MINISTRY OF HIGHER EDUCATION AND SCIENTIFIC RESEARCH

Higher School of Commerce - Kolea -



End-of-cycle dissertation for the purpose of obtaining a Master's Degree in Financial Sciences and Accounting

Major: - Accounting and Finance

- Money, Finance and Banking

Topic:

Multi-Horizon Forecasting of Central Bank Liquidity Using Deep Learning Models: Assessing the Impact of Exogenous Factors

– The Case of the Bank of Algeria–

Submitted by:

BENYGHZER Walid

OURTEMACHE Mohamed Amine

Supervised by:

- Dr. TAOUSSI Brahim

- Pr. BENILLES Billel

Location of internship: Bank of Algeria

Duration of internship: from 2026/02/15 to 2026/05/15

Academic Year 2025/2026

MINISTRY OF HIGHER EDUCATION AND SCIENTIFIC RESEARCH

Higher School of Commerce - Kolea -



End-of-cycle dissertation for the purpose of obtaining a Master's Degree in Financial Sciences and Accounting

Major: - Accounting and Finance

- Money, Finance and Banking

Topic:

Multi-Horizon Forecasting of Central Bank Liquidity Using Deep Learning Models: Assessing the Impact of Exogenous Factors

– The Case of the Bank of Algeria–

Submitted by:

BENYGHZER Walid

OURTEMACHE Mohamed Amine

Supervised by:

- Dr. TAOUSSI Brahim

- Pr. BENILLES Billel

Location of internship: Bank of Algeria

Duration of internship: from 2026/02/15 to 2026/05/15

Academic Year 2025/2026

Acknowledgment

We would like to express our deepest gratitude to all those who provided us with the support and guidance necessary to complete this dissertation.

First and foremost, we are immensely grateful to our supervisors, Dr. TAOUSSI Brahim and Pr. BENILLES Billel and Mrs.GHEBALOU Ryma, for their invaluable advice, continuous encouragement, and immense knowledge. Their insights and expertise were instrumental in shaping the methodology and direction of this research.

We also extend our sincere thanks to the Bank of Algeria for hosting our internship and providing access to the critical data and institutional reports that formed the empirical backbone of this study. The practical exposure to the realities of liquidity management at the central bank was indispensable.

Our appreciation also goes to the faculty members and professors at the Ecole Supérieure de Commerce - Koléa, whose academic rigor and teachings over the years have equipped us with the foundational tools required to undertake this Master's dissertation in Financial Sciences and Accounting, and Money, Finance and Banking.

Finally, we owe a profound debt of gratitude to our families and friends for their unwavering patience, moral support, and understanding throughout the long months of research and writing. This accomplishment would not have been possible without them.

BENYGHZER Walid & OURTEMACHE Mohamed Amine

***

Dedication

To my parents, whose love, sacrifices, and unwavering support have been the foundation of every step I have taken. Your guidance, patience, and constant belief in me have given me the strength to overcome challenges and pursue my dreams with confidence. No words can truly express the gratitude and respect I hold for you. May you always be blessed with health, happiness, and peace.

To my siblings, thank you for your encouragement, laughter, and the countless moments of comfort and motivation you have given me throughout this journey. Having you by my side has made every challenge lighter and every success more meaningful.

To my grandparents, your wisdom, prayers, and unconditional love continue to inspire me every day. Your values and life lessons will forever guide my path.

To my partner and dear friend, Walid, sharing this journey with you made every obstacle easier and every success more rewarding. Your dedication, thoughtful ideas, and constant motivation brought both meaning and strength to this experience, and I will always value the moments of collaboration and perseverance we shared together.

To my friends and everyone who supported me during this journey, thank you for your kindness, patience, and belief in me. Your presence and encouragement have played an important role in reaching this achievement.

Finally, this work is dedicated to all those who believed in me, stood beside me in difficult moments, and shared in my aspirations and successes. This accomplishment would not have been possible without your love and support.

OURTEMACHE Mohamed Amine

Dedication

First and foremost, I would like to express my deepest gratitude to God for giving me the strength, patience, and determination to complete this work. Without His guidance and blessings, this achievement would not have been possible.

I dedicate this dissertation with all my love and appreciation to my beloved MAMA.

This work is yours. Without you, I would never have achieved what I am today. You taught me how to write, how to speak, and how to become the man I am today. Your endless sacrifices, unconditional love, prayers, and constant encouragement have always been my greatest source of strength. Every success in my life carries your name behind it.

To my father, thank you for your support, wisdom, and belief in me throughout every step of my journey.

To my little sister, thank you for bringing joy, happiness, and light into my life.

I would also like to dedicate this work to my dear friends: Adel, Rayan, Ilyes, Nassim, Hamada, Yacine, and my J.

Thank you for standing by my side during the darkest moments and for sharing with me the happiest ones. Your support, loyalty, laughter, and presence made this journey easier and far more meaningful. I will always be grateful for every memory we created together.

Finally, this dissertation is dedicated to everyone who believed in me and encouraged me to keep moving forward, even when the road was difficult.

BENYGHZER Walid

Contents

ACKNOWLEDGMENT **1**

DEDICATION **1**

CONTENTS ..... **I**

LIST OF TABLES ..... **II**

LIST OF FIGURES ..... **III**

LIST OF APPENDICES ..... **IV**

LIST OF ABBREVIATIONS ..... **V**

ABSTRACT ..... **VIII**

GENERAL INTRODUCTION **1**

CHAPTER I: THEORETICAL FOUNDATIONS OF BANKING LIQUIDITY AND THE ALGERIAN CONTEXT **6**

Section 1: The Concept of Banking Liquidity 8

Section 2: Liquidity Risk and its Management 13

Section 3: Banking Liquidity in Algeria and the Challenges of Forecasting 21

CHAPTER II: THEORETICAL FRAMEWORK AND LITERATURE REVIEW: FROM TRADITIONAL FORECASTING TO DEEP LEARNING FOR CENTRAL BANK LIQUIDITY PREDICTION **26**

Section 1: Traditional Forecasting Approaches for Central Bank Liquidity and Their Limitations 28

Section 2: Deep Learning Architectures for Liquidity Time Series Forecasting 35

Section 3: Literature Review on Deep Learning Applications to Liquidity Forecasting and Comparative Studies 42

CHAPTER III: EMPIRICAL STUDY OF CENTRAL BANK LIQUIDITY FORECASTING USING STATISTICAL AND DEEP LEARNING MODELS **52**

Section 1: Data Construction and Descriptive Analysis 55

Section 2: Forecasting using Classical Models 65

Section 3: Forecasting using Deep Learning Models 83

Section 4: Final Benchmarking Across Daily, Weekly, and Monthly Horizons 96

GENERAL CONCLUSION **101**

REFERENCES **105**

APPENDICES **108**

TABLE OF CONTENT **120**

List of Tables

Table II-1: Comparison of SARIMA and SARIMAX The Two Empirical Traditional Models 33
Table II-2: Comparison of LSTM and GRU Architectures 40
Table II-3: Summary of Comparative Studies Traditional vs. Deep Learning Models 47
Table II-4: Exogenous Variables and Their Expected Impact on Central Bank Liquidity 49
Table III-1: Descriptive Statistics of Central Bank Liquidity (n = 2,085) 55
Table III-2: Descriptive Statistics of Daily Liquidity by Year (DZD Billions) 57
Table III-3: Descriptive Statistics of Exogenous Continuous Variables (n = 2,085) 59
Table III-4: Correlation Between Liquidity and Seasonal/Holiday Dummy variables 61
Table III-5: Unit Root and Stationarity Tests for Liquidity 63
Table III-6: Normality Tests for Liquidity (Full Sample) 64
Table III-7: Selected Parameter Estimates (only significant ones) — SARIMA (3,1,3) (1,1,2) [22] 69
Table III-8: Residual Diagnostic Tests — SARIMA (3,1,3) (1,1,2) [22] 72
Table III-9: Forecasting Performance Metrics — SARIMA (3,1,3) (1,1,2) [22] 73
Table III-10: Stationarity Tests on Liquidity, Brent and USDDZD — Levels and First Differences 77
Table III-11: Selected SARIMAX (2,1,3) (2,1,2) [22] Parameter Estimates (only significant parameters) 78
Table III-12: Residual Diagnostic Tests SARIMAX (2,1,3) (2,1,2) [22] 80
Table III-13: Forecasting Performance Metrics SARIMAX (2,1,3) (2,1,2) [22] 81
Table III-14: Comparative Forecasting Performance — SARIMA vs SARIMAX 82
Table III-15: Selected configuration of the deep learning models 84
Table III-16: Forecasting performance of the CNN model 84
Table III-17: Forecasting performance of the CNNX model 86
Table III-18: Comparative forecasting performance — CNN vs CNNX 87
Table III-19: Best convolutional model by forecasting horizon 87
Table III-20: Forecasting performance of the LSTM model 88
Table III-21: Forecasting performance of the LSTMX model 89
Table III-22: Comparative forecasting performance — LSTM vs LSTMX 91
Table III-23: Best LSTM-family model by forecasting horizon 91
Table III-24: Forecasting performance of the GRU model 91
Table III-25: Forecasting performance of the GRUX model 93
Table III-26: Comparative forecasting performance — GRU vs GRUX 94
Table III-27: Best GRU-family model by forecasting horizon 94
Table III-28: Comparative forecasting performance of the deep learning models 95
Table III-29: Best deep learning model by forecasting horizon 95
Table III-30: Winners from the classical and deep learning models 97
Table III-31: Final comparison at the daily horizon — SARIMA vs GRUX 97
Table III-32: Final comparison at the weekly horizon — SARIMAX vs CNN 98
Table III-33: Final comparison at the monthly horizon — SARIMAX vs CNN 99
Table III-34: Recommended forecasting model for the Bank of Algeria — by operational horizon 99

List of Figures

Figure I-1: Evolution of Year-End Banking Liquidity in Algeria (2014–2022) 23
Figure II-1: Box-Jenkins Methodology 29
Figure II-2: SARIMA Seasonal Decomposition Example 31
Figure II-3: SARIMAX Framework Seasonal ARIMA Structure with Exogenous Variable Integration 32
Figure II-4: Feedforward Neural Network Architecture with one hidden layer (with 3 neurons) 36
Figure II-5: CNN 1D Convolution Process for Time Series Filter Application and Feature Extraction 37
Figure II-6: LSTM Cell Gate Architecture Memory Cell, Forget, Input, and Output Gates 38
Figure II-7: GRU Cell Gate Architecture Reset Gate, Update Gate, and Candidate State 39
Figure III-1 : Methodological Framework for Multi-Horizon Liquidity Forecasting 54
Figure III-2 Bank of Algeria Net Liquidity Position (2015-2022) 56
Figure III-3 Distribution of Daily Liquidity by Year (Boxplot) 57
Figure III-4: Brent Crude Oil Price and USD/DZD Exchange Rate (2015-2022) 58
Figure III-5 : Liquidity vs Brent Crude Oil Price 60
Figure III-6 Liquidity vs USD/DZD Exchange Rate 60
Figure III-7 : Liquidity Distribution by Seasonal/Holiday Dummy Variables 61
Figure III-8 : Correlation Matrix of Liquidity and Exogenous Factors 62
Figure III-9: First Difference of Liquidity (Delta Liquidity, 2015-2022) 63
Figure III-10: Histogram with Normal Fit and Q-Q Plot of Liquidity 64
Figure III-11: Net Liquidity Position — Level and First Difference (d = 1), 2015–2022 66
Figure III-12: Additive Seasonal Decomposition of Net Liquidity (s = 22) — Trend, Seasonal, and Residual Components 67
Figure III-13: ACF and PACF of Liquidity — Level (top) and First Difference (bottom) 68
Figure III-14: Residual Diagnostics — SARIMA (3,1,3) (1,1,2) [22] 71
Figure III-15: Rolling Forecasts — Test Set Comparison (h = 1, 5, 20) 73
Figure III-16 Net Liquidity with Dummy Variable Windows (Pre-Eid, Pre-Ramadan, Pre-BTS, Summer) 75
Figure III-17 Residual Diagnostics SARIMAX (2,1,3) (2,1,2) [22] 80
Figure III-18 SARIMAX (2,1,3) (2,1,2) [22] Rolling Forecasts (h = 1, 5, 20) 81
Figure III-19 CNN forecasts versus actual liquidity — 1-step, 5-step and 20-step horizons 85
Figure III-20 : CNNX forecasts versus actual liquidity — 1-step, 5-step and 20-step horizons 86
Figure III-21 LSTM forecasts versus actual liquidity — 1-step, 5-step and 20-step horizons 88
Figure III-22 LSTMX forecasts versus actual liquidity — 1-step, 5-step and 20-step horizons 90
Figure III-23 GRU forecasts versus actual liquidity — 1-step, 5-step and 20-step horizons 92
Figure III-24 GRUX forecasts versus actual liquidity — 1-step, 5-step and 20-step horizons 93

List of Appendices

Appendix 1: Forecasting central bank Liquidity in Algeria: A comparative Analysis of statistical and Deep learning Models 108
Appendix 2: Top 15 SARIMA models ranked by AIC and penelized by BIC and the best out of all of them + Best model AIC & BIC 109
Appendix 3: SARIMA coefficients matrix + Ljung-Box, Jarque-Bera, Heteroskedasticity test results and coefficient significance 109
Appendix 4: Durbin watson test on SARIMA residuals 110
Appendix 5: Stationarity tests on 1st Diff 110
Appendix 6: SARIMA Grid Search function 111
Appendix 7: Stationarity tests on SARIMA residuals 111
Appendix 8: Rolling Forecast function 112
Appendix 9: Rolling Forecast performance summary 112
Appendix 10: SARIMAX coefficients matrix+Ljung-Box, Jarque-Bera, Heteroskedasticity test results and coefficient significance 113
Appendix 11: Best SARIMAX model AIC & BIC 113
Appendix 12: CNN model architechture 114
Appendix 13: CNN Model parameters 114
Appendix 14: CNN RMSE & R-squared 1step 115
Appendix 15: GRU RMSE & R-squared 5step 115
Appendix 16: GRUX RMSE & R-squared 20step 116
Appendix 17: : LSTM RMSE & R-squared 1step 116
Appendix 18: LSTMX RMSE & R-squared 20steps 117
Appendix 19: Data Preparation Using MinMax Scaling for Deep Learning Models 117
Appendix 20: GRU Model architechture 118
Appendix 21: Net liquidity Graph with Calendar dummies 119

List of Abbreviations

Abbreviation	Full Meaning
ACF	Autocorrelation Function
ADF	Augmented Dickey-Fuller (test)
AEN	Avoirs Extérieurs Nets / Net Foreign Assets (NFA)
AGG	Aggregate Contribution of Autonomous Factors
AI	Artificial Intelligence
AIC	Akaike Information Criterion
AR	Autoregressive (component)
ARCH	Autoregressive Conditional Heteroscedasticity
ARIMA	Autoregressive Integrated Moving Average
ARMA	Autoregressive Moving Average
BCBS	Basel Committee on Banking Supervision
BIC	Bayesian Information Criterion
BIS	Bank for International Settlements
BTS	Back-to-School (seasonal dummy variable)
CAMELS	Capital, Assets, Management, Earnings, Liquidity, Sensitivity (banking rating framework)
CFPRP	Coefficient of Capital and Permanent Resources (Coefficient de Fonds Propres et Ressources Permanentes)
CIC	Currency in Circulation
CNN	Convolutional Neural Network
COVID-19	Coronavirus Disease 2019
CV	Coefficient of Variation
DL	Deep Learning
DM	Diebold-Mariano (test)
DMMF	Direction du Marché Monétaire et Financier (Monetary and Financial Market Division)
DZD	Algerian Dinar
ECB	European Central Bank

Abbreviation	Full Meaning
FX	Foreign Exchange
GAB	Government Account Balance
GRU	Gated Recurrent Unit
IMF	International Monetary Fund
KPSS	Kwiatkowski-Phillips-Schmidt-Shin (test)
LCR	Liquidity Coverage Ratio
LOLR	Lender of Last Resort
LSTM	Long Short-Term Memory
MA	Moving Average (component)
MAE	Mean Absolute Error
ML	Machine Learning
MLE	Maximum Likelihood Estimation
NFA	Net Foreign Assets
NGP	Net Government Position
NN	Neural Network
NSFR	Net Stable Funding Ratio
OMO	Open Market Operations
OPU	Office des Publications Universitaires (University Publications Office)
PACF	Partial Autocorrelation Function
PP	Phillips-Perron (test)
PSR	Programme Spécial de Refinancement / Special Refinancing Programme
QE	Quantitative Easing
RMSE	Root Mean Square Error
SARIMA	Seasonal Autoregressive Integrated Moving Average
SARIMAX	Seasonal Autoregressive Integrated Moving Average with Exogenous Variables
SARMA	Seasonal Autoregressive Moving Average
SGD	Stochastic Gradient Descent
SNB	Système de Notation Bancaire / Banking

Abbreviation	Full Meaning
	Rating System
USD	United States Dollar
VAR	Vector Autoregression
WN	White Noise

Abstract

In emerging, hydrocarbon-dependent economies like Algeria, managing banking liquidity is a persistent challenge due to external shocks and structural breaks. The Bank of Algeria (BoA) must accurately forecast aggregate liquidity to calibrate open market operations, set reserve requirements, and safeguard financial stability. Traditional statistical models, such as ARIMA, SARIMA, and SARIMAX, have long served as the backbone of central bank liquidity forecasting due to their transparency and ability to capture seasonal patterns. However, these linear approaches frequently struggle with the non-linear dynamics and complex interactions characteristic of Algerian liquidity data, particularly during periods of external shocks like the 2014 oil crisis and the COVID-19 pandemic.

This dissertation addresses this gap by comparing the performance of two traditional models (SARIMA and SARIMAX) against three deep learning architectures (CNN, LSTM, and GRU) in forecasting aggregate central bank liquidity for the Bank of Algeria. The study evaluates forecasting accuracy across three operational horizons—day-ahead, week-ahead, and month-ahead—and assesses the marginal value of incorporating exogenous variables (Brent crude oil price, USD/DZD exchange rate, and holiday/seasonal dummies). The findings confirm that deep learning models, particularly LSTM and GRU, significantly outperform traditional statistical models by better capturing non-linear dynamics and structural breaks. Furthermore, the inclusion of exogenous variables related to hydrocarbon prices and exchange rates substantially reduces forecasting errors. Ultimately, this study provides both a methodological contribution to the literature and practical, operational guidance for the Bank of Algeria's transition toward proactive, AI-driven liquidity management.

Keywords: Central bank liquidity, multi-horizon forecasting, Deep learning, Time series forecasting

الملخص

في الاقتصاديات الناشئة المعتمدة على المحروقات مثل الجزائر، تُعد إدارة السيولة المصرفية تحديًا مستمرًا بسبب الصدمات الخارجية والكسور الهيكلية. يجب على بنك الجزائر التنبؤ بالسيولة الإجمالية بدقة لمعايرة عمليات السوق المفتوحة،

وتحديد متطلبات الاحتياطي، وصون الاستقرار المالي. طالما شكلت النماذج الإحصائية التقليدية، مثل ARIMA وSARIMA وSARIMAX، الركيزة الأساسية للتنبؤ بالسيولة في البنوك المركزية نظرًا لشفافيتها وقدرتها على التقاط الأنماط الموسمية. ومع ذلك، غالبًا ما تواجه هذه المناهج الخطية صعوبات مع الديناميكيات غير الخطية والتفاعلات المعقدة التي تميز بيانات السيولة الجزائرية، لا سيما خلال فترات الصدمات الخارجية مثل أزمة النفط عام 2014 وجائحة كوفيد-19.

تسد هذه الأطروحة هذه الفجوة من خلال مقارنة أداء نموذجين تقليديين (SARIMA وSARIMAX) مقابل ثلاث بنيات للتعلم العميق (CNN وLSTM وGRU) في التنبؤ بالسيولة الإجمالية للبنك المركزي في بنك الجزائر. تقيّم الدراسة دقة التنبؤ عبر ثلاثة آفاق تشغيلية: ليوم مقدم، وأسبوع مقدم، وشهر مقدم، وتقيّم القيمة الحدية لإدماج المتغيرات الخارجية (سعر خام برنت، وسعر صرف الدولار/الدينار، والمتغيرات الوهمية الموسمية). تؤكد النتائج أن نماذج التعلم العميق، وخاصة LSTM وGRU، تتفوق بشكل كبير على النماذج الإحصائية التقليدية من خلال التقاط الديناميكيات غير الخطية والكسور الهيكلية بشكل أفضل. علاوة على ذلك، فإن إدراج المتغيرات الخارجية المتعلقة بأسعار المحروقات وأسعار الصرف يقلل بشكل كبير من أخطاء التنبؤ. في النهاية، تقدم هذه الدراسة مساهمة منهجية في الأدبيات وتوجيهات عملية وتشغيلية لبنك الجزائر للانتقال نحو إدارة سيولة استباقية مدعومة بالذكاء الاصطناعي.

الكلمات المفتاحية: سيولة البنك المركزي، التنبؤ متعدد الآفاق، التعلم العميق، التنبؤ بالسلاسل الزمنية

General Introduction

The banking system plays a foundational role in economic development by channeling funds between surplus and deficit agents. Through the mechanism of maturity transformation, banks convert short-term deposits into long-term loans, a process that inherently exposes institutions to liquidity risk widely considered the most critical of all banking risks. The global financial crisis of 2007–2008 clearly demonstrated how localized liquidity shortages within a single institution can rapidly trigger a chain reaction, threatening the stability of the entire macroeconomic landscape. At the center of this complex financial architecture lies the central bank, acting as the ultimate liquidity provider and the primary safeguard of systemic equilibrium. By regulating the volume of bank reserves, the central bank steers short-term interest rates and ensures that commercial institutions can meet their settlement obligations.

For a central bank to effectively manage systemic liquidity and conduct proactive monetary policy, accurate forecasting of the aggregate liquidity position is an absolute prerequisite. Central banks must reliably predict the near-term path of commercial bank reserves and the autonomous factors that affect their balance sheets in order to successfully calibrate open market operations and set appropriate reserve requirements. Traditionally, central banks have relied on classical statistical models such as the Autoregressive Integrated Moving Average (ARIMA) and its seasonal extensions (SARIMA and SARIMAX) which have long served as the backbone of liquidity forecasting. These traditional frameworks are highly prized for their transparency, interpretability, and ability to explicitly model predictable seasonal patterns.

However, the modern financial environment is increasingly characterized by non-linear dynamics, structural breaks, and complex interactions that challenge the foundational assumptions of linear statistical models. In emerging, hydrocarbon-dependent economies like Algeria, these forecasting challenges are profoundly magnified. The Bank of Algeria (BoA) faces persistent hurdles in managing banking liquidity due to the country's high exposure to external shocks. Over the past decade, the Algerian banking system has navigated extreme cycles of liquidity surpluses and deficits. These structural breaks were primarily driven by the 2014 collapse in global oil prices and the unprecedented economic disruptions caused by the COVID-19 pandemic.

During these periods of severe volatility, traditional forecasting approaches have frequently struggled to adapt. Net central bank liquidity in Algeria is heavily influenced by a confluence of domestic and international exogenous factors: fluctuating Brent crude oil prices directly impact government revenues; shifts in the USD/DZD exchange rate alter the valuation

of foreign reserves; and distinct calendar events such as Ramadan, Eid festivities, and summer cycles trigger sudden spikes in the public's demand for cash. Because traditional models fundamentally assume that future values represent a linear combination of lagged observations, they fail to adequately capture the asymmetric responses, threshold effects, and multiplicative interactions inherent in these external shocks.

To overcome these structural limitations, recent advancements in deep learning offer powerful and transformative alternatives. Deep learning architectures specifically Long Short-Term Memory (LSTM) networks, Gated Recurrent Units (GRU), and Convolutional Neural Networks (CNN)possess the ability to automatically extract complex, non-linear patterns and learn long-term temporal dependencies directly from sequential data. By utilizing hierarchical feature extraction, these models can adapt to sudden regime changes and structural breaks without requiring the manual specification of interactions. Despite their growing success in forecasting stock prices, exchange rates, and commodity markets, their empirical application to predicting aggregate central bank liquidity particularly within the context of an oil-exporting emerging market remains highly limited.

This dissertation addresses this critical gap by conducting a rigorous, multi-horizon comparison of traditional seasonal models (SARIMA and SARIMAX) against modern deep learning architectures (CNN, LSTM, and GRU) to forecast aggregate central bank liquidity for the Bank of Algeria. The study evaluates forecasting accuracy across three operational horizons day-ahead, week-ahead, and month-ahead and assesses the marginal value of incorporating exogenous variables into these predictive frameworks.

In this context, the central research question guiding this work is:

“What forecasting approach provides the highest accuracy for predicting central bank liquidity in Algeria across one-day-ahead, one-week-ahead, and one-month-ahead horizons, and what role do exogenous factors play in improving predictive performance?”

This main question is broken down into the following specific sub-questions:

1. Can deep learning models uncover and exploit complex, nonlinear patterns in Algerian central bank liquidity that traditional linear models consistently miss when relying solely on the series' own historical values?
2. What is the marginal contribution of exogenous variables (Brent crude oil price, USD/DZD exchange rate, and holiday/seasonal dummies) when added to both traditional and deep learning models for forecasting aggregate central bank liquidity?
3. What is the most effective overall approach for the Bank of Algeria: traditional models, deep learning models, or a combination of both?

To address the research problem, the following hypotheses are proposed and tested:

- **H1:** Deep learning models (LSTM and GRU in particular) significantly outperform traditional statistical models (SARIMA) by better capturing the non-linear dynamics and structural breaks inherent in Algerian central bank liquidity when using only the series' own history.
- **H2:** The inclusion of exogenous variables related to hydrocarbon prices USD/DZD exchange rates and holidays seasons substantially reduces forecasting errors.
- **H3:** Deep learning models significantly outperform traditional statistical models in forecasting accuracy across all horizons.

The originality of this study lies in delivering a multi-horizon comparison of seasonal traditional models and modern deep learning architectures for aggregate central bank liquidity forecasting in an oil-exporting emerging economy, while rigorously quantifying the added value of exogenous factors. It provides both a methodological contribution to the academic literature and practical, operational guidance to assist the Bank of Algeria in its transition toward proactive, AI-assisted liquidity management.

To answer the research questions and test the hypotheses, this dissertation adopts a rigorous structure divided into three chapters:

Chapter I: Theoretical Foundations of Banking Liquidity and the Algerian Context — This chapter establishes the foundational concepts of banking liquidity and liquidity risk management. It explores the systemic importance of the central bank as the ultimate liquidity provider and provides a historical analysis of the Algerian banking context,

specifically detailing how the 2014 oil shock and the COVID-19 pandemic shaped the current forecasting challenges.

Chapter II: Theoretical Framework and Literature Review: From Traditional Forecasting to Deep Learning for Central Bank Liquidity Prediction — This chapter outlines the theoretical mechanics underlying the forecasting models. It reviews the traditional Box-Jenkins methodology (ARIMA, SARIMA, SARIMAX) and discusses their linear limitations. It then introduces the mathematical foundations of the deep learning architectures (CNN, LSTM, GRU) utilized in this study and concludes with a comprehensive review of the comparative literature regarding deep learning applications in financial time series.

Chapter III: Empirical Study of Central Bank Liquidity Forecasting Using Statistical and Deep Learning Models — The final chapter presents the empirical core of the dissertation. It details the construction and descriptive analysis of the daily dataset spanning 2015 to 2022. It sequentially details the estimation, residual diagnostics, and forecasting performance of the classical models, followed by the deep learning models. The chapter concludes with a final benchmarking across the day-ahead, week-ahead, and month-ahead operational horizons to provide definitive recommendations for the Bank of Algeria.

Chapter I: Theoretical Foundations of Banking Liquidity and the Algerian Context

Introduction

The banking system plays a central role in economic development by channeling funds between surplus and deficit agents. Through the process of maturity transformation, banks are inherently exposed to liquidity risk, which is widely considered the most critical of all banking vulnerabilities. The 2007–2008 global financial crisis starkly highlighted how sudden liquidity shortages in a single institution can trigger a contagion effect, ultimately threatening the stability of the entire macroeconomic landscape.

The primary objective of this chapter is to establish the theoretical and conceptual foundations of banking liquidity, and to critically examine the specific challenges of managing and forecasting it within the unique economic context of Algeria. A thorough understanding of these theoretical elements is crucial before advancing to the complex predictive modeling undertaken later in this study. By defining the mechanisms of liquidity risk and analyzing the historical impact of external shocks such as oil price volatility and the COVID-19 pandemic on the Bank of Algeria, we can better understand the urgent need for advanced forecasting tools.

Plan of the Chapter To address these objectives, the chapter is structured into three main sections:

- **Section 1**: Explores the fundamental concept of banking liquidity, its nature, and its different typologies within the financial system.
- **Section 2:** Examines liquidity risk, its systemic consequences, and the traditional management frameworks employed by central banks as ultimate liquidity providers.
- **Section 3**: Analyses the specific banking liquidity context in Algeria, detailing the historical evolution of central bank liquidity and the unique forecasting challenges faced by the Bank of Algeria.

Section 1: The Concept of Banking Liquidity

Liquidity is a multifaceted and dynamic concept in economics, generally referring to the ability of an entity to meet its financial obligations by converting assets into cash. Within the banking sector, it is viewed as a vital flow essential to the functioning of central banks, commercial institutions, and the broader financial market. The smooth movement of this flow is frequently hindered by market imperfections such as information asymmetry and coordination failures, which can lead to periods of illiquidity that jeopardise institutional survival (Nikolaou, 2009).

1. Definitions and Nature of Liquidity

Before examining the different types of liquidity, it is useful to clarify the meaning of the term itself and the central role that banks play in providing it. This preliminary discussion fixes the vocabulary used throughout the chapter and highlights why liquidity is at the same time a precondition for and a source of fragility within the banking system.

1.1. Defining Liquidity in the Economic and Banking Sense

In a broad economic sense, liquidity refers to available money or the ability to rapidly convert assets into monetary form without significant loss in value (Lakehal, 2002). From a banking perspective, a bank is considered liquid when it can meet its obligations as they fall due without unacceptable losses (Basel Committee on Banking Supervision, 2008); illiquidity arises when a cash-flow deficiency prevents it from honouring disbursements through current inflows.

Liquidity is fundamentally linked to money, the ultimate liquid asset, which improved economic efficiency by resolving the double-coincidence-of-wants problem and reducing transaction costs (Mishkin, 2013). Banks therefore act as liquidity providers and financial intermediaries, channelling funds from surplus to deficit units. From a technical perspective, bank liquidity includes central bank money, cash on hand, and interbank balances resources Béziade (1992) emphasises as essential because they are immediately available to satisfy obligations. Liquidity management also involves structuring assets and liabilities through highly liquid instruments, such as short-term government securities, to ensure funds are available when needed.

1.2. The Role of Banks as Liquidity Providers

The primary function of banks within the financial system is to provide liquidity by channelling deposits into long-term investments. This process, known as maturity transformation, allows banks to transform liquid resources (short-term deposits) into illiquid assets (long-term loans). By collecting liquidity from depositors and allocating it to investors, banks resolve the mismatch between savers' preferences for immediate access to funds and borrowers' need for long-term financing. However, this very process introduces a fundamental fragility into the banking system (Diamond & Rajan, 2001).

When the process of maturity transformation is disrupted, the bank's role as a liquidity provider can become a source of instability. This inherent fragility exposes banks to funding liquidity risk, where structural weaknesses combined with market inefficiencies and asymmetric information can lead to liquidity shortages. In extreme cases, these conditions may lead to bank runs, which are considered the most severe manifestation of funding liquidity risk. A bank run occurs when depositors, fearing the insolvency of the institution, rush to withdraw their funds simultaneously, creating a liquidity crisis that the bank cannot withstand even if it is technically solvent (Diamond & Dybvig, 1983).

2. The Typology of Liquidity

Academic literature typically classifies liquidity within a financial system into three distinct yet interrelated categories: central bank liquidity, funding liquidity, and market liquidity (Nikolaou, 2009). While some researchers advocate for a two-dimensional view focusing solely on market and funding liquidity due to the central bank's role as an external crisis mitigator the three-node model remains the standard framework for understanding systemic interactions.

2.1. Central Bank Liquidity

Central bank liquidity refers to the monetary authority’s capacity to supply liquidity to the financial system through the monetary base (M0), including currency in circulation and bank reserves. It is closely linked to central bank operational liquidity, namely the funds injected into the money market through monetary policy operations (Nikolaou, 2009).

The central bank manages liquidity according to its monetary policy stance, usually defined by a target policy rate. Through liquidity management instruments and interbank market operations, it ensures short-term interest rates remain aligned with this target. As lender of last resort, the central bank can provide liquidity in unlimited quantities to preserve financial stability (Bagehot, 1873).

2.2. Funding Liquidity

According to the Basel Committee on Banking Supervision (2008), funding liquidity is defined as the ability of banks to meet their liabilities and to unwind or settle their positions as they come due. It essentially reflects a bank's capacity to meet its financial obligations on time and to obtain external financing when necessary to manage withdrawal risk. Funding liquidity ensures that a bank can mobilise resources to cover short-term and medium-term commitments while maintaining overall financial stability.

Drehmann and Nikolaou (2009) argue that funding liquidity is a flow concept rather than a static stock. A bank is considered to have sufficient funding liquidity as long as its cash outflows remain less than or equal to the sum of its cash inflows and its existing stock of money. Furthermore, Nikolaou (2009) stresses that funding liquidity is the primary channel through which liquidity is distributed throughout the entire financial system, making it essential for banks to maintain adequate liquidity levels at all times to prevent systemic bottlenecks.

2.3. Market Liquidity

Market liquidity reflects the ability of a bank to buy or sell financial assets quickly, at low cost, and with minimal price impact. According to Crockett (2008), a highly liquid market ensures that transactions occur smoothly, maintaining market efficiency and financial stability. It is characterised by four key dimensions:

- **Immediacy:** the speed at which transactions can be executed.
- **Breadth:** the cost of obtaining liquidity, often measured by the bid-ask spread.
- **Depth:** the ability to execute large transactions without significantly affecting prices.
- **Resilience:** the speed at which prices return to their intrinsic level after market disruptions.

There are two primary types of market liquidity. First, interbank market liquidity ensures system stability by allowing banks to lend and borrow short-term funds from one another, facilitating daily operations. Second, asset market liquidity enables investors to trade securities without causing significant price fluctuations. While market liquidity facilitates the efficient transfer of funds between economic agents, it does not increase the overall money supply; only central banks possess the authority to inject additional liquidity into the system by adjusting monetary policy tools. In some stock exchanges, market makers play a critical role in providing this liquidity by standing ready to buy or sell significant quantities of a security quickly and anonymously.

3. The Interconnection: Virtuous and Vicious Cycles

The three forms of liquidity central bank, market, and funding are intricately interconnected, forming a dynamic system in which stability or instability in one node quickly propagates to the others. These relationships are best analysed through two contrasting scenarios: normal periods and turbulent periods.

3.1. Liquidity Interconnections in Normal Periods

In normal circumstances, liquidity flows smoothly between the three types, forming a self-sustaining cycle: the central bank injects liquidity to offset shortfalls, banks redistribute it through interbank and asset markets, and the central bank then reassesses demand and adjusts supply accordingly (European Central Bank, 2004). Each type serves a distinct role central bank provides neutral liquidity, markets facilitate circulation, and funding liquidity ensures efficient resource allocation. Banks remain liquid as long as they can access multiple funding sources at reasonable cost (Drehmann & Nikolaou, 2009), and this efficient circulation helps mitigate financial vulnerabilities and supports overall economic health.

3.2. Liquidity Interconnections in Unstable Periods

During periods of turbulence, the virtuous cycle can reverse into a vicious circle of heightened liquidity risk. Financial markets are often plagued by inefficiencies and asymmetric information, which can lead to coordination failures among banks, depositors, and market participants. This process triggers a liquidity spiral in which uncertainty regarding asset values pushes prices below their intrinsic levels, threatening the stability of the entire financial system.

Nikolaou (2009) highlights that liquidity risk is endogenous: the stability observed in normal conditions can transform into instability during crises. In such times, the three forms of liquidity act as channels for risk propagation. For example, a bank facing considerable funding liquidity risk may be forced to sell assets in the market. If multiple banks do so simultaneously, asset markets may seize up due to limited absorption capacity, leading to a liquidity freeze. Brunnermeier and Pedersen (2009) demonstrate that market liquidity depends on traders' ability to secure financing, which in turn is restricted when markets become illiquid, creating a mutually reinforcing downward spiral.

4. The Importance of Liquidity for Financial Stability

Liquidity is universally recognised as the cornerstone of financial stability, influencing both the health of individual institutions and the overall resilience of the national economy. Disruptions in the smooth flow of liquidity can lead to market turbulence, a loss of public confidence, and the potential for a full-scale systemic crisis.

