: Predictive analytics, ATM cash demand forecasting*

CHAPTER ONE

INTRODUCTION

1.1 Background of the Study

Recently the accelerated advancement of data analytic and artificial intelligence technologies has transformed decision making processes in different industries like banking, logistics, manufacturing and retail. From these technologies, predictive analytics has emerged as essential tool for forecasting future events by analyzing historical data using data mining, machine learning and statistical techniques (Punia & Shankar,2022; Wolinak,2024). Predictive analytics has been broadly applied in demand forecasting, risk management, inventory optimization and operational planning supporting organizations to transition from reactive to proactive decision-making (Seyedan,2024).

Several research in manufacturing, e-commerce, and logistics sectors emphasizes the effectiveness of predictive analytics in demand forecasting and resource optimization (Owen, 2023; Bin Yusof *et al*., 2024; A. Agatić, 2022). These studies suggested that advanced predictive analytics models often achieve high accuracy effective implementation depends not only on algorithms but also on organizational readiness, data quality and organizational support.

In banking industry, predictive analytics has been implemented to improve customer satisfaction, operational efficiency, and strategic planning. Applications include in the areas of fraud detection, credit risk assessment, customer segmentation (Adeniran *et al*.,2024; Wolinak,2024). One functional area is Automated Teller Machine (ATM) cash demand forecasting which has crucial role in securing continuous cash availability while minimizing excess cash holdings and replenishment costs. Effective ATM cash management have direct relation with customer experience, operational cost control and service reliability (Psomakelis *et al*. (2025).

Various international studies illustrate that predictive analytics significantly improves ATM cash demand forecasting accuracy compared to traditional manual approaches. For example, a study by Sarveswararao & Ravi (2020) applied chaos theory and deep learning models to ATM withdrawal data and reported actual improvement in forecasting performance. In the same way Psomakelis *et*

al. (2025) presented that machine learning-based forecasting models enables banks to optimize cash replenishment schedules and reduce ATM cash-out incidents.

Regardless of the demonstrated benefits of predictive analytics, its adoption rate remains disparate specially in developing economies. Studies on analytics claims that technological availability only is insufficient, but organization must have appropriate data governance structure, skilled human resource, supportive leadership and conducive regulatory environment (Wolinak,2024; Adeniran *et al*., 2024).

In Ethiopia, the banking sector is undergoing gradual digital transformation, with increased deployment of core banking systems, electronic payment platforms, and ATM networks. The Commercial Bank of Ethiopia (CBE), as the largest public bank in the country, generates vast volumes of transactional data through its extensive ATM infrastructure. However, existing evidence suggests that ATM cash replenishment decisions are still largely manual and experience-based, with limited use of predictive analytics for demand forecasting. While digital data exist, their use for advanced analytics remains minimal, indicating a possible gap between data availability and analytics adoption.

Based on preliminary observation and informal discussions with bank staff ATM cash replenishment practices are guided by traditional decision-making approach. This study will propose a predictive model for ATM cash demand forecasting using Machine Learning (ML) approach based on data collected from CBE.

1.2 Background of the Company

CBE is the largest and most influential commercial bank in Ethiopia. Officially established in 1963, with historical roots dating back to 1942, CBE dominates the banking sector in terms of asset size, customer base, and branch coverage. As of 30 June 2024, it manages over 1.4 trillion Ethiopian Birr in assets, serves more than 45 million customers, and operates network of 1,942 branches nationwide. The bank also has a presence beyond Ethiopia, including a branch in South Sudan and a representative office in Dubai (CBE, 2024).

CBE has embraced various forms of digital banking, including mobile banking, internet banking, agent banking, and ATMs. With over 3,000 ATMs deployed across the country, CBE has significantly improved access to financial services, especially in urban and semi-urban areas

(CBE,2024). However, cash management at these ATMs is still largely managed manually, without the application of predictive analytics tools for forecasting or optimizing cash replenishment.

This operational gap presents a valuable opportunity to investigate how predictive analytics could be introduced to enhance ATM demand forecasting at CBE. Improving demand forecasting could help reduce cash shortages, improve service availability, and support the bank's broader digital transformation goals.

1.3 Statement of the Problem

Predictive analytics has emerged as transformative tool in different sectors for improving operational efficiency, optimizing resource allocation and enabling proactive decision-making Adekunle et al.,2021; Adeniran et al.,2024. One functional area is ATM cash demand forecasting which have crucial role in securing continuous cash availability while minimizing excess cash holdings and replenishment costs. Effective ATM cash management have direct relation with customer experience, operational cost control and service reliability (Psomakelis et al. (2025). This benefit is particularly important in cash intensive economies where ATM availability is closely related to financial accessibility and customer satisfaction.

Preliminary observation at the CBE of Ethiopia, which have the largest ATM network in the country, indicates that ATM cash replenishment decisions are manual and experience based. Although the bank generates a high volume of transaction data, these data are not systematically analyzed or used for predictive forecasting. This practice may lead to inefficient cash management and increased operational costs, challenges that are also noted in studies of banks operating in similar emerging-market contexts (Seyedan, 2024).

Previous studies demonstrate that predictive analytics can significantly improve ATM cash demand forecasting by enhancing operational efficiency and reducing cash shortage and excess holdings (Wolinak,2024; Adeniran *et al*.,2024). However, in the Ethiopian banking context there is a lack of empirical evidence on machine learning models for ATM cash demand forecasting.

The purpose of this study is to propose predictive model for ATM cash demand forecasting using Machine learning approach that support data-driven decision-making, improves ATM cash replenishment planning and contributes to more efficient cash management practices and also improves service availability for customers. As far as the researcher's knowledge, there is no prior

study that applies machine learning model to ATM cash demand forecasting in Ethiopia. Therefore, this study tries to fill the knowledge gap.

1.4 Research Questions

This research work tries to address the following research questions.

1. What data features most significantly influence ATM cash demand forecasting at CBE?

2. How can machine learning models be applied to ATM transaction data to forecast cash demand?

3. Which machine learning model provides the most accurate and reliable prediction based on evaluation metrics?

1.5 Objective of the Study

1.5.1 General Objective

To propose a predictive machine learning model for ATM cash demand forecasting at CBE, demonstrating how technology can enhance business strategy and organizational efficiency.

1.5.2 Specific Objective

- To review relevant literature.
- To collect and preprocess ATM transaction data for use in predictive modeling.
- To train and evaluate machine learning model for ATM cash demand forecasting.
- To compare model performance using accuracy metrics (RMSE, MAE, MAPE and R2)
- To recommend the most effective predictive model as strategic solution for ATM cash management at CBE.

1.6 Significance of the Study

This research work is significant in many ways. Academically, the findings in this research will contribute to the existing body of knowledge as referral material. Future researchers will find this work beneficial because it is bound to contribute information and enlighten them on the importance of predictive analytics on demand forecasting. By proposing predictive modeling to a large, state-owned bank to implement predictive analytics, the study offers practical implications for banking professionals, policymakers, and IT leaders working to modernize legacy systems

1.7 Scope (Delimitation) of the Study

Geographically, this study is delimited to CBE, focusing on selected ATMs where transaction data is available. Conceptually, it examines transactional temporal and contextual data features that influence ATM cash demand, while integrating business, technological, and organizational perspectives to emphasize the strategic role of predictive analytics. Methodologically, the study employs quantitative approach using historical ATM transaction dataset, applying machine learning models and evaluating them with statistical accuracy metrics.

1.8 Limitation of the Study

First, the study confined to the commercial Bank of Ethiopia, therefore the findings may not be generalized to other banks. Second, the study relies partly on informally obtained information about existing ATM replenishment practices, which may limit the objectivity of the findings. Another limitation was lack of adequate research material conducted in this area within the country.

1.9 Definition of Key Terms

Predictive Analytics Predictive analytics refers to the use of statistical techniques, data mining, and machine learning algorithms to analyze historical data to make forecasts about future events (Owen, 2024).

Automated Teller Machine (ATM) An electronic banking device that enables customers to access financial services such as cash withdrawal, deposits, and balance inquiries without the need for a human teller (Gökçay, 2021).

Machine Learning (ML) A branch of artificial intelligence that focuses on building algorithms that can learn patterns from data and improve their performance at making predictions or decisions without being explicitly programmed (Gökçay,2021).

Evaluation Metrics Quantitative measures used to assess the accuracy of predictive models by comparing predicted values against actual outcomes. They provide a way to judge how well a model performs and are essential for validating forecasting approaches (Gökçay,2021).

1.10 Organization of the Report

This paper is organized into five chapters. The first chapter is introduction which consists background of the study, statement of the problem, research question, objectives, significance of the study, scope and limitation of the study. The second chapter reviews the related literatures which consists theoretical background, review of previous empirical studies and conceptual framework. The third chapter presents the research methodology applied in this study. The fourth chapter presents result and discussion of the study. The fifth chapter presents summary of findings, conclusion, recommendation, limitation and future areas of study.

CHAPTER TWO

REVIEW OF THE RELATED LITERATURE

Introduction

This Chapter presents review of literature that form the basis of this study. The Chapter consist theoretical definitions and empirical reviews of past studies that have been done.

2.1 Theoretical Review

2.1.1 Forecasting

Forecasting is the process of predicting future values based on historical and current data. It is one of the oldest analytics activities in business and management and supports decision where future is uncertain (Hyndman & Athanasopoulos, 2018). Forecasting methods are grouped as two broad families, which are qualitative and quantitative. Qualitative method relies on experiment judgement, market research, and structured opinion and they are used when historical data is limited. Quantitative methods rely on numerical historical data and statistical or machine learning techniques and they are used when adequate data is available to detect patterns. Forecasting can be also classified by time horizon. Short term forecast covers days or weeks and used for operational decisions like inventory and cash management, medium term forecasts cover months and supports planning decisions such as budgeting and long-term forecasts covers years and supports strategic decisions such as expansion.

Forecasting theory has developed from simple statistical models towards machine learning approaches that can capture nonlinear patterns. Classical approaches such as the naïve method, moving averages, and exponential smoothing have been widely used to capture seasonality and trends. While these methods are simple and transparent, they often fail to address the complexity of modern financial systems where demand is influenced by multiple nonlinear factors (Hyndman & Athanasopoulos, 2018). Modern approaches based on machine learning handle this complex pattern better however they need more computational resources.

Recent research has extended forecasting theory into banking operations, particularly ATM cash demand. Styrkársson (2020) highlighted that accurate forecasting is essential to minimize idle cash costs and ensure customer satisfaction. The researcher compared classical statistical methods with

machine learning approaches, demonstrating that neural networks provided superior accuracy in predicting daily withdrawals. Similarly, Fallahtafti *et al*. (2022) emphasized that forecasting models must adapt to external shocks such as the COVID-19 pandemic. Their findings showed that while machine learning models excelled in stable environments, statistical models like ARIMA and SARIMA were more reliable during volatile periods.

Together, these studies illustrate the evolution of forecasting theory from classical statistical approaches to modern adaptive models. Forecasting theory provides the foundation for ATM cash demand prediction, explaining why historical withdrawal patterns, salary cycles, and holiday effects can be leveraged to anticipate future cash requirements.

2.1.2 Demand Forecasting

Demand forecasting is the estimation of future demand for product or service based on historical trends, market conditions and external variables (Adekunle & Chukuma-Eke, 2021). It is a core activity in operation management because it directly affects how organizations plan resources, allocate resource and serve their customers.

Accurate demand forecasting enables organizations to balance supply and demand, reduce costs and improve service reliability. Studies indicate that traditional forecasting approaches based on fixed schedules or simple averages are often insufficient due to demand variability caused by seasonality, customer behavior, and special events Psomakelis *et al*. (2025).

2.1.3 Predictive Analytics

Predictive analytics is branch of analytics that uses statistical techniques, data mining and machine learning algorithms to analyze historical data to forecast future events. (Owen, 2024). By identifying trends and potential fluctuations in demand, predictive analytics enables businesses to make data-driven decisions that minimize risks and enhance operational efficiency. This transforms demand forecasting from reactive process to proactive strategy which help business to make informed decision based on data driven insight (Adekunle & Chukuma-Eke, 2021).

According to Owen, (2024) the process involves data collection, data processing, model development, validation and implementation. As organizations rely on data driven strategies predictive analytics has become important component of operational and strategic decision making

in various industries including finance, logistic and manufacturing. (Owen, 2023; Bin Yusof1 *et al*., 2024; A. Agatić, 2022).