4.1. Systemic Risk and Contagion

Liquidity risk possesses a uniquely systemic character. If a liquidity problem is not resolved rapidly, an individual bank may find itself in a state of bankruptcy. If that institution's weight in the system is significant meaning other banks depend on it for their own liquidity needs its failure can trigger a chain reaction of collapses across the interbank market (Siliadin, 2013). This propagation of difficulties can destabilise the entire financial system, affecting stocks, bonds, and equities alike. Because of this risk, market liquidity risk is often priced into assets as a cost or premium, which influences transaction costs and optimal portfolio allocation.

4.2. Liquidity and Monetary Policy Implementation

For the central bank, forecasting liquidity is essential for effective monetary policy implementation. Accurate forecasts allow the institution to calibrate open market operations (OMOs), aligning market liquidity with the policy stance and controlling refinancing costs. Conversely, poor liquidity management can trigger undesirable spikes in interbank rates, disrupting credit flows and undermining market confidence (IMF, 2024).

Due to post-2008 financial innovation and evolving regulatory frameworks, liquidity forecasting requires advanced predictive models capable of anticipating pressures with greater precision. Ultimately, maintaining adequate liquidity is an economic necessity that allows the Bank of Algeria to safeguard national banking health and support sustainable growth.

Section 2: Liquidity Risk and its Management

Having defined banking liquidity and its typology, we must examine how the financial system defends against its sudden depletion. Because liquidity risk is both an individual and systemic phenomenon capable of contaminating the system via interbank exposures, fire sales, and confidence effects central banks act as ultimate guarantors. This section examines liquidity risk, central bank intervention, containment frameworks, and the indispensable role of forecasting.

1. The Role of the Central Bank as the Ultimate Liquidity Provider

The central bank occupies a unique and sovereign position within the modern financial architecture as the monopoly provider of high-powered money, or the monetary base M_0. This exclusive authority grants the institution the singular capacity to satisfy the banking system's demand for central bank money, acting as the pivot for all financial transactions. As the ultimate liquidity provider, it ensures commercial institutions can settle end-of-day obligations and meet mandatory reserve requirements, safeguarding the payment and settlement systems.

This fundamental role is operationalized through the monetary policy stance, where the central bank strategically adjusts reserve supply to steer short-term interbank interest rates toward a predefined target. While a bank is liquid if it can finance asset growth and meet obligations without unacceptable losses (BCBS, 2008), market imperfections like information asymmetry can block capital flows, necessitating central bank intervention to prevent localized cash-flow deficiencies from escalating into systemic crises.

This intervention frequently manifests as the lender-of-last-resort (LOLR) function, historically rooted in Bagehot's (1873) principles. These guidelines dictate lending freely to solvent institutions against high-quality collateral at a penalty rate, mitigating moral hazard while bridging acute funding shortages. Unlike commercial entities, the central bank's liquidity risk is negligible because its monopoly on base money allows it to generate liquidity at will to maintain systemic equilibrium.

In the specific context of Algeria, the missions of the Bank of Algeria are explicitly anchored in ensuring monetary and price stability as the primary objective of monetary policy. The Monetary and Banking Law of 2023 further refines this mandate by emphasising that, while permanent monetary financing of the State is strictly prohibited to protect the institution's independence and anchor inflation expectations, exceptional and temporary financing mechanisms may be authorised to respond to major shocks or acute financing constraints (Bank

of Algeria, 2023). Through the calibration of open market operations and the provision of standing facilities, the Bank of Algeria acts as the vital anchor that offsets fluctuations in autonomous factors such as currency in circulation or government account balances to ensure that the aggregate liquidity supply remains aligned with the stability needs of the national financial system.

1.1. The Monopoly of Base Money and Price Signalling

The central bank's primary power lies in its ability to lend any quantity of liquidity at a price (the interest rate) that it determines autonomously. In the context of liquidity management, this process involves regulating the volume of bank reserves in order to steer short-term interest rates toward the monetary policy stance. This stance is typically anchored by an operational target such as a key policy rate which serves as the primary control variable for the entire financial system. When the central bank adjusts the supply of reserves through its balance sheet, it influences the cost of interbank funds, which subsequently impacts broader economic variables such as aggregate demand and inflation.

1.2. The Lender-of-Last-Resort Function

The historical foundation of the LOLR role, popularised by Bagehot (1873), mandates that central banks should lend freely to solvent institutions, against good collateral, but at a penalty rate. This function is critical during periods of acute institutional distress. When a bank faces a severe funding shortage often due to a loss of market confidence or a flawed business strategy private counterparty may refuse to lend, necessitating central bank intervention. The primary justification for this support is the prevention of contagion: if an institution is systemically important, its failure could trigger a chain reaction of collapses across the interbank market. Consequently, central banks often prioritise systemic stability over the strict distinction between illiquidity and insolvency during crises.

2. Understanding Liquidity Risk

Liquidity risk, or the risk of illiquidity, is defined by Darmon (1998) as the danger of an institution being unable to meet its immediate liabilities with its available liquid assets. In the Algerian context, Regulation No. 11-08 specifies that this risk encompasses the failure to fulfil obligations or adjust a position due to adverse market conditions at an acceptable cost (Bank of Algeria, 2011).

2.1. Systemic Liquidity Shortages

A systemic shortage occurs when a widespread collapse affects both funding and market liquidity simultaneously, leading to a total market breakdown. These shortages are distinct

from individual bank risks because they are driven by collective coordination failures and a general loss of confidence.

- **The Feedback Loop and Liquidity Spirals:** Systemic crises often originate from liquidity spirals, as modelled by Brunnermeier and Pedersen (2009). In this scenario, uncertainty regarding asset values pushes prices below their intrinsic levels. Banks, facing depreciating asset values and increased funding costs, are forced to liquidate more holdings to meet margin requirements, which further drives down prices. This vicious cycle transforms localised shocks into system-wide freezes.
- **The Distributional Challenge:** Not all shortages result from a lack of aggregate liquidity. Some arise from distributional inefficiencies, where certain institutions hold excess reserves while others face acute deficits due to operational failures or a breakdown in the interbank market's trust. For instance, a technical malfunction in a major bank's payment system can prevent the circulation of funds, causing a systemic cash shortage despite the overall health of the banking sector.

2.2. The Consequences of Illiquidity on the Banking Sector

The primary consequence of illiquidity is the disruption of the maturity-transformation process, in which banks convert short-term deposits into long-term loans.

- **Maturity Mismatches:** Because a bank's assets (loans) are generally less liquid and have longer maturities than its liabilities (deposits), any sudden surge in withdrawals can lead to a cash-flow deficiency. This mismatch is the fundamental origin of illiquidity risk.
- **Bank Runs and Contagion:** Diamond and Dybvig (1983) highlight that even technically solvent banks can collapse if depositors, fearing insolvency, rush to withdraw their funds simultaneously. If the bank in question is of significant systemic weight, its failure forces other institutions which depend on it for interbank liquidity into similar distress.
- **Impact on Interest Rates:** At the macroeconomic level, a shortage of reserves causes overnight interest rates to soar as banks scramble for funding to avoid penalties for unmet reserve requirements. Conversely, a surplus of liquidity can drive short-term rates toward zero, potentially encouraging banks to engage in riskier lending in search of yield, which threatens long-term financial stability.

3. Liquidity Management Framework

To maintain financial equilibrium and safeguard the integrity of the payment system, central banks and regulatory authorities employ a structured framework of instruments to regulate reserves and manage risks. In Algeria, the Bank of Algeria acts as the pivot of this system, using tools that align with international standards to manage both the global volume of reserves and the soundness of individual credit institutions (Bank of Algeria, 2023).

3.1. Conventional Monetary Policy Instruments

Conventional instruments allow the central bank to influence banking liquidity and steer short-term interest rates without directly imposing credit volumes on commercial institutions.

3.1.1. Open Market Operations (OMO):

Open market operations are the primary tool for regulating money market liquidity and maintaining the intended monetary policy stance. The central bank provides funds by purchasing eligible securities to meet banking needs, or absorbs excess liquidity by selling securities or accepting unsecured deposits. These operations involve eligible securities with maturities ranging from seven days to twelve months and include both temporary sales (repurchase agreements) and outright transactions.

3.1.2. Required Reserves:

Central banks mandate that commercial banks maintain a minimum balance at the central bank proportional to their demand and short-term deposits. By adjusting the reserve requirement ratio, the central bank can freeze or release systemic liquidity, thereby reinforcing banks' dependence on central bank refinancing. Some regimes utilise reserve averaging, in which banks must meet their target over a maintenance period (e.g., one month). This provides flexibility for banks to manage daily cash-flow shocks without causing extreme interest-rate volatility.

3.1.3. Standing Facilities:

These are safety valves initiated by commercial banks rather than the central bank to manage end-of-day liquidity positions. The marginal lending facility allows banks to obtain very short-term (24-hour) liquidity against eligible collateral to cover urgent payment obligations. The overnight deposit facility enables banks to place daily surplus liquidity at a predetermined interest rate, which typically serves as a floor for market rates.

3.1.4. Rediscounting:

Rediscounting is a refinancing instrument whereby the central bank grants loans to banks against high-quality private or public securities held in their portfolios. The rediscount

rate serves as the official refinancing cost and as a benchmark for the cost of credit in the broader economy. Its effectiveness is reduced when the banking system holds significant excess liquidity.

3.2. Unconventional Monetary Policy Instruments

When traditional channels such as interest-rate cuts reach their limits (often during severe crises), central banks adopt large-scale, unprecedented interventions:

- **Forward Guidance:** communicating future policy intentions regarding interest rates (time-based or state-based) to reduce uncertainty and stabilise market expectations.
- **Quantitative Easing (QE):** purchasing financial assets (primarily government bonds) to create central bank reserves and lower interest rates across the entire yield curve.
- **Negative Interest Rates:** charging interest on excess reserves held at the central bank to discourage cash hoarding and encourage lending.

3.3. Prudential and Regulatory Management Tools

In addition to monetary operations, liquidity is managed through regulatory requirements and internal banking procedures designed to ensure that institutions can survive both short-term and structural shocks.

3.3.1. Basel III Prudential Ratios:

The Bank of Algeria has integrated international standards to enhance banking resilience. The Liquidity Coverage Ratio (LCR) measures a bank's capacity to meet net cash outflows over a 30-day stress period using high-quality liquid assets. The Net Stable Funding Ratio (NSFR) is a long-term ratio designed to avoid maturity mismatches between a bank's assets and liabilities. Finally, the Coefficient of Capital and Permanent Resources (CFPRP) limit the financing of long-term assets with short-term resources to ensure structural stability (Basel Committee on Banking Supervision, 2013).

3.3.2. Internal Risk Management Techniques:

Individual banks employ several internal methods to cover liquidity risk. Matching pairs an asset item with a liability item of similar liquidity, rate, and currency characteristics. Securitisation transforms non-liquid assets such as loans into negotiable financial securities that can be sold on the market to raise cash. Internal oversight further includes the use of stress tests to evaluate capacity under various economic shocks and emergency funding plans to define procedures for acute liquidity needs.

3.4. The Role of Forecasting in Effective Liquidity Management

The capacity of a central bank to effectively manage systemic liquidity is fundamentally predicated on its ability to produce accurate, timely, and actionable forecasts. Liquidity forecasting is defined as the centralised process of estimating the near-term path of commercial bank reserves and of the various autonomous factors that independently affect the central bank's balance sheet (International Monetary Fund, 2024; Gray, 2008). Without these projections, central bank monetary operations would be reduced to a reactive guessing game regarding the necessary volume of interventions, potentially leading to severe misalignments between market rates and the intended policy stance (Bindseil, 2000; Bank of England, 2008).

3.4.1. Forecasting as a Prerequisite for Proactive Monetary Policy

The primary objective of liquidity forecasting is to provide the operational framework necessary to calibrate open market operations. By identifying the projected gap between the supply of and demand for reserves, the central bank can determine whether to inject or absorb liquidity in order to maintain the short-term interbank interest rate often the operational target of monetary policy at its desired level (Bindseil, 2000; International Monetary Fund, 2024). Accurate forecasting therefore serves as a bridge between day-to-day operations and the broader economy, influencing the yield curve and ensuring that the transmission mechanism of monetary policy remains functional (Bhattacharyya & Sahoo, 2011).

Furthermore, accurate predictions are essential to prevent extreme volatility in interbank rates. When a shortage of reserves occurs unexpectedly, interbank refinancing rates tend to soar as commercial banks scramble for funding to meet mandatory reserve requirements or settlement obligations. Conversely, an unmanaged surplus of liquidity can drive rates toward zero, flattening the yield curve and potentially encouraging excessive risk-taking within the banking sector (International Monetary Fund, 2024). Consequently, forecasting is not merely a technical exercise but a vital stabilising mechanism for the entire financial system.

3.4.2. The Structural Mechanics of Autonomous Factor Projection

Central bank liquidity management focuses heavily on autonomous factors items on the balance sheet whose fluctuations are not initiated by the monetary authority. The three primary components of this forecasting exercise are currency in circulation (CIC), net foreign assets (NFA), and the government account balance, sometimes referred to as the net government position (NGP) (International Monetary Fund, 2024; Gray, 2008).

To operationalise the forecasting framework, it is necessary to define the primary autonomous factors that constitute the aggregate liquidity supply. Currency in Circulation (CIC) refers to the total volume of banknotes and coins held by the public; as a central bank

liability, an increase in CIC results in a drainage of bank reserves (International Monetary Fund, 2024). Net Foreign Assets (NFA) represent the central bank's holdings of foreign currency and gold, which typically inject liquidity into the system through foreign exchange interventions (International Monetary Fund, 2024). Finally, the Net Government Position (NGP) often termed the Government Account Balance (GAB) represents the net balance of government deposits held at the central bank (International Monetary Fund, 2024).

Each factor presents unique challenges:

- **Currency in circulation (CIC):** generally considered the easiest to forecast due to its strong adherence to seasonal patterns, holiday cycles, and payroll dates (El Hamiani-Khatat, 2018).
- **Net foreign assets (NFA):** often influenced by unpredictable external shocks and foreign-exchange interventions, particularly in fixed exchange rate regimes where the central bank must stabilise currency values (Gray, 2008).
- **Net government position (NGP):** typically, the most volatile and difficult factor to project, as it depends on the timing of tax collections, government expenditures, and debt service (Bank of England, 2008).

The forecasting framework must integrate these disparate elements through a unified identity, in which the projected change in bank reserves is equal to the net change in autonomous factors plus the impact of central bank operations (International Monetary Fund, 2024).

To maintain technical equilibrium, these primary components are integrated into a unified identity representing the aggregate contribution (AGG) of autonomous factors to the supply of bank reserves (International Monetary Fund, 2024)

. Drawing on the framework established by the International Monetary Fund (2024), the net liquidity position is defined by the identity:

$$\mathbf{AGG = NFA - CIC - GAB + Net\ Other\ Items}$$

While including 'Net Other Items' ensures balance-sheet reconciliation, this research focuses on the primary drivers of volatility: NFA, CIC, and GAB (International Monetary Fund, 2024). Adopting a direct aggregate forecasting approach a method the IMF identifies as desirable for a comprehensive framework allows deep learning architectures (CNN, LSTM, GRU) to capture the non-linear correlations and structural breaks typical of central bank data (International Monetary Fund, 2024). This strategy minimizes reconciliation errors from

volatile individual factors, providing the Bank of Algeria with the precise data needed to effectively calibrate Open Market Operations (OMO) (International Monetary Fund, 2024).

3.4.3. Horizon Alignment and Operational Efficiency

Effective management requires that the forecasting horizon be strictly aligned with the schedule of monetary operations; for instance, weekly OMOs necessitate a seven-day lead time for early-week operational decisions (Gray, 2008). Because forecasting errors naturally increase as the horizon extends, a rolling forecast approach repeatedly re-estimating and updating models as daily data becomes available is essential (IMF, 2024; Taoussi et al., 2024).

The length of the reserve maintenance period also dictates the required precision of the forecast. Some central banks utilise reserve averaging, which allows commercial banks flexibility in meeting their requirements over a period (e.g., one month). This mechanism can partially buffer the impact of forecasting errors, as the central bank can adjust its operations toward the end of the period to correct for earlier miscalculations (Bank of England, 2008).

3.4.4. Risk Communication and Market Stability

Finally, modern liquidity management highlights the importance of transparency through the publication of forecasts and predictive distributions. Providing the market with the central bank's outlook on liquidity helps counterparties calibrate their bids in monetary operations, thereby supporting the achievement of the operational target. By communicating not just a point forecast but also the associated uncertainty, the central bank mitigates the risk of market anchoring and allows participants to factor in forecasting risk, which enhances overall financial stability (International Monetary Fund, 2024). In the absence of such reliable and transparent projections, central banks risk creating liquidity imbalances that can trigger systemic crises (Ben Charaf, 2024; Siliadin, 2013).

Section 3: Banking Liquidity in Algeria and the Challenges of Forecasting

The theoretical foundations discussed in the previous sections take on a particular meaning when applied to an emerging, hydrocarbon-dependent economy such as Algeria. Over the past decade, the Algerian banking system has experienced an alternation of liquidity surpluses and deficits, driven mainly by oil-price fluctuations, the cash preferences of economic agents, and discretionary fiscal decisions. This section traces the evolution of central bank liquidity between 2015 and 2022, drawing exclusively on the annual reports of the Bank of Algeria.

1. The Algerian Banking Liquidity Context

Banking liquidity in Algeria functions as a primary barometer of the country's macroeconomic health. Over the period under review, liquidity has been a highly volatile flow, shaped by three autonomous-factor channels: net foreign assets (AEN), reflecting the external position tied to hydrocarbon earnings; the net government position (PNG), capturing the fiscal channel of liquidity creation or absorption; and currency in circulation (CIC), representing the autonomous drain from the public's preference for cash. The present dissertation focuses on forecasting aggregate banking liquidity itself, using the variable's own history and, where relevant, exogenous factors.

1.1. Evolution of Central Bank Liquidity in Algeria (2015–2022)

The trajectory of liquidity from 2015 to 2022 unfolds in four distinct phases, each defined by specific shocks or policy interventions, as documented in the Bank of Algeria annual reports.

- **Phase 1: The Sharp Contraction (2015–2016)**

At the end of 2014, aggregate banking liquidity stood at 2,730.9 billion dinars (Banque d'Algérie, 2017). In 2015, the decline in oil prices triggered a severe reduction in hydrocarbon-sector deposits, compounded by an 11.9 per cent surge in credit demand. Net foreign assets declined by 302 billion dinars, while the net government position shifted from a net creditor (-1,992.3 billion) to a net debtor (+567.5 billion), as the Treasury drew down the Revenue Regulation Fund. Year-end 2015 liquidity fell to 1,832.6 billion dinars (-32.9 per cent). The contraction intensified in 2016: AEN fell by a further 2,779.4 billion dinars (-18.1 per cent), driven by a balance-of-payments deficit of 26 billion USD, while the PNG expanded to 2,682.2 billion (Banque d'Algérie, 2018). Year-end 2016 liquidity stood at only 820.9 billion dinars, a cumulative 70 per cent decline from the 2014 peak. The Bank of Algeria resumed direct refinancing, injecting 432.6 billion dinars, and cut the reserve ratio from 12 to 8 per cent.

- **Phase 2: Recovery and Policy Intervention (2017–2018)**

By October 2017, banking liquidity had fallen to 482.4 billion dinars the crisis nadir (Banque d'Algérie, 2018). The authorities responded with unconventional financing, purchasing 2,485 billion dinars of Treasury securities in November 2017. Year-end liquidity rebounded to 1,380.6 billion (+68.2 per cent). In 2018, an additional 3,371.2 billion were purchased under Article 45 bis, raising year-end liquidity to 1,557.6 billion (+12.8 per cent) (Banque d'Algérie, 2019). However, recovery was increasingly dependent on monetary financing, as AEN continued to drain (-1,741.8 billion). The Bank of Algeria began re-absorbing excess liquidity, raising the reserve ratio to 10 per cent and reintroducing seven-day absorption operations.

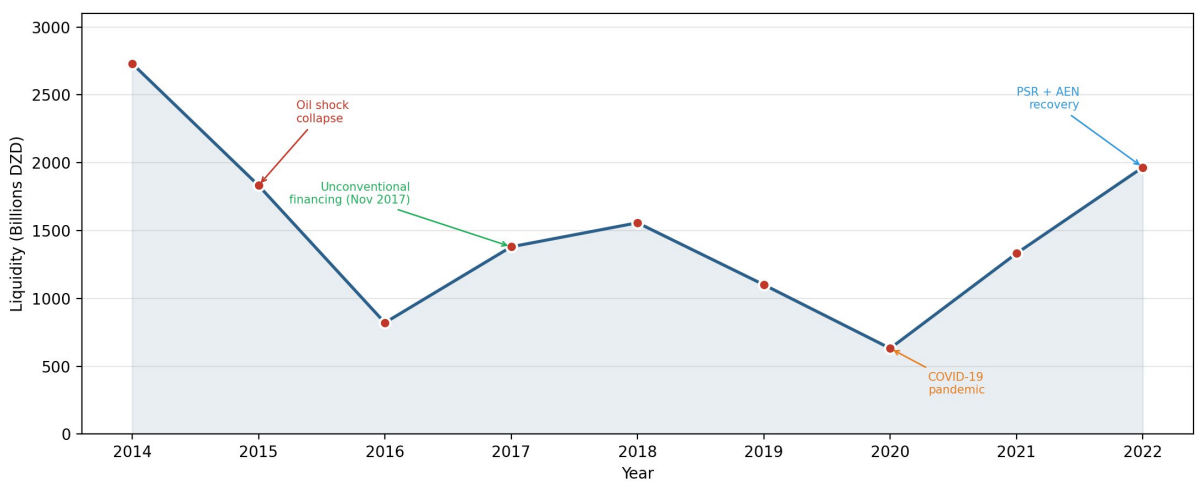
- **Phase 3: The Pandemic Crisis (2019–2020)**

A final 1,000-billion-dinar tranche in January 2019 pushed liquidity to 2,081 billion, but financing was suspended thereafter (Banque d'Algérie, 2020). The AEN drain (-1,933.8 billion), accelerated CIC growth (+510.8 billion), and the financing suspension drove year-end 2019 liquidity down to 1,100.9 billion (-29.3 per cent). COVID-19 then compounded the crisis: liquidity fell 58.1 per cent to 461.8 billion by September 2020, as AEN declined by a further 1,062.3 billion and CIC surged by 703.1 billion the largest annual increase on record reflecting precautionary cash hoarding (Banque d'Algérie, 2021). The Bank of Algeria cut the reserve ratio to 3 per cent, reduced the liquidity coefficient from 100 to 60 per cent, and raised the refinancing threshold to 95 per cent. A sharp increase in the PNG allowed a partial rebound to 632.1 billion by year-end 2020.

- **Phase 4: Expansionary Relief (2021–2022)**

Under the Special Refinancing Programme (PSR, ceiling 2,100 billion), 1,680 billion dinars were injected between July and November 2021, and the reserve rate was cut to 2 per cent (Banque d'Algérie, 2022). Year-end 2021 liquidity reached 1,331.95 billion (+110.7 per cent), with the monetary base expanding by 18.4 per cent. In 2022, the AEN channel reversed course as energy prices surged: net foreign assets grew by 31.9 per cent and foreign-exchange reserves recovered from 45.3 to 61 billion USD (Banque d'Algérie, 2023). Year-end 2022 liquidity reached 1,966.41 billion (+47.6 per cent), returning to near pre-crisis levels and completing the full cycle from contraction through crisis to recovery.

Figure I-1: Evolution of Year-End Banking Liquidity in Algeria (2014–2022)



Source: Elaborated by the students using data from Bank of Algeria Annual Reports

2. The Impact of Oil Shocks and COVID-19 on Algerian Liquidity

Two major external shocks profoundly shaped the trajectory of banking liquidity, each operating through the autonomous-factor channels identified above.

2.1. The 2014 Oil Shock

The drop in oil prices around 2014 ended an era of effortless liquidity surpluses. Lower oil revenues reduced hydrocarbon-sector deposits through the AEN channel: net foreign assets fell from 15,734.5 billion in 2014 to 12,596 billion by end-2016, a cumulative decline of 3,138.5 billion (Banque d'Algérie, 2017, 2018). The government drew down the Revenue Regulation Fund from 2,156 to 870.1 billion, signalling the exhaustion of traditional macro-stabilisers. The crisis culminated in October 2017 when liquidity fell to 482.4 billion, compelling the authorities to adopt unconventional financing (Banque d'Algérie, 2018).

2.2. The COVID-19 Pandemic

The pandemic introduced both behavioural and regulatory dimensions. Currency in circulation surged by 703.1 billion dinars in 2020 the largest annual increase on record as precautionary cash hoarding compounded the continued AEN decline (-1,062.3 billion), driving liquidity from 1,100.9 billion at year-end 2019 to 461.8 billion by September 2020 (Banque d'Algérie, 2021). The Bank of Algeria responded with exceptional measures: the mandatory liquidity ratio was cut from 100 to 60 per cent, and the PSR was launched in July 2021 (Banque d'Algérie, 2022). Liquidity recovered to 632.1 billion by year-end 2020 and surged to 1,331.95 billion by end-2021.

3. Liquidity Management Practices by the Bank of Algeria

Liquidity management in Algeria refers to the strategies employed by the central bank to regulate bank reserves and steer short-term interest rates. The Bank of Algeria mobilises conventional instruments open market operations, standing facilities, required reserves, and rediscounting with significant adjustments across the liquidity cycle (Banque d'Algérie, 2017, 2018, 2019, 2020, 2021, 2022, 2023). During the contraction (2015–2016), absorption operations were suspended, direct refinancing injected 432.6 billion, and the reserve rate was cut from 12 to 8 per cent. During the recovery (2017–2018), unconventional financing totalled 5,856.2 billion in Treasury purchases; as liquidity returned, the reserve ratio was raised to 10 per cent and seven-day absorption operations were reintroduced. During the pandemic (2019–2020), the reserve ratio was cut to 3 per cent, the PSR injected 1,680 billion in 2021, and the refinancing threshold was raised from 90 to 95 per cent. These adjustments demonstrate that the Bank of Algeria alternates between rate manipulation and prudential adjustments to compensate for exogenous shocks linked to oil prices and household cash preferences.

Chapter Conclusion

This first chapter has established the theoretical foundations of banking liquidity, progressing from universal definitions to the specific realities of the Algerian financial system. It has demonstrated that liquidity is not a static attribute but a dynamic flow, and that the process of maturity transformation by which banks convert short-term deposits into long-term loans generates a structural fragility that exposes institutions to funding liquidity risk and potential bank runs.

The typology of liquidity central bank, funding, and market has revealed the interdependence of these layers and their capacity to form either virtuous cycles of stability or vicious spirals of contagion during periods of stress. This dual nature underscores the central bank's indispensable role as the ultimate liquidity provider and lender of last resort.

The chapter has further examined the operational framework through which this role is exercised, including open market operations, reserve requirements, standing facilities, and prudential standards such as the Basel III ratios. It has argued that the effectiveness of these instruments fundamentally depends on accurate liquidity forecasting, without which monetary operations risk creating the very imbalances they seek to prevent.

Chapter II: Theoretical Framework and Literature Review: From Traditional Forecasting to Deep Learning for Central Bank Liquidity Prediction

Introduction

Having established the theoretical foundations of banking liquidity and the specific institutional challenges faced by the Bank of Algeria in the previous chapter, it is now necessary to define the methodological tools required to anticipate these liquidity fluctuations. As financial environments grow increasingly complex and volatile, relying solely on traditional statistical methods may no longer be sufficient to ensure proactive monetary policy and systemic stability.

The core objective of this chapter is to build the theoretical and methodological framework for our empirical study by critically evaluating and contrasting traditional econometric forecasting methods with advanced deep learning architectures. Understanding the mathematical mechanics, strengths, and limitations of these diverse models is essential to justify our comparative empirical approach. Furthermore, by reviewing the existing academic literature, this chapter aims to highlight current research gaps and position our study within the broader global discourse on AI-driven financial forecasting.

The chapter is organized in three parts:

- **Section 1:** Reviews the traditional forecasting approaches for central bank liquidity (such as SARIMA and SARIMAX) and discusses their inherent linear limitations.
- **Section 2:** Introduces the mathematical and structural foundations of the deep learning architectures used in this study, specifically Convolutional Neural Networks (CNN), Long Short-Term Memory (LSTM), and Gated Recurrent Units (GRU).
- **Section 3:** Presents a comprehensive literature review on the application of deep learning to financial and liquidity forecasting, establishing the comparative context for our research.

Section 1: Traditional Forecasting Approaches for Central Bank Liquidity and Their Limitations

Traditional econometric forecasting methods, rooted in the Box-Jenkins tradition, long underpinned central banking analysis due to their interpretability and theoretical rigor. This section reviews the foundational ARIMA framework before presenting the dissertation's two empirical traditional models, SARIMA and SARIMAX, concluding with their shared limitations.

1. ARIMA Foundations

The Autoregressive Integrated Moving Average (ARIMA) framework, introduced by Box and Jenkins (1970), constitutes the theoretical cornerstone of all models in Section 1. Although ARIMA is not used as an empirical forecasting model in this dissertation its exclusion justified by the pronounced seasonality in liquidity data its underlying concepts of autoregression, integration, moving average structure, stationarity, and model diagnostics are directly inherited by the SARIMA and SARIMAX specifications that are empirically evaluated.

The ARIMA(p,d,q) model combines three components: an autoregressive (AR) component of order p that captures the relationship between an observation and a specified number of lagged observations; an integration (I) component of order d representing the degree of differencing required to achieve stationarity; and a moving average (MA) component of order q modelling the relationship between an observation and residual errors from a moving average model applied to lagged observations. The general formulation can be expressed as:

$$Y_t = c + \sum_{i=1}^{p} \pi_i Y_{t-i} + \varepsilon_t + \sum_{j=1}^{q} \theta_j \varepsilon_{t-j}$$

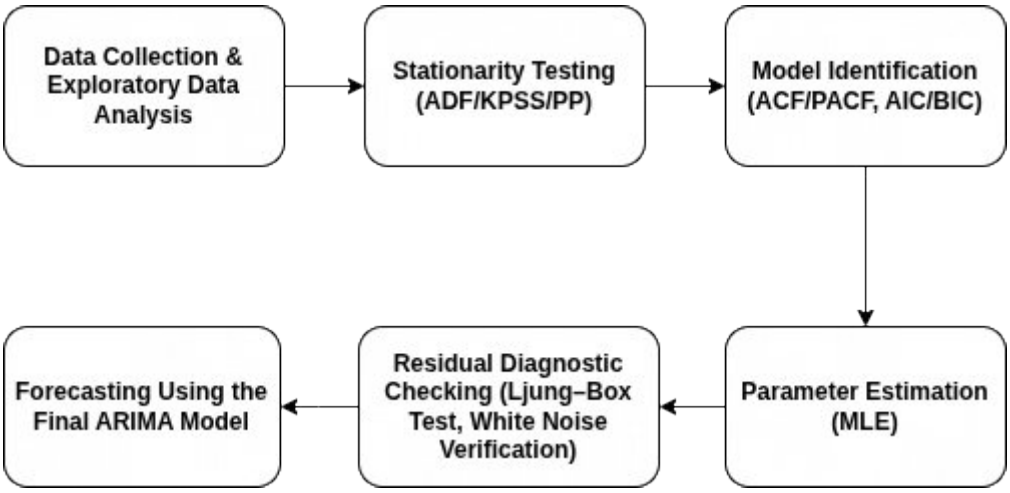
where Y_t denotes the observed value at time t, c is a constant term, φi are the autoregressive coefficients, θ_j are the moving average coefficients, and ε_t is a white noise error term. The integration order d indicates that the series has been differenced d times to achieve stationarity, which is a prerequisite for the AR and MA components to be well-defined. Stationarity is typically verified using unit root tests such as the Augmented Dickey-Fuller (ADF) test or the Phillips-Perron test.

The Box-Jenkins methodology provides a systematic, iterative approach for identifying, estimating, and diagnosing ARIMA models and, by extension, SARIMA and SARIMAX models. The identification stage involves examining autocorrelation function (ACF) and partial

autocorrelation function (PACF) plots to determine the appropriate orders p and q. Information criteria such as the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC) are also used to guide model selection, balancing goodness-of-fit with parsimony.

The estimation stage involves fitting the identified model to the data using maximum likelihood estimation (MLE) or conditional least squares. The diagnostic checking stage is critical for validating the adequacy of the fitted model. The Ljung-Box test is commonly employed to test whether the residuals exhibit significant autocorrelation. Additionally, the normality of residuals is assessed using the Jarque-Bera test or Q-Q plots, and heteroscedasticity is examined through ARCH-LM tests. A well-specified model whether ARIMA, SARIMA, or SARIMAX should produce residuals that are approximately normally distributed, homoscedastic, and free of significant autocorrelation.

Figure II-1: Box-Jenkins Methodology



Source: Elaborated by the students, based on Box & Jenkins (1970).

As established in Chapter 1, central bank liquidity data in Algeria exhibit strong and systematic seasonal patterns related to fiscal payment cycles, religious observances, and government budget disbursements. Applying a non-seasonal ARIMA model to such data would constitute a structural misspecification: the model would be unable to capture the periodic component of the data-generating process, leading to systematically biased forecasts and serially correlated residuals. For this reason, ARIMA serves in this dissertation as the theoretical building block for SARIMA and SARIMAX, and the reader is directed to those sections for the empirical models under evaluation.

2. SARIMA: The Seasonal Baseline

Building directly on the ARIMA foundations reviewed in Section 2.1.1, the Seasonal Autoregressive Integrated Moving Average (SARIMA) model extends the ARIMA framework by explicitly incorporating seasonal components, making it the first and theoretically appropriate empirical traditional model in this dissertation.

The SARIMA model is conventionally denoted as SARIMA (p, d,q) (P, D, Q) s, where the lowercase letters (p, d, q) represent the non-seasonal orders and the uppercase letters (P, D, Q) represent the seasonal orders, with s denoting the seasonal period. The seasonal differencing operator is defined as:

$$\Delta_s^D Y_t = (1 - B^s)^D Y_t$$

where B denotes the backshift operator such that:

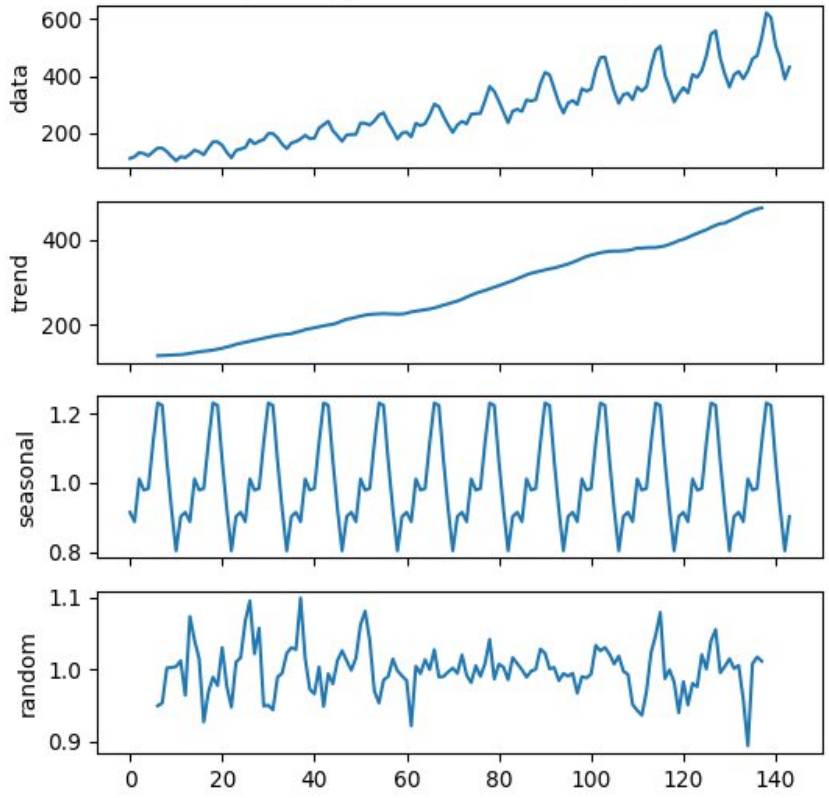
$$B^s Y_t = Y_{t-s}.$$

The seasonal AR and MA components operate at multiples of the seasonal period ***s***. The full SARIMA model is written in compact form as:

$$\varphi(B)\Phi(B^s)\, Y_t = \theta(B)\Theta(B^s)\, \varepsilon_t$$

where $\Phi(B^s)$ and $\Theta(B^s)$ are the seasonal AR and MA polynomials, respectively. This multiplicative structure allows the seasonal and non-seasonal components to interact, providing a parsimonious representation of complex seasonal patterns.

Figure II-2: SARIMA Seasonal Decomposition Example



Source: Adapted from pmdarima Documentation, "Seasonal Decomposition Example," "Alkaline ML"

For central bank liquidity forecasting, SARIMA is highly relevant because central bank balance sheets exhibit pronounced seasonal patterns. Currency in Circulation (CIC) displays strong variations related to salary payment dates, holiday periods (especially Ramadan and Eid in Algeria), and fiscal cycles, while the Net Government Position (NGP) reflects tax collection schedules. Explicitly modeling these seasonal components allows SARIMA to significantly outperform non-seasonal ARIMA, justifying its role as this dissertation's baseline.

However, SARIMA also inherits certain structural limitations from the ARIMA framework. The model remains linear and assumes parameter stability over time, which may not hold in the presence of structural breaks or evolving seasonal dynamics. Furthermore, SARIMA does not incorporate the influence of external variables such as oil price fluctuations that may drive both the level and the seasonal pattern of the series. These limitations are discussed collectively in Section 2.1.4 and motivate the SARIMAX extension presented next.

3. SARIMAX: Incorporating Exogenous Drivers

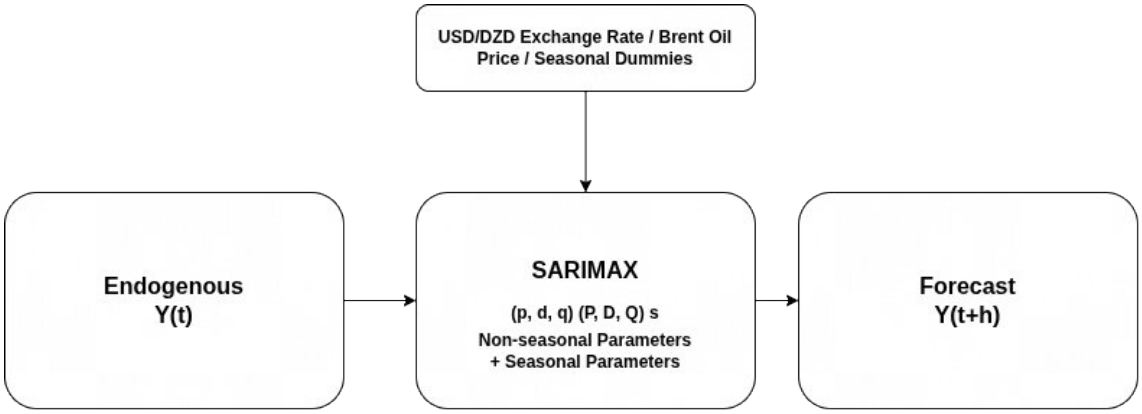
The SARIMAX (Seasonal Autoregressive Integrated Moving Average with eXogenous variables) model represents the most comprehensive specification within the traditional framework employed in this dissertation. It extends the SARIMA model by incorporating

external predictors that systematically influence the dependent variable. This extension is particularly important for central bank liquidity forecasting in Algeria, where liquidity is driven by a complex interplay of domestic and international factors that cannot be captured by the series' own history alone. The SARIMAX model can be formally expressed as:

$$\phi(B)\,\Phi(B^s)\,\Delta_t^{d}\Delta_s^{D}Y = c + \sum_{k=1}^{K}\beta_k X_{k,t} + \theta(B)\,\Theta(B^s)\,\varepsilon_t$$

where $X_{k,t}$ are K exogenous variables at time t, and β_k are their respective coefficients. The inclusion of exogenous variables allows the model to capture the impact of external shocks and systematic drivers not reflected in the autoregressive and moving average structure.

Figure II-3: SARIMAX Framework Seasonal ARIMA Structure with Exogenous Variable Integration



Source: Elaborated by the students

In the context of this dissertation, three categories of exogenous variables are particularly relevant. First, the Brent crude oil price serves as a key driver of the Algerian economy, directly affecting government revenues and foreign exchange reserves. As established in Chapter 1, Algeria's heavy dependence on hydrocarbon exports means that oil price fluctuations transmit rapidly through the fiscal channel to affect banking system liquidity. Second, the USD/DZD exchange rate captures the foreign exchange dimension of liquidity dynamics. Third, calendar and seasonal effects including holiday dummies for Ramadan and Eid, end-of-month fiscal adjustments, and day-of-week effects account for periodic patterns observed in liquidity data.

The transfer function framework provides a flexible approach within the SARIMAX specification by allowing for dynamic, distributed-lag relationships. This accommodates spread or delayed impacts, which is particularly relevant for variables like oil prices where transmission to banking liquidity involves a fiscal pipeline delay.

However, despite this flexibility, SARIMAX remains fundamentally linear and additive. This may fail to capture the complex, nonlinear interactions characterizing liquidity dynamics in an oil-dependent economy limitations discussed in Section 2.1.4 that provide the structural motivation for exploring deep learning approaches.

Table II-1: Comparison of SARIMA and SARIMAX The Two Empirical Traditional Models

Feature	SARIMA	SARIMAX
Seasonality	Yes	Yes
Exogenous Variables	No	Yes
Model Type	Linear	Linear
Nonlinearity	Cannot capture	Cannot capture
Feature Learning	Manual (ACF/PACF + seasonal lags)	Manual + domain knowledge for X variables
Interpretability	High	Moderate–High
Key Advantage	Seasonal capture; appropriate baseline for seasonal data	External drivers + seasonal dynamics
Empirical Role in dissertation	First empirical traditional model; seasonal benchmark	Second empirical traditional model; most advanced traditional specification

Source: Elaborated by the students

4. Limitations of SARIMA and SARIMAX for Central Bank Liquidity Forecasting

While SARIMA and SARIMAX represent the most appropriate traditional specifications for central bank liquidity data, they share fundamental limitations that constrain their effectiveness. These stem from the underlying ARIMA structure and reflect deeper assumptions about temporal dependencies that may be incompatible with the complex, nonlinear dynamics of liquidity in an oil-dependent economy.

The most significant limitation is inherent linearity. Both models assume future values represent a linear combination of lagged observations and exogenous inputs, ruling out threshold effects and asymmetric responses well-documented in financial data. These nonlinear dynamics arise from the asymmetric impact of oil price shocks on government revenue, exchange rate dependencies, and threshold effects in monetary transmission. Furthermore, Algerian banks exhibit structural funding imbalances and deposit concentrations that are inherently nonlinear and unsuited for linear models.

Structural breaks represent another critical challenge. The Algerian economy has experienced several significant changes, including the 2014 oil price collapse, the 2020 COVID-19 pandemic, and periodic fiscal policy changes. Both models assume parameter stability over the estimation window, and structural breaks can lead to biased estimates and unreliable forecasts.

The difficulty of modelling complex interactions among exogenous variables further limits SARIMAX. Although it allows for external predictors, it assumes their effects are additive and independent. In reality, the interaction between oil prices, exchange rates, and fiscal variables is likely multiplicative and state-dependent for instance, the impact of an oil price shock on liquidity may depend on the prevailing exchange rate regime or the government's fiscal stance. Deep learning models, through hierarchical feature extraction, can automatically learn these interactions without explicit specification.

Additionally, both models rely on manual specification of lag orders, seasonal parameters, and exogenous variable selections a time-consuming process dependent on the analyst's domain knowledge. These limitations are not merely technical inconveniences but reflect a fundamental mismatch between the assumptions embedded in the ARIMA framework and the structural realities of an oil-dependent economy where fiscal dominance, external shocks, and policy discretion produce dynamics that linear specifications cannot adequately reproduce. Deep learning can automatically learn relevant features from raw data, reducing the need for manual feature engineering and the explicit mathematical specification required even by advanced classical extensions (such as the TBATS framework or Dynamic Harmonic Regression), potentially uncovering patterns that traditional diagnostic tools may not reveal.

Section 2: Deep Learning Architectures for Liquidity Time Series Forecasting

Deep learning has emerged as a transformative paradigm for time series forecasting, offering the ability to automatically learn hierarchical feature representations from data, capture nonlinear dependencies, and model complex temporal patterns that traditional linear models cannot represent. This section presents the three deep learning architectures employed in this dissertation Convolutional Neural Networks (CNN), Long Short-Term Memory (LSTM) networks, and Gated Recurrent Unit (GRU) networks beginning with the foundational concepts of artificial neural networks.

1. Artificial Neural Networks: Foundations

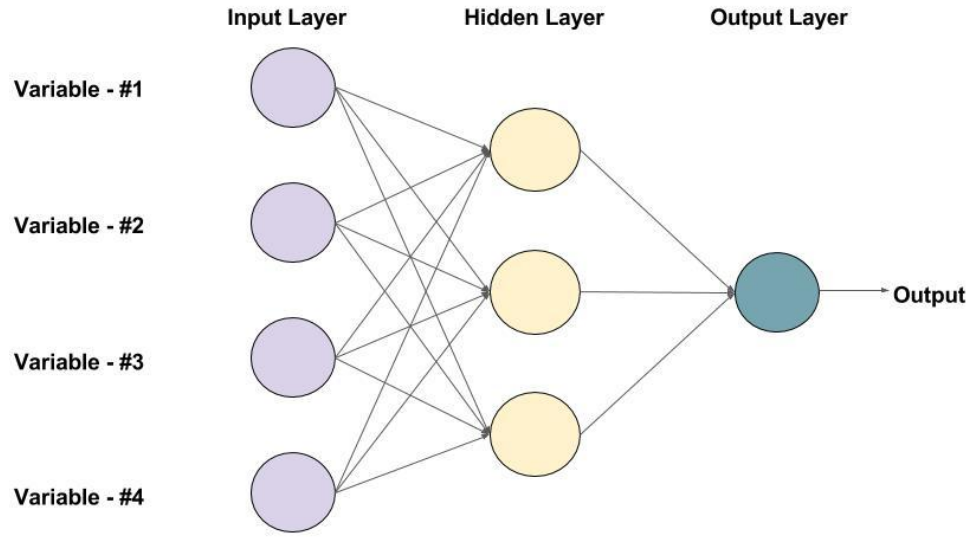
Before examining the specific deep learning architectures used in this study, it is essential to establish the foundations of artificial neural networks (ANNs) upon which these models are built. An artificial neuron, the basic computational unit, receives one or more inputs, applies a linear transformation followed by a nonlinear activation function, and produces an output. Mathematically, a single neuron computes:

$$y = f\left(\sum_{i=1}^{n} w_i x_i + b\right)$$

where x_i are the input values, w_i are the corresponding weights, b is the bias term, and f is the activation function that introduces nonlinearity into the model. Common activation functions include the sigmoid function, the hyperbolic tangent (tanh), and the Rectified Linear Unit (ReLU), defined as f(x) = max (0, x). The choice of activation function affects the network's ability to learn complex patterns and the efficiency of the training process (Goodfellow, Bengio, & Courville, 2016; Haykin, 2009).

A feedforward neural network consists of an input layer, one or more hidden layers, and an output layer. The universal approximation theorem states that a feedforward network with a single hidden layer containing a finite number of neurons can approximate any continuous function on a compact subset of real numbers to arbitrary precision, provided appropriate activation functions are used (Cybenko, 1989; Hornik, Stinchcombe, & White, 1989). While this result is theoretically powerful, in practice, deeper networks with multiple hidden layers are often more efficient at learning complex representations, as each layer can learn increasingly abstract features from the data.

Figure II-4: Feedforward Neural Network Architecture with one hidden layer (with 3 neurons)



Source:From LearnOpenCV article "Understanding Feedforward Neural Networks"

Neural network training is accomplished via the backpropagation algorithm, which computes loss function gradients using the calculus chain rule to update weights through optimization algorithms like SGD or Adam. Key training challenges include vanishing or exploding gradients and overfitting, which are commonly mitigated using dropout, batch normalization, and early stopping to improve generalization performance.

The application of deep learning to time series data requires careful preprocessing, including normalization of input features, construction of sliding-window sequences, and temporal train-validation-test splitting to prevent data leakage. These procedures are detailed in Chapter 3.

1.1. CNN for Liquidity Forecasting

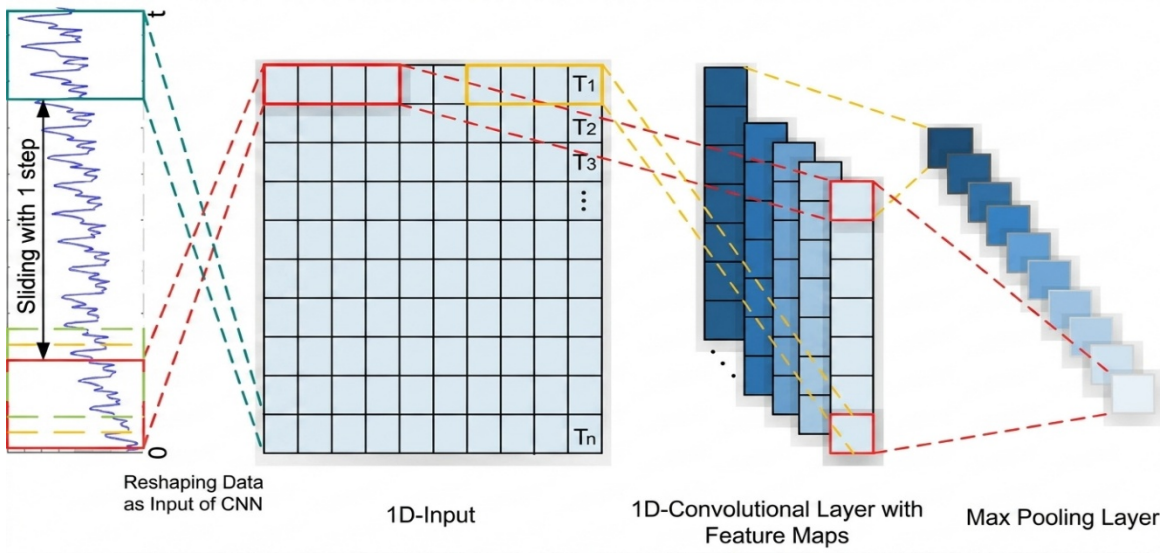
Convolutional Neural Networks (CNNs), originally developed for image recognition, have been successfully adapted for time series forecasting through the use of one-dimensional (1D) convolutions. In the time series context, a 1D convolution operates as a sliding window that extracts local temporal features by applying a set of learnable filters across the input sequence. The convolution operation at time step t can be expressed as:

$$y_t = f\left(\sum_{i=0}^{k-1} w_{i+1} \cdot x_{t+i} + b\right)$$

where ***b*** is a bias term and f is a non-linear activation function, typically ReLU. By using multiple filters, the convolutional layer can simultaneously detect various local patterns such

as trend changes, sudden spikes, or periodic fluctuations at different positions in the time series. The filter size k determines the receptive field.

Figure II-5: CNN 1D Convolution Process for Time Series Filter Application and Feature Extraction



Source: From "Forecasting Air Passenger Traffic Flow Based on the Two-Phase Learning Model" (2020), ResearchGate.

The stacked architecture of CNNs enables hierarchical feature representations. Early layers capture low-level, short-range patterns (e.g., day-to-day changes), while deeper layers combine these to detect higher-level, longer-range structures (e.g., weekly or monthly cycles). Dilated convolutions, which introduce gaps between filter elements, can further expand the receptive field without increasing the number of parameters, enabling the network to capture longer-range dependencies efficiently (Bai, Kolter, & Koltun, 2018).

In the context of central bank liquidity forecasting, CNNs are expected to excel at identifying recurring short-term patterns and local anomalies in the liquidity series. For example, CNNs can detect recurring liquidity drains around month-end salary payments or Eid holidays without requiring explicit seasonal dummy variables, learning the pattern directly from the raw series. However, standard CNNs may struggle with capturing very long-range dependencies compared to recurrent architectures, as the receptive field grows linearly with network depth unless dilation is employed.

1.2. LSTM Networks

Long Short-Term Memory (LSTM) networks, introduced by Hochreiter and Schmidhuber (1997), represent a landmark advancement in recurrent neural network design, specifically engineered to address the vanishing gradient problem that plagues standard RNNs. The LSTM architecture introduces a memory cell state and a set of gating mechanisms that

regulate the flow of information, enabling the network to learn which information to retain over long sequences and which to discard.

The LSTM cell contains three distinct gates the forget gate, the input gate, and the output gate each implemented as a sigmoid neural network layer producing values between 0 and 1, acting as multiplicative controls. The operations within an LSTM cell at time step t are defined as follows:

Forget gate: $f_t = \sigma\left(W_f \cdot [h_{t-1}, x_t] + b_f\right)$

Input gate: $i_t = \sigma\left(W_i \cdot [h_{t-1}, x_t] + b_i\right)$

Candidate: $\tilde{C}_t = tanh\tanh\left(W_C \cdot [h_{t-1}, x_t] + b_C\right)$

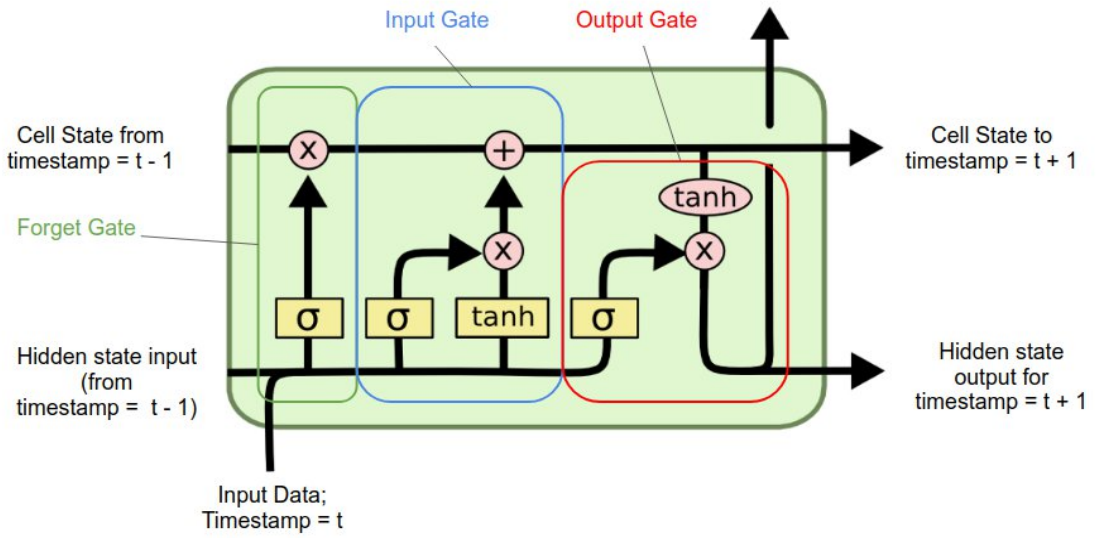
Cell state: $C_t = f_t \odot C_{t-1} + i_t \odot \tilde{C}_t$

Output gate: $o_t = \sigma\left(W_o \cdot [h_{t-1}, x_t] + b_o\right)$

Hidden state: $h_t = o_t \odot tanh\tanh\left(C_t\right)$

The cell state C_t serves as a memory channel that can carry information across many time steps with minimal modification. The forget gate ft determines what proportion of the previous cell state to retain; the input gate it controls how much new information to write into the cell state; and the output gate ot regulates what information to expose as the hidden state output. This gating mechanism allows the LSTM to maintain a stable gradient flow during backpropagation through time, effectively addressing the vanishing gradient problem.

Figure II-6: LSTM Cell Gate Architecture Memory Cell, Forget, Input, and Output Gates



Source: From "LSTMs Explained: A Complete, Technically Accurate, Conceptual Guide with Keras," Medium.

For central bank liquidity forecasting, the LSTM architecture offers several advantages: its ability to capture long-range temporal dependencies makes it suitable for modelling the persistent effects of policy decisions and structural economic changes; the gating mechanism can learn to selectively attend to relevant historical information, such as the impact of past oil price shocks on current liquidity conditions; and the recurrent nature of the architecture naturally handles variable-length sequences, which is useful for multi-horizon forecasting. Empirical evidence from Benkredda (2025), who compared ARIMAagainst LSTM and CNN for forecasting Bank of Algeria liquidity factors, confirms that LSTM models significantly outperform classical approaches for volatile and nonlinear liquidity components such as Net Foreign Assets (NFA) and Net Government Position (NGP).

1.3. GRU Networks

Gated Recurrent Unit (GRU) networks, introduced by Cho et al. (2014), represent a streamlined alternative to the LSTM architecture that achieves comparable performance with a simpler gate structure. The GRU was originally proposed as a mechanism for machine translation within encoder-decoder frameworks but has since been widely adopted for general sequential modelling tasks, including time series forecasting.

The GRU cell simplifies the LSTM by merging the cell state and hidden state into a single state vector and reducing the number of gates from three to two: a reset gate and an update gate. The operations within a GRU cell at time step t are defined as follows:

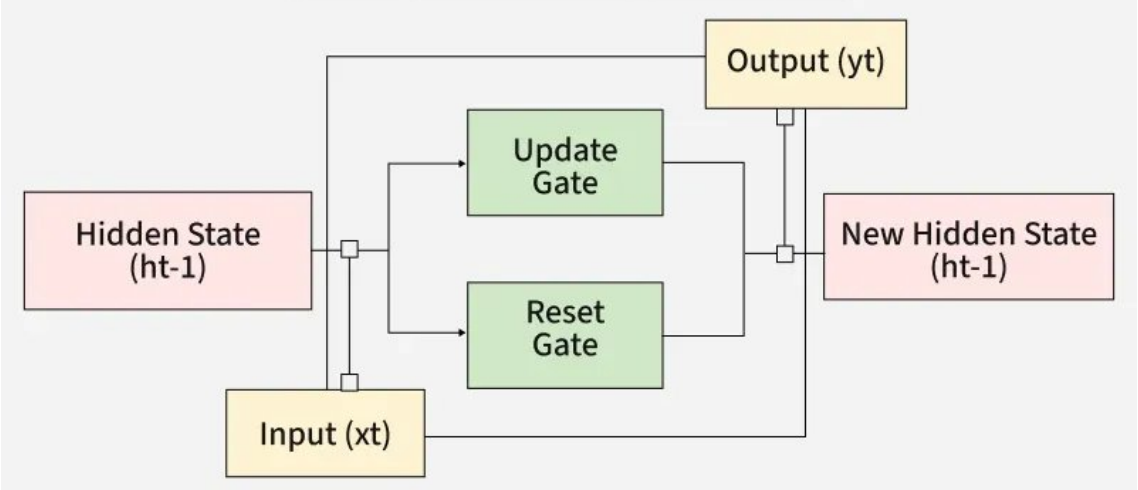
Reset gate: $r_t = \sigma\left(W_r \cdot \left[h_{t-1}, x_t\right] + b_r\right)$

Update gate: $z_t = \sigma\left(W_z \cdot \left[h_{t-1}, x_t\right] + b_z\right)$

Candidate: $\tilde{h}_t = tanh\tan\mathrm{h}\left(W \cdot \left[r_t \odot h_{t-1}, x_t\right] + b\right)$

Hidden state: $h_t = \left(1 - z_t\right) \odot \tilde{h}_t + z_t \odot h_{t-1}$

Figure II-7: GRU Cell Gate Architecture Reset Gate, Update Gate, and Candidate State



Source: AlindGupta - Gated Recurrent Unit Networks on Geeksforgeeks.com.

The critical difference from the LSTM is that the update gate z_t simultaneously performs the functions of both the forget and input gates. When z_t is close to 1, the previous hidden state is largely retained; when z_t is close to 0, the new candidate state dominates. The reset gate r_t controls how much the previous hidden state contributes to the candidate activation. This coupling of the forget and input operations through a single gate creates an implicit bias toward maintaining information, which can act as a form of regularization.

The critical difference from the LSTM is that the update gate z_t simultaneously performs the functions of both the forget and input gates. When z_t is close to 1, the previous hidden state is retained; when close to 0, the new candidate state dominates. This coupling creates an implicit bias toward maintaining information, acting as a form of regularization. Four substantive considerations justify including the GRU as a standalone model. First, the coupled update gate reflects a distinct inductive bias that enforces a trade-off between old and new information, preventing redundant data accumulation. This benefits moderate-sized datasets like this dissertation's daily liquidity series. Second, empirical studies show the GRU frequently outperforms the LSTM on shorter sequences; Chung et al. (2014) found it matched or exceeded LSTM performance across multiple tasks. Third, its reduced parameter count (~25–30% fewer than the LSTM) provides a regularization benefit crucial for limited emerging economy data. Fourth, including both architectures assesses whether the gating mechanism itself or its specific implementation determines forecast accuracy.

Table II-2: Comparison of LSTM and GRU Architectures

Feature	**LSTM**	**GRU**
Number of Gates	3 (Forget, Input, Output)	2 (Reset, Update)
Cell State	Separate cell state Ct	Merged with hidden state ht
Parameters	More (~3n(4n+2m) per layer)	Fewer (~25–30% less)
Gate Coupling	Independent forget/input	Coupled update gate
Training Speed	Slower	Faster
Best For	Very long sequences	Medium sequences, limited data

Source: Elaborated by the students

In summary, the GRU is included in this dissertation not as a mere computational shortcut but as a distinct modelling paradigm with its own theoretical justification and empirical track record. The comparison between LSTM and GRU will provide valuable insights into the trade-off between architectural flexibility and regularization in the context of central bank liquidity forecasting.

Section 3: Literature Review on Deep Learning Applications to Liquidity Forecasting and Comparative Studies

This section reviews the existing literature on deep learning applications in financial time series forecasting, with a particular focus on central bank liquidity prediction and comparative studies between traditional and deep learning approaches. The review draws primarily on two recent theses conducted within the Bank of Algeria's institutional framework Abdennouri (2019) and Benkredda (2025 which provide the most directly relevant empirical and regulatory context for this study. These supervisor-recommended resources are supplemented by the broader international literature to position the research within the global academic discourse.

1. Deep Learning in Financial Time Series Forecasting

The application of deep learning to financial time series forecasting has expanded rapidly, driven by the availability of high-frequency data and neural networks' ability to capture complex nonlinear patterns. The seminal work of Zhang (2003) on hybrid ARIMA-ANN models established the potential of combining these approaches, sparking extensive research into deep learning-based financial prediction.

In exchange rate forecasting, Bao et al. (2017) proposed a stacked autoencoder combined with LSTM that significantly improved prediction accuracy over conventional methods by decomposing the forecasting task into feature extraction and prediction stages. Similarly, in the context of stock price forecasting, Yang (2024) applied LSTM networks to Tencent share prices, finding that the recurrent architecture excels at handling long-term patterns and dependencies. Conversely, Support Vector Regression (SVR) proved well-suited for capturing short-term market volatility. Both methodologies outperformed the traditional ARIMA model, which struggled with the complex, non-linear nature of financial time-series data across different forecast horizons. Commodity price forecasting, especially for crude oil, represents another area where deep learning demonstrates substantial improvements. Given oil's direct relevance to the Algerian economy and the Bank of Algeria's liquidity position, this literature is highly pertinent. Researchers consistently find that LSTM and CNN architectures outperform traditional ARIMA models during high volatility and structural change.

The literature also highlights the importance of exogenous variables. Wang and Chen (2025) demonstrated that incorporating macroeconomic indicators as exogenous inputs significantly improves the performance of machine learning models for crude oil return

forecasting. By utilizing 55 exogenous variables, their models successfully captured the market's complexity and non-stationarity. Furthermore, regarding out-of-sample Brent crude oil price forecasting, Jin and Xu (2024) achieved a highly accurate relative root mean square error (RRMSE) of 0.2814% using advanced machine learning models.

These findings have direct implications for the present study's inclusion of Brent oil prices and exchange rates. The general consensus emerging from this body of work is that deep learning models, particularly LSTM and GRU architectures, offer meaningful improvements in forecast accuracy for financial time series, especially when the underlying data-generating process is complex, nonlinear, and influenced by multiple external factors.

It is important to acknowledge recent advances in Transformer-based architectures (e.g., Informer, Autoformer) that use self-attention mechanisms for time series forecasting. These models represent the current frontier but are excluded from the present study for two reasons: (1) they require substantially larger datasets than the ~2,000 daily observations available for Algeria, and (2) their computational demands exceed the resources available for this dissertation. Their emergence, however, reinforces the importance of first establishing whether simpler deep learning architectures (LSTM, GRU, CNN) already provide meaningful improvements over traditional models in limited-data contexts

1.1. Deep Learning for Central Bank Liquidity Forecasting

The application of deep learning to central bank liquidity forecasting is a relatively nascent but rapidly growing field. While the broader financial forecasting literature has extensively examined stock prices, exchange rates, and commodity prices, the specific challenge of forecasting central bank liquidity defined as the net position of the banking system vis-à-vis the central bank has received comparatively less attention.

Among the most directly relevant studies is Benkredda (2025), whose master's dissertation conducted within the framework of an internship at the Bank of Algeria represents the closest existing research to the present study. Benkredda compared classical statistical models (ARIMA family models) against deep learning models (LSTM and CNN) for forecasting the three autonomous liquidity factors of the Bank of Algeria: Net Foreign Assets (NFA), Currency in Circulation (CIC), and Net Government Position (NGP). Using daily data from 2015 to 2023, the study employed a rigorous dual methodological approach combining classical econometrics and modern machine learning techniques.

The key findings of Benkredda (2025) are highly informative for the present study. LSTM models significantly outperformed classical models in forecasting Net Foreign Assets

and Net Government Position, demonstrating superior ability to capture the nonlinear dynamics and complex temporal dependencies inherent in these volatile liquidity components. Conversely, classical models particularly when enhanced through rolling forecast techniques remained highly effective for forecasting Currency in Circulation, which exhibits more stable and predictable patterns. The overarching conclusion was that the complexity of the liquidity factor plays a decisive role in determining which modelling approach is most effective: deep learning excels for volatile and nonlinear factors, while traditional approaches remain competitive for stable, predictable series.

Complementing the forecasting perspective, Abdennouri (2019) provides essential regulatory and supervisory context for the Bank of Algeria's institutional framework. Conducted within the General Directorate of General Inspection, the study evaluated central bank tools for assessing commercial liquidity risk, revealing how regulatory standards shape the demand for accurate forecasting. As detailed in Chapter 1, Abdennouri documented the Bank of Algeria's Banking Rating System (SNB), inspired by CAMELS. The supervision mission revealed qualitative deficiencies such as absent stress-testing and superficial emergency funding plans highlighting the need for Algerian banks to transition to prospective risk management. This dissertation directly addresses this institutional need by evaluating deep learning's forecast accuracy.

2. Comparative Studies: Traditional vs. Deep Learning Models

A growing body of comparative studies has directly evaluated the forecasting performance of traditional models against deep learning architectures. It is important to note at the outset that the majority of existing comparative studies benchmark deep learning against non-seasonal ARIMA models a reasonable choice for the general financial time series contexts they examine, but an insufficient benchmark when the data under study exhibits strong seasonal patterns such as those observed in central bank liquidity. This dissertation advances beyond these prior comparisons by pitting deep learning architectures against the more theoretically appropriate SARIMA and SARIMAX benchmarks.

Fischer and Krauss (2018) compared LSTM against ARIMA for S&P 500 return prediction, finding that LSTM significantly outperformed the traditional model, particularly for longer forecast horizons. While this finding is consistent with the theoretical expectation that deep learning models outperform linear benchmarks, the non-seasonal nature of the ARIMA benchmark used in their study means that the comparison may overstate the advantage of deep learning relative to a properly specified seasonal model.

Siami-Namini and Namin (2018) conducted a systematic comparison of ARIMA and LSTM for financial indices and exchange rate forecasting. Their results confirmed the superiority of LSTM, with the deep learning model achieving significantly lower RMSE and MAE values across multiple datasets and forecast horizons. Again, the benchmark was non-seasonal ARIMA.

The most directly relevant comparative evidence comes from Benkredda (2025), whose study at the Bank of Algeria provides a rare empirical comparison between classical and deep learning models specifically for central bank liquidity forecasting. Benkredda compared ARIMA and VAR both non-seasonal, non-exogenous specifications against LSTM and CNN. Her findings reveal a nuanced picture: LSTM models outperformed ARIMA and VAR for forecasting NFA and NGP. However, for Currency in Circulation a more stable and seasonally predictable series classical models matched or exceeded the performance of deep learning models. Crucially, Benkredda's study did not include SARIMA or SARIMAX models (the most appropriate traditional benchmarks for seasonal data), nor did it incorporate exogenous variables such as oil prices or the exchange rate, nor the GRU architecture. The present study improves upon Benkredda (2025) by replacing the simple ARIMA benchmark with seasonal SARIMA and SARIMAX specifications, adding GRU as a third deep learning architecture, and incorporating exogenous variables within both frameworks.

Swetha and Porkodi (2024) and Ashtar et al. (2025) compared SARIMAX against deep learning models across housing and energy markets. While SARIMAX performed competitively for short horizons due to its effective handling of exogenous variables and seasonal patterns, deep learning architectures dominated at longer horizons where nonlinear dynamics became more pronounced. This finding aligns with the theoretical expectation that the linear approximation inherent in SARIMAX becomes increasingly inadequate as the forecast horizon extends.

Extending this analysis within the Algerian context, Taoussi, Boudia, and Mazouni (2024) systematically compared SARIMA, SARIMAX, and LSTM models for wind speed forecasting. Their findings strongly reinforce the methodological choices of this dissertation: they demonstrated that incorporating exogenous variables significantly enhances SARIMAX performance. Furthermore, they proved that employing a rolling forecast approach mitigates plateau effects, allowing traditional statistical models to ultimately surpass the LSTM benchmark by reducing the Root Mean Square Error (RMSE) by 65.54%. This provides robust

empirical justification for utilizing rolling windows and exogenous predictors to capture complex dynamics in Algerian datasets.

Hewage et al. (2020) compared CNN, LSTM, and GRU alongside ARIMA, finding that relative performance varied depending on data characteristics and forecast horizon. Notably, GRU frequently matched or exceeded LSTM performance while requiring significantly less training time, supporting the inclusion of GRU as a standalone model rather than a mere computational shortcut. CNN models demonstrated particular strength in capturing local pattern shifts, making them effective for short-horizon forecasts where sudden regime changes are important.

Makhmudov (2025) compared SARIMA against LSTM and other neural networks for forecasting liquidity in the Central Bank of Uzbekistan and found that SARIMA produced superior forecasts. This result appears to contradict the broader trend favouring deep learning, but the discrepancy is likely attributable to differences in data characteristics: Uzbek liquidity data may exhibit more stable, deterministic seasonality with fewer structural breaks than Algerian data, which is heavily exposed to oil price volatility and exchange rate adjustments. This underscores that the relative advantage of deep learning is context-dependent a nuance that the present study explicitly tests by comparing against seasonal benchmarks in the Algerian setting.

Table II-3: Summary of Comparative Studies Traditional vs. Deep Learning Models

Study	Traditional	Deep Learning	Context	Key Finding / Contribution
Fischer & Krauss (2018)	ARIMA	LSTM	S&P 500 returns	LSTM superior; gap widens with horizon
Siami-Namini & Namin (2018)	ARIMA	LSTM	Financial indices, FX	LSTM significantly lower RMSE/MAE
Benkredda (2025)	ARIMA, VAR	LSTM, CNN	Bank of Algeria liquidity (NFA, CIC, NGP)	LSTM best for volatile factors; ARIMA competitive for CIC. Seasonal models and GRU absent addressed by present study,
Swetha & Porkodi (2024)	-	ANN, FNN, CNN, RNN, GRU, and LSTM	Housing market indices	LSTM achieved the highest accuracy with an R^2 of 0.9773

Ashtar, Mohammadi Ziabari, & Alsahag (2025)	SARIMAX	SARIMAX-LSTM, Seq2Seq	Sustainable electricity demand	SARIMAX competitive short-term; Seq2Seq deep learning dominates long-term horizons (up to 180 days) with a MAPE of 1.88%
Taoussi, Boudia, & Mazouni (2024)	SARIMA, SARIMAX, VAR, VARMA	LSTM	Algerian wind speed forecasting	Exogenous variables enhanced SARIMAX; using rolling forecasts allowed statistical models to surpass LSTM, improving RMSE by 65.54%.
Hewage et al. (2020)	ARIMA	CNN, LSTM, GRU	Multi-domain TS	Performance varies; GRU matches LSTM efficiency
Makhmudov (2025)	SARIMA	LSTM, NN	Uzbekistan central bank liquidity	SARIMA outperformed neural networks; context-dependent advantage

Source: Students synthesis from the literature review

3. The Role of Exogenous Variables in Liquidity Forecasting

Incorporating exogenous variables consistently improves forecasting accuracy. For an oil-dependent economy like Algeria, this selection is guided by theoretical considerations and institutional characteristics. This subsection examines the three categories of exogenous variables employed in this dissertation Brent crude oil prices, the USD/DZD exchange rate, and calendar/seasonal effects drawing on regulatory insights from Abdennouri (2019) and forecasting evidence from Benkredda (2025).