2.1.4 Machine Learning

Machine learning theory provides the foundation for algorithms that learn patterns from historical data and improve predictive accuracy without being explicitly programmed. Unlike classical statistical models, which rely on fixed assumptions about linearity and stationarity, machine learning methods adapt to complex, nonlinear relationships and can capture hidden structures in large datasets (Mitchell, 1997; Hyndman & Athanasopoulos, 2018).

Mitchell (1997) defines machine learning as the study of computer algorithms that improve automatically through experience. In forecasting applications, this means models can continuously refine their predictions as more data becomes available. For ATM cash demand, machine learning theory justifies the use of algorithms such as Random Forest, XGBoost, and neural networks, which can handle high-dimensional data and nonlinear withdrawal patterns (Sarveswararao & Ravi, 2020; Fallahtafti *et al*., 2022).

Recent studies empasizes the effectiveness of machine learning in financial forecasting. Gökçay (2020) compared traditional ARIMA models with machine learning approaches such as LSTM and neural networks, finding that machine learning methods achieved higher accuracy in predicting ATM withdrawals. Similarly, Sarveswararao and Ravi (2020) demonstrated that ensemble methods like Random Forest and deep learning architectures such as LSTM and CNN outperform classical models when chaos and exogenous variables (e.g., day-of-week effects) are incorporated.

2.1.5 Time Series Analysis

Time series is a sequence of observations recorded at regular interval of time. A time series is described by four components which are the trend, seasonal components, cyclical component and the noise components. The trend captures the long run movement, the seasonal component captures repeating calendar cycles, the cyclical components capture long term fluctuations that are not tied to calendar, and the noisy component captures the random variation that cannot be explained (Box et al., 2015).

A core concept in time series is autocorrelation which is the relationship between an observation and earlier observation in the same series. If today value is closely related to yesterday's value or

to the value seven days ago the series is called to be autocorrelated because demand on any day is shaped by demand patterns from recent past. This is why lag features and rolling mean features useful as inputs to a forecasting model (Hyndman & Athanasopoulos, 2018).

The other property is stationarity which means that statistical properties of the series such as its mean and variance do not change over time. Classical models assume stationarity while machine learning models can often work with non-stationary data when the features are designed to catch the changing trend (Box et al., 2015).

2.1.6 Linear Regression

Linear regression is one of the oldest and simplest machine learning method. The method was structured by Gauss and Legendre in the early 19th century and continues to be used widely because of its simplicity and interpretability (Gokçay, 2020). It works by finding straight line relationship between the input feature and target variable and estimates a weight for each feature (James et al., 2013).

The model works by giving a coefficient to each input features so that the predicted value is the sum of the features multiplied by their coefficients plus an intercept term. The coefficients are approximated by minimizing the squared difference between predicted and actual values on the training data. Once trained the model create a single equation which connects the inputs to the predicted output. This model is fast and easy to explain but it can not capture complex pattern or when the relationship with variables is not linear.

2.1.7 Random Forest

The Random Forest model proposed by L. Breiman in 2001 belongs to the decision tree family. It works by building many decision trees in different random samples of the data then averaging their predictions to generate the final output. This averaging process is known as bootstrap aggregation or bagging and it reduces the variance of the individual tree, so the combined model is more stable and accurate than any single tree.

Random Forest use two layers of randomness to keep the trees different from each other. First each tree is trained on a random subset of the rows of the training data which is drawn with substitute. Second at each split inside a tree the algorithm considers only a random subset of a feature rather than all the features. This two-design choice stops the trees from becoming too similar and prevent

the model from depending heavily on any single feature. So the final prediction is the average of the predictions made by all trees in the forest.

The strength of Random Forest is its ability to catch nonlinear patterns, its resistance to overfitting when properly configured and the feature importance score it crates that helps to explain model decision.

2.1.8 Extreme Gradient Boosting XGBoost

It is one of tree-based ML model, it was first introduced by Chen and Guestrin in 2016. This method builds tree one after another, where each new tree corrects the error made by the previous tree. This step-by-step correction is known as gradient boosting, and it allows the model to gradually reduce its prediction error round by round.

This method has become widely in applied machine learning because it integrates accuracy with computational efficiency. XGBoost includes various built-in features which improves performance and control overfitting. This includes shrinkage, that scales down the contribution of each tree by a learning rate, row and feature subsampling at each tree to add randomness and L1 and L2 regularization terms that penalize complex trees. These features make the algorithm more reliable on noisy datasets than basic gradient boosting model.

2.1.9 Model Evaluation and Generalization

Building a machine learning model is the half work after building judging how well the model will perform on new unseen data, which is the question that model evaluation answer (Hastie et al., 2009). The most common evaluation is design train-test split in which part of the data is used to fit the model and another part is held back to measure how well it predicts. For time series data the split must be in time order with earlier data used for training and later for test to avoid leakage of future information in to training process. A strong version of this approach is time series cross validation that the model is trained and tested across many moving windows to give a more stable estimate of performance (Hyndman & Athanasopoulos, 2018).

The central concept in evaluation is the bias-variance trade-off. Which means a model with high bias is too simple and cannot capture the real patter in the data so it under fits and a model with high variance is too complex and follows the noise in the training data so it over fits. The skill in

model development is to find a balance which gives low error on both training set and test set, which is the practical meaning of good generalization (Hastie et al., 2009)

2.1.10 ATM Cash Management

ATM cash management is the duty of keeping a machine stocked with enough cash to serve customers meanwhile avoiding too much idle cash. Poor cash management have two costs. One is the cost of cash out incident due to machine runs out of cash which results in customer dissatisfaction that harms service quality. The other cost is idle cash due to overloading the machine with cash this makes the bank to lose interest it could have earned, also faces higher transport and insurance cost (Simutis et al., 2008). Good cash management balances these two costs.

Forecasting is core to this balance. If a bank can predict daily demand of each ATM, it can plan replenishment precisely and reduce idle cash and cash outs (Castro, 2009). Studies that integrate demand forecasting with replenishment planning shows lower operating cost and fewer service failures (Ekinci et al., 2015; Psomakelis et al., 2025). Beside cost optimization recent studies shows that ATM cash management performance also influenced by organizational and managerial factors. Assaf et al. (2025) found that technological readiness, managerial competence and regulatory and environmental support greatly influence ATM operational performance and this operational performance play's role of mediator that converts these organizational assets into improved customer satisfaction.

2.1.11 Decision Support System

A decision support system (DSS)is a computer-based system that helps mangers to make decision by combining data, analytical models, and user-friendly interface (Sharda et al., 2020). DSS are generally built around three main components. The first is the data component that stores relevant historical and current information. The second is the model component which contains the analytical method that turns data into recommendation and the third is the user interface that allows the manager to interact with the system, ask questions and view result in understandable way (Arnott & Pervan, 2014).

2.1.12 Big Data Analytics in Banking

Big data analytics is the analysis of very large, diverse and fast-moving datasets to extract patterns and support decisions (McAfee & Bryniolfsson, 2012). Banks are within the most data rich organizations in any economy because each transaction, login payments and customer interactions generate record. The challenge banks faces are turning collected data into useful outcomes. Three properties are commonly used to describe big data which are volume refers to the sheer amount of data produced, velocity refers to the speed at which new data arrives and variety referring to the different forms the data takes including text, images and structured records.

2.1.13 Ethiopian Banking and Digital Transformation Context

Ethiopian banking sector is regulated by national Bank of Ethiopia (NBE) that licenses banks, set payments policy and supervises this financial system. CBE is the major player in the sector. According to Geda (2022) CBE accounts for about 60 percent of national deposit and lending in recent years however this dominance has declined from earlier as private commercial banks have grown. This concentration means any operational improvement in CBE has significant effect on the wider banking system.

Digital transformation in Ethiopia quickened over the past decade. The Federal Government of Ethiopia adopted A national Digital Strategy in 2020 and NBE followed with a National Digital Payments Strategy (NDPS) for the period 2021 to 2024 which sets out goals for a secure, competitive and inclusive digital payment system (Geda,2022). Studies shows that Ethiopian banking is progressing toward digital operation but still maturing regulatory framework and shows that the gain from digitalization depends on organizational readiness and supportive regulation not only on technology alone (Abera,2025).

2.1.14 Data Governance and Organizational Readiness for Analytics in Banking

Data governance describes that the policies, processes and standards that guarantee data is managed accurately, consistently and securely throughout its lifecycle. Mishra (2019) argued that data governance becomes a major driver for retail banks because it supports operational efficiency, customer experience and regulatory compliance. Strong data governance enables banks to consolidate customer data, reduce duplication and error and feed reliable input to advanced analytics for risk detection and fraud prevention. This same study identifies that the main barrier

to adoption are organizational silos, aging legacy system and regulatory landscape, without this foundations analytics outputs cannot be trusted and cannot support decision making.

Organizational Readiness explains the extent to which a banks structure, skills, leadership and culture are prepared to understand an act of analytics. Aykanat et al. (2025) findings showed that the barrier to big data analytics in banking affect organizational preparation which shapes the business analytics adoption and strength of a data driven culture. This shows that banks need to build internal readiness including staff skills and management support before analytics tools deliver business value.

2.1.15 Business IT Alignment

Business IT alignment is the extent to which the goals, plans and activities of an organization's information technology are matched to its business goals (Henderson & Venkatraman, 1993). Alignment is strong when technology investment supports the strategy and operations of the business and weak when the organization invests on technology that does not produce business value and useful technologies stay underused.

Henderson and Venkatraman (1993) explained alignment in two dimensions. The first is the strategic fit between business strategy and IT strategy. The second is the functional integration between business operation and IT operations. Strong alignment requires both dimensions to be in place at the same time, that technology that fit the strategy but cannot be operationalized or operational technology that does not support the strategy fails to create value.

2.2 Empirical Review

Al-zeir (2024) conducted research on ATM availability forecasting model using machine learning algorithms. The purpose of this study was to build a model that forecasts the total cash withdrawal amount for the next hour across the Bank of Palestine ATM network. The study used large dataset of 9 million ATM transaction from 172 ATMs collected during 2017 and 2018. The raw data was cleaned; preprocessed and new features were engineered using different approaches. Three machine learning models were trained and tuned, Random Forest, XGBoost and LSTM, and they were evaluated using R2 and MAPE.

The result shows that LSTM gained the best result in two tasks with R2 of 91% and MAPE of 0.37 to predict the next hour withdrawal for 172 ATMs grouped into hourly time series and considered as one ATM. And R2 of 96%-99% and MAPE of 0.28-0.08 to predict the rolling mean transaction amount of the next even hours. The study also found that largest demand peaks happened on salary transfer days. For the limitation, the study didn't group ATMs by location.

Fallahtafti *et al.* (2022) aimed to forecast ATM cash demand for periods before and during the COVID-19 pandemic. To achieve this, first ATMs were categorized based on accessibility and surrounding environmental factors. Then different statistical and machine learning models employed and tested on each group, including parametric models such as ARIMA and SARIMA and non-parametric models like Random Forest, Support vector machine and k-nearest neighbors. To provide the feature matrix for machine learning models some influential variables are added, and new fitness measure were proposed that combined error size and the direction of change.

The result indicates that grouping ATMs by category improved forecasting quality. Before COVID-19 and in times where there is minor disturbance in withdrawal pattern the machine learning models performs better but during COVID-19 and there is a sudden shock in demand and massive volatility in withdrawal pattern the statistical models gave better short-term forecasts. The study also showed that sophisticated models will not always outperform simpler models.

Gökcay (2020) conducted research on ATM cash stock prediction using different machine learning approaches. The purpose of the study was to predict the daily amount of cash withdrawn from ATMs with the lowest possible error, using past withdrawal data. The researcher-built models using Linear regression, support vector machine, artificial neural networks, LSTM, and the statistical model ARIMA using the same data. The models were trained using cross validation method and tested for both weekly and monthly prediction period and compared across eight different feature sets.

The result indicates that machine learning models in comparison with ARIMA have higher accuracy and among the machine learning models, LSTM gives the most accurate prediction and use less features compared to other models.