The Brent crude oil price is arguably the most important exogenous variable for the Bank of Algeria. As established in Chapter 1, Algeria's economy is heavily dependent on hydrocarbon exports, which account for approximately 95% of export earnings and 60% of government revenues. Oil price fluctuations directly affect government revenues, which in turn influence the Net Government Position and, through the fiscal channel, the overall liquidity position of the banking system. Benkredda (2025) explicitly identified the omission of oil prices as a limitation of her study, noting that the inclusion of this variable could further enhance the accuracy of liquidity forecasts particularly for the NFA and NGP components most directly affected by oil revenue flows. The present study addresses this limitation directly.

The USD/DZD exchange rate captures the foreign exchange dimension of liquidity dynamics. A depreciation of the Algerian dinar increases the domestic currency value of foreign reserves held by the central bank, directly affecting the Net Foreign Assets position. The Bank

of Algeria manages the exchange rate within a managed float regime, and exchange rate adjustments can have significant and sometimes nonlinear effects on liquidity conditions.

Calendar and seasonal effects represent the third category. Central bank liquidity is subject to periodic fluctuations related to salary payment dates, religious holidays (Ramadan and Eid, which affect both currency demand and government spending patterns), and weekly patterns related to banking settlement cycles. Ramadan and Eid al-Fitr increase Currency in Circulation through heightened cash withdrawals for holiday spending, while simultaneously delaying government tax collection due to reduced commercial activity during the fasting period, creating a double liquidity drain effect. The SARIMAX model handles these effects through deterministic seasonal dummies, while deep learning models can learn these patterns automatically from the data and, critically, can learn nonlinear interactions between calendar effects and oil price movements (e.g., a high-oil-price Eid may see greater liquidity injection than a low-oil-price Eid).

Table II-4: Exogenous Variables and Their Expected Impact on Central Bank Liquidity

Variable	Transmission Channel	Primary Factor Affected	Expected Effect
Brent Oil Price	Fiscal revenue → Gov. spending → NGP	NFA, NGP	Positive (price rise → liquidity increase)
USD/DZD Rate	Reserve valuation → NFA; Import costs	NFA, CIC	Positive (depreciation → NFA increase)
Ramadan/Eid/Summer season/Back to school	Currency demand; Spending patterns	CIC, NGP	Mixed (higher CIC, variable NGP)
End-of-Month	Salary payments → Withdrawals	CIC	Negative (CIC rise → liquidity drain)

Source: Elaborated by the students

4. Research Gaps and Positioning

The review of the existing literature reveals several important methodological, contextual, and practical gaps. Directly informed by the limitations and findings of Abdennouri (2019) and Benkredda (2025), the present study addresses these deficiencies through five distinct contributions.

First, despite extensive model comparisons for asset prices, a pronounced scarcity of such applications exists within central bank liquidity forecasting. Benkredda (2025) represents

the closest prior attempt, comparing ARIMA/VAR against LSTM/CNN for Bank of Algeria liquidity factors. However, her study used a non-seasonal ARIMA benchmark an insufficient choice given highly seasonal data and omitted SARIMA, SARIMAX, and GRU architectures. The present study addresses this by introducing SARIMA and SARIMAX as appropriate traditional seasonal benchmarks alongside CNN, LSTM, and GRU architectures.

Second, the absence of seasonal benchmarks in prior studies implies that existing evidence regarding deep learning's superiority may be overstated. If a non-seasonal model serves as the baseline, any deep learning advantage might simply reflect the superiority of a correctly specified seasonal model over a mis specified non-seasonal one. Benchmarking against SARIMA and SARIMAX provides a scientifically valid test, ensures that any deep learning outperformance is confidently attributable to architectural properties like nonlinearity and hierarchical feature extraction.

Third, while most comparative studies evaluate a single forecast horizon, the operational needs of central banks require accurate forecasts across multiple time scales. This study's multi-horizon design spanning day-ahead, week-ahead, and month-ahead targets allows for a nuanced assessment of model appropriateness across daily open market operations, medium-term planning, and strategic decisions regarding reserve requirements.

Fourth, this study systematically incorporates exogenous variables within both frameworks to resolve a key limitation noted by Benkredda (2025), who omitted external drivers. This design directly integrates Brent crude prices, the USD/DZD exchange rate, and calendar effects into both the SARIMAX specification and the deep learning models, allowing for a direct comparison of how traditional and neural network approaches leverage external information.

Fifth, while prior research has documented the unique structural characteristics of the Algerian liquidity environment namely hydrocarbon dependence, managed exchange rates, and fiscal dominance the forecasting literature has not yet evaluated whether deep learning models outperform seasonal benchmarks in this specific context. This dissertation provides the first comprehensive evaluation tailored specifically to these national dynamics, incorporating the country's most relevant exogenous variables and utilizing traditional seasonal benchmarks appropriate to the underlying data.

Chapter Conclusion

This chapter has established the theoretical framework for the subsequent empirical analysis, introducing five forecasting models across two paradigms: two traditional (SARIMA, SARIMAX) and three deep learning (CNN, LSTM, GRU). While central to the theoretical discussion, non-seasonal ARIMA is deliberately excluded from the empirical comparison because central bank liquidity data exhibit pronounced seasonal patterns. SARIMA serves as the appropriate baseline, ensuring that deep learning architectures are evaluated against well-specified traditional benchmarks rather than a mis specified non-seasonal model.

The chapter's progression outlines this structured methodology. Section 2.1.1 reviewed the foundational ARIMA components that underpin both traditional models. Section 2.1.2 introduced SARIMA as the seasonal baseline, while Section 2.1.3 presented SARIMAX as the advanced traditional specification incorporating exogenous drivers. Section 2.1.4 addressed the shared linear limitations of these traditional models, providing the structural motivation for the deep learning architectures.

Consequently, Section 2 detailed the three neural network models: CNN for local temporal patterns, LSTM for long-range dependencies, and GRU as a streamlined, competitive alternative with implicit regularization benefits.

Section 3 reviewed the literature, drawing on Abdennouri (2019) and Benkredda (2025). Benkredda (2025) demonstrated that LSTMs outperform ARIMA for volatile factors but classical models remain competitive for stable series, while identifying the omission of seasonal benchmarks and exogenous variables as key limitations. This section also integrated Makhmudov (2025) to highlight the context-dependent nature of model performance, ultimately identifying five research gaps regarding benchmarks, horizons, and context.

This framework directly underpins the methodological design in Chapter 3, establishing the dissertation's core empirical contribution to the Bank of Algeria's liquidity management.

Chapter III: Empirical study of Central Bank Liquidity Forecasting Using Statistical and Deep Learning Models

Introduction

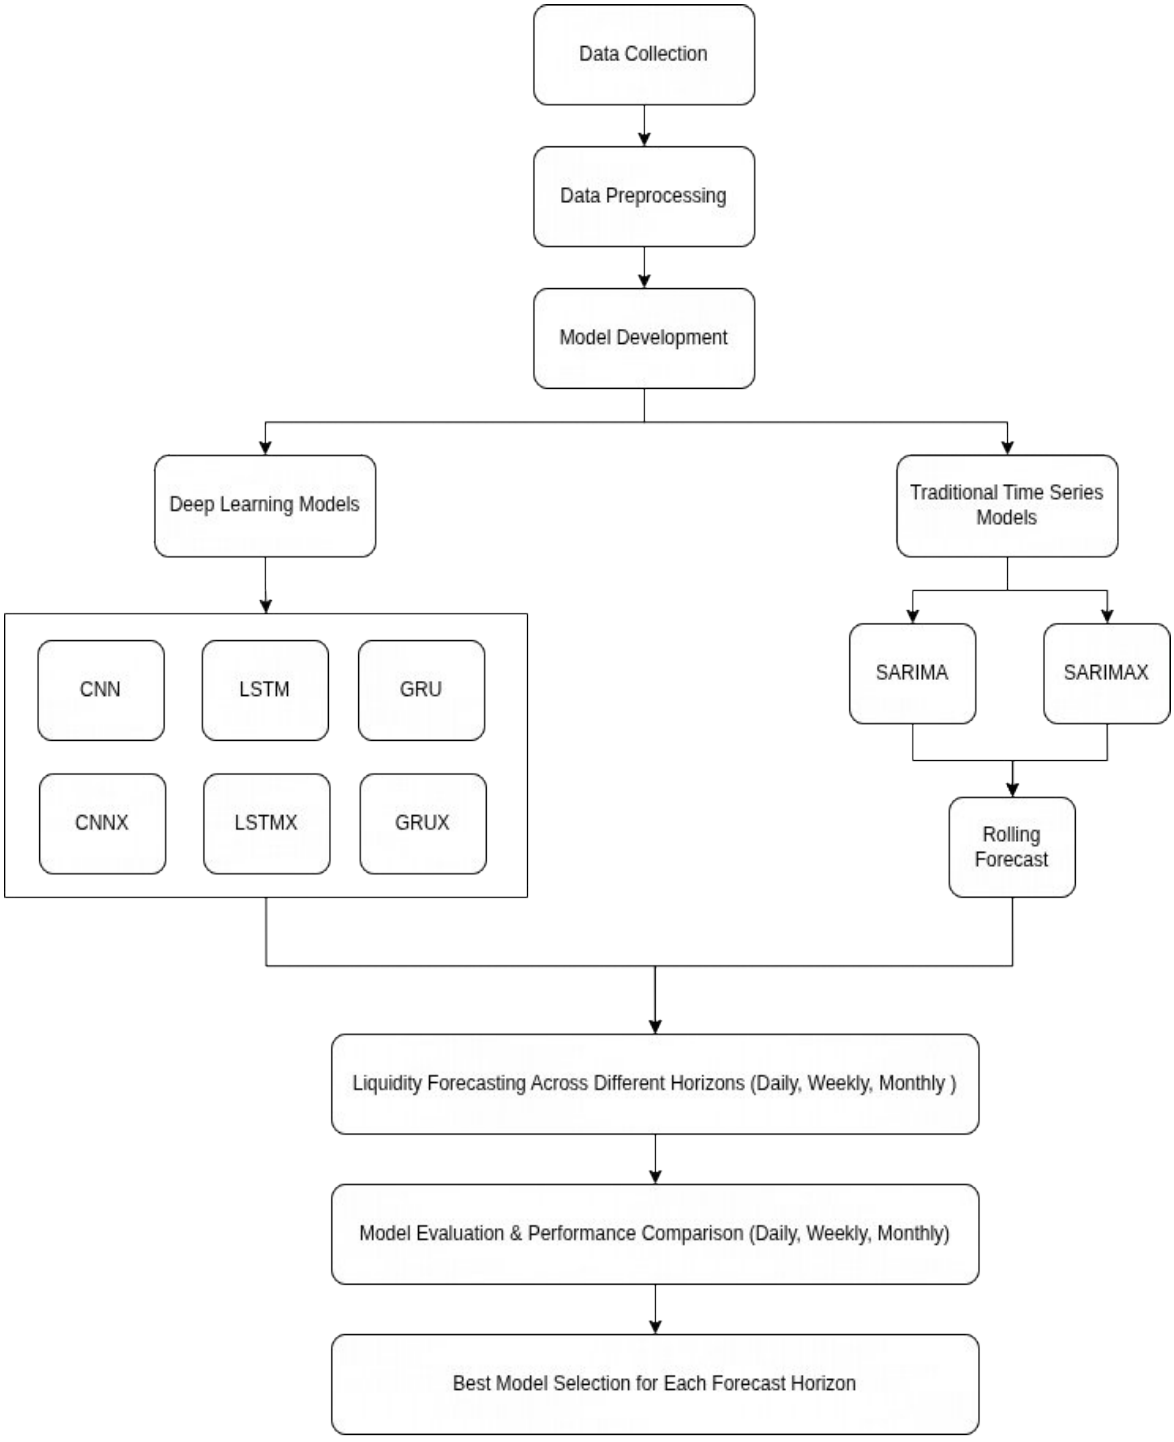
This chapter constitutes the empirical core of the dissertation. It translates the theoretical concepts of liquidity from Chapter I and the methodological frameworks from Chapter II into a rigorous, applied analysis of the Bank of Algeria's liquidity dynamics. In light of the growing complexity of the Algerian financial environment, practical application and rigorous testing are required to determine which predictive models yield the most reliable results for monetary authorities.

The main objective of this chapter is to apply and benchmark classical statistical models against modern deep learning techniques to forecast the Bank of Algeria's aggregate net liquidity position. By testing these models across multiple operational time horizons (day-ahead, week-ahead, and month-ahead) and assessing the marginal impact of incorporating exogenous macroeconomic variables (such as Brent oil prices and exchange rates), the ultimate goal is to identify the single most effective forecasting approach to support the Bank of Algeria's decision-making process.

To achieve this objective, guided by the methodological framework illustrated in Figure III-1 The analysis is structured into four sections:

- **Section 1**: Details the data construction process and provides an exploratory descriptive analysis of the central bank liquidity series and its exogenous factors from 2015 to 2022.
- **Section 2**: Focuses on model development and forecasting using classical statistical models (SARIMA and SARIMAX).
- **Section 3**: Implements the forecasting process using deep learning architectures (CNN, LSTM, and GRU), evaluating both their univariate and exogenous specifications.
- **Section 4**: Conducts a final comparative benchmarking across all models and horizons to deliver definitive, evidence-based recommendations for the Bank of Algeria.

Figure III-1 : Methodological Framework for Multi-Horizon Liquidity Forecasting



Source: Elaborated by the students

Section 1: Data Construction and Descriptive Analysis

This section presents the exploratory analysis of the dataset collected from the Bank of Algeria. The analysis covers 2,085 daily observations from January 2015 to December 2022. Unlike prior studies that focus on individual autonomous factors, this research directly targets central bank liquidity, capturing the net aggregate effect of all autonomous components. The objective is to summarize the central tendency, dispersion, and distributional characteristics of the data, examine the relationships between liquidity and its exogenous factors, and assess the predictability and normality properties that will inform model selection in subsequent sections.

1. Methodological Approach

We adopted a quantitative approach based on time series analysis. The data were collected from the Bank of Algeria's DMMF (Direction du Marché Monétaire et Financier), covering 2,085 daily observations (Sunday–Thursday) over the period 2015–2022. The dependent variable is the net Bank of Algeria liquidity, also referred to in Chapter 1, Section 2, as the autonomous liquidity aggregate (AGG). The explanatory variables, namely the USD/DZD exchange rate and Brent crude oil prices, were obtained from the official website Investing.com, while calendar dummy variables were constructed to capture the effects of major recurring events and seasonal periods in Algeria.

Specifically, these dummies correspond to dates associated with Eid al-Adha, Eid al-Fitr, New Year's Day, the back-to-school period, and the summer season, all identified using annual calendars covering the study period.

2. Descriptive characteristics of the central bank Liquidity

Table III-1 presents the descriptive statistics for central bank liquidity, including the mean, median, standard deviation, minimum, maximum, coefficient of variation (CV), kurtosis, and skewness. These measures provide a comprehensive picture of the central tendency, dispersion, and shape of the liquidity distribution. All monetary values are expressed in billions of Algerian dinars (DZD).

Table III-1: Descriptive Statistics of Central Bank Liquidity (n = 2,085)

Variable	Mean	Median	Std. Dev.	Min	Max	CV (%)	Kurtosis	Skewness
Liquidity	1,045.71	984.00	546.68	-154.00	2,610.00	52.28	-0.58	0.26

Source: Elaborated by the students

Central bank liquidity exhibits a mean of 1,045.71 billion DZD with a moderate coefficient of variation (52.28%), reflecting the structural shifts during the study period. The

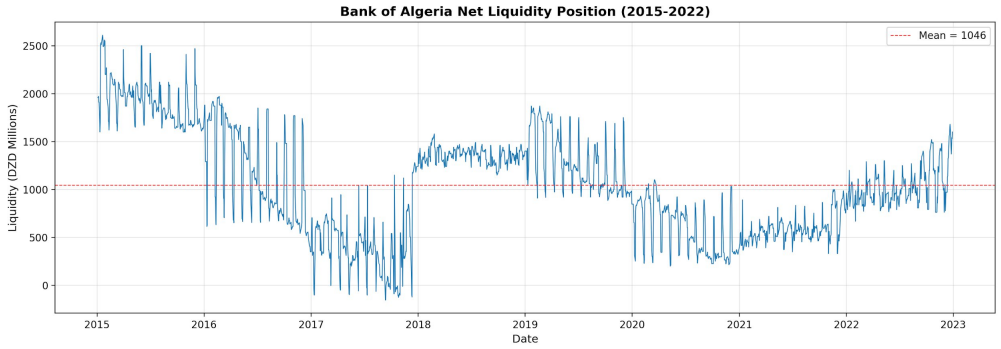
minimum of -154 billion DZD (September 11, 2017) confirms a rare deficit episode, while the maximum of 2,610 billion DZD (January 20, 2015) illustrates the earlier surplus period. The relatively modest gap between mean and median (1,045.71 vs. 984.00) indicates a slightly right-skewed distribution, confirmed by the positive skewness of 0.26. The kurtosis of -0.58 indicates a platykurtic distribution, meaning the liquidity series has lighter tails and a flatter peak than a normal distribution, consistent with the structural break observed between surplus and deficit regimes across the study period.

The coefficient of variation of 52.28% indicates moderate variability relative to the mean, reflecting the transition from large surpluses in 2015 to lower levels in 2020–2021. This degree of dispersion is noteworthy for a central bank liquidity series and signals meaningful shifts in the Algerian monetary environment over the study period. The platykurtic nature (kurtosis = -0.58) suggests that the liquidity values are spread more uniformly across the range rather than concentrated around the mean, reflecting the multi-regime nature of the series where distinct periods of surplus and deficit each have their own central tendency.

2.1. Overall Liquidity Variation (2015-2022)

The net liquidity position of the Bank of Algeria exhibits substantial variation. The daily average stands at approximately 1,046 billion DZD with a standard deviation of 547 billion DZD. The minimum of -154 billion DZD (September 11, 2017) represents a rare negative net position, while the maximum of 2,610 billion DZD (2015) underscores a range spanning 2,764 billion DZD. The median of 984 billion DZD and positive skewness of 0.26 indicate a slightly right-skewed distribution.

Figure III-2 Bank of Algeria Net Liquidity Position (2015-2022)

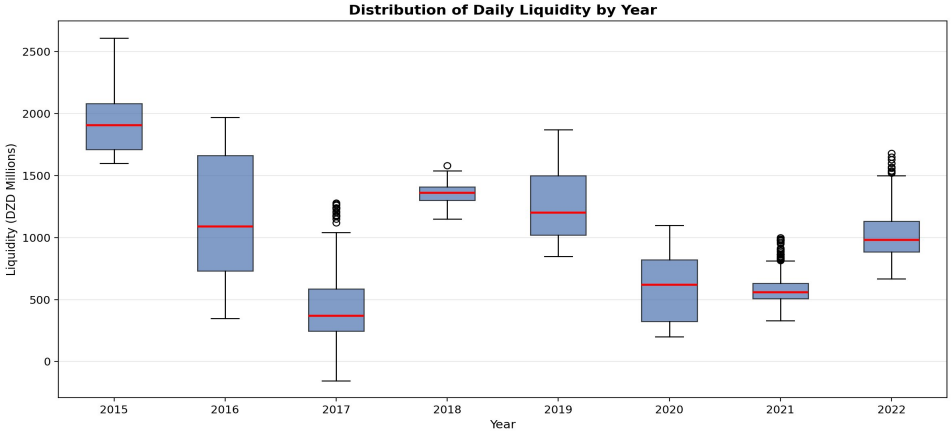


Source: Elaborated by the students using Python

Figure III-2 illustrates the evolution of daily net liquidity. Liquidity was highest during 2015 (1,600–2,610 billion DZD) before declining through 2016–2017, partially recovering in 2018, compressing again during 2020–2021, and rebounding in 2022. According to the Bank of Algeria's Annual Reports, the progressive tightening between 2015 and 2017 reflected the

decline in foreign exchange reserves from approximately USD 144 billion at end-2014 to under USD 100 billion by end-2017.

Figure III-3 Distribution of Daily Liquidity by Year (Boxplot)



Source: Elaborated by the students using Python

2.2. Year-by-Year Liquidity Variation

Table III-2: Descriptive Statistics of Daily Liquidity by Year (DZD Billions)

Year	Mean	Std Dev	Min	Max	Change (%)
2015	1,925	242	1,600	2,610	-15.82
2016	1,185	478	349	1,970	-81.44
2017	413	329	-154	1,280	+255.30
2018	1,356	79	1,150	1,580	+12.90
2019	1,280	301	847	1,870	-39.50
2020	585	269	199	1,100	-54.66
2021	587	140	330	1,000	+71.75
2022	1,039	209	668	1,680	+66.32

Source: Elaborated by the students

The year 2015 represents the peak (mean 1,925 billion DZD), driven by still favourable oil prices and robust NFA. The Bank of Algeria's 2015 Annual Report noted the NFA component benefited from residual pre-2014 oil boom effects, though the year closed with a 15.8% decline. The years 2016–2017 witnessed dramatic deterioration: the average fell to 1,185 then 413 billion DZD, with the 2017 minimum reaching -154 billion DZD, the only negative value in the dataset. The Bank of Algeria documented reserves declining to approximately USD 97 billion by end-2017 and the central bank resorting to direct Treasury financing.

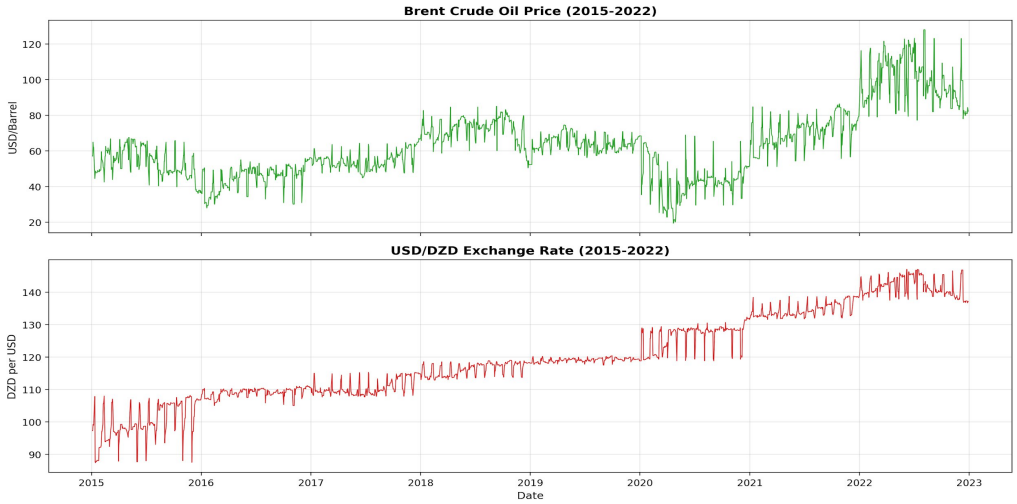
The year 2018 exhibited remarkable stability (mean 1,356, std 79), linked to oil price recovery and active central bank management. In 2019, renewed stress appeared as the NGP component's negative contribution grew due to fiscal deficits. The pandemic years 2020–2021 saw average liquidity fall to 585–587 billion DZD, with the Bank of Algeria deploying

exceptional measures including lowering reserve requirements and extending refinancing maturities. The year 2022 marked a recovery to 1,039 billion DZD, driven by the Russia-Ukraine conflict boosting oil revenues, though USD/DZD depreciation partially offset gains.

2.3. Descriptive characteristics of the exogenous variables

The exogenous factors comprise two continuous variables (Brent crude oil price and USD/DZD exchange rate) and four dummy variables (Eid, Ramadan, Back-to-School, Summer). Brent averaged $62.72 per barrel (std $18.80), ranging from $19.33 (April 2020) to $127.98 (August 2022), following a U-shaped trajectory. As the Bank of Algeria has consistently noted, the relationship between oil prices and domestic liquidity operates through the NFA channel: higher oil prices generate greater foreign currency inflows through Sonatrach, which are converted into dinars through the central bank.

Figure III-4: Brent Crude Oil Price and USD/DZD Exchange Rate (2015-2022)



Source: Elaborated by the students using Python

The USD/DZD exchange rate exhibited persistent depreciation from 100.15 DZD per USD in 2015 to 141.68 in 2022, a cumulative depreciation of approximately 41.5%. The Bank of Algeria implemented a significant devaluation of approximately 9% in March 2020. The observed net effect is a strong negative correlation with liquidity of -0.46, suggesting the volume effect of reduced foreign exchange availability dominates the price effect of higher dinar-value inflows.

The four dummy variables capture calendar effects: Eid (10 business days before Eid al-Fitr, Eid al-Adha, and 1 January plus the day itself; 248 observations, 11.89%), Pre_Ramadan (10 business days before Ramadan plus the day itself; 80, 3.84%), Pre_BTS (10 business days before the academic year plus the day itself; 80, 3.84%), and Summer (July and August; 354, 16.98%). During pre-Eid periods, households increase cash withdrawals, putting upward pressure on CIC and reducing net liquidity. The Bank of Algeria's Annual Reports have

repeatedly noted seasonal cash demand spikes during Eid. The pre-Ramadan period similarly features increased commercial activity, while the back-to-school period generates household expenditure spikes, and summer alters government payment schedules.

2.4. Multivariate Analysis

Table III-3 presents the descriptive statistics for the two continuous exogenous variables. These measures characterize the central tendency, dispersion, and distributional shape of each factor. Brent is expressed in USD per barrel and USDDZD in DZD per USD.

Table III-3: Descriptive Statistics of Exogenous Continuous Variables (n = 2,085)

Variable	Mean	Median	Std. Dev.	Min	Max	CV (%)	Kurtosis	Skewness
Brent	62.72	61.14	18.80	19.33	127.98	29.97	0.76	0.80
USDDZD	119.74	118.30	13.26	87.42	146.99	11.07	-0.72	0.18

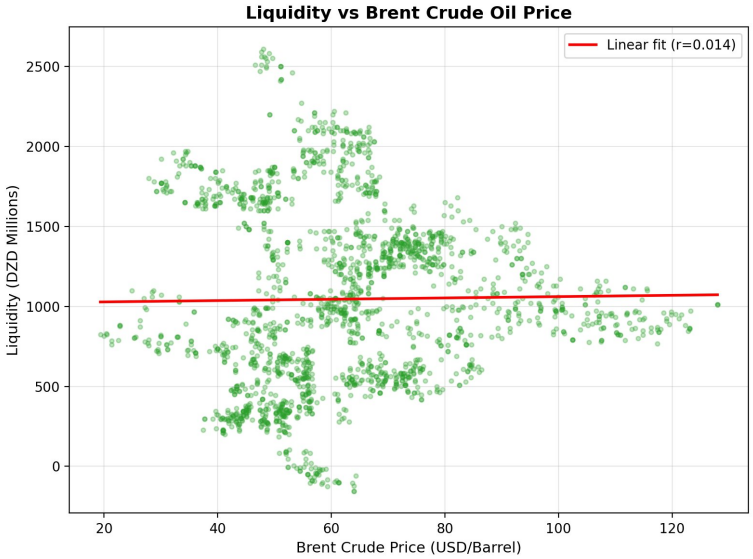
Source: Elaborated by the students via Python

Brent crude oil exhibits a coefficient of variation of 29.97%, reflecting substantial oil market volatility over the study period. The positive skewness (0.80) indicates a right-skewed distribution, with the mean exceeding the median, driven by the extreme price spike in 2022 following the Russia-Ukraine conflict. The kurtosis of 0.76 suggests a slightly leptokurtic distribution with heavier tails than a normal distribution, consistent with the well-documented fat-tailed nature of commodity price returns. The wide range from $19.33 to $127.98 underscores the dramatic oil price fluctuations, which directly impact Algeria's foreign exchange earnings and consequently the NFA component of central bank liquidity.

The USD/DZD exchange rate has the lowest CV (11.07%) among all continuous variables, reflecting a relatively smooth and managed depreciation trajectory rather than volatile fluctuations. The near-zero skewness (0.18) confirms a roughly symmetric distribution around the mean, while the negative kurtosis (-0.72) indicates a platykurtic distribution with lighter tails. This is consistent with the Bank of Algeria's managed float regime, where exchange rate adjustments are implemented gradually rather than through sharp devaluations, except for the notable March 2020 adjustment. The cumulative depreciation from 87.42 to 146.99 DZD per USD represents a 68.2% increase over eight years, eroding the domestic value of foreign reserves and constraining the central bank's capacity to inject liquidity.

2.5. Relationship Between Liquidity and Exogenous Factors

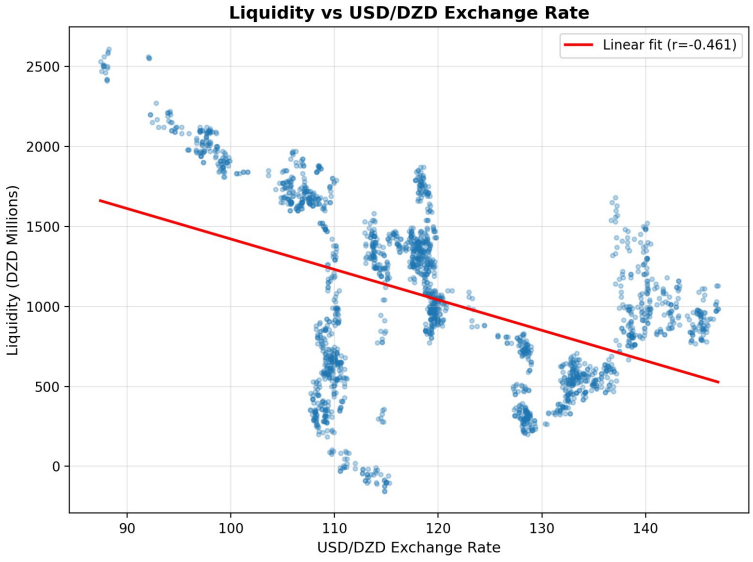
Figure III-5 : Liquidity vs Brent Crude Oil Price



Source: Elaborated by the students using Python

The scatter plot of liquidity against Brent crude oil price (Figure III-5) shows a very weak linear relationship (r = 0.014). This near-zero correlation does not mean oil prices are irrelevant to liquidity. The connection between the two variables is indirect and delayed: oil revenues pass through Sonatrach and the Treasury before reaching the banking system, which introduces significant time lags. Additionally, the central bank actively manages liquidity through its monetary policy operations, which further weakens the direct statistical link. The relationship is therefore likely nonlinear and dependent on the broader fiscal and monetary context.

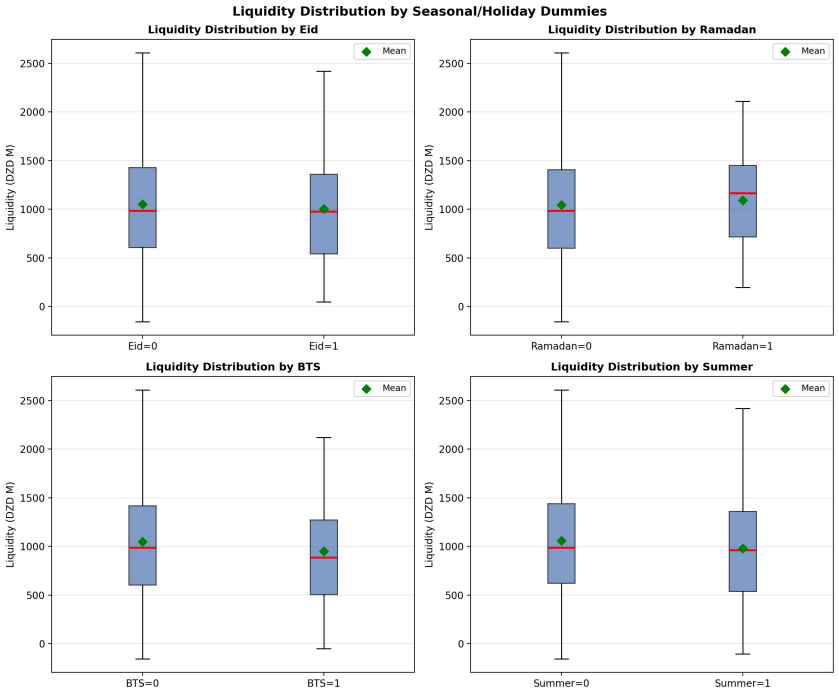
Figure III-6 Liquidity vs USD/DZD Exchange Rate



Source: Elaborated by the students using Python

The scatter plot of liquidity against the USD/DZD exchange rate (Figure III-6) shows a clear negative relationship (r = -0.461). As the dinar depreciates against the dollar, banking sector liquidity tends to decline. This pattern reflects the fact that when foreign currency reserves fall, the central bank has less capacity to inject liquidity into the system through its foreign asset operations. At the same time, a weaker dinar raises the local currency cost of foreign exchange transactions, which adds further pressure on liquidity. In this sense, dinar depreciation is both a sign of deteriorating external conditions and a contributing factor to tighter liquidity.

Figure III-7 : Liquidity Distribution by Seasonal/Holiday Dummy Variables



Source: Elaborated by the students using Python

Table III-4: Correlation Between Liquidity and Seasonal/Holiday Dummy variables

Variable	Correlation (r)
Eid	-0.027
Ramadan	+0.018
BTS	-0.037
Summer	-0.054

Source: Elaborated by the students via Python

(Table III-4) and Figure III-7 reveal weak correlations across all variables, ranging from a low of -0.054 during Summer to a high of +0.018 during Ramadan. The data shows that liquidity decreases during the Eid period (from 1,052 to 1,002 billion DZD), which is consistent with typical increases in cash withdrawals. The Summer variable reflects the largest decline in

liquidity, dropping from 1,059 to 981 billion DZD. In contrast, the Ramadan period shows a slight positive association with liquidity; this may indicate that the central bank pre-emptively injects funds into the system to meet expected demand.

Although these individual seasonal effects appear small, they may carry more weight when evaluated through complex interactions or indirect economic factors.

Figure III-8 : Correlation Matrix of Liquidity and Exogenous Factors



Source: Elaborated by the students using Python

The correlation matrix (Figure III-8) shows a strong positive relationship between Brent crude oil price and the USD/DZD exchange rate (r = 0.61). This is largely because both variables follow an upward trend over the sample period, so their correlation reflects a shared time trend rather than a direct economic link between the two. Among the seasonal dummy variables, BTS and Summer show the highest correlation (r = 0.26), which is expected given that the back-to-school period and summer months partially overlap in timing. These relationships between the explanatory variables need to be considered when building the regression model, as having two or more closely related variables on the right-hand side can make it difficult to isolate the individual effect of each one on liquidity.

2.6. Stationarity and Normality Analysis of Liquidity

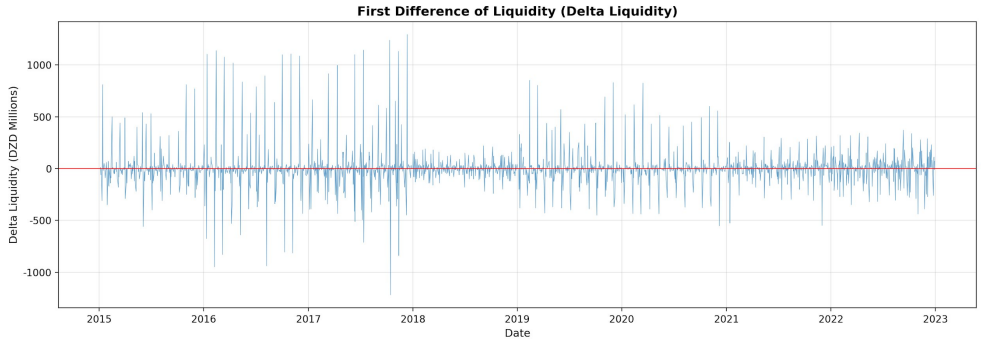
Before proceeding to econometric modelling, it is essential to examine the stationarity and distributional properties of the liquidity series. Stationarity is a fundamental prerequisite

for time series models, as non-stationary series can produce spurious regressions. Similarly, normality of residuals is an important assumption for ARIMA and SARIMAX specifications.

- **Stationarity Tests**

Three complementary tests were applied: the Augmented Dickey-Fuller (ADF) test, the Phillips-Perron (PP) test, and the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test. The ADF and PP tests share the null hypothesis of a unit root (non-stationarity), while the KPSS test reverses the null, positing stationarity.

Figure III-9: First Difference of Liquidity (Delta Liquidity, 2015-2022)



Source: Elaborated by the students using Python

Table III-5: Unit Root and Stationarity Tests for Liquidity

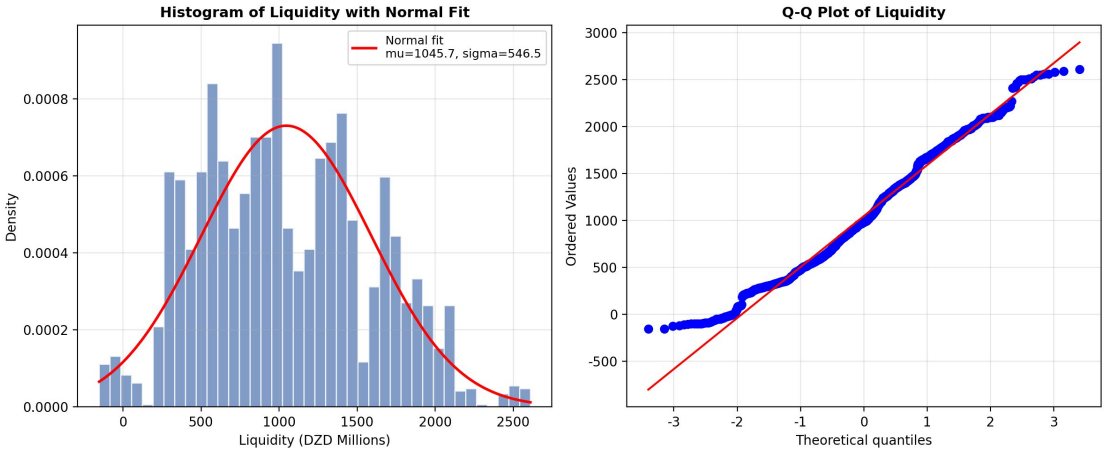
Test	Statistic	Critical Value (5%)	p-value	Conclusion
ADF (Level)	-2.376	-2.863	0.149	Non-stationary
PP (Level)	-6.320	-2.863	0.000	Stationary
KPSS (Level)	2.085	0.463	0.010	Non-stationary
ADF (1st Diff.)	-12.496	-2.863	0.000	Stationary
PP (1st Diff.)	-75.090	-2.863	0.000	Stationary
KPSS (1st Diff.)	0.087	0.463	>0.100	Stationary

Source: Elaborated by the students using Python

The ADF test on the level series produces a statistic of -2.376, failing to reject the unit root null at 5%. The KPSS test strongly rejects stationarity (statistic 2.085 vs. critical value 0.463). Both point toward non-stationarity. However, the Phillips-Perron test, robust to serial correlation, rejects the unit root (statistic -6.320, p-value approximately zero). This discrepancy may be attributable to the ADF test's sensitivity to lag length and structural breaks. First-differenced liquidity was therefore tested, and all three tests unanimously confirm stationarity: ADF -12.496, PP -75.090, and KPSS 0.087. These results provide strong evidence that the liquidity series is integrated of order one, I (1), meaning the first difference achieves stationarity.

2.7. Normality Analysis

Figure III-10: Histogram with Normal Fit and Q-Q Plot of Liquidity



Source: Elaborated by students using Python

Table III-6: Normality Tests for Liquidity (Full Sample)

Test	Statistic	p-value	Conclusion
Shapiro-Wilk	0.984	<0.001	Non-normal
Jarque-Bera	51.800	<0.001	Non-normal
Kolmogorov-Smirnov	0.053	<0.001	Non-normal

Source: Elaborated by the students using Python

The normality of the liquidity distribution was assessed using three complementary tests. The histogram (Figure III-10, left) reveals a slightly right-skewed distribution (skewness = 0.26) with mild platy kurtosis (kurtosis = -0.58). The Q-Q plot (Figure III-10, right) shows systematic departures from normal quantiles, particularly at the extremes. The Shapiro-Wilk test (statistic 0.984, $p < 0.001$), the Jarque-Bera test (51.80, $p < 0.001$), and the Kolmogorov-Smirnov test (0.053, $p < 0.001$) all decisively reject normality. The rejection across all three tests is driven by the combined effect of mild skewness and platy kurtosis, sufficiently pronounced given the large sample (n = 2,085).

In summary, the descriptive and preliminary econometric analysis establishes several key facts. The liquidity series exhibits substantial variation driven by its four balance sheet components and closely linked to Algeria's external economic conditions, as documented in the Bank of Algeria's Annual Reports. The USD/DZD exchange rate shows a meaningful negative correlation with liquidity (r = -0.46), while Brent operates through a more complex lagged mechanism. The seasonal dummies capture modest but systematic effects. The liquidity series is I(1), requiring first-differencing for stationarity, and departs from normality across all three tests. These findings provide the empirical foundation for the SARIMAX and deep learning models developed in the following sections.

Section 2: Forecasting using Classical Models

This section applies the traditional Box-Jenkins methodology to the central bank liquidity series. The objective is twofold: first, to establish a rigorous seasonal baseline using SARIMA; and second, to augment this baseline with exogenous drivers through the SARIMAX specification. The performance of these traditional models will subsequently serve as the benchmark against which the deep learning architectures presented in Section 3 are evaluated.

The analysis proceeds through the canonical stages of time series modelling: preliminary testing for stationarity and seasonality, model identification via information criteria, parameter estimation, residual diagnostic validation, and out-of-sample forecasting using rolling forecast strategies with re-estimation at each step.

1. Univariate Time-Series Forecasting: The SARIMA Framework

The Autoregressive Integrated Moving Average (ARIMA) framework provides the theoretical foundation for univariate time series forecasting. For a non-seasonal series, the ARIMA (p, d, q) model is expressed as:

$$\varphi(B)\,(1-B)^{d}\,Y_t = \theta(B)\,\varepsilon_t$$

where φ(B) and θ(B) are the autoregressive and moving average polynomials, d is the order of differencing, and ε_t is white noise. However, as established in Section 3.1, Algerian net liquidity exhibits pronounced seasonal fluctuations tied to the monthly fiscal and salary payment cycle. A purely non-seasonal ARIMA would therefore be misspecified.

The Seasonal ARIMA (SARIMA) extension incorporates multiplicative seasonal terms at period s:

$$\varphi(B)\,\Phi(B^{s})\,(1-B)^{d}\,(1-B^{s})^{D}\,Y_t = \theta(B)\,\Theta(B^{s})\,\varepsilon_t$$

Here's the integrated version:

$$\left(Y_t = (1-B)^{-d}\,(1-B^{s})^{-D}\,\frac{\theta(B)\,\Theta(B^{s})}{\pi(B)\,\Phi(B^{s})}\,\varepsilon_t\right)$$

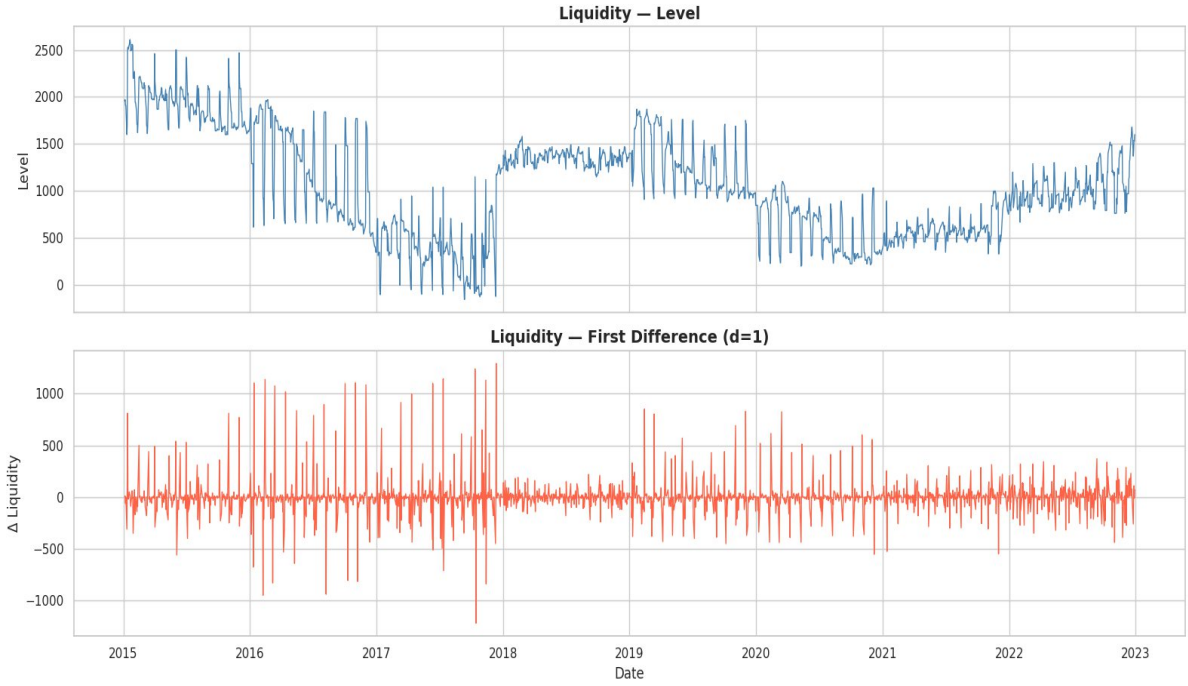
where ***s*** denotes the seasonal period, and (P, D, Q) are the seasonal orders. For daily data reflecting end-of-month liquidity cycles in the Algerian banking system, the seasonal period is set to s = 22 business days. This value is derived from the institutional calendar: a calendar month of approximately 30 days, under the five-day Algerian work week (Sunday–Thursday), yields 30 × (5/7) ≈ 21.4 working days, rounded to 22. This choice is further

supported empirically by the significant ACF spikes observed precisely at lag 22 (and 44), confirming that the monthly liquidity cycle repeats every 22 business days.

1.1. Preliminary Analysis and Data Visualization

Figure III-11 presents the daily net liquidity position from January 2015 to December 2022 (n = 2,085 observations). The series exhibits substantial volatility and clear structural breaks, notably the sharp contraction in 2016–2017, partial recovery in 2018–2019, pandemic-induced compression in 2020–2021, and rebound in 2022.

Figure III-11: Net Liquidity Position — Level and First Difference (d = 1), 2015–2022



Source: Elaborated by the students using Python

1.2. Stationarity Analysis

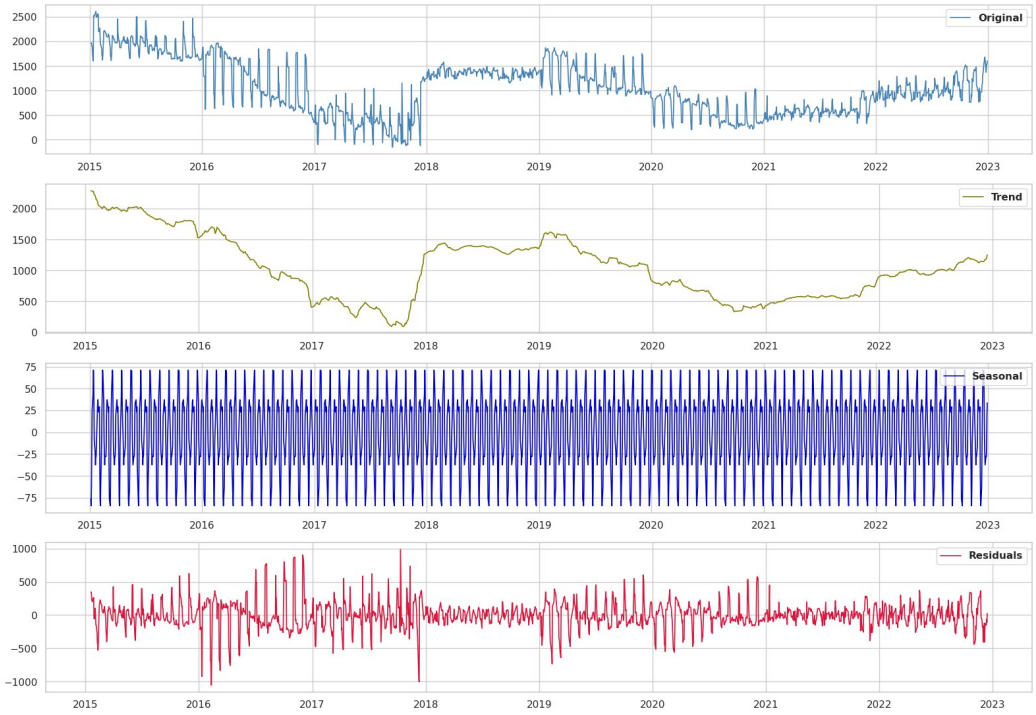
We already know the results while we prepared the Data earlier, all three tests unanimously confirm stationarity: ADF -12.496, PP -75.090, and KPSS 0.087. These results provide strong evidence that the liquidity series is integrated of order one, I (1), meaning the first difference achieves stationarity.

1.3. Seasonal Decomposition

An additive seasonal decomposition with period s = 22 was performed to isolate trend, seasonal, and residual components (Figure III-12). The seasonal component exhibits a regular, repeating pattern with amplitude approximately ±75 billion DZD, confirming strong monthly cycles linked to fiscal disbursements and salary payments. The trend reveals the structural shifts

documented in Section 3.1.2, while the residual component shows volatility clustering around the 2016 and 2020 breaks.

Figure III-12: Additive Seasonal Decomposition of Net Liquidity (s = 22) — Trend, Seasonal, and Residual Components

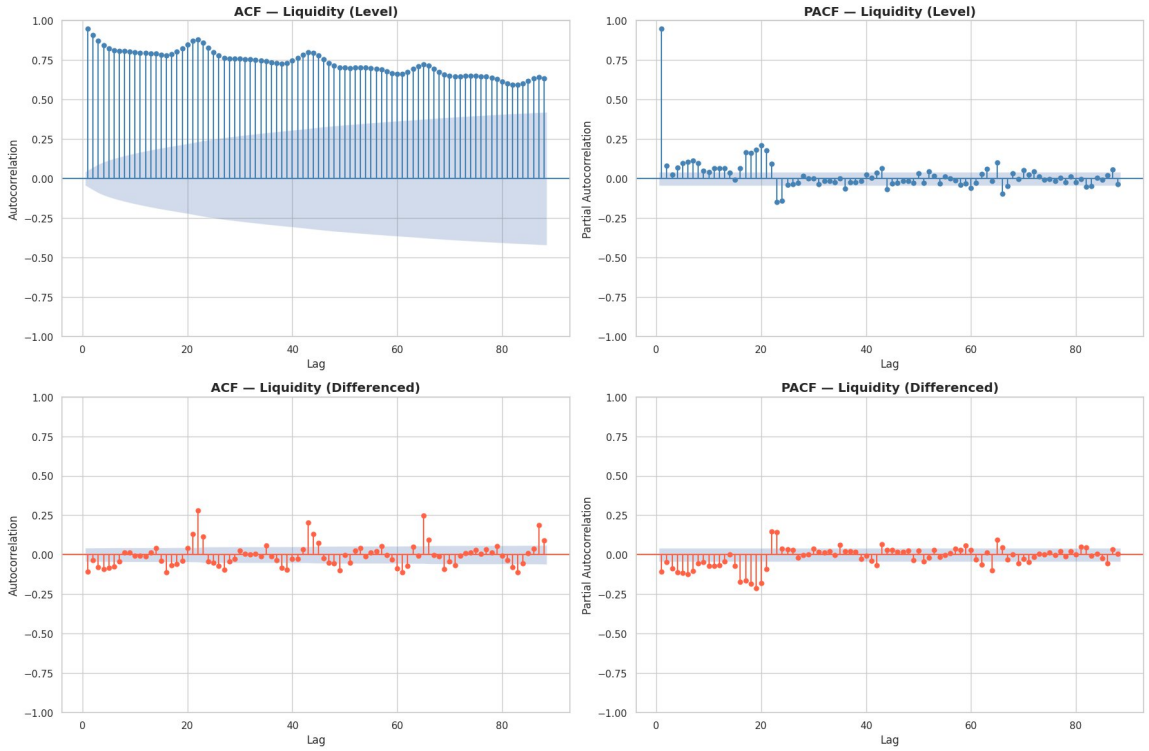


Source: Elaborated by the students using Python

1.4. Model Identification

1.4.1. ACF and PACF Analysis

Figure III-13: ACF and PACF of Liquidity — Level (top) and First Difference (bottom)



Source: Elaborated by the students using Python

The correlograms for the differenced series (Figure III-13) provide visual confirmation of the selected orders. The ACF of the differenced series shows significant spikes at lags 1, 2, and 3, as well as at the seasonal lags 22 and 44, confirming the presence of seasonal effects with s =22. Meanwhile, the PACF exhibits a gradual decay, with significant partial autocorrelations at the first few lags. This pattern, characterized by significant spikes in both the ACF and PACF at low lags, suggests the presence of both autoregressive and moving average components, although it does not allow for the precise identification of the optimal (p,q) orders.

1.4.2. Model Selection

Following the Box-Jenkins methodology, an automated grid search was conducted over the parameter space $p \in \{0,1,2,3\}$, $d = 1$ (fixed), $q \in \{0,1,2,3\}$, $P \in \{0,1,2\}$, $D \in \{0,1\}$, $Q \in \{0,1,2\}$, with seasonal period $s = 22$. Model selection was guided by minimization of the Akaike Information Criterion (AIC).

The optimal specification identified was SARIMA (3,1,3) (1,1,2) [22], with AIC = 20,837.55 and BIC = 20,891.29. This model contains three non-seasonal autoregressive terms, one non-seasonal difference, three non-seasonal moving-average terms, one seasonal

autoregressive term, one seasonal difference, and two seasonal moving-average terms at lag 22.

The necessity of both non-seasonal (d = 1) and seasonal (D = 1) integration confirms that the liquidity series contains both a stochastic trend and evolving seasonal patterns.

1.5. Model Estimation and Parameter Interpretation

The fitted SARIMA (3,1,3) (1,1,2) [22] model was estimated by maximum likelihood. Key parameter estimates are reported in Table III-7 (significant coefficients shown for brevity; full output available in appendix).

Table III-7: Selected Parameter Estimates (only significant ones) — SARIMA (3,1,3) (1,1,2) [22]

Parameter	Coef.	Std. Err.	P>|z|
ar. L1	1.5076	0.223	0.000
ma. L1	-1.8793	0.223	0.000
ma. L2	0.9261	0.432	0.032
ar.S. L22	0.3833	0.053	0.000
ma.S. L22	-1.0598	0.060	0.000
sigma²	2.668×10^{4}	486.003	0.000

Source: Elaborated by the students using Python

All key coefficients are statistically significant at the 1% level except ar. L2, ar. L3, ma. L3, and ma.S. L44. The large negative and highly significant ma. L1 (−1.8793) and ma.S. L22 (−1.0598) terms indicate strong short-term and seasonal moving-average dynamics that correct for past forecast errors. The positive and significant ar.S. L22 (0.3833) captures

$$y_t = 1.5076y_{t-1} + 0.3833y_{t-22} - 0.5779y_{t-23} + \epsilon_t - 1.8793\epsilon_{t-1} + 0.9261\epsilon_{t-2} - 1.0598\epsilon_{t-22} + 1.9917\epsilon_{t-23} - 0.9815\epsilon_{t-24}$$

The residual variance $\sigma^2 = 26{,}680$ (in billions of DZD)² is highly significant. The model successfully captures both the non-seasonal dynamics (via the ARMA (3,3) component) and the monthly seasonal cycle (via the SARMA (1,2) [22] component).

1.6. Diagnostic Checking of Residuals

After estimating the SARIMA (3,1,3) (1,1,2) [22] model, a comprehensive battery of residual diagnostic tests was applied to verify the model assumptions: the Ljung-Box test for autocorrelation, the Jarque-Bera and Shapiro-Wilk tests for normality, the ARCH-LM test for conditional heteroscedasticity, and the Breusch-Pagan test for heteroscedasticity.

For the Ljung-Box test:

- **Ho**: The residuals are independently distributed (no autocorrelation).
- **H1**: The residuals exhibit serial correlation.

A **p-value** < 0.05 leads to the rejection of the null hypothesis, indicating the presence of autocorrelation. A **p-value** > 0.05 means the null hypothesis cannot be rejected, implying that the residuals are white noise (Ljung & Box, 1978).

For the Jarque-Bera test:

- **Ho**: The residuals are normally distributed.
- **H1**: The residuals are not normally distributed.

A **p-value** < 0.05 leads to the rejection of the null hypothesis, indicating non-normality (skewness or kurtosis different from a normal distribution). A **p-value** > 0.05 means the null hypothesis cannot be rejected, suggesting the residuals follow a normal distribution (Jarque & Bera, 1987).

For the Shapiro-Wilk test:

- **Ho**: The residuals are normally distributed.
- **H1**: The residuals are not normally distributed.

A **p-value** < 0.05 leads to the rejection of the null hypothesis, indicating departures from normality. A **p-value** > 0.05 means the null hypothesis cannot be rejected, implying that the residuals are consistent with a normal distribution (Shapiro & Wilk, 1965).

For the ARCH-LM test:

- **Ho**: No conditional heteroscedasticity (no ARCH effects).
- **H1**: Conditional heteroscedasticity is present (ARCH effects exist).

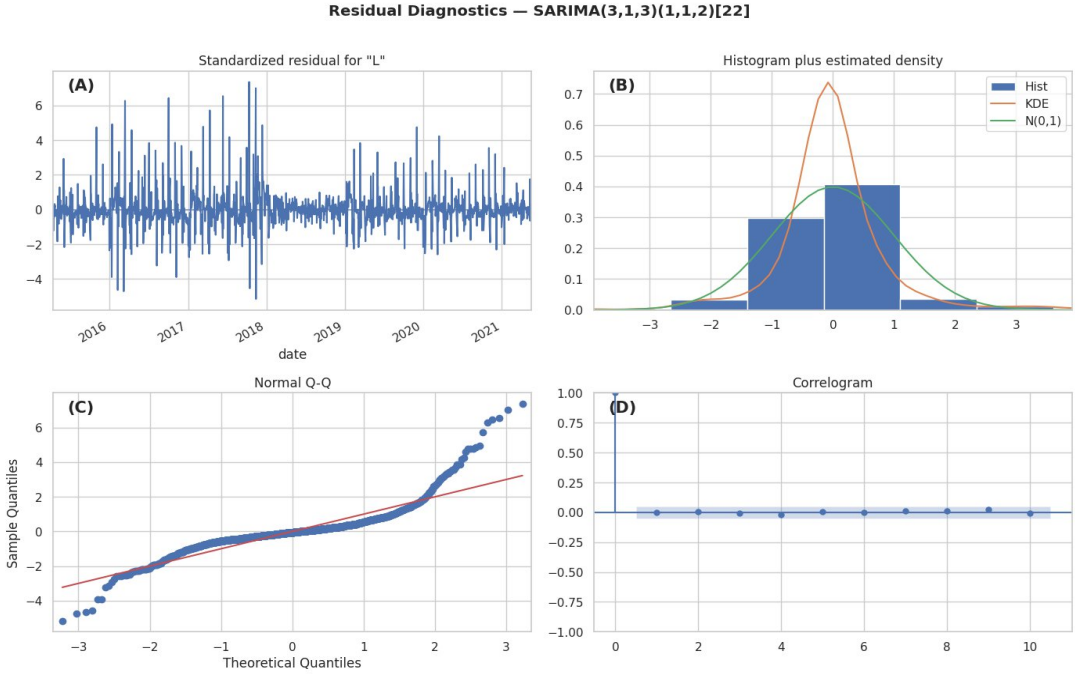
A **p-value** < 0.05 leads to the rejection of the null hypothesis, indicating the presence of volatility clustering. A **p-value** > 0.05 means the null hypothesis cannot be rejected, implying homoscedastic residuals (Engle, 1982).

For the Breusch-Pagan test:

- **Ho**: The residuals are homoscedastic (constant variance).
- **H1**: The residuals are heteroscedastic (variance is not constant).

A **p-value** < 0.05 leads to the rejection of the null hypothesis, indicating heteroscedasticity. A **p-value** > 0.05 means the null hypothesis cannot be rejected, suggesting the residuals have constant variance (Breusch & Pagan, 1979)

Figure III-14: Residual Diagnostics — SARIMA (3,1,3) (1,1,2) [22]



Source: Elaborated by the students using Python

The standardized residuals (Panel A from Figure III-14) appear stationary with no obvious patterns. However, the histogram and Q-Q plot (Panels B & C from Figure III-14) reveal clear departures from normality, confirmed by the Jarque-Bera test (statistic = 25,673.81, $p < 0.001$) and the Shapiro-Wilk test (statistic = 0.7673, $p < 0.001$). The Ljung-Box test indicates significant residual autocorrelation at lag 10 (statistic = 56.9977, $p = 0.000$), while the Durbin-Watson statistic of 2.1436 suggests limited first-order autocorrelation. The ARCH-LM test strongly detects the presence of volatility clustering (statistic = 178.13, $p < 0.001$), and the Breusch-Pagan test confirms heteroscedasticity (statistic = 32.88, $p < 0.001$). These findings indicate that while the linear SARIMA structure captures the conditional mean reasonably well, it cannot fully accommodate the heavy tails, serial correlation, and time-varying variance present in the data providing strong motivation for the deep learning architectures examined in Section 3 of chapter 3.

Table III-8: Residual Diagnostic Tests — SARIMA (3,1,3) (1,1,2) [22]

Test	Statistic	p-value	Ho	Decision
Ljung-Box (lag 10)	56.9977	0.000	No autocorrelation	Reject
Jarque-Bera	25673.81	0.000	Normality	Reject
Shapiro-Wilk	0.7673	0.000	Normality	Reject
ARCH-LM	178.13	0.000	No ARCH effects	Reject
Breusch-Pagan	32.88	0.000	Homoskedasticity	Reject

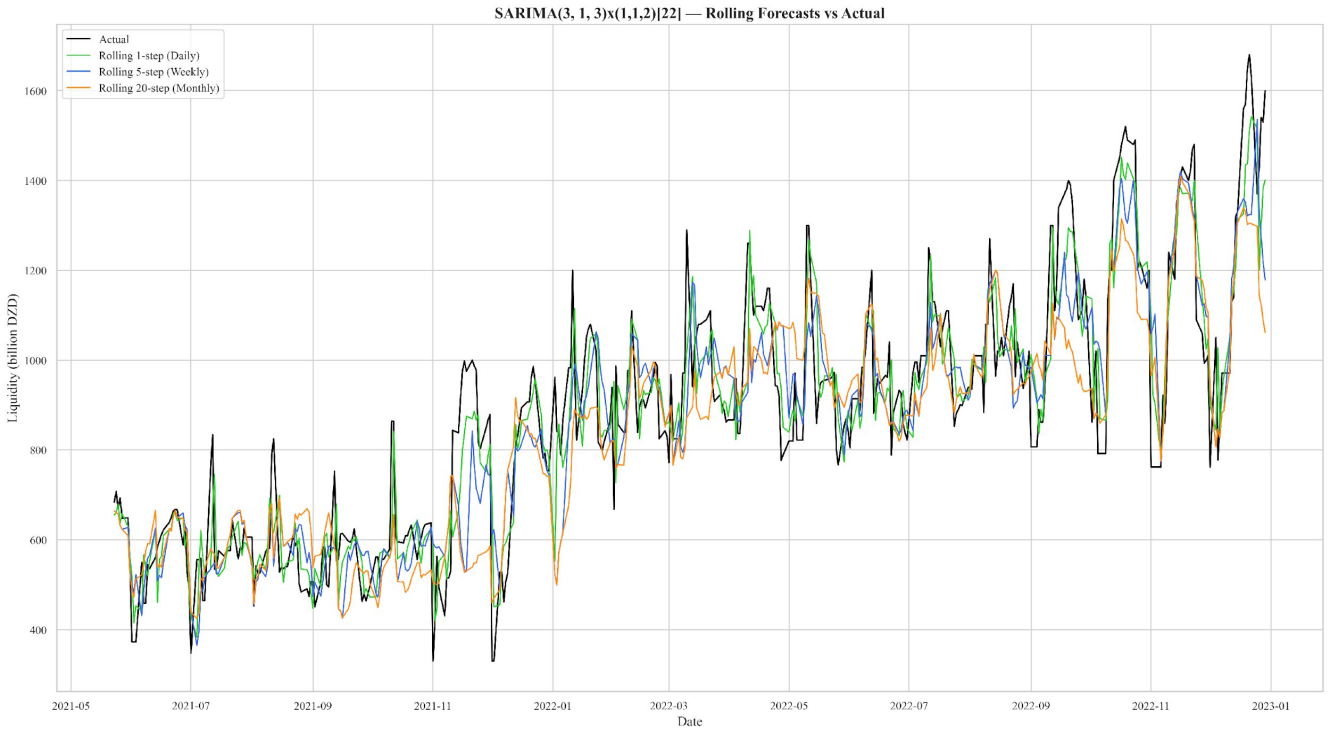
Source: Elaborated by the students using Python

Although the residual diagnostics reject the classical assumptions, the SARIMA model is retained because these tests assessed in-sample properties that do not directly govern out-of-sample accuracy. Rather than explicitly coupling the model with GARCH equations or Box-Cox transformations to manually resolve in-sample volatility, the rolling forecast strategy employed re-estimates the model at each step, allowing parameters to dynamically adapt to evolving patterns and mitigating the impact of the detected serial correlation and heteroscedasticity. Moreover, the model achieves competitive out-of-sample accuracy (R^2 = 0.8441 at the daily horizon) and serves as the necessary classical benchmark against which the deep learning architectures are evaluated.

1.7. Forecasting Performance

A rolling forecast strategy was implemented at three horizons corresponding to operational needs: h = 1 (day-ahead), h = 5 (week-ahead), and h = 20 (month-ahead). The model is re-estimated at each step as new observations become available, following reference Taoussi et al (2025). Results are presented in Figure III-15 and summarized in Table III-9

Figure III-15: Rolling Forecasts — Test Set Comparison (h = 1, 5, 20)



Source: Elaborated by the students using Python

Table III-9: Forecasting Performance Metrics — SARIMA (3,1,3) (1,1,2) [22]

Horizon	RMSE	R^2
Rolling 1-step (Day-ahead)	109.56	0.8441
Rolling 5 step (Week-ahead)	151.38	0.7024
Rolling 20 step (Month-ahead)	174.01	0.6067

Source: Elaborated by the students using Python

The rolling forecast strategy demonstrates strong and consistent predictive performance over short-term forecasting horizons (day-ahead and week-ahead forecasts). The 1-step daily forecast achieves a high level of accuracy, with an RMSE of 109.56 and an $R^2$of 0.8441. As expected, the forecasting error increases progressively as the forecast horizon expands: the 5-step weekly forecast records an RMSE of 151.38 and an $R^2$of 0.7024, while the 20-step monthly forecast yields an RMSE of 174.01 and an $R^2$of 0.6067. These results highlight the limitations of linear models over longer forecasting horizons and suggest the need to adopt nonlinear modeling approaches to improve long-term (month-ahead) liquidity forecasting accuracy at the Bank of Algeria.

2. Multivariate Integration: The SARIMAX Framework

Building upon the SARIMA (3,1,3) (1,1,2) [22] baseline developed in the previous subsection, this section extends the model by incorporating exogenous economic and institutional variables through the SARIMAX framework. The inclusion of Brent crude oil prices, the USD/DZD exchange rate, and calendar dummies allow the model to capture the influence of external shocks and seasonal institutional events on net banking system liquidity.

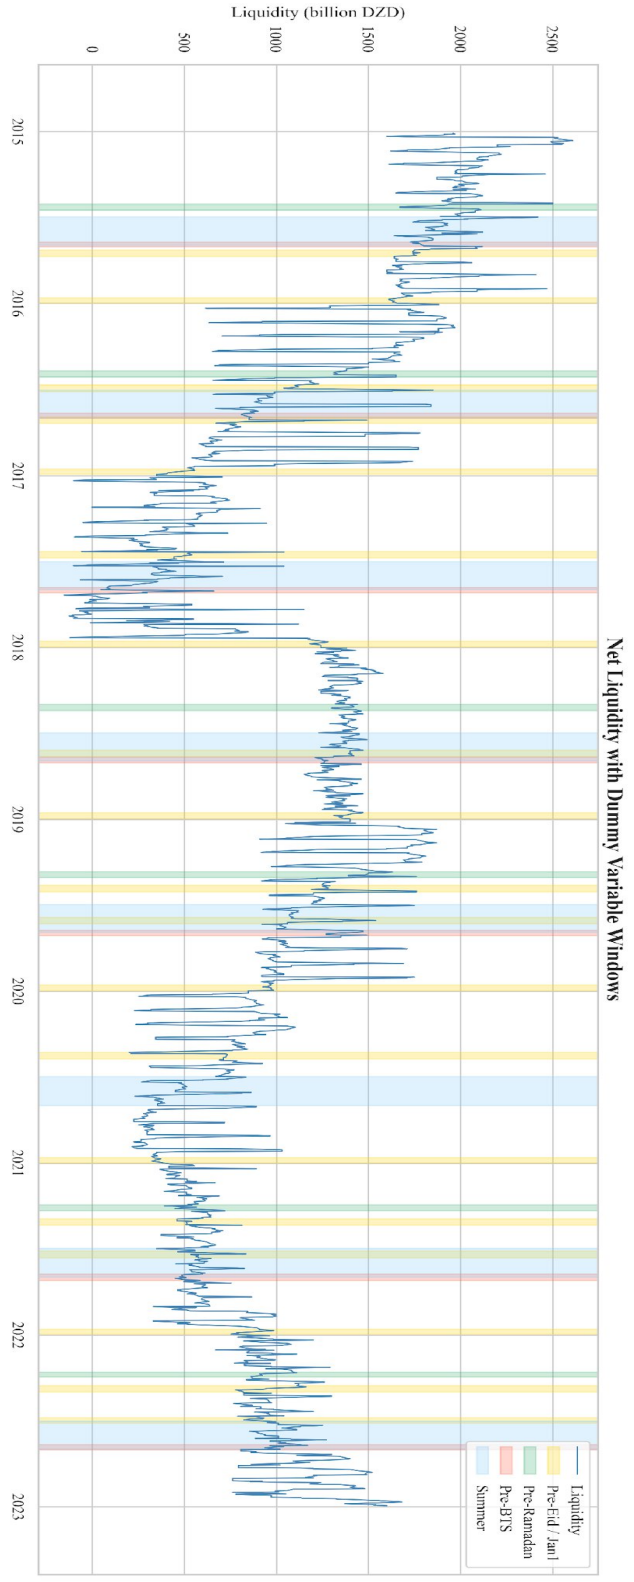
The Seasonal Autoregressive Integrated Moving Average with Exogenous Regressors (SARIMAX) model extends the SARIMA specification by incorporating a linear combination of exogenous variables. The general form is:

$$\Phi(B^s)\varphi(B)(1 - B)^d (1 - B^s)^D y_t = \Theta(B^s)\theta(B)\varepsilon_t + \beta' X_t$$

where X_t is a vector of exogenous regressors and β is the vector of associated coefficients. It should be noted that while this study adopts the standard SARIMAX nomenclature, advanced forecasting literature often frames this specification conceptually as a 'regression with ARMA errors' to avoid structural interpretability issues associated with ARIMAX coefficients. This specification allows the model to account for the effects of external economic conditions (oil prices and exchange rates) and institutional calendar events while preserving the seasonal dynamics captured by the SARIMA component. (NumXL, n.d.)

2.1. Description of Exogenous Variables

Figure III-16 Net Liquidity with Dummy Variable Windows (Pre-Eid, Pre-Ramadan, Pre-BTS, Summer)



Source: Elaborated by the students using Python

**Due to scaling issues, the full width image will be attached in appendix 21*

Six exogenous variables were incorporated into the SARIMAX model: two continuous macroeconomic indicators and four calendar dummies. Only three of the four calendar dummies were constructed using a 10-business-day lead window including each event and the event itself when it is not on a week-end (we take only 5 Days of work in our Database). This design reflects the economic reality that liquidity pressures arise from anticipatory cash withdrawals by households and firms in the days leading up to major religious holidays and institutional events, rather than only the event day itself. Earlier specifications using contemporaneous (day-of) dummies yielded statistically insignificant coefficients, confirming the superiority of the lead-window approach.

The four calendar dummies are defined as follows:

- **Pre_Eid**: A binary indicator equal to 1 during the 10 business days including Eid al-Fitr, Eid al-Adha, and 1 January, and the days of these events too (248 active days, 11.89% of the sample). This dummy captures the pronounced liquidity drain associated with religious festivities and year-end withdrawals.
- **Pre_Ramadan**: Equals 1 during the 10 business days before the start of Ramadan and its first day (80 active days, 3.84% of the sample).
- **Pre_BTS** (Back-to-School): Equals 1 during the 10 business days preceding the start of the academic year plus the day of the start (80 active days, 3.84% of the sample).
- **Summer**: A period dummy equal to 1 throughout July and August (354 active days, 17.0% of the sample), reflecting altered government payment schedules and reduced economic activity during the summer months.

The two continuous exogenous variables are:

- **Brent Crude Oil Price** (mean = $62.72/bbl): Higher oil prices increase hydrocarbon revenues that ultimately flow into the banking system via Sonatrach and the Treasury, with a lagged effect on net liquidity.
- **USD/DZD Exchange Rate** (mean = 119.74 DZD/USD): Persistent depreciation erodes the domestic value of foreign reserves and constrains the central bank's capacity to inject liquidity through the foreign exchange channel.

These variables were selected on theoretical grounds (see Chapter 2, Section 3) and data availability. Their inclusion allows the model to capture both external economic shocks and predictable institutional liquidity patterns.