Psomakelis *et al*. (2025) conducted study on ATM supply scheduling optimization. The purpose of the study was to build a unified pipeline that predicts the optimal time for ATM to be resupplied and then plans the supply route with the lowest cost. This study used data from National Bank of Greece including 863 ATMs at first and keeping 643 after cleaning. Linear regression, deep learning and other machine learning models were used for prediction step and Integer linear programming were used for routing step. The records were enhanced with external features like day of the week, holidays and if it is workday. The result of the study showed that decision tree model gave the best accurate prediction, while deep learning also showed promising result. The limitation in this study is it uses only 24 ATM test set which is small share of the full network.

Styrkársson (2020) conducted research on Forecasting demand for ATM replenishment. The objective of the study was to create forecasting model using machine learning and data mining techniques to predict daily cash demand transaction in a single ATM. The linear regression model was compared with a more complex neural network for predicting daily demand and three classical methods, naïve method, moving average and exponential smoothing are used for predicting monthly demand. The models are then compared using three different error measures which are the mean absolute error, the squared error and the mean absolute percentage error. The result of the study showed with the forecasting models an order policy using the EOQ model is created.

Suder *et al.* (2024) aimed to test the forecasting accuracy of recently implemented econometric tools compared to traditional models when predicting ATM cash demand and to verify if COVID-19 driven market change conditions impacted the predictive power of the tested models. The data used for this study consists daily withdrawals from 61 ATMs of one of the largest European ATM networks operating in Krakow, Poland and covered the period between January 2017 and April 2021. In this study SARIMA model was compared with two recently implemented methods, XGBoost and Bayesian vector autoregression model. The result indicates that the recently implemented models were more accurate than SARIMA, and with XGBoost providing the most accurate forecasts in most cases. XGBoost is also the least affected by the changes that came with the pandemic, although it was it was not the best model in every periods.

Riabykha *et al.* (2022) conducted research on ATM cash flow prediction at one of the largest commercial banks in Russia focusing on handling the nonstationarity of real-world ATM data. The

purpose of this study was to build a data preprocessing pipeline that addresses nonstationarity and to compare local and global modeling approaches for forecasting customer cash demand. This study used ATM transaction records from VTB Bank which operates a large network of machine across Russian market. Various forecasting methods were proposed under two paradigms called the local model paradigm where a separate model is trained for each ATM and the global model paradigm where a single model is trained across many ATMs at once. These methods were compared against classical time-series approaches and against earlier methods used specifically for cash demand forecasting and the result shows that proposed methods showed better forecasting performance.

Kiyaei and Kiaee (20205) conducted study on smart ATM cash replenishment planning using deep reinforcement learning motivated by the increased importance of stable cash access during the COVID-19 lockdown period. The purpose of this study was to design a real-time replenishment system which minimizes ATM funding and security cost under uncertain cash demand and that accounts for the higher security fees that apply when replenishment falls on holidays and weekends. The study presented the problem as a sequential decision problem and used a Double Deep Q-Network (DQN). The recommended methods were compared with several prediction based neural network approaches including a clustered deep neural network (CDNN), a recurrent neural network (RNN), LSTM and a shallow neural network (SNN). The result of the study shows that Double DQN reduced the total replenishment cost more effectively than the prediction-based methods because prediction only systems rely on the forecast alone without optimizing the replenishment action itself. The authors mentioned that LSTM can reach competitive cost in some contexts, but reinforcement learning provides more direct way to learn cost minimizing actions. The contribution of this study shows that the value of forecasting in cash management increases when the forecast is connected to decision making layer.

Ekstrand and Mattsson (2025) conducted research on time series forecasting for ATM cash demand with a specific focus on adapting models to anomalous events. The purpose of the study was to develop forecasting model which remains accurate when unusual event disrupts demand and to test methods which allows the model to adapt swiftly to such disruptions. The data used in this study covers the period from 2023 to 2024 of ATM transaction data enhanced by exploratory analysis of trends, seasonality, local events, operational disruptions and weather as exogenous

variable. State of the art forecasting models were implemented and then adjusted in two ways called fine-tuning on data from the anomaly window and constrastive learning to better separate normal and abnormal patterns. The study was during the disruption around the period of the Russian invasion of Ukraine. The result shows that the adaption method improved performance during anomalous periods and that clustering ATMs to groups of similar machines also improved forecast quality. The contribution of this study is it offers practical methods for handling shocks in ATM demand which is relevant for a bank that operates in volatile environment.

Oflazoğlu and Rençber (2023) applied Random Forest and XGBoost to predict bank profitability. The purpose of the study was to test if ensemble machine learning models can correctly predict bank profitability indicators namely Return on Assets (ROA) and Return on Equity (ROE) and to compare them with traditional Multiple Linear Regression baseline. This study used financial ratio data from Turkish deposit banks covering the period from 2010 up to 2020, this period includes episodes of currency depreciation and high inflation. In the study three models were trained, Multiple Linear Regression, Random Forest and XGBoost and evaluated using R2 and RMSE.

The result shows that Random Forest produced the lowest RMSE on ROA, followed by XGBoost and then linear baseline while R2 values stayed around 40 percent across the three models. The authors concluded that XGBoost and Random Forest are usable for predicting financial outcomes in banking and outperform a linear baseline even when the underlying data is noisy. The contribution of the study to present research is methodological, it provides that Random Forest and XGBoost are reliable tools for banking analytics that supports the choice of these two models for ATM cash demand forecasting at CBE. Although prior studies shows that machine learning can forecast ATM cash demand no publishable study has yet applied specific to CBE. Which is the gap this study addresses.

CHAPTER THREE

RESEARCH METHODOLOGY

The research methodology serves as systematic approach to address research problems by outlining the process employed in this study. This Chapter consists of Research Design, Research Approach, Population, Sampling techniques, Sample Size, Data Source, Data Collection Procedure, Data Analysis, Reliability and Validity Test and Ethical consideration of the research.

3.1 Research Approach

Research approach can be classified into three approaches. This are quantitative, qualitative and mixed. in This study the researcher used a quantitative method to achieve the objective of the study. This approach was chosen because the research questions are about predicting numerical values namely the daily ATM cash demand and quantitative methods are the most appropriate way to answer this questions. The data used in the study are also numerical and structured consisting of daily ATM transaction records and contextual variables such as day of the week, the month, salary days, holidays and ATM identity.

The approach was predictive rater than descriptive or explanatory. A descriptive approach would only summarize what happened in the past and explanatory would test why something happened. This study used past data to predict future demand which is why it is called predictive. The approach was also applied, since the goal was to propose a machine learning model.

3.2 Research Design

The study used a quantitative predictive modeling design. This design fit the objective of the study which was to build a machine learning models that forecast ATM cash demand. Under this design historical ATM transaction data and contextual variables were preprocessed and used as inputs and daily total transaction amount in ETB was used as the output to be predicted.

Three machine learning models were trained and compared under a same condition. Linear Regression were used as a baseline model, Random Forest and XGBoost were used as main predictive models. Model performance was measured using four evaluation metrics called MAE, RMSE, MAPE and R2. The design followed a standard machine learning practice for time-series

data including a time-based split between training and test sets to prevent data leakage and give a realistic estimate of how the models would perform of new unseen days.

3.3 Population

The population of this study is all ATM cash withdrawal handled by CBE during the study period from May 2024 up to May 2026. Each transaction represents a unit of analysis and the aggregated withdrawal amount over specific time intervals form the basis for predictive modeling. While the study does not directly analyze individual customers, the population indirectly reflects the behavior of bank clients who utilize ATM services.

3.4 Sampling Size

The sample size in this study is defined by the total number of daily ATM cash withdrawal transactions records collected over a two-year period, from May 2024 to May 2026. The sample covers 15 ATMs and 15 districts that covers different areas, with one daily record per ATM giving a total of 9,831 raw daily records. This timeframe was selected to ensure that the dataset captures complete seasonal and cyclical variations, including salary cycles, holidays, and annual fluctuations in customer behavior and the total number of records is large enough to allow a time-based train and test spilt that leaves a sizeable test set for evaluation. The per ATM record depth is enough to support the lag and rolling features built from past demand that requires several weeks of history before they became valid.

3.5 Sampling Technique

A purposive sampling technique was applied to select the 15 ATMs included for the study. Selection criteria include traffic volume (high-traffic, medium traffic and low-traffic machines). The traffic level was only used as guide selecting a balanced sample and was not used as input variable in the forecasting models this is because an ATMs traffic level is derived from its own demand and including it would bias the model. This sampling approach ensures that the models is trained and tested on a representative mix of ATM types rather tan only quiet or only busy machines.

3.6 Data Source and Type

This study used secondary data obtained from the bank's internal ATM transaction records. The data is quantitative in nature, consisting of numerical and categorical variables suitable for predictive modeling.

3.7 Data collection procedure

The data used in this study was obtained as secondary data from the bank. While the bank's electronic transaction logging system is the original source, the dataset was provided to the researcher in processed and aggregated form by the bank, after personal customer information was removed.

3.8 Data Analysis

The data analysis was carried out using the Python programming language. Before modeling, the raw transaction data was cleaned. The transaction dates were converted to a proper date format, text fields were standardized, and the data was checked for missing and invalid values. Outliers in the demand values were treated using the interquartile range method and capped rather than deleted, so that real high demand days were kept. New input features were then created from the cleaned data through feature engineering. These included calendar features such as the day of the week and the month, event features for salary days and religious holidays, and lag and rolling features built from past demand and past transaction counts.

Three machine learning models were trained and compared. Linear Regression was used as a baseline, because any useful model should improve on a simple linear approach. Random Forest (Breiman, 2001) and Extreme Gradient Boosting, known as XGBoost (Chen & Guestrin, 2016), were used as the main models, since both handle nonlinear patterns in structured data well. All three models were trained on the same data so that the comparison would be fair.

The data was split into a training set and a test set using a per-ATM time-based split. For each ATM, the earliest 80 percent of days were used for training and the latest 20 percent for testing. For the Linear Regression model, the features were standardized using a scaler that was fitted on the training data only.

Model performance was measured using four evaluation metrics which are Mean Absolute Error (MAE), Root Mean Square Error (RMSE), Mean Absolute Percentage Error (MAPE) and the Coefficient of Determination (R2). Each metric captures a different aspect of accuracy. MAE gives the average error in Ethiopian Birr, RMSE gives more weight to large errors, MAPE expresses the error as a percentage, and R2 shows how much of the variation in demand the model explains. MAE and R2 were treated as the primary metrics, because MAPE can be unreliable when actual demand is close to zero. In addition, feature importance was calculated from the tree-based models to identify which variables most influenced the predictions, which adds interpretability and practical insight for decision makers.

3.9 Evaluation Metrics

Model performance was measure using four metrics which are Mean Absolute Error (MAE), Root Mean Square Error, Mean Absolute Percentage Error and Coefficient of Determination(R2). Each metrics catches different sides of accuracy and reading them together gives a complete picture than any one alone (Hyndman & Koehler, 2006).

MAE gives the average prediction error in Ethiopian birr. RMSE express the error in ETB but squares the error before averaging so this gives more weigh to large misses that matters when a single cash-out is more costly than several small error. MAPE states the error as a percentage of the actual demand which allows comparison across ATMs of different size. And R2 shows the share variation in daily cash demand which the model explains.

This evaluation metrics are calculated. Let n denote the number of observations in the test set, y_i the actual ATM cash demand on day i, $\hat{y}_i$ the demand predicted by the model on day i, and $\bar{y}$ the mean of the actual demand values in the test set.

Mean Absolute Error (MAE) measures the average size of the prediction error in Ethiopian Birr. It is the simple average of the absolute differences between actual and predicted values.

$$MAE = (1/n) \times \Sigma |y_i - \hat{y}_i|$$

Root Mean Square Error (RMSE) also measures the average error in Ethiopia Birr, but it squares the error before averaging.

$$RMSE = \sqrt{[(1/n) \times \Sigma (y_i - \hat{y}_i)^2]}$$

Mean Absolute Percentage Error (MAPE) shows the error as a percentage of the actual demand which enables comparison across ATMs of different sizes. The metric is multiplied by 100 so that result is read directly as a percentage.

$$MAPE = (100/n) \times \Sigma \left| (y_i - \hat{y}_i) / y_i \right|$$

Coefficient of Determination (R2) measures the share of the variation in daily cash demand that the model explains. A value of 1 means as a perfect fit and a value of 0 means the model performs no better than predicting the mean and a negative value means the model performs worse than the mean.

$$R^2 = 1 - [\Sigma (y_i - \hat{y}_i)^2 / \Sigma (y_i - \bar{y})^2]$$

In all four formulas the summation Σ over the *n* observations in the test set.

In this study MAE and R2 were considered as the primary evaluation metrics for selecting the best model. MAPE was used as a supporting metric because it can become unstable when the actual demand is close to zero and RMSE was reported as check on the size of the largest error.

3.10 Tools and Software

The analysis was carried using python programming language (Van Rossum & Drake, 2009). Open-source libraries were used, the pandas library was used for data handling and feature engineering (McKinney, 2010). The scikit-learn library was used for Linear Regression, Random Forest, data scaling and evaluation metrics (Pedregosa et al., 2011). And the XGBoost library was used for gradient boosting model (Chen & Guestrin, 2016). These tools are chosen because they are widely used in research and are open source which makes the analysis reproducible by a reader who has the data and scripts.

3.11 Validity

To ensure validity, the study uses transaction data directly obtained from the bank's internal records. This guarantees that the dataset accurately reflects actual ATM withdrawal behavior. The inclusion of contextual variables such as salary cycles and holiday calendars further strengthens construct validity by capturing factors known to influence withdrawal demand.

For external validity the findings should be interpreted with the size of the sample in mind. This study used only 15 ATMs out of more than 2500 ATMs that CBE operates overall. So, the result of this study is indicative of ATMs with similar profiles rather than generalizable to the full network. A larger and diverse sample of ATMs are needed before the finding could be claimed to apply to all CBE machines

3.12 Reliability

Reliability is ensured through the consistency of the bank's electronic transaction logging system, which automatically records each withdrawal with standardized fields such as amount, timestamp, and ATM identifier. As the data is generated and maintained by the bank's operational systems, it is free from researcher bias and replicable under similar conditions.

The reliability of the study was also helped by steps which prevent data leakage and this makes the analysis easy to repeat. The training set and test set were split by time for each ATM. And the full process which includes cleaning preprocessing, modeling and evaluation was saved in set of numbered scripts that can be run in order. A fixed random seed was used for both Random Forest and XGBoost. This erases the small changes that random parts of these models would normally cause from one run to the next. Therefore, anyone who runs the same scripts on the same data will get the same result shown in Chapter 4.

A stronger reliability check uses time-series- cross-validation in which the models are repeatedly trained on earlier windows and tested on later windows that follows them which gives a more stable estimates of model performance in different periods. Time-series cross validation was not implemented in this study because the goal was to establish baseline feasibility under a clean time-based split.

3.13 Ethical Consideration

This study used secondary data provided by the bank. All transaction records were anonymized and aggregated prior to analysis to ensure customer confidentiality. No personally identifiable information will be accessed or used. The data was obtained with the bank's consent and used solely for academic purposes. Ethical approvals were sought in line with institutional requirements, and the study adheres to principles of responsible data handling, including accuracy, transparency, and respect for privacy.

Beyond data privacy this study takes cautious position on how the model should be used. The forecasting model is recommended as decision support tool that helps cash management staff not fully automated system that replace human judgement, this is because human review of model output is treated as one part of responsible deployment, specially on unusual days when banks operational knowledge differs from what the model predicts.

CHAPTER FOUR

DATA ANALYSIS AND RESULT DISCUSSION

Introduction

This chapter presents the data analysis and models result discussion. Which includes description of dataset, data cleaning, feature engineering, model development and comparison of the three machine learning models. The three models used are Linear Regression, Random Forest and Extreme Gradient Boosting (XGBoost). The analysis was undertaken with the help of Python programming language.

4.1 Dataset Description

The dataset was acquired from Commercial Bank of Ethiopia, the largest government bank in the country. The data period covers from May 2024-May 2026 aggregated daily transaction, with 9831 records. It includes data from 15 ATMs selected based on transaction volume. The data was provided under formal data sharing agreement, and no personal customer information is included.

Table 4.1 Statistical summary

	Information
Study period	May 2024 to May 2026
Number of ATMs	15
Total records	9831
Mean daily demand	957,807
Median daily demand	880,900
Standard deviation of demand	582,747
Minimum daily demand	200
Maximum daily demand	2,504,025

4.2 Data cleaning

The dataset was checked for data quality issues. Three main cleaning steps were undertaken. The steps are data type correction, checking for impossible values and outlier treatment. Each step is discussed below.

4.2.1 Data Type Correction and Standardization

It was found that the transaction date was stored as text string and needed to be converted to a proper date format. Transaction amount and number of transactions were confirmed to already be in integer format. Text fields for terminal ID and district name were standardized to avoid inconsistencies from mixed case values. One column was removed because it does not provide additional information beyond what it is already captured by terminal ID which uniquely identifies each ATM.

4.2.2 Impossible Values

The dataset was verified no record where transaction amount or transaction count was zero or negative. This shows all records represents real and completed transaction. No record was removed in this step.

4.2.3 Outlier Detection and Treatment

Outlier detection was applied to the total transaction amount and number of transactions using the interquartile range (IQR) method. The IQR was calculated as the difference between the 75^{th} and 25^{th} percentiles. The upper bound was set at Q3 plus 1.5 times the IQR, which gave a value of 2,504,025 ETB. A total of 169 records or 1.7% of the dataset had values above this bound for the total transaction amount and 31 outlier records or 0.3% were found and capped for the number of transactions.

These records were treated using winsorization, which means the extreme values were capped at the upper bound rather than deleted. This approach is consistent with the recommendation of Fallahtafti et al. (2022) to preserve outlier records that represent real demand pattern.

4.3 Exploratory Data Analysis

Exploratory Data Analysis (EDA) was undertaken in this step to understand the patterns in ATM cash demand before building the models. This helps to identify which features are likely to be useful for prediction. Three key patterns were examined, specifically demand by day of week, demand by month, and the effect of special events such as salary days and holidays.

4.3.1 Demand by Day of Week

Figure 4.1 shows the average ATM cash demand for each day of the week. Saturday has the highest average demand at 1,102,967 which is 30 percent above Sunday that have lowest average demand at 880,344. Weekday demand is stable ranging from 880,742 on Wednesday to 949,048 on Tuesday. This result is consistent with the finding of Gokçay (2020), who also found day of week is strong predictor of ATM demand.

Figure 4.1 Average ATM Cash Demand by Day of Week

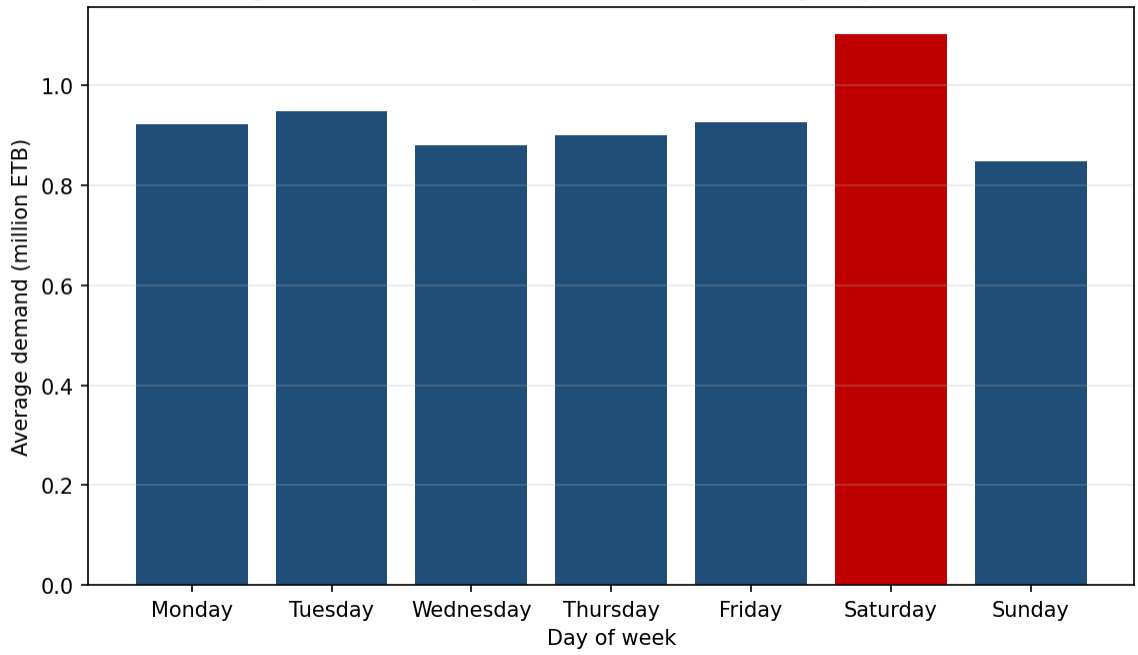


Figure 4.1: Average ATM Cash Demand by Day of Week (ETB)

4.3.2 Demand by Month

Figure 4.2 shows average demand for each month. April shows the highest monthly demand at 998,273, while August shows the lowest at 861,394. Demand tends to be highest in the first half of the year. The variation across months shows that the month of the year is useful feature for the model.

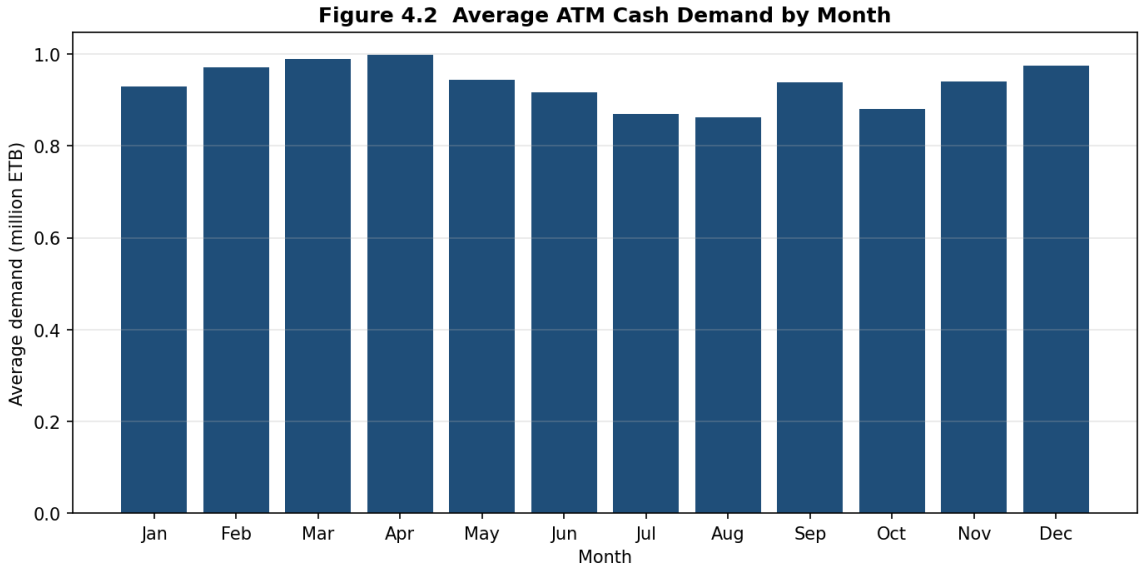


Figure 4.2: Average ATM Cash Demand by Month (ETB)

4.3.3 Effect of Key Events

Figure 4.3 shows how demand changes on salary days, preholiday days and holiday (religious holidays) days compared to normal days. Salary days shows a 15.3% increase above the normal daily average of 916,000. This implies the behavior of CBE customers which receive their salary between 25th to 30th of each Ethiopian calendar month. This pattern was also observed by Al-zeir (2024), where salary transfer days were found the most important driver of ATM withdrawals.

Pre-holiday days show the highest demand increase with 20.1% above norma daily average demand. This increase may be related to customers withdraw for holiday celebration before the exact holiday day. Religious holiday days themselves shows a 16.5% decrease in demand

compared to normal days. This makes sense because customers have already withdrawn cash on the days before. These patterns justify including both holiday and pre-holiday features in the holiday.