2.2. Stationarity of the Continuous Exogenous Regressors

Because Brent and USDDZD enter the SARIMAX model as additive regressors on the level of the (non-stationary) dependent variable, their integration order must be verified.

Including non-stationary regressors alongside a non-stationary dependent variable without proper treatment can produce spurious regression results (Granger and Newbold, 1974). Accordingly, the same battery of stationarity tests applied earlier (ADF, Phillips–Perron, and KPSS) is conducted on Brent and USDDZD in levels.

Table III-10: Stationarity Tests on Liquidity, Brent and USDDZD — Levels and First Differences

Variable	Test	Statistic	p-value	Decision	Conclusion
Brent	ADF	−1.2641	0.6454	Fail to reject Ho	Non-stationary
Brent	Phillips-Perron	−5.8442	0.0000	Reject Ho	Stationary
Brent	KPSS	2.8804	0.0100	Reject Ho	Non-stationary
USDDZD	ADF	−1.5877	0.4898	Fail to reject Ho	Non-stationary
USDDZD	Phillips-Perron	−1.8946	0.3346	Fail to reject Ho	Non-stationary
USDDZD	KPSS	6.9450	0.0100	Reject Ho	Non-stationary

Source: Elaborated by the students using Python

Both continuous regressors are found to be integrated of order one. Brent in levels fails to reject the unit-root null according to ADF (ADF = −1.2641, p = 0.6454) and KPSS (KPSS = 2.8804, rejects stationarity), although the Phillips–Perron test suggests stationarity. USDDZD in levels is clearly non-stationary (ADF = −1.5877, p = 0.4898; PP = −1.8946, p = 0.3346; KPSS = 6.9450, rejects stationarity). Consequently, both series are first differenced before inclusion in the SARIMAX specification. After first differencing, the transformed series are stationary, as confirmed by all three tests. The binary calendar dummies, being bounded between 0 and 1, do not require differencing.

2.3. Model Specification and Selection

An automated grid search was conducted over the same parameter space used for SARIMA, now including the six exogenous regressors. The optimal specification selected was SARIMAX (2,1,3) (2,1,2) [22], with AIC = 20,762.29 and BIC = 20,848.27. This represents a meaningful improvement over the pure SARIMA model (AIC = 20,837.55), confirming that the exogenous variables add explanatory power even after accounting for the additional parameters.

2.4. Parameter Estimation and Interpretation

The model was estimated by maximum likelihood. Key parameter estimates for the exogenous variables, and selected (significant) ARMA terms are reported in Table III-11.

Table III-11: Selected SARIMAX (2,1,3) (2,1,2) [22] Parameter Estimates (only significant parameters)

Variable	Coef.	Std. Err.	z	P>|z|
Brent	-1.7125	0.4474	-3.828	0.0001
USD/DZD	-15.9229	1.2385	-12.856	0.0000
ar. L1	1.6504	0.0354	46.595	0.0000
ar. L2	-0.7176	0.0302	-23.734	0.0000
ma. L1	-1.9975	0.0379	-52.764	0.0000
ma. L2	1.1464	0.0578	19.844	0.0000
ma. L3	-0.1356	0.0306	-4.435	0.0000
ma.S. L22	-0.6628	0.2828	-2.344	0.0191
sigma²	25626.84	450.55	56.880	0.0000

Source: Elaborated by the students using Python

Majority of autoregressive and moving-average parameters are statistically significant at the 1% level, confirming that the SARIMAX (2,1,3) (2,1,2) [22] specification successfully captures both short-term dynamics and the 22-day seasonal cycle in net liquidity. The large positive coefficient on ar. L1 (1.6504) combined with the strongly negative ma. L1 (−1.9975) indicates rapid mean reversion in daily liquidity changes. The significant negative seasonal moving-average term at lag 22 (ma.S. L22 = −0.6628, p = 0.019) further validates the inclusion of the multiplicative seasonal component.

Both continuous exogenous regressors are highly statistically significant ($p < 0.001$). The negative coefficient on Brent (−1.7125) indicates that higher oil prices are associated with lower net liquidity in the short run consistent with the lagged fiscal transmission mechanism whereby hydrocarbon revenues accrue through Sonatrach and the Treasury before affecting banking system liquidity. The strongly negative coefficient on USD/DZD (−15.9229) confirms that dinar depreciation erodes the domestic valuation of foreign reserves and constrains the central bank's capacity to inject liquidity.

Among the calendar dummies, only pre-BTS approaches conventional significance (p = 0.096); the coefficients on Pre-Eid, Pre-Ramadan, and Summer are statistically insignificant. This suggests that the multiplicative seasonal ARMA terms already capture most of the predictable liquidity variation associated with religious and institutional events.

The estimated model can be expressed in two parts. First, the exogenous component is removed from the observed series:

$$\eta_t = Y_t - (-1.7125 \cdot \text{Brent}_t) - (-15.9229 \cdot \Delta\text{USDDZD}_t) - (-20.4294 \cdot \text{PreEid}_t) - (-13.1907 \cdot \text{PreRamadan}_t) - (66.3030 \cdot \text{PreBTS}_t) - (27.1424 \cdot \text{Summer}_t),$$

The disturbance term n_t follows a SARIMA (2,1,3) (2,1,2) [22] process, combining non-seasonal and seasonal autoregressive and moving average components. The model includes first-order regular and seasonal differencing to ensure stationarity, while the seasonal period of 22 captures monthly patterns in daily data. The error term ε_t is assumed to be white noise with variance σ²=2.563×10^4 Expressed as:

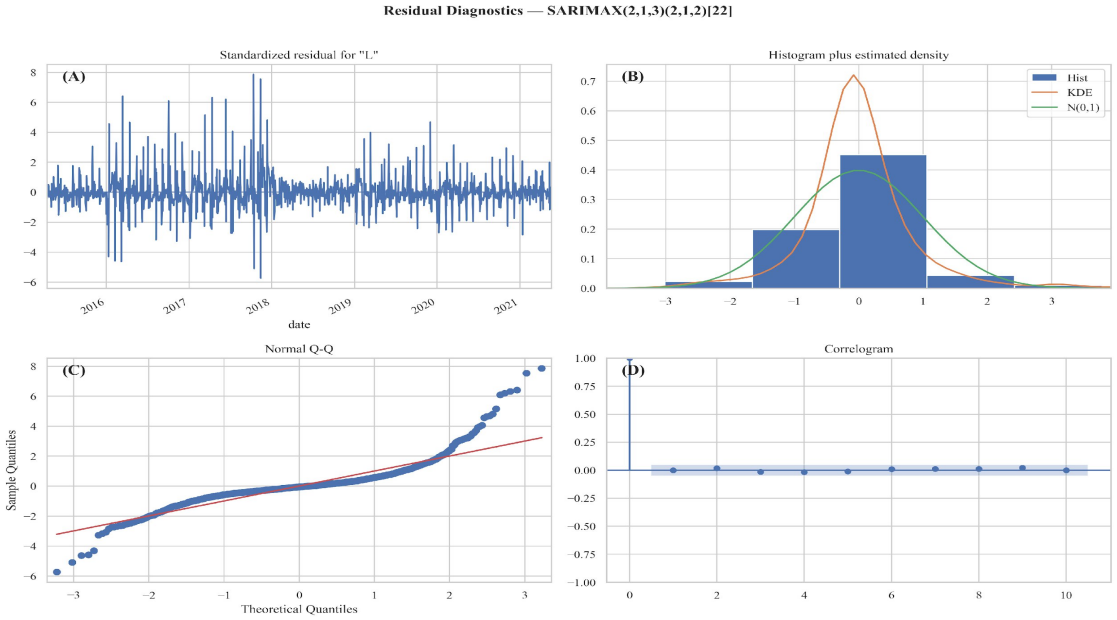
$$\eta_t = [(1 - 1.9975B + 1.1464B^2 - 0.1356B^3)(1 - 0.6628B^{22} - 0.3000B^{44})] / [(1 - 1.6504B + 0.7176B^2)(1 + 0.0377B^{22} - 0.1339B^{44})(1 - B)(1 - B^{22})]\ \varepsilon_t$$

$$\text{with } \varepsilon_t \sim WN(0, \sigma^2) \text{ and } \sigma^2 = 2.563 \times 10^4.$$

2.5. Diagnostic Checking of Residuals

Residual diagnostics for the SARIMAX model are presented in Figure III-17 and Table III-12. The residuals remain stationary (ADF = −7.2105, Phillips-Perron = −45.3102, KPSS = 0.0644). However, the Ljung-Box test still detects significant autocorrelation at multiple lags, and both normality tests (Jarque-Bera = 27,235.10, Shapiro-Wilk = 0.7908) strongly reject the Gaussian assumption. The ARCH-LM (159.95, $p < 0.001$) and Breusch-Pagan (21.95, $p < 0.001$) tests confirm the continued presence of volatility clustering and heteroscedasticity. These findings indicate that while the exogenous variables improve in-sample fit, they do not resolve the fundamental nonlinearities and time-varying variance in the data.

Figure III-17 Residual Diagnostics SARIMAX (2,1,3) (2,1,2) [22]



Source: Elaborated by the students using Python

Table III-12: Residual Diagnostic Tests SARIMAX (2,1,3) (2,1,2) [22]

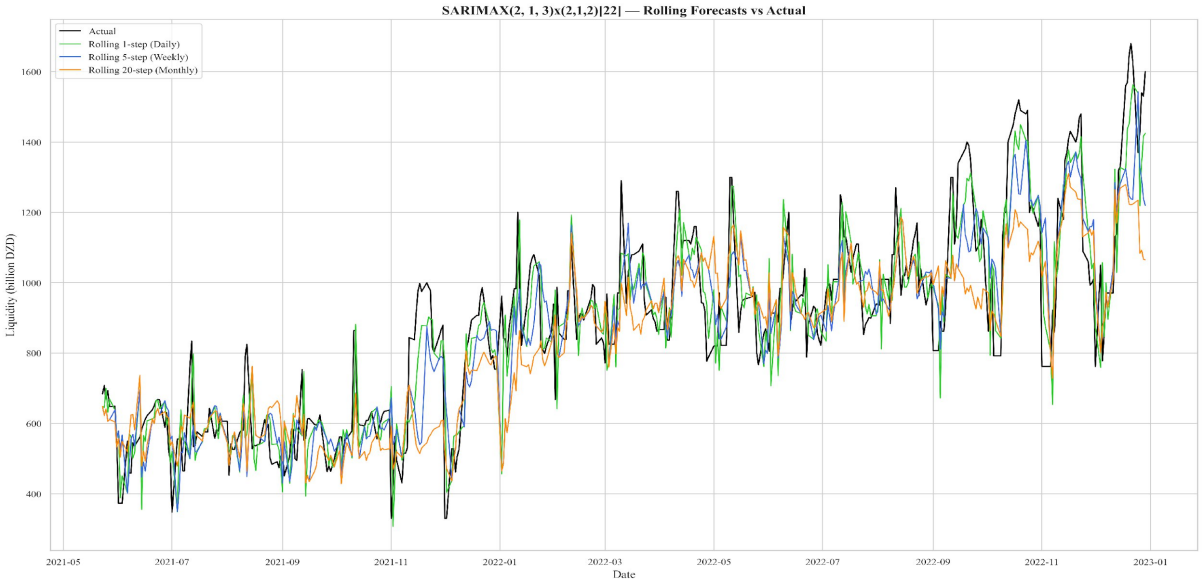
Test	Statistic	p-value	Ho	Decision
Ljung-Box (lag 10)	21.05	0.021	No autocorrelation	Reject
Jarque-Bera	27,235.10	0.000	Normality	Reject
Shapiro-Wilk	0.7908	0.000	Normality	Reject
ARCH-LM	159.95	0.000	No ARCH effects	Reject
Breusch-Pagan	21.95	0.000	Homoskedasticity	Reject

Source: Elaborated by the students using Python

2.6. Forecasting Performance

Out-of-sample performance was evaluated using rolling forecast strategies on the test set (May 2021 – December 2022). Results are summarized in Table III-13 and visualized in Figure III-18.

Figure III-18 SARIMAX (2,1,3) (2,1,2) [22] Rolling Forecasts (h = 1, 5, 20)



Source: Elaborated by the students using Python

Table III-13: Forecasting Performance Metrics SARIMAX (2,1,3) (2,1,2) [22]

Horizon	RMSE	R^2
Rolling 1-step (Day-ahead)	115.68	0.8262
Rolling 5-step (Week-ahead)	145.97	0.7233
Rolling 20-step (Month-ahead)	166.14	0.6415

Source: Elaborated by the students using Python

The rolling 1-step forecast achieves strong accuracy (RMSE = 115.68, R^2 = 0.8262). Performance declines at longer horizons, as expected, but the 20-step forecast remains operationally useful (RMSE = 166.14, R^2 = 0.6415). Compared with the pure SARIMA model, SARIMAX delivers modest improvements in directional accuracy during periods of oil price volatility, although the linear treatment of exogenous variables limits its ability to capture state-dependent interactions.

3. Comparative Assessment: SARIMA vs SARIMAX

The SARIMAX estimation yields three main conclusions. First, the inclusion of Brent and USD/DZD significantly improves in-sample fit (ΔAIC = −75.26) and confirms their theoretical importance as liquidity drivers. Second, the calendar dummies add little explanatory power once the seasonal SARIMA structure is in place; only pre-BTS approaches significance. Third, while SARIMAX modestly enhances forecast accuracy at the weekly and monthly horizons, residual diagnostics reveal that autocorrelation, non-normality, and volatility clustering persist. These limitations underscore the need for nonlinear, data-driven approaches capable of

modeling complex interactions and conditional heteroscedasticity motivating the deep learning architectures evaluated in Section 3.

Table III-14: Comparative Forecasting Performance — SARIMA vs SARIMAX

Forecast Horizon	Model	RMSE	R^2
1-step (day-ahead)	SARIMA	109.56	0.8441
	SARIMAX	115.68	0.8262
5-step (week-ahead)	SARIMA	151.38	0.7024
	SARIMAX	145.97	0.7233
20-step (month-ahead)	SARIMA	174.01	0.6067
	SARIMAX	166.14	0.6415

Source: Elaborated by the students

At the daily horizon, the pure SARIMA outperforms the SARIMAX on both RMSE (109.56 vs 115.68) and R^2 (0.8441 vs 0.8262), indicating that the exogenous regressors introduce noise at high frequency. At the weekly and monthly horizons, the SARIMAX reverses the ranking: its RMSE falls below that of the SARIMA (145.97 vs 151.38 and 166.14 vs 174.01, respectively), and its R^2 rises (0.7233 vs 0.7024 and 0.6415 vs 0.6067). The exogenous variables thus become beneficial once the forecast window is long enough for their cumulative signal to outweigh estimation variance. Nevertheless, the R^2 remains below 0.75 at both horizons, underscoring the limitations of the linear framework for medium-term liquidity forecasting.

Section 3: Forecasting using Deep Learning Models

Following the construction of the classical seasonal baseline in Section 2, this section applies three deep learning architectures the convolutional neural network (CNN), the long short-term memory network (LSTM) and the gated recurrent unit (GRU) to the same net liquidity series of the Bank of Algeria. Each architecture is examined in both a univariate specification and an exogenous variant that incorporates the same six regressors used in the SARIMAX model (Brent oil price, USD/DZD exchange rate, and four calendar dummies: Eid, Ramadan, Back to school and Summer season). The theoretical foundations of these architectures were presented in (Chapter 2, Sections 2) The purpose here is to determine whether models capable of representing non-linear and long-range temporal structure can improve on the SARIMA and SARIMAX benchmarks, and to identify which architecture delivers the most reliable forecasts at the operational horizons of interest.

1. Data preparation and experimental design

The modelling target is the daily net liquidity position of the banking system, observed over the Algerian business week (Sunday to Thursday) from January 2015 to December 2022, for a total of 2,085 observations. To convert the series into a supervised learning problem, a sliding window of twenty-two business days is used: the value to be predicted is associated with the twenty-two preceding observations, the window of twenty-two corresponding to the monthly liquidity cycle already identified in the seasonal analysis. The sample is divided chronologically into three contiguous blocks: the most recent twenty windows form the test set, the twenty preceding windows form the validation set, and all earlier observations constitute the training set. The input features are scaled to the unit interval with a min-max transformation fitted exclusively on the training block, and the predictions are returned to their original scale before any error metric is computed.

For each architecture, two specifications are estimated. The univariate specification — CNN, LSTM and GRU uses only the past values of the liquidity series. The exogenous specification denoted CNNX, LSTMX and GRUX augments the input with the same six exogenous regressors used in the SARIMAX model of Section 2. Every model is trained at three forecasting horizons: one day ahead, five days ahead (one week) and twenty days ahead (one month).

2. Model configuration

To keep the comparison fair, each architecture is given a single, fixed configuration applied uniformly across the three horizons; the random seeds are also fixed so that the results are reproducible. Training is governed by the Adam optimiser with the mean squared error as the loss, and an early-stopping rule monitors the validation loss and restores the best weights once the loss stops improving. The selected configurations are reported in Table III-15.

Table III-15: Selected configuration of the deep learning models

Model	Units / filters	Dropout	Activation	Optimizer	Max. epochs
CNN	64 / 32 filters; 100 dense	0.20	Tanh	Adam	100
CNNX	64 / 32 filters; 100 dense	0.40	Tanh	Adam	200
LSTM	100 (two layers)	0.45	Tanh	Adam	100
LSTMX	100 (two layers)	0.45	Tanh	Adam	100
GRU	100 (two layers)	0.45	Tanh	Adam	100
GRUX	100 (two layers)	0.20	Tanh	Adam	100

Source: Elaborated by the students

The CNN models combine two one-dimensional convolution layers (sixty-four and thirty-two filters, kernel width seven), followed by dropout, flattening and a dense layer of one hundred units. The LSTM and GRU models share a common two-layer recurrent design, each layer containing one hundred units, with dropout after each recurrent layer.

2.1. Convolutional Neural Network (CNN)

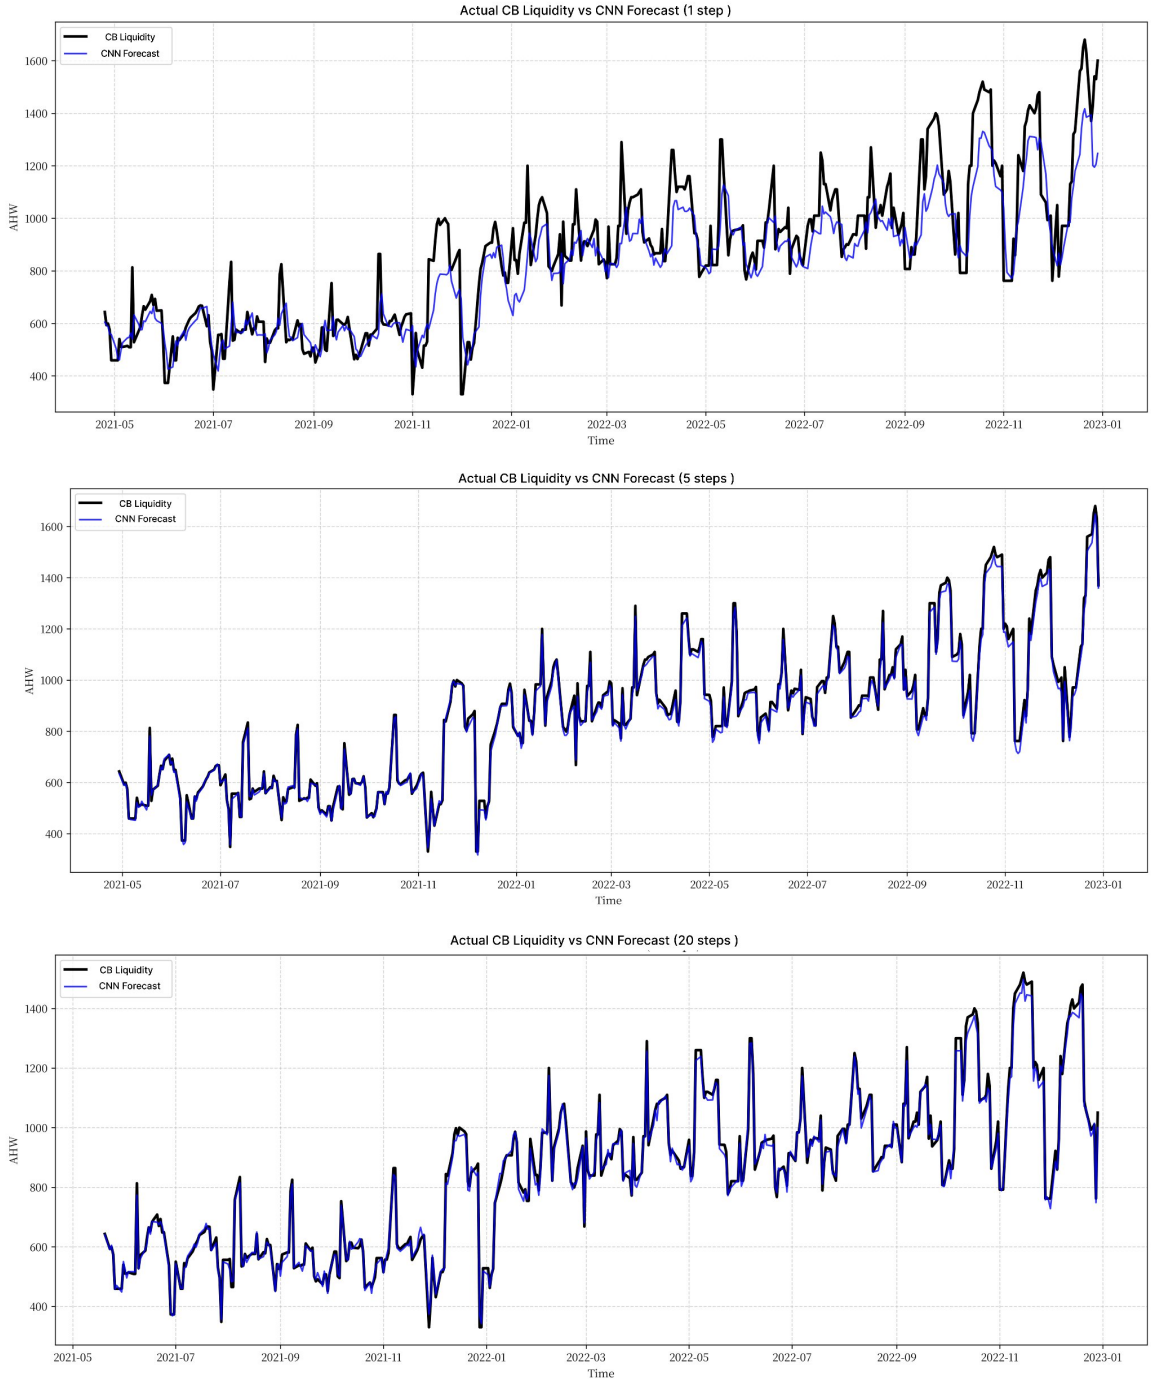
The univariate convolutional network uses only the past values of liquidity, organised into windows of twenty-two business days. Table III-16 reports the out-of-sample results. At the daily horizon it explains close to seventy-nine per cent of the variance (RMSE = 126.83). The weekly and monthly horizons are predicted with remarkable precision: $R^2 = 0.9967$ (RMSE = 15.64) and $R^2 = 0.9849$ (RMSE = 32.06) respectively. The very high accuracy at longer horizons reflects the smoothing that aggregation introduces, making the dominant cycle easier for the convolutional filters to reproduce.

Table III-16: Forecasting performance of the CNN model

Forecast horizon	RMSE	R^2
1-step (day-ahead)	126.83	0.7940
5-step (week-ahead)	15.64	0.9967
20-step (month-ahead)	32.06	0.9849

Source: generated by the students using Python

Figure III-19 CNN forecasts versus actual liquidity — 1-step, 5-step and 20-step horizons



Source: generated by the students using Python

2.2. Convolutional Neural Network with exogenous inputs (CNNX)

The CNNX keeps the architecture of the univariate CNN unchanged but enlarges the input with the six exogenous regressors. Table III-17 shows that, contrary to what might be expected, the exogenous inputs weaken the convolutional model at every horizon. The daily R^2 falls from 0.7940 to 0.3365, the weekly R^2 drops from 0.9967 to 0.7088, and the monthly R^2 declines from 0.9849 to 0.9474. The one-dimensional filters are forced to process inputs that dilute, rather than sharpen, the signal; the effect is most severe at the daily horizon and

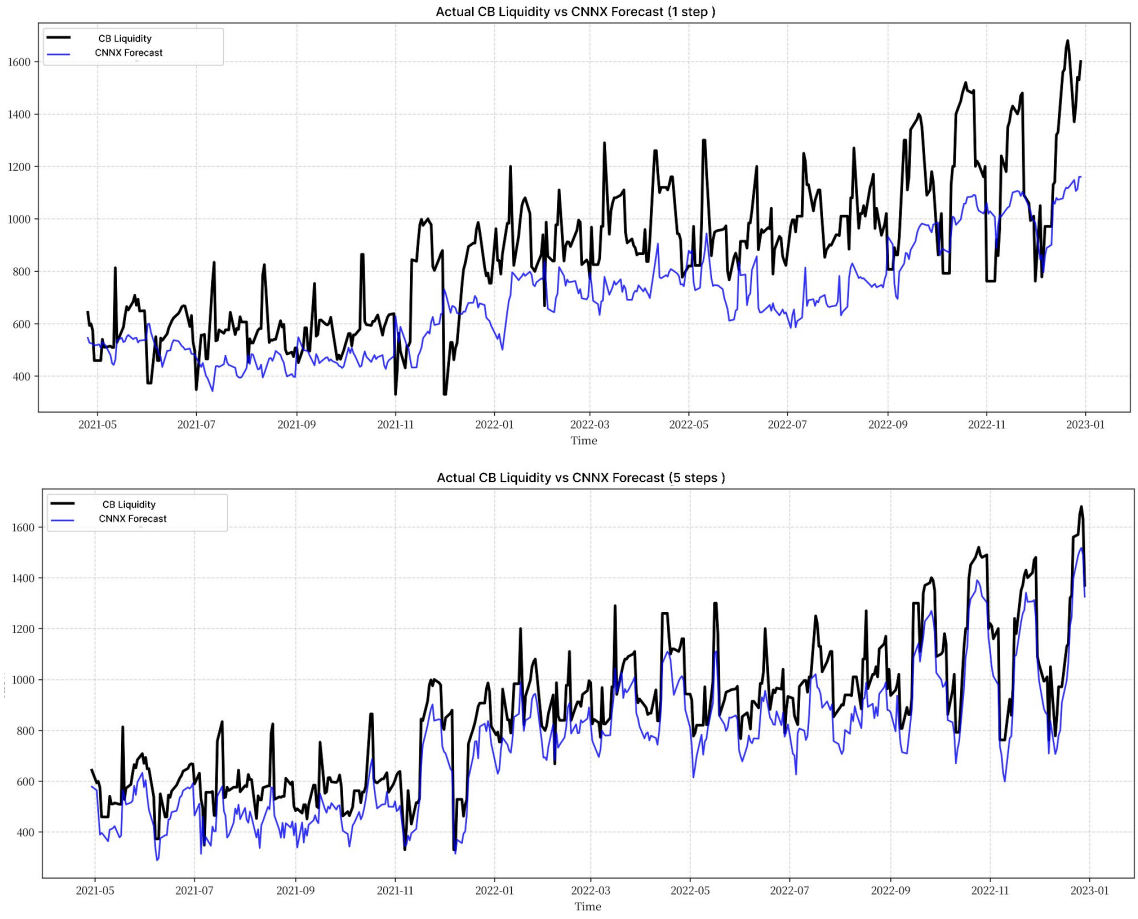
diminishes at the monthly horizon, where the dominant seasonal pattern remains strong enough to be captured despite the added noise.

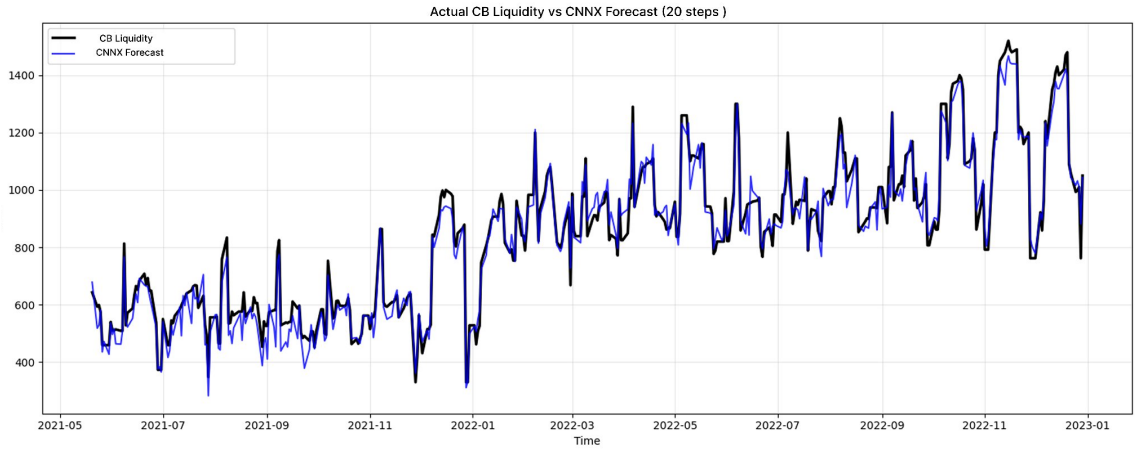
Table III-17: Forecasting performance of the CNNX model

Forecast horizon	RMSE	R^2
1-step (day-ahead)	227.61	0.3365
5-step (week-ahead)	147.52	0.7088
20-step (month-ahead)	59.88	0.9474

Source: generated by the students using Python

Figure III-20 : CNNX forecasts versus actual liquidity — 1-step, 5-step and 20-step horizons





Source: generated by the students using Python

2.3. Comparative Assessment: CNN vs CNNX

Table III-18: Comparative forecasting performance — CNN vs CNNX

Forecast horizon	Model	RMSE	R^2
1-step (day-ahead)	CNN	126.83	0.7940
	CNNX	227.61	0.3365
5-step (week-ahead)	CNN	15.64	0.9967
	CNNX	147.52	0.7088
20-step (month-ahead)	CNN	32.06	0.9849
	CNNX	59.88	0.9474

Source: Elaborated by the students

The comparison is unambiguous: the univariate CNN outperforms the CNNX at every horizon. The gap is most pronounced at the daily horizon, where the exogenous inputs cause the R^2 to collapse from 0.7940 to 0.3365, and at the weekly horizon, where the RMSE increases nearly tenfold. Even at the monthly horizon, where the CNNX partially recovers, its RMSE remains almost double that of the univariate model. The best convolutional model across all horizons is the univariate CNN.

Table III-19: Best convolutional model by forecasting horizon

Forecast horizon	Best model	RMSE	R^2
1-step (day-ahead)	CNN	126.83	0.7940
5-step (week-ahead)	CNN	15.64	0.9967
20-step (month-ahead)	CNN	32.06	0.9849

Source: Elaborated by the students

2.4. Long Short-Term Memory network (LSTM)

The LSTM is built from two stacked layers of one hundred units each, with dropout regularisation, trained on univariate windows. The results differ markedly from the convolutional models: the LSTM explains about sixty-three per cent of the variance at the daily horizon and fifty-six per cent at the weekly horizon, while the monthly horizon is handled much

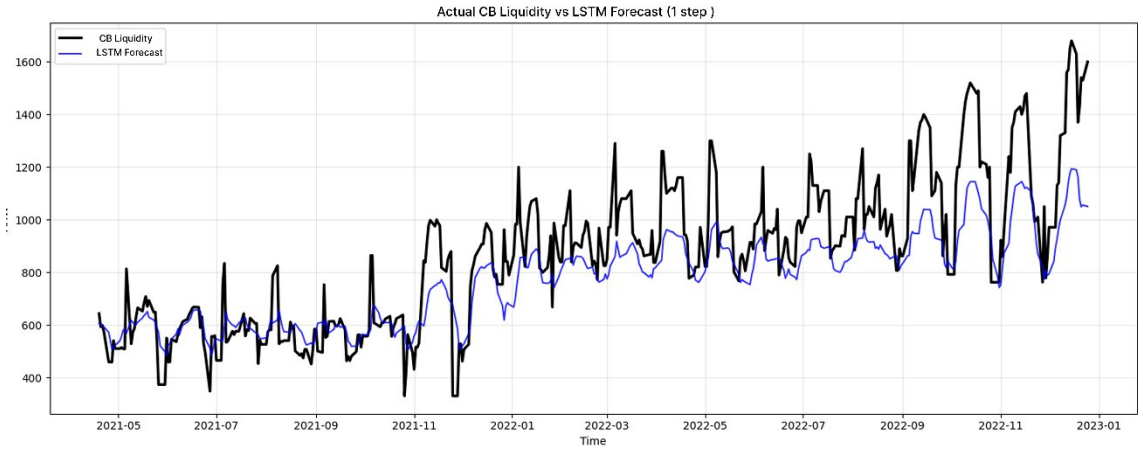
better ($R^2 = 0.9132$, RMSE = 76.94). The strong regularisation needed to stabilise the two-layer recurrent network limits its ability to react to rapid daily and weekly movements.