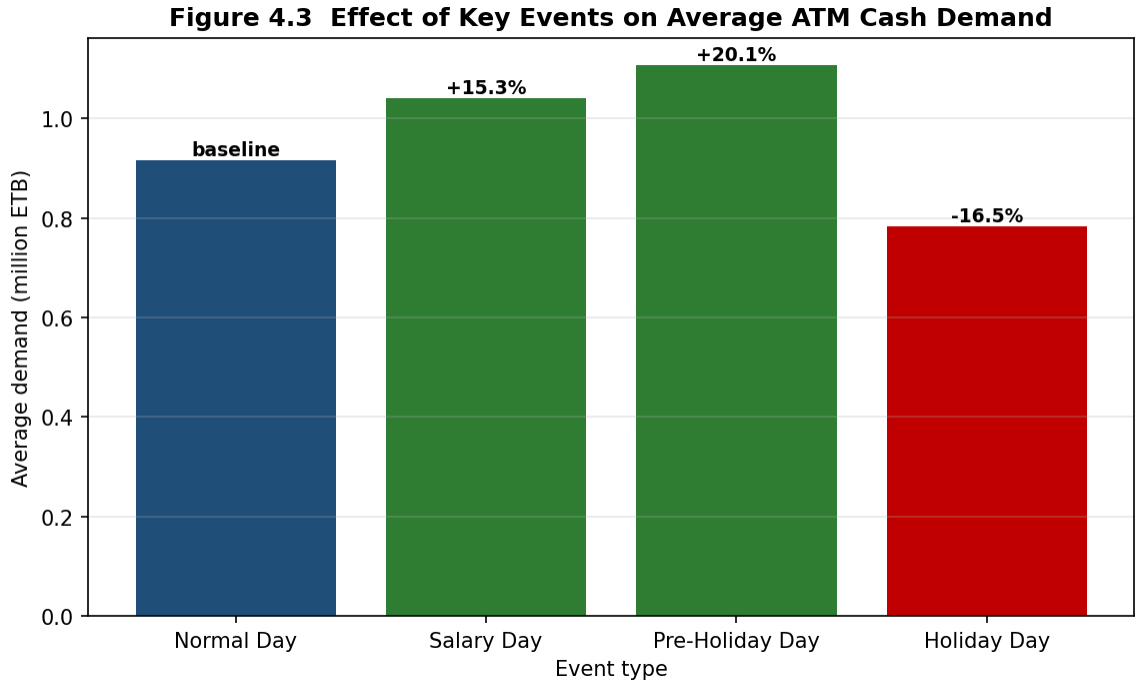


Figure 4.3: Effect of Key Events on Average ATM Cash Demand (ETB)

4.3.4 Additional Data Visualization

Five additional charts are presented in this section that give deeper view of the data before the models were built. The charts support the patterns described above and the feature choices explained in the next section.

Figure 4.4 shows the average daily ATM cash demand throughout all ATMs over the study period. The fine line shows the daily values while the bold line shows the 30-day moving average. Demand moves up and down rapidly from day to day and the moving average shows a clear up ward trend rising from 0.9 million ETB in early 2024 to about 1.3 million ETB by late 2026 which means average demand grew over the period. The upward trend is significant for the modeling because the test set covers the later higher demand days so that the models are tested on demand levels that are higher than the ones they were trained.

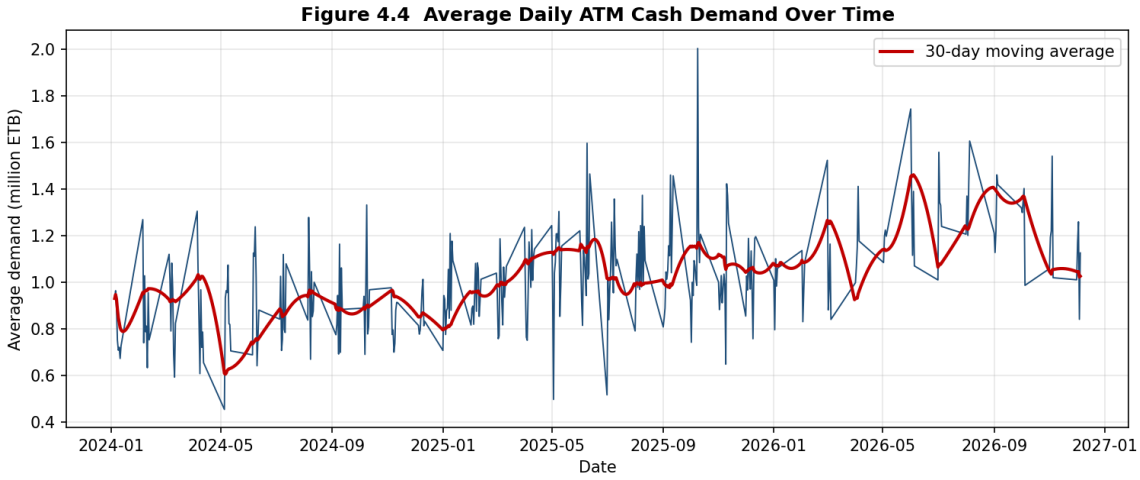


Figure 4.4 Average Daily ATM Cash Demand Over Time

Figure 4.5 divides the daily demand into three parts, which are the trend, the weekly seasonal pattern and the residual. The trend part shows a gradual rise over the period which shows the average demand growth over time. The seasonal part repeats every 7 days which verifies the weekly pattern in figure 4.1 where Saturday is high and Sunday is low. And the residual part is what is left after removing the trend and the season and it represents the day-to-day randomness that the model can not easily predict. This division shows that demand follows a rising trend together with repeating weekly patterns which is one reason the lag and rolling features work well.

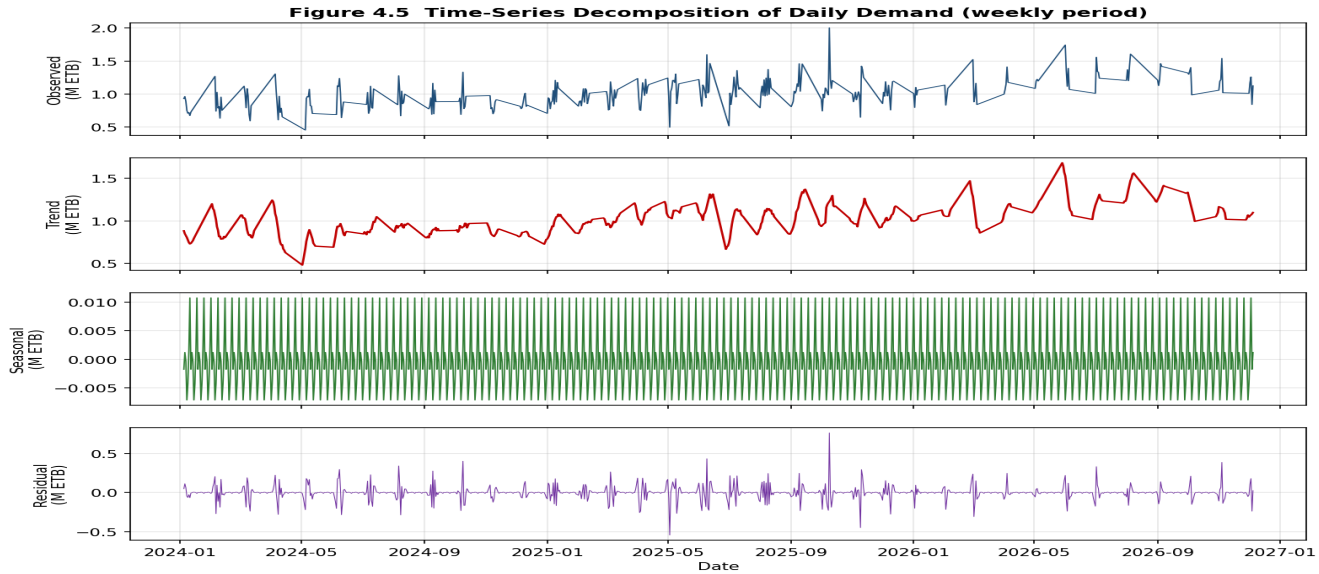


Figure 4.5 Time-series Decomposition of Daly Demand

Figure 4.6 shows the full distribution of demand for each day of the week not the average. The box shows the middle half values and the line inside each box is median. Saturday has the highest median demand and widest box this means demand is both higher and more variable in Saturday. And Sunday has the lowest median. The weekday boxes are similar to each other. This view adds more detail than the average in figure 4.1 because it shows that some days are higher on average and harder to predict.

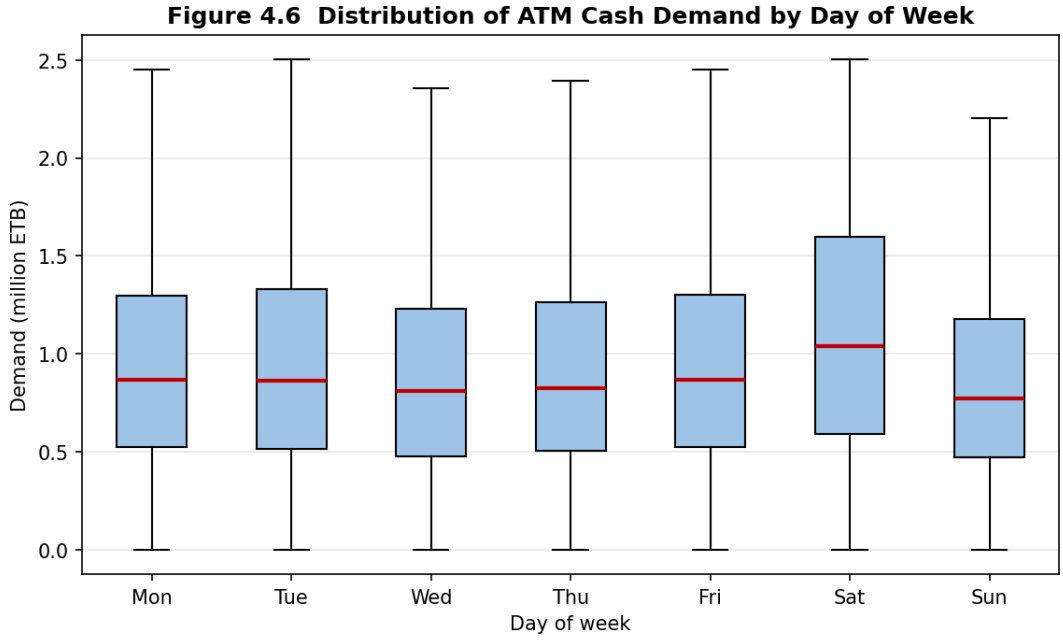


Figure 4.6 Distribution of ATM Cash Demand by Day of Week

Figure 4.7 shows the average demand for each day of the Gregorian month. Demand is higher in the first third of the month and lower in the second half. The 25th of the Gregorian month which is expected to be high is not higher than the other days. This is because the data shows most salary period are based on Ethiopian calendar which starts from 25th Ethiopian calendar month that falls on different Gregorian dates through the year. This pattern in this figure supports the decision to build the salary feature on the Ethiopian calendar rather than the Gregorian as explained in section 4.5.2.

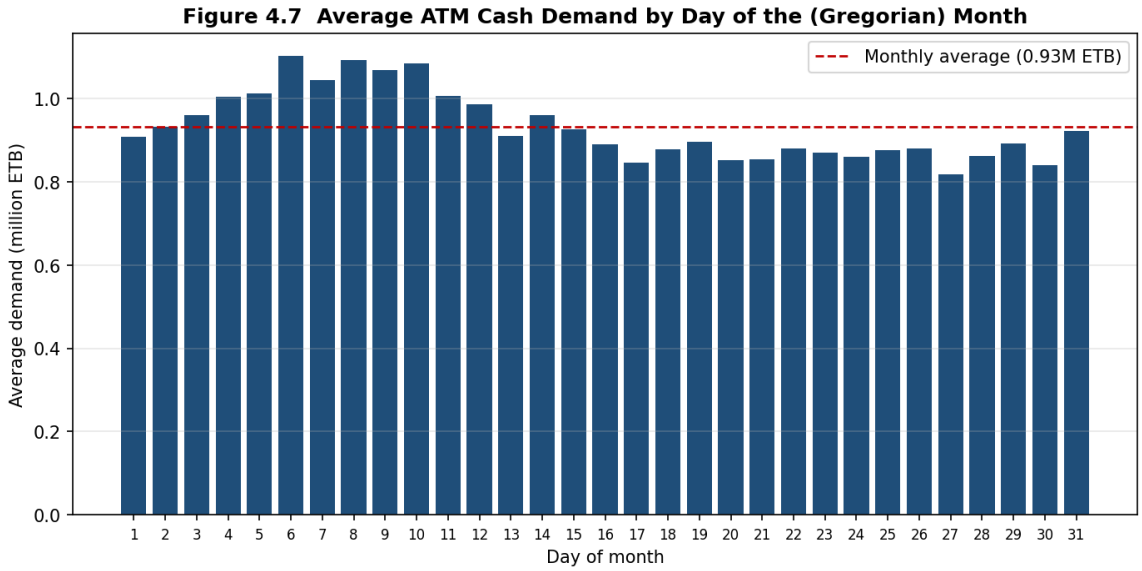


Figure 4.7 Average ATM Cash Demand by Day of the Month

4.4 Feature Engineering

New features were added and grouped in to four categories, namely calendar features, salary and holiday features, categorical encoding and lag features. Summary of all features created are provided in Table 4.2 below.

Table 4.2: New Features Description

	Feature Name	Type	Description
1	day_of_week	Calendar	Day number (0=Monday,6=Sunday)
2	month	Calendar	Month number (1 to 12)
3	day_of_month	Calendar	Day of month (1 to 31)
4	quarter	Calendar	Quarter of year (1 to 4)
5	week_of_year	Calendar	Week number of the year
6	days_to_month_end	Calendar	Days remaining until month end
7	is_weekend	Binary	1 if Saturday or Sunday, else 0
8	is_salary	Binary	1 if within salary window, else 0
9	is_holiday	Binary	1 if religious holiday, else 0
10	is_pre_holiday	Binary	1 if 1 to 3 days before holiday

11	DISTRICT_ENCODED	Encoded	District name as number (0 to 14)
12	TERMINAL_ENCODED	Encoded	ATM terminal ID as number (0 to 14)
13	lag_1	Lag	Demand from the previous day
14	lag_7	Lag	Demand from same day last week
15	Rolling_mean_7	Lag	7-day rolling average of past demand
16	count_lag_1	Lag	Transaction count from the previous day
17	count_lag_7	Lag	Transaction count from same day last week
18	count_rolling_7	Lag	7-day rolling average of past count

4.4.1 Calendar Features

Seven features were generated from the transaction date. They are day of week, month, day of month, quarter, week of year, a weekend flag and days remaining until the end of the month. This is due to cash demand builds progressively towards month end not immediately on one day. Fallahtafti et al. (2022) showed that the position of a day within the month is important driver of ATM withdrawal behavior.

4.4.2 Salary Cycle Feature

A 7-day salary forward window was used starting from the salary day itself through the following 6 days. The window is chosen because that ATM demand stay elevated for several days after payday. This feature shows a 15.3% increase over non salary days verifying its predictive value.

The Gregorian salary cycle was removed from the model because initial testing showed that the Gregorian and Ethiopian salary windows overlapped covering the same calendar days. Including both would brings multicollinearity and cause the model to count the salary signal twice.

4.4.3 Holiday Features

Two holiday related features were generated. The first is binary flag for religious holiday days called is_holiday. This feature points major religious holidays observed in Ethiopia. A total of 214 records were flagged as holiday days.

The second feature is a pre-holiday flag called is_pre_holiday. Which covers the 3 days before any religious holiday. The data proved that demand rises by 20.1% on pre-holiday days as customers withdraw in advance. The feature obtained 661 records across the study period. Both features are supported by Al-zeir (2024) findings that shows events related to salary and holidays are the strongest contextual predictors of ATM cash demand.

4.4.4 Lag features

Six lag features were created to provide the model memory of past activity. The three features are based on past demand that is lag_1 which is demand from the previous calendar day for the same ATM, lag_7 which is demand from the same day of the previous week and 7-day rolling mean of past demand. The remaining three are based on past transaction count, namely count_lag_1, count_lag_7 and 7-day rolling mean of past counts. These lag features were computed on a full calendar grid to assure accuracy.

4.4.5 Feature Correlation

Figure 4.8 shows the correlation between the main features and ATM cash demand. The strongest positive correlation is with lag and rolling features that are built from past demand. The 7-day rolling mean of demand and the previous days demand have the highest correlation with current demand. This confirms the feature importance results in section 4.7where the same feature ranked highest. The calendar features show weak correlation in their own because their effect depends on the other factors. This heatmap also shows that demand and transaction count are closely related, since more transaction means usually more cash withdrawn.

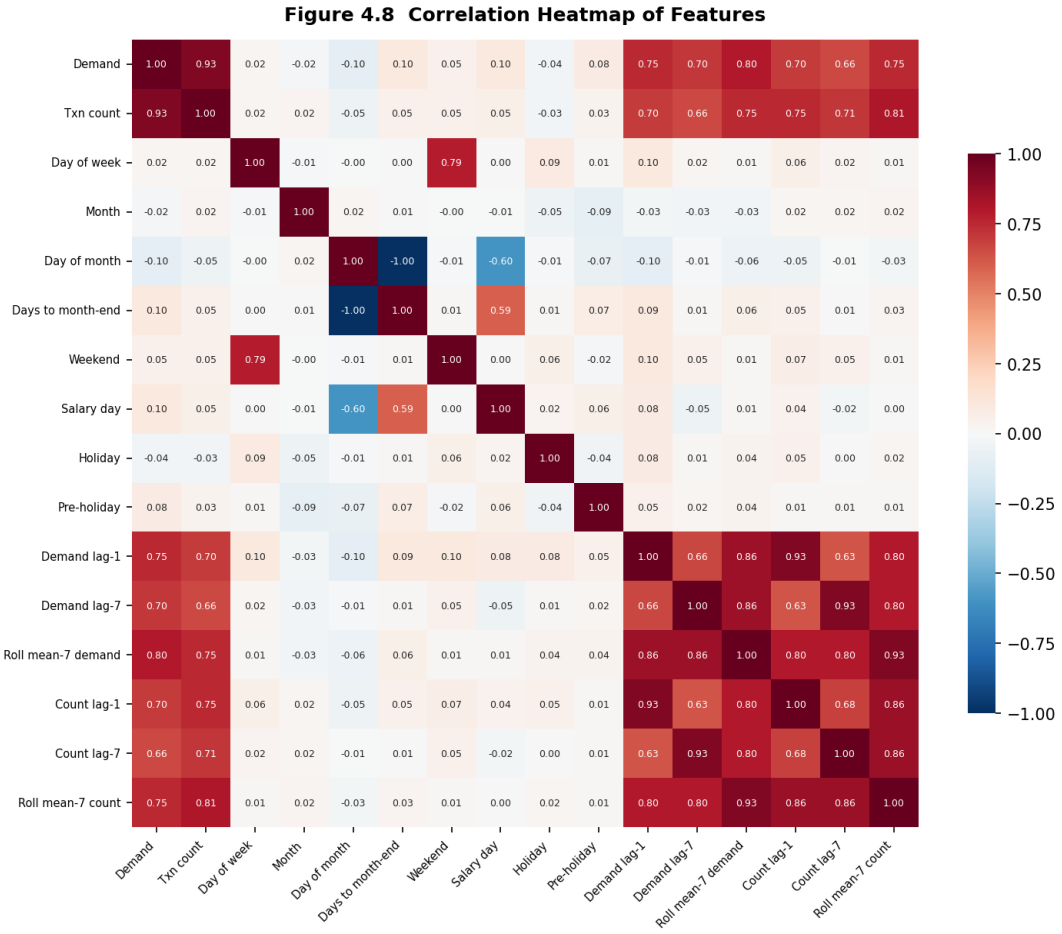


Figure 4.8 Correlation Heatmap of Features

4.5 Data Preprocessing

4.5.1 Removing Incomplete Records

After feature engineering is done 820 records were removed because their lag features contained missing values. These records relate to the first few days of each ATM in the dataset and to days that followed a missing record. After the removal 9,011 records remained for modeling that represents 91.7% of the original data which is large enough sample to train and test the models reliably. These Missing values are because of some days are missing due to machine offline or that day was not recorded. The lag features were built on complete calendar grid any day that fell right after a missing day has no valid value and was left blank. The blank vale was not filled by zero because a missing day means no data was recorded, not the demand was zero so replacing it with zero would distort the lag features.

4.5.2 Train and Test Split

The data was spilt to training and test set using per ATM time based spilt.so for each ATM its earliest 80% of days were used for training and its latest 20% of days were used for testing. This time-based split was used instead of random spilt because ATM demand data has time order. So, splitting by time guarantees the model is always trained on earlier days and tested on later. A random split would made the model see future data during training which is form of data leakage and could produce unrealistic good results (Gokçay, 2020)

4.5.3 Feature Scaling

Feature scaling was used to the linear regression model only using StandardScaler, that transforms each feature to have mean of zero and standard deviation of one. The scaler is fitted on the training data only and then applied to both training and test sets. This prevents data leakage where information from the test set leaks to the training process. Random forest and XGBoost does not require scaling because tree-based models are not sensitive to the scale of input feature.

4.6 Model Development and Hyperparameter Selection

The three models in this study were developed and compared under the same condition to keep the comparison fair. All the models were trained on the same training set and tested on the same test set. Linear Regression was used as the baseline which has no setting to tune and Random Forest and XGBoost have hyperparameters which are settings that control how each mode learns from the data.

The hyperparameter was selected through initial testing where the default settings were tried first and then adjusted to reduce overfitting. The aim was to keep the models simple enough to generalize to new days rather than memorize the training data. For Random Forest 300 trees were used with minimum of five samples per leaf and the square root of the number of features considered at each split. And for XGBoost the number of boosting rounds was set to 200, the maximum tree depth to 4, the learning rate to 0.05 and L1 and L2 regularization were applied. A fixed random seed was used so that the results can be reproduced.

4.7 Model Result and comparison

All three models were evaluated using four evaluation metrics. Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), Mean Absolute Percentage Error (MAPE) and R-squared (R2). In this study MAE and R2 are treated as primary metrics because MAPE can be unreliable when demand values are close to zero.

4.7.1 Linear Regression Result

Linear Regression reached an MAE of 287,254, an RMSE OF 394,607 and R2 of 0.6281 on the test set. This means the model explains 62.8 percent of the variance in ATM cash demand using simple linear approach. The average prediction error is 25.3 percent of the mean daily demand of 1,135,366 ETB on the test set. Figure 4.9 shows the actual versus predicted values of this model. The model follows the general trend of demand but misses sharp peaks and sudden drops.

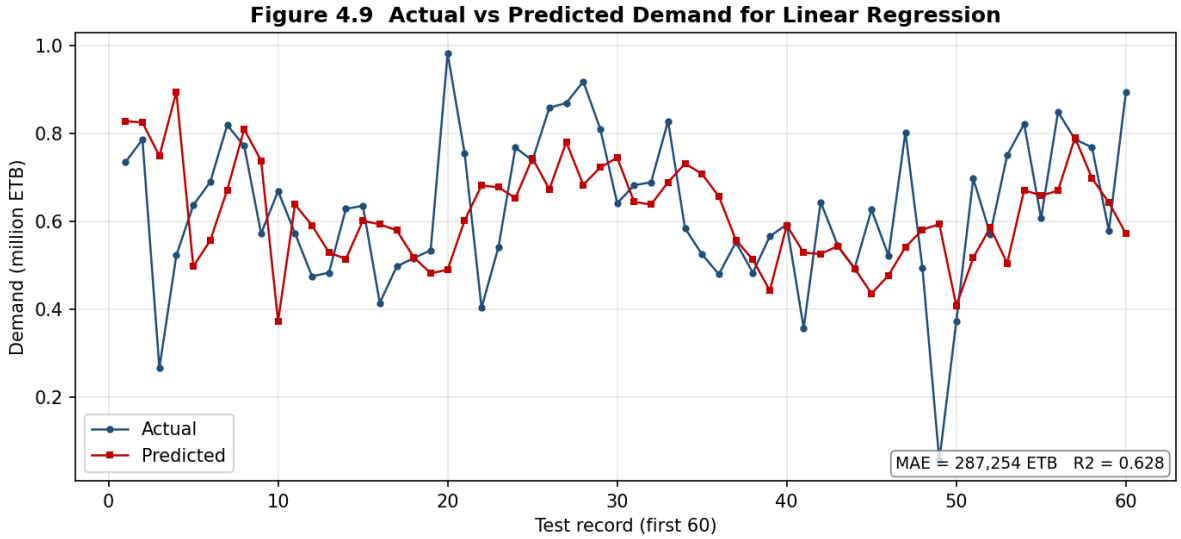


Figure 4.9 Actual vs Predicted Demand for Linear Regression (first 60 test records)

4.7.2 Random Forest Result

Random Forest reached an MAE of 277,939 and RMSE of 381,979 and R2 of 0.6515. This means the model explains 65.2 percent of mean daily demand. The average error is 24.5 of mean daily demand which shows better result than the Linear Regression baseline. Figure 4.10 shows the

actual versus predicted values of this model and shows that Random Forest prediction follows actual demand more closely than Linear Regression.