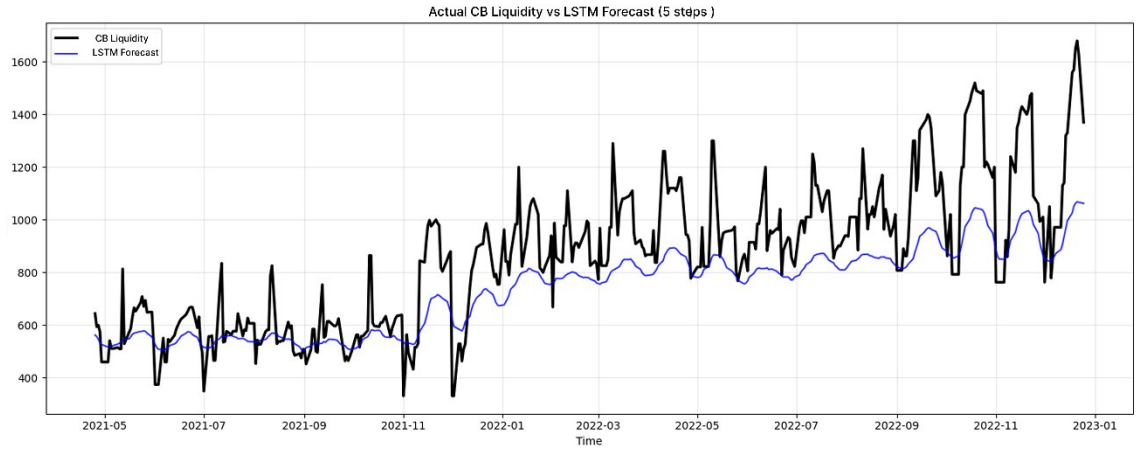
Table III-20: Forecasting performance of the LSTM model

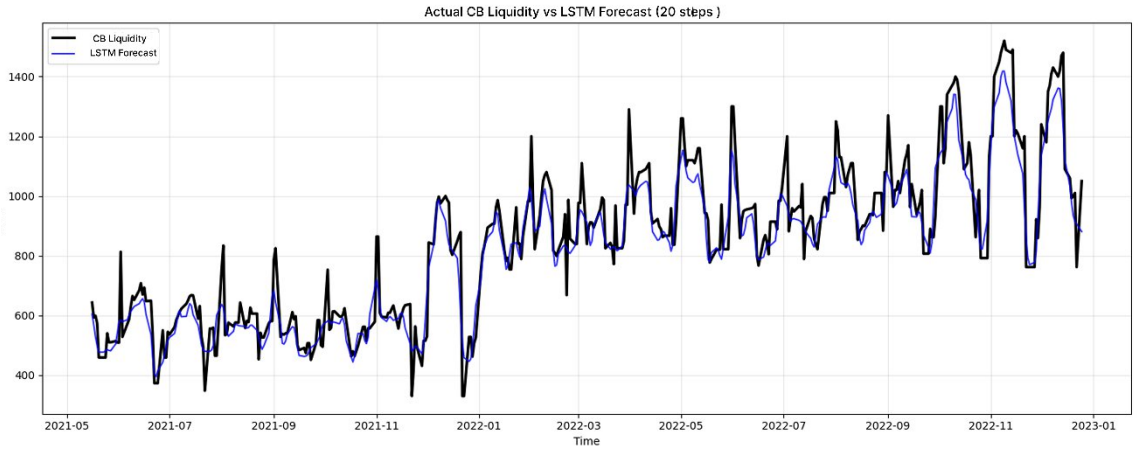
Forecast horizon	RMSE	R^2
1-step (day-ahead)	170.33	0.6284
5-step (week-ahead)	180.45	0.5643
20-step (month-ahead)	76.94	0.9132

Source: generated by the students using Python

Figure III-21 LSTM forecasts versus actual liquidity — 1-step, 5-step and 20-step horizons







Source: generated by the students using Python

2.5. Long Short-Term Memory network with exogenous inputs (LSTMX)

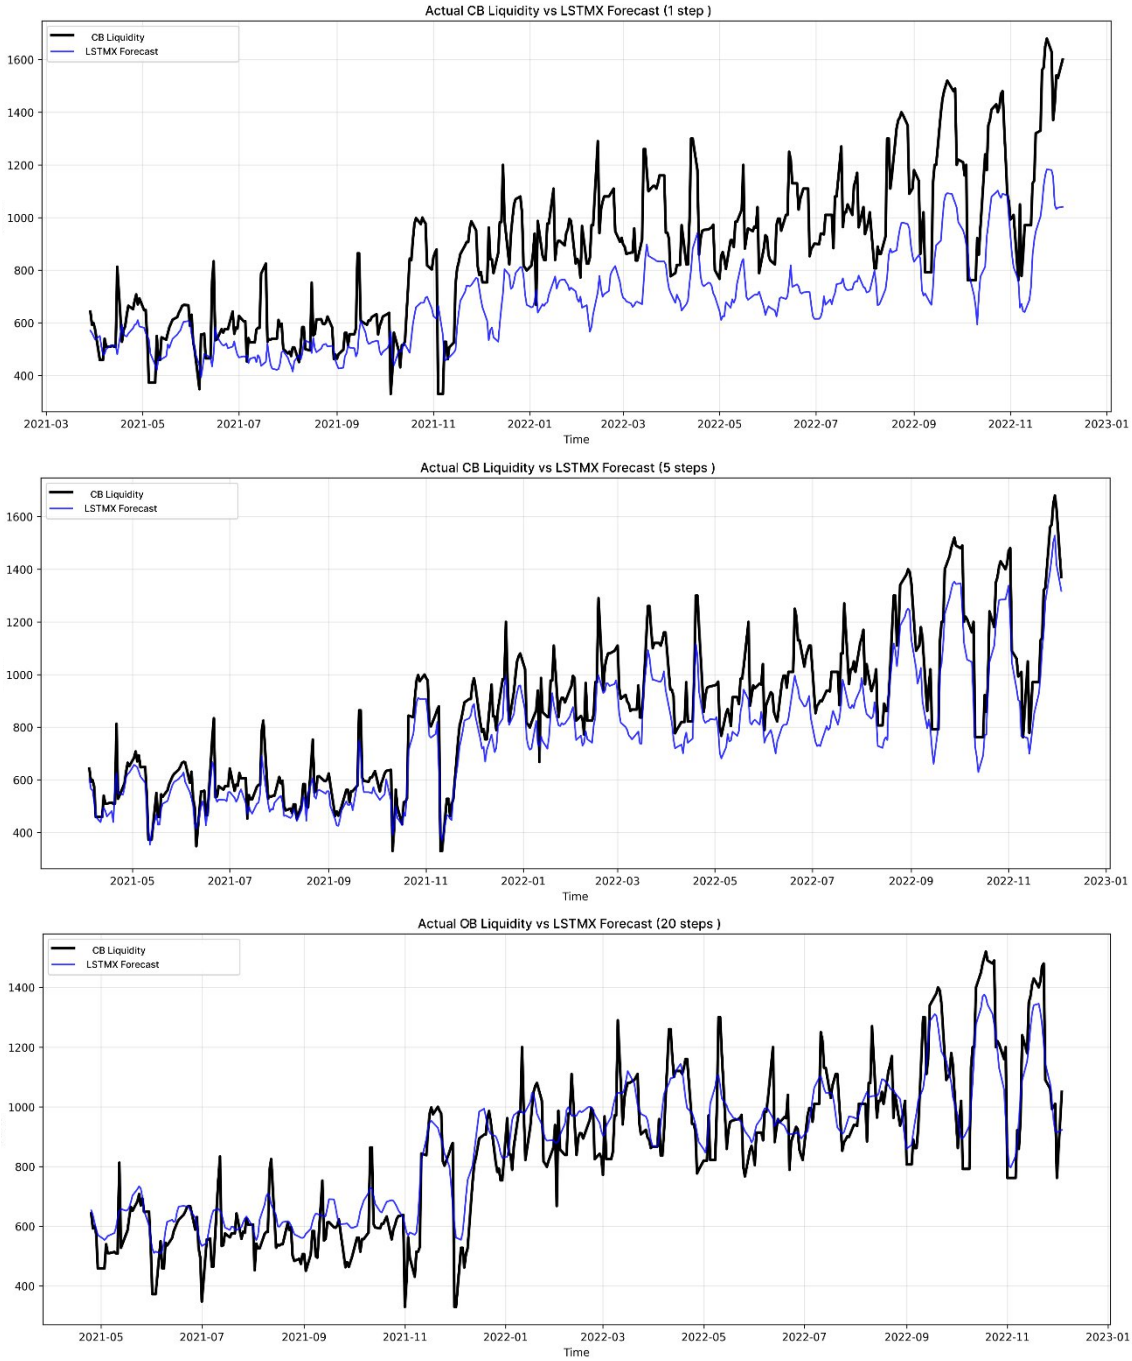
The LSTMX adds the six exogenous regressors to the recurrent architecture. The exogenous inputs do not improve the model: at the daily horizon the R^2 falls to 0.2501 (RMSE = 241.97), a clear deterioration. The weekly horizon improves modestly (R^2 = 0.6033 vs 0.5643), while the monthly horizon weakens marginally (R^2 = 0.8615). The additional regressors introduce more noise than signal, particularly at the daily horizon where the model must allocate capacity to inputs whose short-term relationship with liquidity is weak.

Table III-21: Forecasting performance of the LSTMX model

Forecast horizon	RMSE	R^2
1-step (day-ahead)	241.97	0.2501
5-step (week-ahead)	172.19	0.6033
20-step (month-ahead)	97.20	0.8615

Source: Generated by the students using Python

Figure III-22 LSTMX forecasts versus actual liquidity — 1-step, 5-step and 20-step horizons



Source: Generated by the students using Python

2.6. Comparative Assessment: LSTM vs LSTMX

Table III-22: Comparative forecasting performance — LSTM vs LSTMX

Forecast horizon	Model	RMSE	R^2

1-step (day-ahead)	LSTM	170.33	0.6284
	LSTMX	241.97	0.2501
5-step (week-ahead)	LSTM	180.45	0.5643
	LSTMX	172.19	0.6033
20-step (month-ahead)	LSTM	76.94	0.9132
	LSTMX	97.20	0.8615

Source: Elaborated by the students

The exogenous inputs are predominantly detrimental to the LSTM. The daily R^2 drops from 0.6284 to 0.2501; the weekly horizon shows only a modest improvement (R^2 = 0.6033 vs 0.5643); and the monthly horizon reverts to deterioration. On balance, the univariate LSTM is the more reliable specification, and the best model of the LSTM family at the daily and monthly horizons. The LSTMX retains a slight edge only at the weekly horizon.

Table III-23: Best LSTM-family model by forecasting horizon

Forecast horizon	Best model	RMSE	R^2
1-step (day-ahead)	LSTM	170.33	0.6284
5-step (week-ahead)	LSTMX	172.19	0.6033
20-step (month-ahead)	LSTM	76.94	0.9132

Source: Elaborated by the students

2.7. Gated Recurrent Unit (GRU)

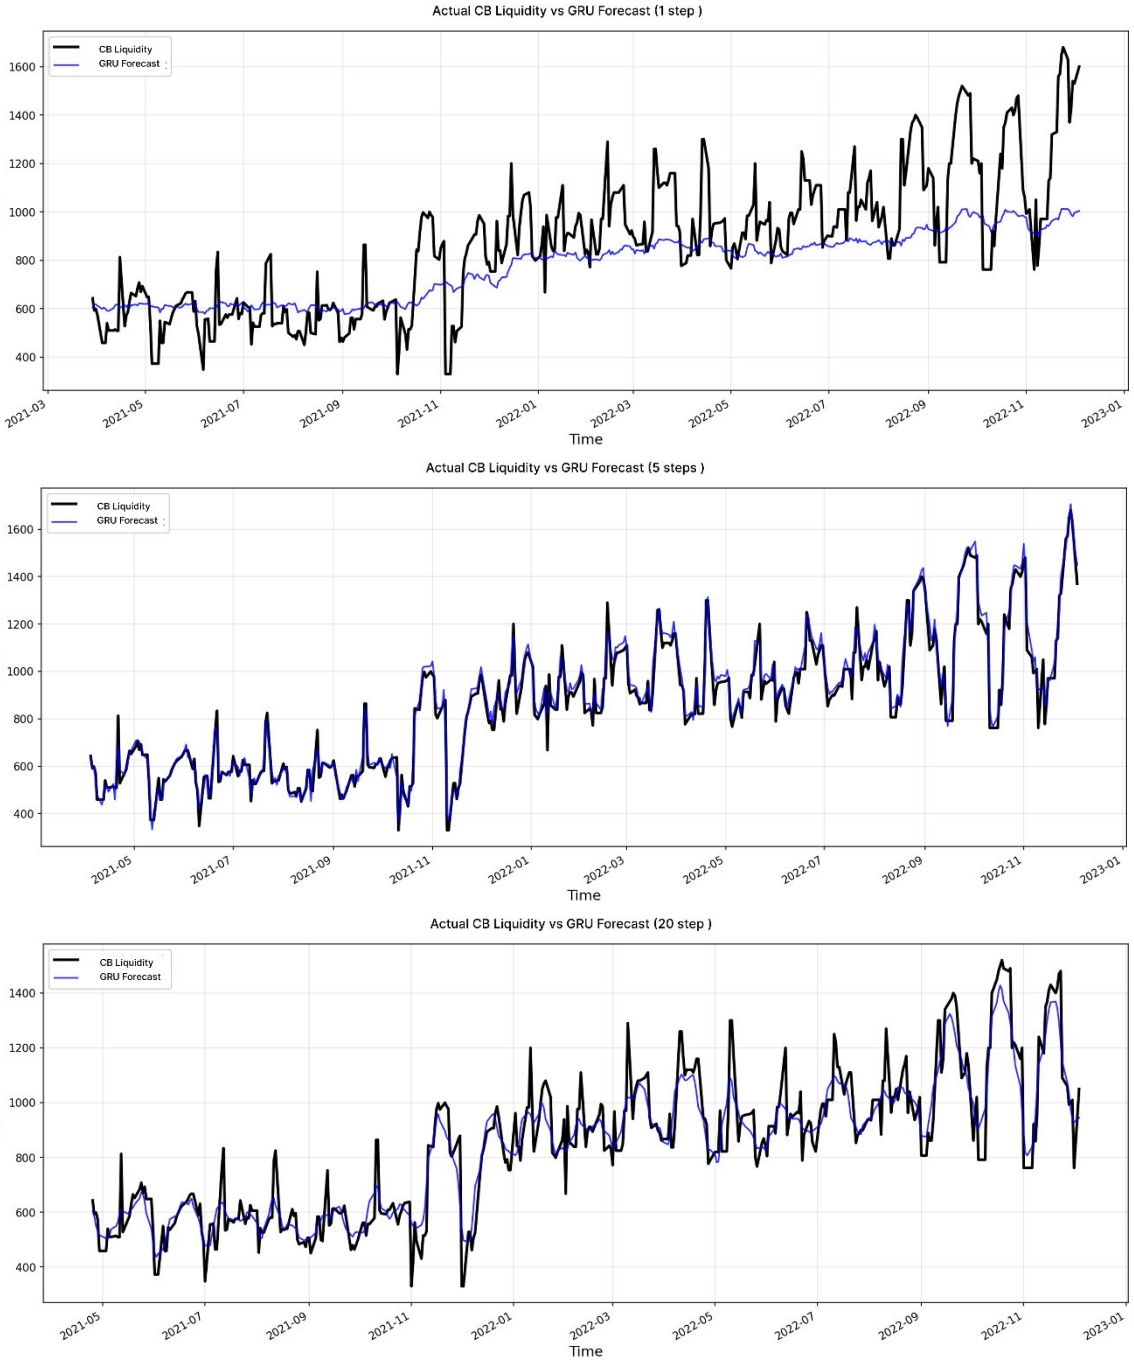
The GRU shares the two-layer, one-hundred-unit design of the LSTM but relies on the simpler GRU cell with fewer parameters. Performance is mixed: the daily horizon is the weakest (R^2 = 0.4529, RMSE = 206.68), the weekly horizon is predicted very accurately (R^2 = 0.9823, RMSE = 36.37), and the monthly horizon also performs well (R^2 = 0.9072). The univariate GRU thus behaves like a strong medium-horizon model, capturing the weekly and monthly structure of liquidity convincingly but struggling with day-to-day variation.

Table III-24: Forecasting performance of the GRU model

Forecast horizon	RMSE	R^2
1-step (day-ahead)	206.68	0.4529
5-step (week-ahead)	36.37	0.9823
20-step (month-ahead)	79.55	0.9072

Source: generated by the students using Python

Figure III-23 GRU forecasts versus actual liquidity — 1-step, 5-step and 20-step horizons



Source: generated by the students using Python

2.8. Gated Recurrent Unit with exogenous factors (GRUX)

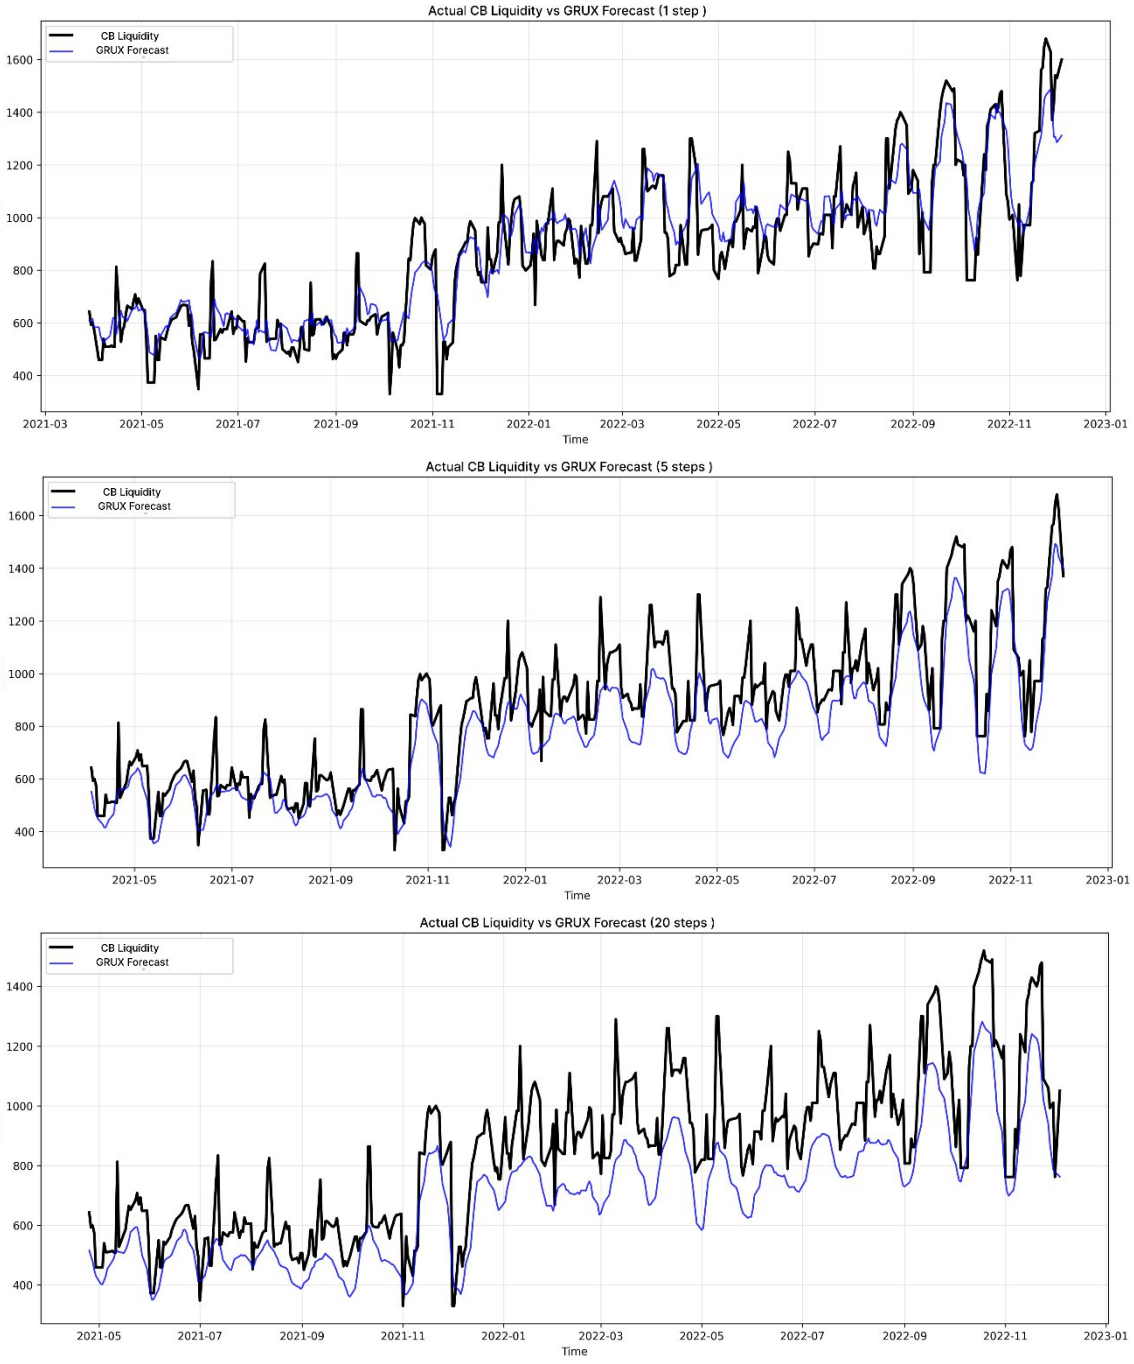
The GRUX combines the gated recurrent architecture with the six exogenous regressors. Unlike the CNNX and LSTMX, the exogenous inputs reshape the model's behaviour in a mixed way: the daily horizon improves substantially ($R^2 = 0.8236$, RMSE = 117.37), confirming that oil-price and calendar information is valuable for short-term prediction. However, the weekly horizon weakens ($R^2 = 0.8239$) and the monthly horizon deteriorates clearly ($R^2 = 0.6743$, RMSE = 149.06).

Table III-25: Forecasting performance of the GRUX model

Forecast horizon	RMSE	R^2
1-step (day-ahead)	117.37	0.8236
5-step (week-ahead)	114.71	0.8239
20-step (month-ahead)	149.06	0.6743

Source: generated by the students using Python

Figure III-24 GRUX forecasts versus actual liquidity — 1-step, 5-step and 20-step horizons



Source: generated by the students using Python

2.9. Comparative Assessment: GRU vs GRUX

Table III-26: Comparative forecasting performance — GRU vs GRUX

Forecast horizon	Model	RMSE	R^2
1-step (day-ahead)	GRU	206.68	0.4529
	GRUX	117.37	0.8236
5-step (week-ahead)	GRU	36.37	0.9823
	GRUX	114.71	0.8239
20-step (month-ahead)	GRU	79.55	0.9072
	GRUX	149.06	0.6743

Source: Elaborated by the students

Unlike the CNN and LSTM families, the GRU architecture presents a genuine trade-off. At the daily horizon, the GRUX achieves a dramatic improvement (R^2 rising from 0.4529 to 0.8236). However, at the weekly and monthly horizons, the univariate GRU is clearly superior. The best model from the GRU family therefore depends on the horizon: the GRUX is preferred for daily forecasting, while the univariate GRU dominates at the weekly and monthly horizons.

Table III-27: Best GRU-family model by forecasting horizon

Forecast horizon	Best model	RMSE	R^2
1-step (day-ahead)	GRUX	117.37	0.8236
5-step (week-ahead)	GRU	36.37	0.9823
20-step (month-ahead)	GRU	79.55	0.9072

Source: Elaborated by the students

3. Synthesis: Comparison of Deep Learning Models

Having identified the best model within each architecture family, this subsection brings together all six deep learning models to determine the best approach at each operational horizon. Table III-28 gathers the RMSE and R^2 at the three forecasting horizons.

Table III-28: Comparative forecasting performance of the deep learning models

Model	Metric	1-step (day-ahead)	5-step (week-ahead)	20-step (month-ahead)
CNN	RMSE	126.83	15.64	32.06
CNN	R^2	0.7940	0.9967	0.9849
CNNX	RMSE	227.61	147.52	59.88
CNNX	R^2	0.3365	0.7088	0.9474
LSTM	RMSE	170.33	180.45	76.94
LSTM	R^2	0.6284	0.5643	0.9132
LSTMX	RMSE	241.97	172.19	97.20
LSTMX	R^2	0.2501	0.6033	0.8615
GRU	RMSE	206.68	36.37	79.55

GRU	R^2	0.4529	0.9823	0.9072
GRUX	RMSE	117.37	114.71	149.06
GRUX	R^2	0.8236	0.8239	0.6743

Source: Elaborated by the students

At the daily horizon, the GRUX achieves the highest R^2 (0.8236) among the deep learning models, followed by the univariate CNN (R^2 = 0.7940). At the weekly horizon, the univariate CNN dominates decisively (R^2 = 0.9967, RMSE = 15.64), far ahead of the univariate GRU (R^2 = 0.9823). At the monthly horizon, the univariate CNN again leads (R^2 = 0.9849, RMSE = 32.06), followed by the CNNX, LSTM and GRU.

No single deep learning architecture dominates at every horizon. The univariate CNN is the strongest model overall, most accurate at the weekly and monthly horizons and reasonable at the daily horizon. The effect of adding exogenous inputs is consistently negative for the convolutional and LSTM architectures; the GRUX is the sole exception, improving only the daily forecast while losing accuracy at longer horizons. Table III-29 summarises the best deep learning model by horizon.

Table III-29: Best deep learning model by forecasting horizon

Forecast horizon	**Best model**	**RMSE**	**R^2**
1-step (day-ahead)	GRUX	117.37	0.8236
5-step (week-ahead)	CNN	15.64	0.9967
20-step (month-ahead)	CNN	32.06	0.9849

Source: Elaborated by the students

The univariate convolutional network is therefore the strongest deep learning approach for forecasting the net liquidity of the Bank of Algeria, dominating clearly at the weekly and monthly horizons. The GRUX is the most accurate deep learning model at the daily horizon, but none of the exogenous variants matches the univariate CNN over the medium and long horizons. These conclusions are carried forward into Section 4, where the deep learning models are compared directly with the classical SARIMA and SARIMAX benchmarks.

Section 4: Final Benchmarking Across Daily, Weekly, and Monthly Horizons

Sections 2 and 3 each delivered a horizon-by-horizon contest. In Section 2, the classical Box-Jenkins family was examined: the SARIMA (3,1,3) (1,1,2) [22] baseline was compared with its exogenous extension SARIMAX (2,1,3) (2,1,2) [22], and the winner was identified at each of the three operational horizons (daily, weekly, monthly). In Section 3, the same exercise was repeated within the deep learning family, with three architectures (CNN, LSTM, GRU) tested in both their univariate and exogenous specifications; here too, a single best model was retained at every horizon. The present section closes the comparative pipeline by bringing the classical winners and the deep learning winners head-to-head, so as to identify the single best forecasting model for the Bank of Algeria's net liquidity at each operational horizon.

The evaluation criteria remain unchanged. Two complementary metrics are used: the coefficient of determination (R^2), which measures the proportion of variance in the actual liquidity series that the model is able to explain, and the root mean squared error (RMSE), expressed in billions of dinars, which quantifies the average magnitude of the forecast errors in the original units of the series. A model is considered superior at a given horizon when it simultaneously achieves the higher R^2 and the lower RMSE; in the rare case where the two metrics disagree, the decision is discussed explicitly.

1. Models Retained from Sections 2 and 3

Before conducting the final comparison, it is useful to recall the models that have survived the intra-family contests. Table III-30 summarises the best classical model and the best deep learning model at each of the three horizons, as established in Sections 2 and 3 respectively.

Table III-30: Winners from the classical and deep learning models

Forecast horizon	Family	Best model (RMSE)	Best model (R^2)
1-step (day-ahead)	Classical	SARIMA — 109.56	SARIMA — 0.8441
1-step (day-ahead)	Deep learning	GRUX — 117.37	GRUX — 0.8236
5-step (week-ahead)	Classical	SARIMAX — 145.97	SARIMAX — 0.7233
5-step (week-ahead)	Deep learning	CNN — 15.64	CNN — 0.9967

Forecast horizon	Family	Best model (RMSE)	Best model (R^2)
20-step (month-ahead)	Classical	SARIMAX — 166.14	SARIMAX — 0.6415
20-step (month-ahead)	Deep learning	CNN — 32.06	CNN — 0.9849

Source: Elaborated by the students

Three observations can be drawn from this recap. First, within the classical family, the inclusion of the exogenous regressors helps only at longer horizons: the pure SARIMA remains preferable at the daily horizon, while the SARIMAX takes over at the weekly and monthly horizons, where the explanatory power of oil prices and the exchange rate become visible. Second, within the deep learning family, the picture is markedly different: a single architecture (the convolutional network in its univariate specification) dominates both the weekly and the monthly horizons, while the daily horizon is best handled by the gated recurrent network in its exogenous form (GRUX). Third, the two families do not converge on the same answer at any horizon, which makes the final benchmarking exercise genuinely informative.

1.1. Final Comparison at the Daily Horizon (h = 1)

The day-ahead horizon is the most demanding because it must capture the high-frequency noise and the abrupt swings that characterise daily liquidity movements. Table III-31 places the classical winner (SARIMA) and the deep learning winner (GRUX) side by side.

Table III-31: Final comparison at the daily horizon — SARIMA vs GRUX

Model	RMSE	R^2
SARIMA (3,1,3) (1,1,2) [22]	109.56	0.8441
GRUX	117.37	0.8236

Source: Elaborated by the students

Both metrics point to the same conclusion: the SARIMA model is the better forecaster at the daily horizon. It achieves a higher coefficient of determination (R^2 = 0.8441 against 0.8236) and a lower root mean squared error (RMSE = 109.56 against 117.37). The gap is modest but consistent across the two metrics, which removes any ambiguity. Substantively, the daily liquidity series is dominated by short-term linear dynamics, by the 22-day seasonal cycle and by transient autoregressive shocks, all of which the seasonal ARMA structure of the SARIMA captures very efficiently when combined with a rolling re-estimation strategy. The GRUX, although it benefits from the exogenous regressors at this horizon, is penalised by the stochastic nature of its training and by the very limited size of the test window for high-

frequency prediction. The classical benchmark therefore remains the most reliable model for day-ahead operational forecasting at the Bank of Algeria.

1.2. Final Comparison at the Weekly Horizon (h = 5)

At the five-day (week-ahead) horizon, the classical winner is the SARIMAX model, while the deep learning winner is the univariate CNN. Table III-32 reports the head-to-head results.

Table III-32: Final comparison at the weekly horizon — SARIMAX vs CNN

Model	RMSE	R^2
SARIMAX (2,1,3) (2,1,2) [22]	145.97	0.7233
CNN (univariate)	15.64	0.9967

Source: Elaborated by the students

The verdict at the weekly horizon is unambiguous and stands in sharp contrast to the daily result. The univariate convolutional network outperforms the SARIMAX by an extraordinary margin on both metrics: the R^2 rises from 0.7233 to 0.9967, while the RMSE falls from 145.97 to 15.64, a reduction of more than ninety per cent. The classical model explains roughly seventy-two per cent of the variance, whereas the CNN explains virtually all of it. This dramatic gap reflects two complementary phenomena. On the one hand, aggregation over five days smooths the high-frequency noise that hurt the deep learning models at the daily horizon, leaving a cleaner signal dominated by the monthly seasonal cycle, a pattern that the one-dimensional convolutional filters extract with remarkable precision. On the other hand, the SARIMAX, although it captures the seasonal structure correctly, suffers from the cumulative effect of multi-step forecast errors and from the linear treatment of the exogenous regressors, which cannot accommodate the non-linear interactions between oil prices, the exchange rate and liquidity. At the weekly horizon, the deep learning approach is therefore clearly superior.

1.3. Final Comparison at the Monthly Horizon (h = 20)

The twenty-day (month-ahead) horizon pits the same two finalists against each other: the SARIMAX on the classical side and the univariate CNN on the deep learning side. Table III-33 summarises the results.

Table III-33: Final comparison at the monthly horizon — SARIMAX vs CNN

Model	RMSE	R^2
SARIMAX (2,1,3) (2,1,2) [22]	166.14	0.6415

Model	RMSE	R^2
CNN (univariate)	32.06	0.9849

Source: Elaborated by the students

The pattern observed at the weekly horizon is confirmed and even reinforced at the monthly horizon. The CNN once again dominates the SARIMAX on both metrics: the R^2 jumps from 0.6415 to 0.9849, while the RMSE drops from 166.14 to 32.06, a reduction of more than eighty per cent. The classical model explains only about sixty-four per cent of the variance at this horizon, a level that is operationally unsatisfactory for medium-term liquidity planning. The CNN, by contrast, retains almost all of its explanatory power. The reason is the same as at the weekly horizon, only more pronounced: the further the forecast extends, the more the linear SARIMAX accumulates errors, while the convolutional network continues to exploit the dominant 22-day seasonal pattern that aggregation makes increasingly visible. The deep learning approach is therefore the clear winner at the monthly horizon as well.

2. Final Synthesis and Recommended Model by Horizon

The three head-to-head comparisons converge on a clear, horizon-dependent recommendation. Table III-34 consolidates the final verdict.

Table III-34: Recommended forecasting model for the Bank of Algeria — by operational horizon

Horizon	Recommended model	Family	RMSE	R^2
1-step (day-ahead)	SARIMA (3,1,3) (1,1,2) [22]	Classical	109.56	0.8441
5-step (week-ahead)	CNN (univariate)	Deep learning	15.64	0.9967
20-step (month-ahead)	CNN (univariate)	Deep learning	32.06	0.9849

Source: Elaborated by the students

The empirical evidence supports a hybrid forecasting strategy for the Bank of Algeria. At the day-ahead horizon, the seasonal ARMA structure of the SARIMA, combined with rolling re-estimation, remains the most accurate tool for capturing short-term liquidity fluctuations. From the weekly horizon onward, however, the univariate convolutional network takes a decisive lead, exploiting the smoother, more cyclical signal that emerges once aggregation attenuates daily noise. Two broader conclusions follow. First, no single model dominates across all horizons; the optimal forecasting tool depends on the operational use to which the prediction will be put. Second, the addition of exogenous regressors is not systematically beneficial: it

helps the SARIMAX at the medium and long horizons but is unnecessary, and even detrimental, in the best deep learning configuration at those same horizons. These findings provide a clear, evidence-based foundation for the policy recommendations developed in the conclusion of this dissertation.

General conclusion

This dissertation investigated a critical operational challenge facing the Bank of Algeria: the accurate multi-horizon forecasting of aggregate central bank liquidity in a volatile, hydrocarbon-dependent economy. The primary aim was to systematically compare the forecasting performance of traditional statistical models (SARIMA and SARIMAX) with state-of-the-art deep learning architectures (CNN, LSTM, and GRU) across day-ahead, week-ahead, and month-ahead horizons, while rigorously assessing the marginal contribution of exogenous variables.

The first chapter established the theoretical foundations of banking liquidity, progressing from universal definitions to the specific realities of the Algerian financial system. It highlighted how central bank liquidity functions as a dynamic flow shaped by net foreign assets, the net government position, and currency in circulation, all heavily exposed to external shocks such as oil price fluctuations and the COVID-19 pandemic. The chapter underscored the indispensable role of the central bank as the ultimate liquidity provider and the necessity of proactive liquidity risk management.

The second chapter provided the theoretical framework for forecasting, moving from traditional econometric approaches to modern neural networks. We examined the structural limitations of linear models, such as their inability to capture the non-linear dynamics and structural breaks inherent in Algerian liquidity data. By contrasting the theoretical underpinnings of SARIMA and SARIMAX with deep learning architectures, this chapter emphasized the need for advanced predictive models to calibrate open market operations and mitigate systemic risks.

The third chapter was dedicated to the empirical forecasting of aggregate central bank liquidity using daily data from 2015 to 2022. We evaluated classical models and deep learning approaches across three operational horizons. The comparative analysis revealed that no single model dominates across all horizons. At the day-ahead horizon, the classical SARIMA model outperformed deep learning approaches, demonstrating efficiency in capturing short-term linear dynamics and seasonal cycles. However, at the week-ahead and month-ahead horizons, the univariate CNN emerged as the standout model, achieving exceptional accuracy by effectively exploiting smoother, cyclical signals.

Based on these results, Hypothesis H1 which stated that deep learning models, particularly LSTM and GRU, would significantly outperform traditional statistical models relying solely on the series' own history is partially validated. While deep learning architectures

demonstrated superior ability to model the inherent dynamics of Algerian liquidity data at medium and long horizons, it was the univariate CNN, rather than the recurrent networks (LSTM/GRU), that delivered the most robust performance. Furthermore, traditional statistical models maintained their superiority for high-frequency, day-ahead forecasting.

Hypothesis H2, regarding the marginal contribution of exogenous variables (Brent crude oil price, USD/DZD exchange rate, and calendar dummies), is largely confirmed but with important nuances. The inclusion of these macroeconomic indicators substantially reduced forecasting errors in traditional models, particularly SARIMAX at the weekly and monthly horizons. However, their impact on deep learning models was mixed: while beneficial for the GRU architecture at the daily horizon (GRUX), adding exogenous inputs generally degraded the performance of the CNN, confirming that univariate deep learning models are often more robust for this specific dataset.

Consequently, Hypothesis H3 which stated that deep learning models significantly outperform traditional statistical models in forecasting accuracy across all horizons is partially supported. Deep learning models delivered markedly superior forecasting accuracy at the weekly and monthly horizons, with the univariate CNN achieving an R^2 of 0.9967 and 0.9849, respectively. However, the hypothesis is rejected for the daily horizon, where the classical SARIMA model (RMSE 109.56, R^2 0.8441) outperformed the best deep learning model, GRUX (RMSE 117.37, R^2 0.8236).

This study provides contributions across three main axes: theoretical, methodological, and managerial.

On a theoretical level, this work deepens our understanding of central bank liquidity in an oil-exporting emerging economy. It highlights the complex, non-linear transmission channels of external shocks such as hydrocarbon price volatility and exchange rate adjustments on domestic liquidity, providing a comprehensive framework that bridges macroeconomic theory with advanced predictive modelling.

Regarding the methodological aspect, this study provides one of the first comprehensive, multi-horizon evaluations of AI-based aggregate central bank liquidity forecasting. By utilizing theoretically appropriate seasonal benchmarks (SARIMA/SARIMAX) rather than mis-specified non-seasonal models and by systematically integrating exogenous variables, this dual-method strategy ensures a rigorous and robust comparison that advances the forecasting literature.

Finally, on the managerial level, this research provides practical, operational guidance for the Bank of Algeria's transition toward proactive liquidity management. By identifying the optimal forecasting tool for each specific operational horizon SARIMA for daily open market operations, and CNN for medium-term planning decision-makers can better calibrate interventions, set reserve requirements, and maintain financial stability in Algeria.

Despite these contributions, it is important to acknowledge the limitations of this research. The relatively limited historical depth of the daily dataset (2015–2022) constrains the training and generalization capabilities of the deep learning architectures. Additionally, the rapid structural shifts and policy changes within the Algerian banking sector may affect the stability of model parameters over time, requiring continual model updating and validation.

For future research, it is recommended to explore more advanced attention-based architectures, such as Transformer models, which may better handle the integration of exogenous variables without suffering from the signal dilution observed in our CNN and LSTM variants. Additionally, incorporating other high-frequency macroeconomic indicators, such as hydrocarbon export volumes or interbank market rates, could further refine the predictive precision.

Based on these findings, we conclude that upgrading forecasting methodologies is essential for effective liquidity management at the Bank of Algeria. The successful integration of these advanced tools balancing the strengths of traditional models and deep learning architectures will ultimately strengthen monetary policy implementation, reinforce financial stability, and support Algeria's economic resilience against global uncertainties.

References

Books

- Bagehot, W. (1873). *Lombard Street: A Description of the Money Market*. London: Henry S. King & Co.
- Box, G. E. P., & Jenkins, G. M. (1970). *Time Series Analysis: Forecasting and Control*.
- Darmon, R. (1998). Le risque de liquidité. In J. P. Patat (Ed.), *Risque et gestion de banque* (pp. 87-103). Paris: Dunod.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.
- Haykin, S. (2009). *Neural Networks and Learning Machines* (3rd ed.). Pearson.
- Lakehal, M. A. (2002). *La gestion de la liquidité bancaire*. Algiers: Editions OPU.
- Mishkin, F. S. (2013). *The Economics of Money, Banking, and Financial Markets* (10th ed.). Boston: Pearson.