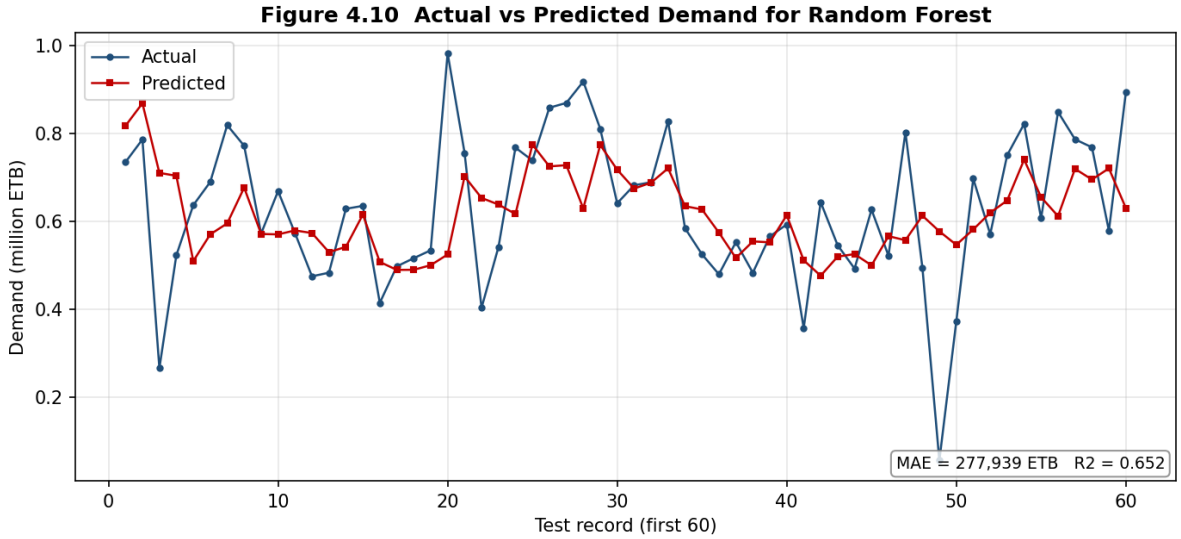


Figure 4.10 Actual vs Predicted Demand for Random Forest (first 60 test records)

Figure 4.11shows the feature importance scores from Random Forest model. The best important feature is the 7- day rolling mean of past demand at 25.0%, followed by the 7-day rolling mean of past transaction counts at 18.7%, the previous days demand lag_1 at 14.8% and last weeks demand lag_7 at 10.4%. The same day transaction count was excluded because it is not known when forecast is made and its lagged version is used instead. Both past demand and past count features explains most of the model's decision. This shows that recent history is the strongest predictor of future demand. Which is consistent with the findings of Sarveswararao and Ravi (2020). The terminal identity and district feature also add smaller contribution which shows that ATM-specific characteristics also matter.

The salary and holiday features show lower importance scores in the model however the exploratory analysis showed that they have a clear effect on demand. This is because the effect of this feature is already captured by the recent demand features. This means demand rises around salary days and before holidays so the previous days demand and the 7-day rolling mean already carry that signal and the model does not need the salary and holiday flags to detect it.

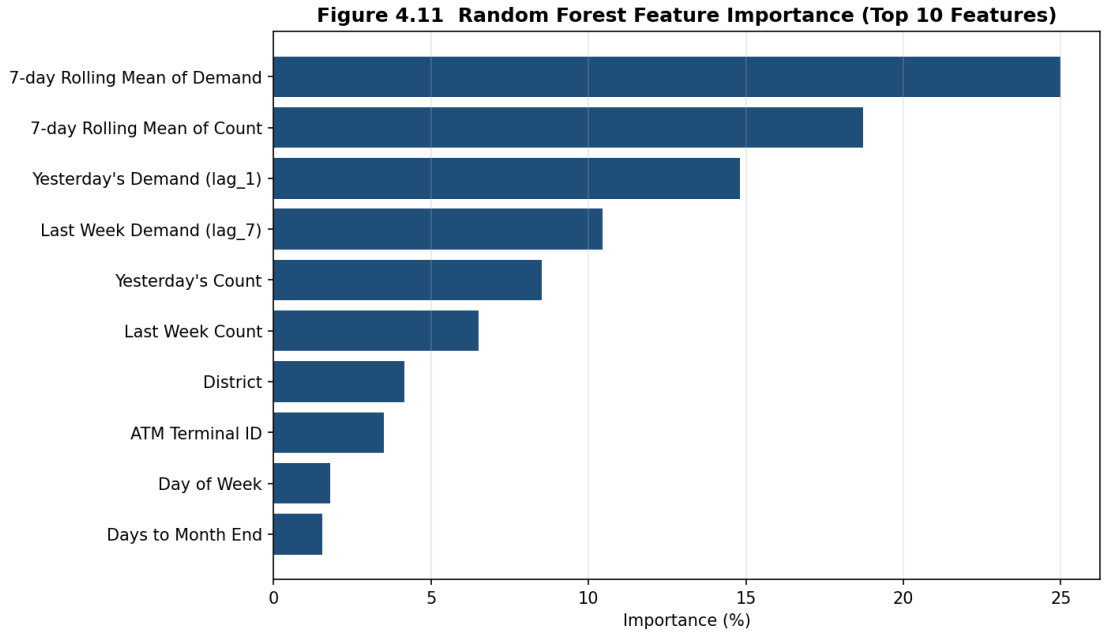


Figure 4.11 Random Forest feature Importance, Top 10 Features

4.7.3 XGBoost Result

XGBoost achieved the best results among the three models. It reached MAE of 268,821 and RMSE of 372,723 and R2 of 0,6682. This means the model explains 66.8 percent of the variance in ATM cash demand. The average error is 23.7 of mean daily demand. Figure 4.12 Shows the actual versus predicted demand for XGBoost also shows that XGBoost prediction track actual demand than the other two models.

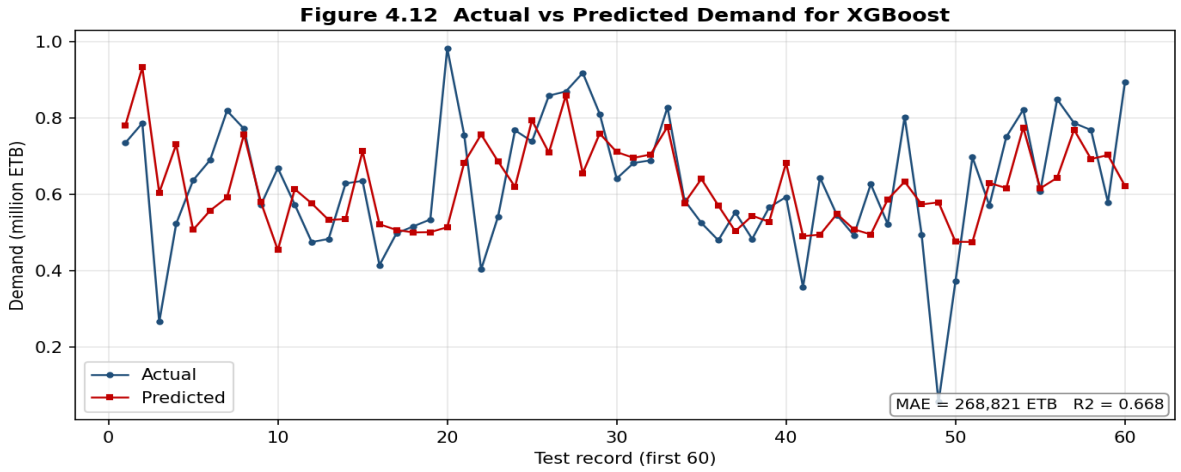


Figure 4.12 Actual vs Predicted Demand for XGBoost (first 60 test records)

4.7.4 Model Comparisons

Table 4.3 shows the full comparison of the three models on the test set. XGBoost ranks first on both MAE and R2, Random Forest second and linear Regression is the baseline. The three models result are close to each other which shows that ATM demand follows pattern that even a simple model can capture well.

Table 4.3 Model Performance Comparison on Test Set

Model	MAE(ETB)	RMSE(ETB)	MAPE	R2
Linear Regression (Baseline)	287,254	394,607	356.46%	0.6281
Random Forest	277,939	381,979	330.87%	0.6515
XGBoost	268,821	372,723	358.39%	0.6682

Figure 4.13 shows the MAE and R2 values of the three models in chart form. The Mae chart shows that XGBoost produces smaller average error. And the R2 chart shows that the three models explain between 63 and 67 percent of demand variance which prove that machine is well suited to ATM cash demand forecasting.

Figure 4.13 Model Performance Comparison - All Three Models

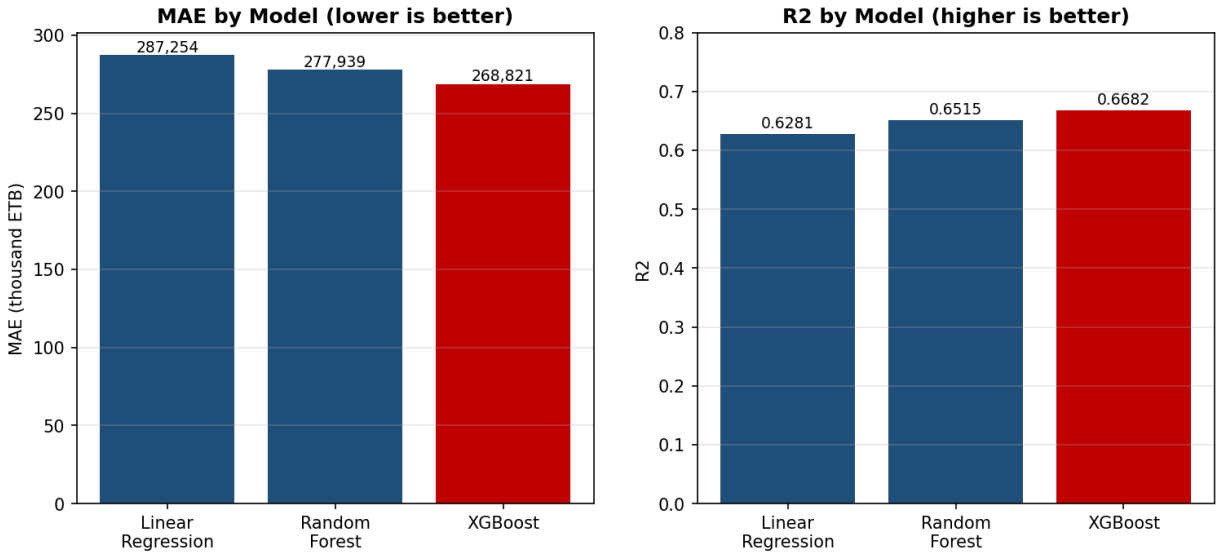


Figure 4.13 Model performance Comparison of MAE and R2 for The Three models

About MAPE all three models show high values and XGBoost shows 358.4%. this looks alarming but this is a known limitation of MAPE on datasets that contain very low demand records. In this study the median percentage error for XGBoost is only about 19% and the high average is caused by small number of very low demand days. Because MAE and R2 are not distorted in this way they give a more accurate picture of model performance.

CHAPTER FIVE

SUMMARY, CONCLUSION AND RECOMMENDATION

Introduction

This Chapter presents summary of findings, conclusion, recommendation obtained from the analysis and interpretation made at previous chapter and limitation of the study and direction for future works.

5.1 Summary of Findings

The main objective of the study was to propose a predictive machine learning model for ATM cash demand forecasting at CBE, demonstrating how technology can enhance business strategy and organizational efficiency. The focus is not just on model accuracy but on how information from the model can be used to improve organizational planning, reduce operational costs and improve service quality. Accordingly, the following are major findings of the study.

- The analysis shows that the most significant features which influence ATM cash demand prediction are the 7-day rolling mean of past demand, the 7-day rolling mean of transaction counts and previous day demand are top three, with recent history features dominating overall.
- The study shows a practical and repeatable process for applying machine learning model to ATM transaction data. This can be applied based on existing data that the bank already generates through its daily ATM operations. The raw transaction records are cleaned, enriched with calendar and event features and with lagged demand and transaction count features and then used to train a model that predicts the next days cash demand for each ATM.
- This study found that salary days and pre-holiday days demand highly increases above average daily demand.
- XGBoost was the best model based on the two main evaluation metrics with an MAE of 268,821ETB and R2 of 0.6682.

5.2 Conclusion

This study aims to examine how predictive analytics and machine learning can support ATM cash demand forecasting at the Commercial Bank of Ethiopia where cash replenishment is currently decided manually from staff experience. Using two years of daily transaction data from 15 ATMs across 15 districts, the study cleaned the data, engineered calendar, event, and lag features, and trained and compared three models which are Linear Regression, Random Forest, and XGBoost. XGBoost gave the best results, predicting daily demand with an average error of about 268,821 ETB and explaining about 67 percent of the variation, with Random Forest a close second. The analysis showed that recent demand history is the strongest driver of future demand, while salary days, pre-holiday days, and holidays produce consistent and predictable shifts.

Collectively, these results show that machine learning can forecast ATM cash demand in the Ethiopian context using data the bank already collects, and that the model is accurate enough to support better replenishment planning, especially on detectable high-demand days. From a Business Information Systems perspective, the main contribution is the finding that the barrier to data-driven cash management at CBE is organizational rather than technical because the data and the methods are already available. The study also fills a gap in the literature, since it is the first to apply machine learning to ATM cash demand forecasting in Ethiopia. Based on the researcher's knowledge. The remaining task is for the bank to build the processes, skills, and management support needed to act on the model's forecasts.

5.3 Recommendation

- The researcher recommends integrating predictive modeling to ATM cash management as decision support tool. According to used method and data XGBoost model is recommended.
- Establishing a formal data governance policy that defines standards for ATM transaction data completeness, timeliness and quality is recommended. Because in this study 820 of expected ATM-day record was missing from the dataset that missing data reduces forecast accuracy in and undermines trust in any predictive system.

- Technology alone does not deliver business value. For predictive model to improve operation the bank has to invest in training for its operation and data teams in areas like data analytics, reporting and the use of decision support systems.

5.4 Limitation and Future area of study

The two years study period relatively short for capturing long term economic cycles so, future work should use longer time to capture long term economic cycles.

Future research should study the organizational adoption process. A study of how bank staff use and trust model outputs, what barriers exist to changing replenishment practices and how management can support the transition to data driven operations would add important knowledge beyond model accuracy.

References

Adekunle, B. I., Chukwuma-Eke, E. C., Balogun, E. D., & Ogunsola, K. O. (2021). Predictive analytics for demand forecasting: Enhancing business resource allocation through time series models. *Journal of Frontiers in Multidisciplinary Research, 2(1).*

Adeniran, I. A., Efunniyi, C. P., Osundare, O. S., & Abhulimen, A. O. (2024). Integrating business intelligence and predictive analytics in banking: A framework for optimizing financial decision-making. *Finance & Accounting Research Journal, 6(8), 1517–1530*. https://doi.org/10.51594/farj.v6i8.1505

Agatić, A., Tijan, E., Hess, S., & Poletan Jugović, T. (2019). Advanced data analytics in logistics demand forecasting. *Scientific Journal of Maritime Research, 33(1), 28–36.*

Al-zeir, L. R. O. (2024). *Automated Teller Machines availability forecasting model using machine learning algorithms* (Master's thesis, Arab American University, Palestine).

Aykanat, Z., Yıldız, T., & Çelik, A. K. (2025). Organizational readiness for big data analytics, business analytics adoption and data-driven culture: The case of Turkish banking sector. Journal of Organizational and End User Computing.

Balakrishnan, T. V., & Kalaiarasi, R. (2021). A literature survey on automated teller machine cash demand analysis and prediction in financial sector. Ilkogretim Online – Elementary Education Online, 20(6), 2146–2161. https://doi.org/10.17051/ilkonline.2021.06.199

Breiman, L. (2001). Random forests. Machine Learning, 45(1), 5-32. https://doi.org/10.1023/A:1010933404324

Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In Proceeding of the 22nd ACMSIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 785-794). https://doi.org/10.1145/2939672.2939785

Commercial Bank of Ethiopia. (2024). *Annual report 2023/24*. https://www.combanketh.et/cbeapi/uploads/CBE_ANNUAL_REPORT_2023_24_e30fd53ce9.pdf

Cossu, G. A. G. (2023). A predictive analytics approach for demand forecasting to support capacity planning (Master's thesis). Eindhoven University of Technology.

Fallahati, A., Aghaaminiha, M., Akbarghandari, S., & Weckman, G. R. (2022). Forecasting ATM cash demand before and during the COVID-19 pandemic using an extensive evaluation of statistical and machine learning models. SN Computer Science, 3(164). https://doi.org/10.1007/s42979-021-01000-0

Gorodetskaya, O., Gobareva, Y., & Koroteev, M. (2022). A machine learning pipeline for forecasting time series in the banking sector. Economies, 10(5), 101. https://doi.org/10.3390/economies10050101

Gökçay, D. E. (2020). *ATM cash stock prediction using different machine learning approaches* (Master's thesis, Sabancı University, Turkey).

Hyndman, R. J., & Athanasopoulos, G. (2018). *Forecasting: Principles and practice* (2nd ed.). OTexts. https://otexts.com/fpp2/

Mitchell, T. M. (1997). *Machine learning*. McGraw-Hill.

Owen, A. (2024). Predictive analytics for demand forecasting in manufacturing. Unpublished manuscript.

Psomakelis, E., Tserpes, K., & Stamati, T. (2025). ATM supply scheduling optimization: Demand prediction to intelligent route planning. *Journal of Intelligent Transportation Systems*.

Punia, S., & Shankar, S. (2022). Predictive analytics for demand forecasting: A deep learning-based decision support system. Knowledge-Based Systems, 258, 109956. https://doi.org/10.1016/j.knosys.2022.109956.

Sarveswararao, V., & Ravi, V. (2020). ATM cash demand forecasting in an Indian bank with chaos and deep learning. *arXiv preprint arXiv:2008.10365*. **https://arxiv.org/abs/2008.10365**

Seyedan, S. (2024). Development of predictive analytics for demand forecasting and inventory management in supply chain using machine learning techniques (Master's thesis). Concordia Institute for Information Systems Engineering.

Styrkársson, V. P. (2020). Forecasting demand for ATM replenishment: How much to order and when (Master's thesis, Reykjavík University).

Suder, M., Gurgul, H., Barbosa, B., Machona, A., & Lach, L. (2024). Effectiveness of ATM withdrawal forecasting and Social Change, 200, 123089. https://doi.org/10.1016/j.techfore.2024.123089

Tornatzky, L. G., & Fleisher, M. (1990). *The processes of technological innovation*. Lexington Books.

Wolniak, R. (2024). Forecasting demand – Utilizing business analytics in Industry 4.0 environments. Scientific Papers of Silesian University of Technology: Organization and Management Series, (196).

Yusof, Z. B. (2024). Analyzing the role of predictive analytics and machine learning techniques in optimizing inventory management and demand forecasting for e-commerce. Malaysia University of Science and Technology.

Appendices

Appendix 1: Sample of ATMs Used in the Study

The study used 15 ATMs selected through purposive sampling across 15 districts, covering high, medium, and low traffic locations. The traffic level was used only to select a balanced sample and was not used as an input variable in the models.

No.	Terminal ID	District	Traffic Level
1	AKF00101	KOLFE	High
2	AKL00171	KALITI	High
3	ANS00233	NIFAS SILK	High
4	ASE00042	SHIRE	High
5	AWT00134	WOLAITA SODO	High
6	ADM00002	DEBRE MARKOS	Medium
7	AGD00091	GONDER	Medium
8	AJJ00019	JIJIGA	Medium
9	AMG00203	MEGENAGNA	Medium
10	ASM00048	SEMERA	Medium
11	AAR00002	ARADA	Low
12	ABD00132	BAHIR DAR	Low
13	ABL00118	BOLE	Low
14	ADD00189	DIRE DAWA	Low
15	AWL00008	WOLDIA	Low

Appendix 2: Full List of Features Used in Modeling

Eighteen features were created from the raw transaction data and used to train the models.

No.	Feature	Type	Description
1	day_of_week	Calendar	Day number (0 = Monday, 6 = Sunday)
2	month	Calendar	Month number (1 to 12)
3	day_of_month	Calendar	Day of month (1 to 31)
4	quarter	Calendar	Quarter of year (1 to 4)
5	week_of_year	Calendar	Week number of the year
6	days_to_month_end	Calendar	Days remaining until month end
7	is_weekend	Binary	1 if Saturday or Sunday, else 0
8	is_salary	Binary	1 if within the salary window, else 0
9	is_holiday	Binary	1 if a religious holiday, else 0
10	is_pre_holiday	Binary	1 if 1 to 3 days before a holiday

11	DISTRICT_ENCODED	Encoded	District as a number (0 to 14)
12	TERMINAL_ENCODED	Encoded	ATM terminal as a number (0 to 14)
13	lag_1	Lag	Demand from the previous day
14	lag_7	Lag	Demand from the same day last week
15	rolling_mean_7	Lag	7-day rolling average of past demand
16	count_lag_1	Lag	Transaction count from the previous day
17	count_lag_7	Lag	Transaction count from the same day last week
18	count_rolling_7	Lag	7-day rolling average of past transaction count

Appendix 3: Model Hyperparameters

The hyperparameters below were selected through initial testing to reduce overfitting. A fixed random seed (42) was used so the results can be reproduced.

Random Forest

Hyperparameter	**Value**
n_estimators (number of trees)	300
max_depth	None (unrestricted)
min_samples_leaf	5
max_features	sqrt
random_state	42

XGBoost

Hyperparameter	**Value**
n_estimators (boosting rounds)	200
learning_rate	0.05
max_depth	4
min_child_weight	20
gamma	1.0
subsample	0.7
colsample_bytree	0.6
reg_alpha (L1)	2.0
reg_lambda (L2)	5.0
random_state	42

Appendix 4: Full Feature Importance Scores

Random Forest

Feature	Importance (%)
rolling_mean_7	24.99
count_rolling_7	18.72
lag_1	14.82
lag_7	10.44
count_lag_1	8.53
count_lag_7	6.51
DISTRICT_ENCODED	4.16
TERMINAL_ENCODED	3.49
day_of_week	1.79
days_to_month_end	1.55
day_of_month	1.53
week_of_year	1.08
month	0.68
is_weekend	0.47
is_salary	0.34
is_pre_holiday	0.32
quarter	0.30
is_holiday	0.29

XGBoost

Feature	Importance (%)
rolling_mean_7	41.90
lag_1	13.00
count_rolling_7	8.76
lag_7	6.68
is_holiday	4.04
count_lag_1	3.41
day_of_week	3.20
is_pre_holiday	2.56
DISTRICT_ENCODED	2.50
TERMINAL_ENCODED	2.24
days_to_month_end	2.04
is_weekend	2.00
is_salary	1.97
count_lag_7	1.62
day_of_month	1.46
week_of_year	1.05
quarter	0.81
month	0.77

Appendix 5: Analysis Scripts

The analysis was carried out in Python using the following scripts, run in order.

Script	Purpose
01_data_understanding.py	Initial exploration of the raw transaction data
02_data_cleaning.py	Data cleaning, outlier treatment, and ATM traffic classification
03_feature_engineering.py	Creating calendar, salary, holiday, and lag features and encoding
04_preprocessing.py	Removing incomplete records, the train and test split, and scaling
05_model_linear_regression.py	Training and evaluating the Linear Regression baseline
06_model_random_forest.py	Training and evaluating the Random Forest model
07_model_xgboost.py	Training and evaluating the XGBoost model
08_model_comparison.py	Comparing the three models and producing the final results