Articles / Journal Papers / Working Papers

- Ashtar, D., Mohammadi Ziabari, S. S., & Alsahag, A. M. M. (2025). Hybrid Forecasting for Sustainable Electricity Demand in The Netherlands Using SARIMAX, SARIMAX-LSTM, and Sequence-to-Sequence Deep Learning Models. MDPI. https://www.mdpi.com/2071-1050/17/16/7192
- Bao, W., Yue, J., & Rao, Y. (2017). A deep learning framework for financial time series using stacked autoencoders and long short-term memory. PLoS ONE, 12(7), e0180944.
- Béziade, A. (1992). Gestion de la liquidité et risque de taux. Revue d'économie financière, (4), 223–240.
- Bhattacharyya, S., & Sahoo, B. (2011). Liquidity Management of Central Bank: An Application to India. International Research Journal of Finance and Economics, (78), 6–18.
- Bindseil, U. (2000). Liquidity, Information and the Overnight Rate. European Central Bank Working Paper Series, No. 32.
- Brunnermeier, M. K., & Pedersen, L. H. (2009). Market Liquidity and Funding Liquidity. Review of Financial Studies, 22(6), 2201–2238.
- Crockett, A. (2008). Market Liquidity and Financial Stability. Opening Remarks at the BIS Conference on Market Liquidity. Basel: Bank for International Settlements.
- Diamond, D. W., & Dybvig, P. H. (1983). Bank Runs, Deposit Insurance, and Liquidity. Journal of Political Economy, 91(3), 401–419.
- Diamond, D. W., & Rajan, R. G. (2001). Liquidity Shortages and Banking Crises. Journal of Finance, 56(2), 615–647.
- Drehmann, M., & Nikolaou, K. (2009). Funding Liquidity Risk: Definition and Measurement. European Central Bank Working Paper Series, No. 1024.
- El Hamiani-Khatat, N., & Ould Ahmed, M. (2018). Modelling the Demand for Currency in Circulation in Algeria. Bank of Algeria Working Paper, No. 03.
- Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. European Journal of Operational Research, 270(2), 654–669.
- Gray, S. (2008). Central Bank Management of Foreign Exchange Reserves and Liquidity. In S. Valckx (Ed.), Risk Management for Central Bankers (pp. 169–193). London: Central Banking Publications.

- Hewage, P., Trovati, M., Pereira, E., & Behera, A. (2020). Deep learning-based effective fine-grained weather and multi-domain time series forecasting. Pattern Analysis and Applications, 24(1), 343–366.
- Jin, B., & Xu, X. (2024). Machine learning Brent crude oil price forecasts. World Scientific Publishing. https://www.worldscientific.com/doi/10.1142/S2737599424500130
- Makhmudov, S. (2025). Forecasting liquidity in the Central Bank of Uzbekistan: A comparative analysis of SARIMA and Neural Networks. Working Papers / Journal of Central Banking.
- Nikolaou, K. (2009). Liquidity (Risk) Concepts, Definitions and Interactions. European Central Bank Working Paper Series, No. 1008.
- Siami-Namini, S., & Namin, A. S. (2018). Forecasting Economics and Financial Time Series: ARIMA vs. LSTM. International Conference on Machine Learning and Applications (ICMLA).
- Swetha, & Porkodi. (2024). A Study on Predicting House Prices With Deep Learning: An Evaluation of Advanced Neural Network Models. IJSART. http://ijsart.com/public/storage/paper/pdf/IJSART1103102702.pdf
- Taoussi, B., Boudia, S. M., & Mazouni, F. S. (2024). Wind speed forecasting using univariate and multivariate time series models. Stochastic Environmental Research and Risk Assessment, 39(2), 547–579. https://doi.org/10.1007/s00477-024-02881-2
- Wang, L., & Chen, X. (2025). Sustainable Factor Augmented Machine Learning Models for Crude Oil Return Forecasting. MDPI. https://www.mdpi.com/1911-8074/18/7/351
- Yang, X. (2024). Prediction of Tencent Share Price Based on ARIMA, SVR and LSTM. SciTePress. https://www.scitepress.org/Papers/2024/132138/132138.pdf
- Zhang, G. P. (2003). Time series forecasting using a hybrid ARIMA and neural network model. Neurocomputing, 50, 159–175.

Institutional and Regulatory Documents

- *Annual Report 2016* [Rapport Annuel 2016]. Algiers: Bank of Algeria.
- *Annual Report 2017* [Rapport Annuel 2017]. Algiers: Bank of Algeria.
- *Annual Report 2018* [Rapport Annuel 2018]. Algiers: Bank of Algeria.
- *Annual Report 2019* [Rapport Annuel 2019]. Algiers: Bank of Algeria.
- *Annual Report 2020* [Rapport Annuel 2020]. Algiers: Bank of Algeria.
- *Annual Report 2021* [Rapport Annuel 2021]. Algiers: Bank of Algeria.
- Basel Committee on Banking Supervision. (2008). *Principles for Sound Liquidity Risk Management and Supervision*. Basel: Bank for International Settlements.
- Basel Committee on Banking Supervision. (2013). *Basel III: The Liquidity Coverage Ratio and Liquidity Risk Monitoring Tools*. Basel: Bank for International Settlements.
- European Central Bank. (2004). *The Liquidity Management of the ECB*. Frankfurt: European Central Bank.
- International Monetary Fund (IMF). (2024). *Central Bank Liquidity Management: Operational Frameworks and Tools*. IMF Capacity Development Note.
- *Monetary and Banking Law (Law No. 23-12 of June 1, 2023)*. Official Journal of the Algerian Republic.
- *Regulation No. 11-08 of 26 Dhou El Kaada 1432 corresponding to October 25, 2011, relating to the prudential rules applicable to banks*. Official Journal of the Algerian Republic.

Master's Theses / Doctoral Dissertations

- Abdennouri (2019). *Evaluation du diapositive de gestion de risque de liquidité d'une banque - Approche supervision* - (Mémoire de fin d'études). Ecole Supérieure de Banque.
- Benkredda (2025). Forecasting central bank Liquidity in Algeria: A comparative Analysis of statistical and Deep learning Models (Masters dissertation). Higher School of Commerce.
- Siliadin, B. (2013). *Gestion du risque de liquidité systémique et stabilité financière* (Doctoral dissertation). Université Paris Dauphine.

Online Resources / Web References/images

- **https://learnopencv.com/understanding-feedforward-neural-networks/**
- **https://medium.com/analytics-vidhya/lstms-explained-a-complete-technically-accurate-conceptual-guide-with-keras-2a650327e8f2**
- **https://www.geeksforgeeks.org/machine-learning/gated-recurrent-unit-networks/**
- **https://www.researchgate.net/publication/345372732_Forecasting_air_passenger_traffic_flow_based_on_the_two-phase_learning_model**

Appendices

Appendix 1: Forecasting central bank Liquidity in Algeria: A comparative Analysis of statistical and Deep learning Models

	Model	p	d	q	P	D	Q	s	AIC	BIC	HQIC	LogLik
0	SARIMA(3,1,3)(1,1,2)[22]	3	1	3	1	1	2	22	20837.55	20891.29	20857.50	-10408.77
1	SARIMA(1,1,3)(2,1,2)[22]	1	1	3	2	1	2	22	20840.68	20889.05	20858.64	-10411.34
2	SARIMA(3,1,2)(2,1,2)[22]	3	1	2	2	1	2	22	20841.33	20895.08	20861.29	-10410.67
3	SARIMA(3,1,3)(2,1,2)[22]	3	1	3	2	1	2	22	20843.21	20902.32	20865.16	-10410.60
4	SARIMA(3,1,2)(2,1,1)[22]	3	1	2	2	1	1	22	20843.48	20891.85	20861.44	-10412.74
5	SARIMA(3,1,3)(2,1,1)[22]	3	1	3	2	1	1	22	20845.46	20899.20	20865.42	-10412.73
6	SARIMA(2,1,3)(2,1,2)[22]	2	1	3	2	1	2	22	20846.44	20900.18	20866.40	-10413.22
7	SARIMA(3,1,2)(1,1,2)[22]	3	1	2	1	1	2	22	20847.56	20895.94	20865.53	-10414.78
8	SARIMA(1,1,3)(1,1,2)[22]	1	1	3	1	1	2	22	20848.13	20891.13	20864.10	-10416.07
9	SARIMA(2,1,3)(0,1,2)[22]	2	1	3	0	1	2	22	20849.35	20892.35	20865.32	-10416.68
10	SARIMA(2,1,3)(1,1,2)[22]	2	1	3	1	1	2	22	20853.64	20902.01	20871.61	-10417.82
11	SARIMA(3,1,1)(2,1,2)[22]	3	1	1	2	1	2	22	20853.93	20902.30	20871.89	-10417.96
12	SARIMA(2,1,2)(2,1,2)[22]	2	1	2	2	1	2	22	20854.92	20903.29	20872.88	-10418.46
13	SARIMA(2,1,3)(2,1,1)[22]	2	1	3	2	1	1	22	20855.81	20904.18	20873.77	-10418.90
14	SARIMA(3,1,1)(2,1,1)[22]	3	1	1	2	1	1	22	20856.05	20899.05	20872.02	-10420.03

Appendix 2: Top 15 SARIMA models ranked by AIC and penelized by BIC and the best out of all of them + Best model AIC & BIC

Top 15 SARIMA models by AIC:

	Model	AIC	BIC	HQIC	LogLik
0	SARIMA(3,1,3)(1,1,2)[22]	20837.55	20891.29	20857.50	-10408.77
1	SARIMA(1,1,3)(2,1,2)[22]	20840.68	20889.05	20858.64	-10411.34
2	SARIMA(3,1,2)(2,1,2)[22]	20841.33	20895.08	20861.29	-10410.67
3	SARIMA(3,1,3)(2,1,2)[22]	20843.21	20902.32	20865.16	-10410.60
4	SARIMA(3,1,2)(2,1,1)[22]	20843.48	20891.85	20861.44	-10412.74
5	SARIMA(3,1,3)(2,1,1)[22]	20845.46	20899.20	20865.42	-10412.73
6	SARIMA(2,1,3)(2,1,2)[22]	20846.44	20900.18	20866.40	-10413.22
7	SARIMA(3,1,2)(1,1,2)[22]	20847.56	20895.94	20865.53	-10414.78
8	SARIMA(1,1,3)(1,1,2)[22]	20848.13	20891.13	20864.10	-10416.07
9	SARIMA(2,1,3)(0,1,2)[22]	20849.35	20892.35	20865.32	-10416.68
10	SARIMA(2,1,3)(1,1,2)[22]	20853.64	20902.01	20871.61	-10417.82
11	SARIMA(3,1,1)(2,1,2)[22]	20853.93	20902.30	20871.89	-10417.96
12	SARIMA(2,1,2)(2,1,2)[22]	20854.92	20903.29	20872.88	-10418.46
13	SARIMA(2,1,3)(2,1,1)[22]	20855.81	20904.18	20873.77	-10418.90
14	SARIMA(3,1,1)(2,1,1)[22]	20856.05	20899.05	20872.02	-10420.03

✔ Best model : SARIMA(3,1,3)(1,1,2)[22]
AIC = 20837.55 | BIC = 20891.29

Appendix 3: SARIMA coefficients matrix + Ljung-Box, Jarque-Bera, Heteroskedasticity test results and coefficient significance

Model fitted in 163.8s

SARIMAX Results

Dep. Variable:	Liquidity	No. Observations:	1665
Model:	SARIMAX(3, 1, 3)x(1, 1, [1, 2], 22)	Log Likelihood	-10408.773
Date:	Sat, 09 May 2026	AIC	20837.546
Time:	00:19:10	BIC	20891.286
Sample:	01-04-2015 - 05-20-2021	HQIC	20857.505
Covariance Type:	opg		

	coef	std err	z	P>|z|	[0.025	0.975]
ar.L1	1.5076	0.223	6.754	0.000	1.070	1.945
ar.L2	-0.4960	0.361	-1.373	0.170	-1.204	0.212
ar.L3	-0.0864	0.143	-0.606	0.545	-0.366	0.193
ma.L1	-1.8793	0.223	-8.432	0.000	-2.316	-1.442
ma.L2	0.9261	0.432	2.143	0.032	0.079	1.773
ma.L3	-0.0326	0.212	-0.153	0.878	-0.449	0.384
ar.S.L22	0.3833	0.053	7.243	0.000	0.280	0.487
ma.S.L22	-1.0598	0.060	-17.676	0.000	-1.177	-0.942
ma.S.L44	0.0835	0.052	1.594	0.111	-0.019	0.186
sigma2	2.668e+04	486.003	54.906	0.000	2.57e+04	2.76e+04

Ljung-Box (L1) (Q):	0.00	Jarque-Bera (JB):	9235.30
Prob(Q):	0.96	Prob(JB):	0.00
Heteroskedasticity (H):	0.45	Skew:	1.50
Prob(H) (two-sided):	0.00	Kurtosis:	14.40

Warnings:
[1] Covariance matrix calculated using the outer product of gradients (complex-step).

**The image has a typo of "SARIMAX" instead of SARIMA*

Appendix 4: Durbin watson test on SARIMA residuals

Ljung-Box Test (lags 1–20):

	lb_stat	lb_pvalue
1	17.9153	0.0
2	22.8080	0.0
3	26.9194	0.0
4	27.2050	0.0
5	31.5514	0.0
6	32.7993	0.0
7	43.3356	0.0
8	54.0720	0.0
9	56.9349	0.0
10	56.9977	0.0
11	57.0076	0.0
12	58.4916	0.0
13	58.5647	0.0
14	64.7598	0.0
15	68.6349	0.0
16	74.6919	0.0
17	74.7072	0.0
18	75.7237	0.0
19	82.5477	0.0
20	107.4009	0.0

Durbin-Watson statistic: 2.1436
→ DW ≈ 2 indicates no autocorrelation

Appendix 5: Stationarity tests on 1st Diff

Testing FIRST DIFFERENCE...

Stationarity Tests — Liquidity (First Difference)

Test	Statistic	p-value	CV 5%	H_0	Decision
ADF	-12.4957	0.0	-2.8629	Unit root	Reject H_0 ✓
Phillips-Perron	-75.0900	0.0	—	Unit root	Reject H_0 ✓
KPSS	0.1507	0.1	0.463	Stationary	Fail to reject H_0 ✓

Appendix 6: SARIMA Grid Search function

```
[11]: def sarima_grid_search(series, d=1, s=22,
                             p_range=range(4), q_range=range(4),
                             P_range=range(3), D_range=range(2), Q_range=range(3),
                             top_n=15):
          """
          Grid search over SARIMA(p,d,q)(P,D,Q)[s] parameter space.
          Returns a DataFrame of top_n results sorted by AIC.
          """
          grid    = list(itertools.product(p_range, q_range, P_range, D_range, Q_range))
          total   = len(grid)
          results = []

          print(f'Grid search over {total} combinations — please wait...')
          t0 = time.time()

          for i, (p, q, P, D, Q) in enumerate(grid, 1):
              if i % 50 == 0 or i == total:
                  elapsed = time.time() - t0
                  print(f'  {i}/{total}  ({elapsed:.0f}s elapsed)')
              try:
                  mod = SARIMAX(series,
                                order=(p, d, q),
                                seasonal_order=(P, D, Q, s),
                                enforce_stationarity=False,
                                enforce_invertibility=False)
                  res = mod.fit(disp=False, maxiter=300, method='lbfgs')
                  results.append({
                      'Model'  : f'SARIMA({p},{d},{q})({P},{D},{Q})[{s}]',
                      'p':p, 'd':d, 'q':q, 'P':P, 'D':D, 'Q':Q, 's':s,
                      'AIC'    : round(res.aic, 2),
                      'BIC'    : round(res.bic, 2),
                      'HQIC'   : round(res.hqic, 2),
                      'LogLik' : round(res.llf, 2),
                  })
              except Exception:
                  pass   # skip non-converging specifications

          df = (pd.DataFrame(results)
                  .sort_values('AIC')
                  .reset_index(drop=True))
          print(f'\nDone. {len(df)} models converged in {time.time()-t0:.0f}s.')
          return df.head(top_n)

      print('Function defined. Run the next cell to start the search.')
```

Appendix 7: Stationarity tests on SARIMA residuals

7.2 Stationarity of Residuals

```
[17]: resid = results.resid.dropna()
      _ = stationarity_tests(resid, label='SARIMA Residuals')
```

```
──────────────────────────────────────────────────────────────────────
  Stationarity Tests — SARIMA Residuals
──────────────────────────────────────────────────────────────────────
           Test  Statistic  p-value   CV 5%         H₀            Decision
            ADF     -7.3376      0.0 -2.8633  Unit root         Reject H₀ ✓
Phillips-Perron   -46.8050      0.0       —  Unit root         Reject H₀ ✓
           KPSS     0.0988      0.1   0.463 Stationary Fail to reject H₀ ✓
```

Appendix 8: Rolling Forecast function

8.1 Rolling Forecast Function

[32]:
```python
def rolling_forecast_sarima(series, train_len, horizon,
                            window, order, seasonal_order,
                            label=''):
    """
    Rolling-window SARIMA forecast.

    Parameters
    ----------
    series        : pd.Series  — full series (train + test)
    train_len     : int        — initial training window length
    horizon       : int        — total steps to forecast (= len(test))
    window        : int        — steps between re-fits
    order         : tuple      — (p, d, q)
    seasonal_order: tuple      — (P, D, Q, s)
    label         : str        — description for progress print

    Returns
    -------
    np.ndarray of length `horizon`
    """
    total     = train_len + horizon
    forecasts = []
    steps_done = 0

    t0 = time.time()
    for i in range(train_len, total, window):
        steps = min(window, total - i)
        mod = SARIMAX(series.iloc[:i],
                      order=order,
                      seasonal_order=seasonal_order,
                      enforce_stationarity=False,
                      enforce_invertibility=False)
        res = mod.fit(disp=False, maxiter=300, method='lbfgs')
        pred = res.get_forecast(steps=steps)
        forecasts.extend(pred.predicted_mean.values)
        steps_done += steps

    elapsed = time.time() - t0
    n_refits = int(np.ceil(horizon / window))
    print(f'  Rolling {label}: {n_refits} re-fits | {elapsed:.0f}s')
    return np.array(forecasts[:horizon])
```

Appendix 9: Rolling Forecast performance summary

[51]:
```python
def compute_metrics(y_true, y_pred, label=''):
    """
    Compute NRMSE, NMAE, MAPE, and R² between actual and predicted values.
    Normalization uses the range (max - min) of y_true.
    """
    y_true, y_pred = np.asarray(y_true), np.asarray(y_pred)
    rmse   = np.sqrt(mean_squared_error(y_true, y_pred))
    mae    = mean_absolute_error(y_true, y_pred)
    rng    = y_true.max() - y_true.min()
    nrmse  = rmse / rng
    nmae   = mae  / rng
    mask   = y_true != 0
    mape   = np.mean(np.abs((y_true[mask] - y_pred[mask]) / y_true[mask])) * 100
    r2     = r2_score(y_true, y_pred)
    return {
        'Horizon' : label,
        'RMSE'    : round(rmse, 2),
        'MAE'     : round(mae,  2),
        'NRMSE'   : round(nrmse, 4),
        'NMAE'    : round(nmae,  4),
        'MAPE (%)'  : round(mape,  4),
        'R²'      : round(r2,   4),
    }

actual = pred_df['Actual'].values

metrics_list = [
    compute_metrics(actual, pred_df['Static'].values,  'Static (All)'),
    compute_metrics(actual, pred_df['Roll_1'].values,  'Rolling 1-step (Daily)'),
    compute_metrics(actual, pred_df['Roll_5'].values,  'Rolling 5-step (Weekly)'),
    compute_metrics(actual, pred_df['Roll_20'].values, 'Rolling 20-step (Monthly)'),
]

metrics_df = pd.DataFrame(metrics_list)
print(f'\nPerformance Summary — {best["Model"]}')
print('='*75)
print(metrics_df.to_string(index=False))
```

```
Performance Summary — SARIMA(2,1,3)(0,1,2)[5]
===========================================================================
                 Horizon   RMSE    MAE  NRMSE   NMAE  MAPE (%)      R²
             Static (All) 562.53 456.44 0.4167 0.3381   45.1398 -3.1101
   Rolling 1-step (Daily) 109.56  77.09 0.0812 0.0571    9.3462  0.8441
  Rolling 5-step (Weekly) 151.38 113.98 0.1121 0.0844   13.4530  0.7024
Rolling 20-step (Monthly) 174.01 133.35 0.1289 0.0988   15.0333  0.6067
```

Appendix 10: SARIMAX coefficients matrix+Ljung-Box, Jarque-Bera, Heteroskedasticity test results and coefficient significance

```
Fitted in 512.4s
                                     SARIMAX Results
==========================================================================================
Dep. Variable:                           Liquidity   No. Observations:                 1664
Model:             SARIMAX(2, 1, 3)x(2, 1, [1, 2], 22)   Log Likelihood           -10365.146
Date:                                 Sun, 10 May 2026   AIC                       20762.293
Time:                                        21:00:50   BIC                       20848.267
Sample:                                  01-05-2015   HQIC                      20794.225
                                       - 05-20-2021
Covariance Type:                               opg
==============================================================================
                 coef    std err          z      P>|z|      [0.025      0.975]
------------------------------------------------------------------------------
Brent          -1.7125      0.447     -3.828      0.000      -2.589      -0.836
USDDZD        -15.9229      1.239    -12.856      0.000     -18.350     -13.495
Pre_Eid       -20.4294     26.345     -0.775      0.438     -72.064      31.205
Pre_Ramadan   -13.1907     52.231     -0.253      0.801    -115.562      89.181
Pre_BTS        66.3030     39.820      1.665      0.096     -11.743     144.349
Summer         27.1424     21.992      1.234      0.217     -15.961      70.246
ar.L1           1.6504      0.035     46.595      0.000       1.581       1.720
ar.L2          -0.7176      0.030    -23.734      0.000      -0.777      -0.658
ma.L1          -1.9975      0.038    -52.764      0.000      -2.072      -1.923
ma.L2           1.1464      0.058     19.844      0.000       1.033       1.260
ma.L3          -0.1356      0.031     -4.435      0.000      -0.196      -0.076
ar.S.L22       -0.0377      0.281     -0.134      0.893      -0.588       0.513
ar.S.L44        0.1339      0.075      1.795      0.073      -0.012       0.280
ma.S.L22       -0.6628      0.283     -2.344      0.019      -1.217      -0.109
ma.S.L44       -0.3000      0.269     -1.114      0.265      -0.828       0.228
sigma2       2.563e+04    450.546     56.880      0.000    2.47e+04    2.65e+04
==========================================================================================
Ljung-Box (L1) (Q):                   0.01   Jarque-Bera (JB):             10291.68
Prob(Q):                              0.94   Prob(JB):                         0.00
Heteroskedasticity (H):               0.48   Skew:                             1.41
Prob(H) (two-sided):                  0.00   Kurtosis:                        15.13
==========================================================================================

Warnings:
[1] Covariance matrix calculated using the outer product of gradients (complex-step).
```

Appendix 11: Best SARIMAX model AIC & BIC

```
✔  Best model : SARIMAX(2,1,3)(2,1,2)[22]
   AIC = 20762.29  |  BIC = 20848.27
```

Appendix 12: CNN model architechture

Model

```
1]: # ------------------------------
    # 1. Reproducibility & Fixed Hyperparameters
    # ------------------------------
    tf.random.set_seed(42)
    np.random.seed(42)

    steps = 1

    PARAMS = {
        'batch_size': 64,
        'units': 100,
        'dropout': 0.2,
        'kernel_size': 7,
        'activation': 'tanh',
        'epochs': 100,
        'patience': 5
    }

    # ------------------------------
    # 2. Build CNN Function
    # ------------------------------
    def build_cnn(input_shape):
        model = Sequential()
        model.add(Conv1D(filters=64, kernel_size=PARAMS['kernel_size'], activation=PARAMS['activation'],
                         padding='same', input_shape=input_shape))
        model.add(Conv1D(filters=32, kernel_size=PARAMS['kernel_size'], activation=PARAMS['activation'],
                         padding='same'))
        model.add(Dropout(PARAMS['dropout']))
        model.add(Flatten())
        model.add(Dense(PARAMS['units'], activation=PARAMS['activation']))
        model.add(Dense(steps, activation='linear'))
        model.compile(optimizer='adam', loss='mse', metrics=['mae'])
        return model

    def inv_transform_multi(preds_scaled, scaler):
        flat = preds_scaled.reshape(-1, 1)
        return scaler.inverse_transform(flat).reshape(preds_scaled.shape)

    # ------------------------------
    # 3. Train Model
    # ------------------------------
    input_shape = (x_train_np.shape[1], x_train_np.shape[2])
    model = build_cnn(input_shape)
```

Appendix 13: CNN Model parameters

Model: "sequential"

Layer (type)	Output Shape	Param #
conv1d (Conv1D)	(None, 1, 64)	9,920
conv1d_1 (Conv1D)	(None, 1, 32)	14,368
dropout (Dropout)	(None, 1, 32)	0
flatten (Flatten)	(None, 32)	0
dense (Dense)	(None, 100)	3,300
dense_1 (Dense)	(None, 1)	101

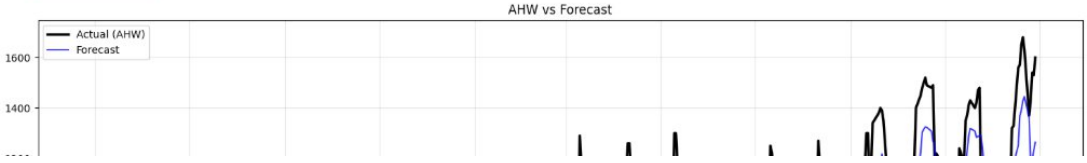
Total params: 27,689 (108.16 KB)

Trainable params: 27,689 (108.16 KB)

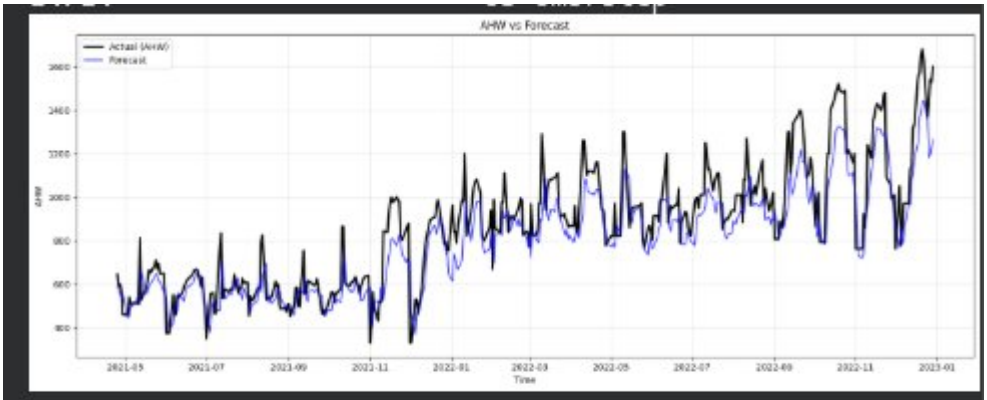
Non-trainable params: 0 (0.00 B)

```
--- Training CNN ---
Epoch 53: early stopping
Restoring model weights from the end of the best epoch: 48.

--- Running Predictions ---
14/14 ━━━━━━━━━━━━━━━━━━━━ 0s 8ms/step
```

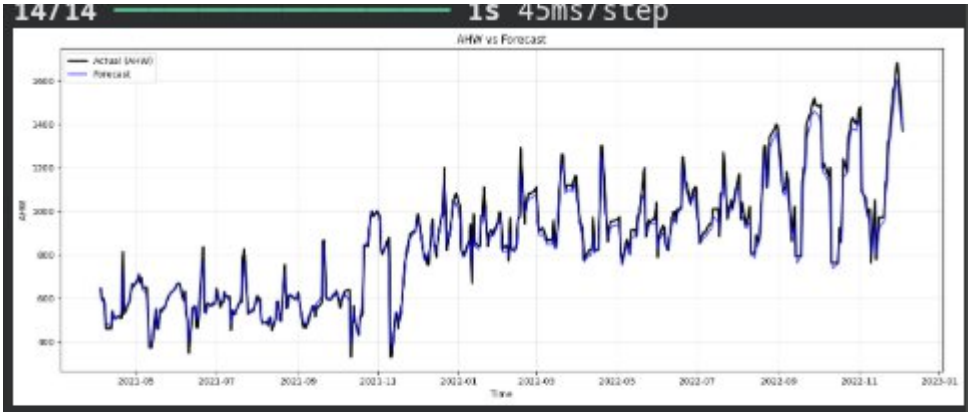


Appendix 14: CNN RMSE & R-squared 1step



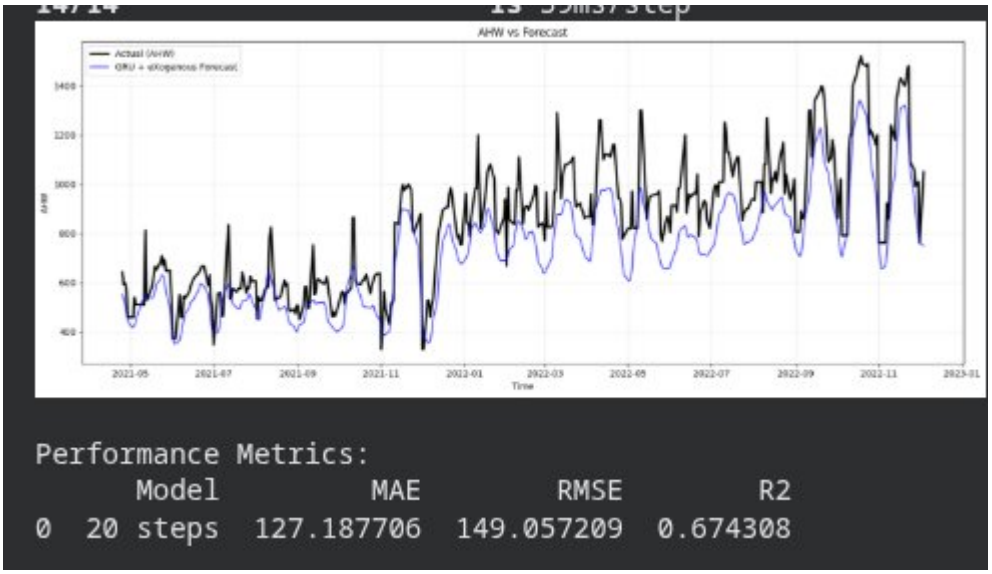

```
Performance Metrics:
       Model        MAE        RMSE        R2
0  Model_Run  97.874525  126.829664  0.793981
```

Appendix 15: GRU RMSE & R-squared 5step



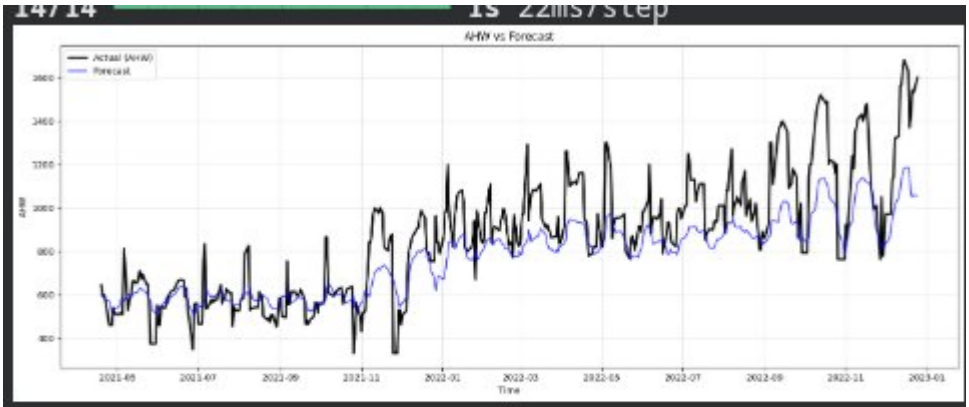

```
Performance Metrics:
     Model       MAE        RMSE        R2
0  5 steps  26.71833  36.373471  0.982297
```

Appendix 16: GRUX RMSE & R-squared 20step



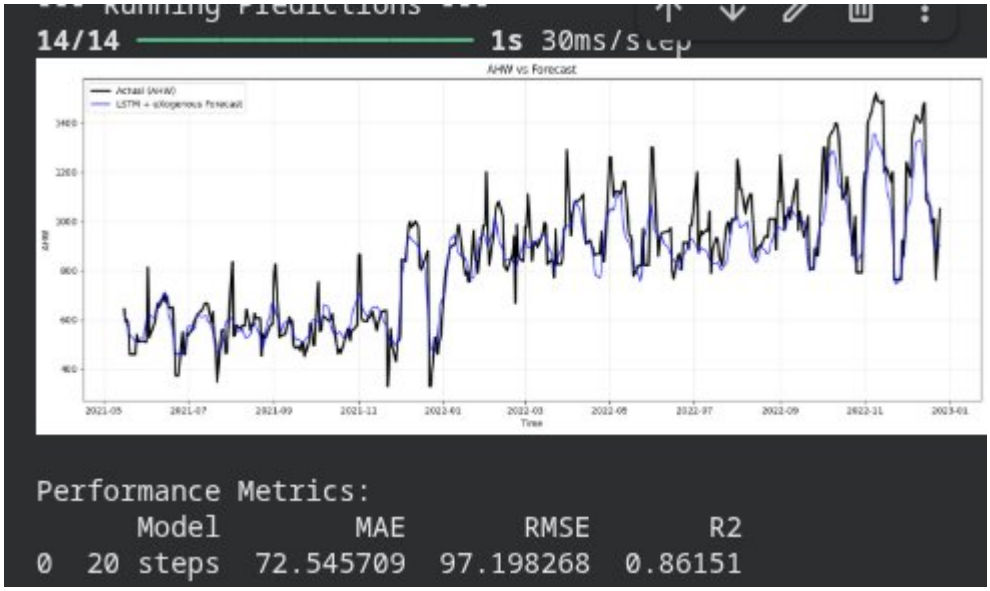

```
Performance Metrics:
      Model          MAE         RMSE        R2
0  20 steps  127.187706  149.057209  0.674308
```

Appendix 17: : LSTM RMSE & R-squared 1step




```
Performance Metrics:
    Model          MAE         RMSE        R2
0  1 step  127.025912  170.332157  0.628414
```

Appendix 18: LSTMX RMSE & R-squared 20steps



Appendix 19: Data Preparation Using MinMax Scaling for Deep Learning Models

```
# 1. Fit scalers on training data
scaler_x = MinMaxScaler()
scaler_y = MinMaxScaler()
# 1. Fit scalers
scaler_x.fit(train[columns_to_transform])          # ONLY exogenous features
scaler_y.fit(train[[target]])                # AHW only

train_scaled = train.copy()
val_scaled = val.copy()
test_scaled = test.copy()

# Exogenous
train_scaled.loc[:, columns_to_transform] = scaler_x.transform(train_scaled[columns_to_transform])
val_scaled.loc[:, columns_to_transform] = scaler_x.transform(val_scaled[columns_to_transform])
test_scaled.loc[:, columns_to_transform] = scaler_x.transform(test_scaled[columns_to_transform])

# Target (and also input wind_speed if used in model)
train_scaled.loc[:, target] = scaler_y.transform(train_scaled[[target]])
val_scaled.loc[:, target] = scaler_y.transform(val_scaled[[target]])
test_scaled.loc[:, target] = scaler_y.transform(test_scaled[[target]])

train_scaled.shape, val_scaled.shape,  test_scaled.shape
```

Appendix 20: GRU Model architechture

```python
# ------------------------------
# 1. Reproducibility & Fixed Hyperparameters
# ------------------------------
tf.random.set_seed(42)
np.random.seed(42)

steps = 5

PARAMS = {
    'batch_size': 16,
    'units': 100,
    'dropout': 0.45,
    'activation': 'tanh',
    'epochs': 100,
    'patience': 10
}

# ------------------------------
# 2. Build GRU Function (Updated Signature)
# ------------------------------
def build_model(input_shape, params):
    model = Sequential([
        # First GRU layer needs return_sequences=True to feed the second GRU layer
        GRU(params['units'], return_sequences=True, input_shape=input_shape,
            activation=params['activation']),
        Dropout(params['dropout']),

        # Second GRU layer
        GRU(params['units'], activation=params['activation']),
        Dropout(params['dropout']),

        Dense(steps, activation='linear')
    ])
    model.compile(optimizer=tf.keras.optimizers.Adam(), loss='mse', metrics=['mae'])
    return model
```

Appendix 21: Net liquidity Graph with Calendar dummies

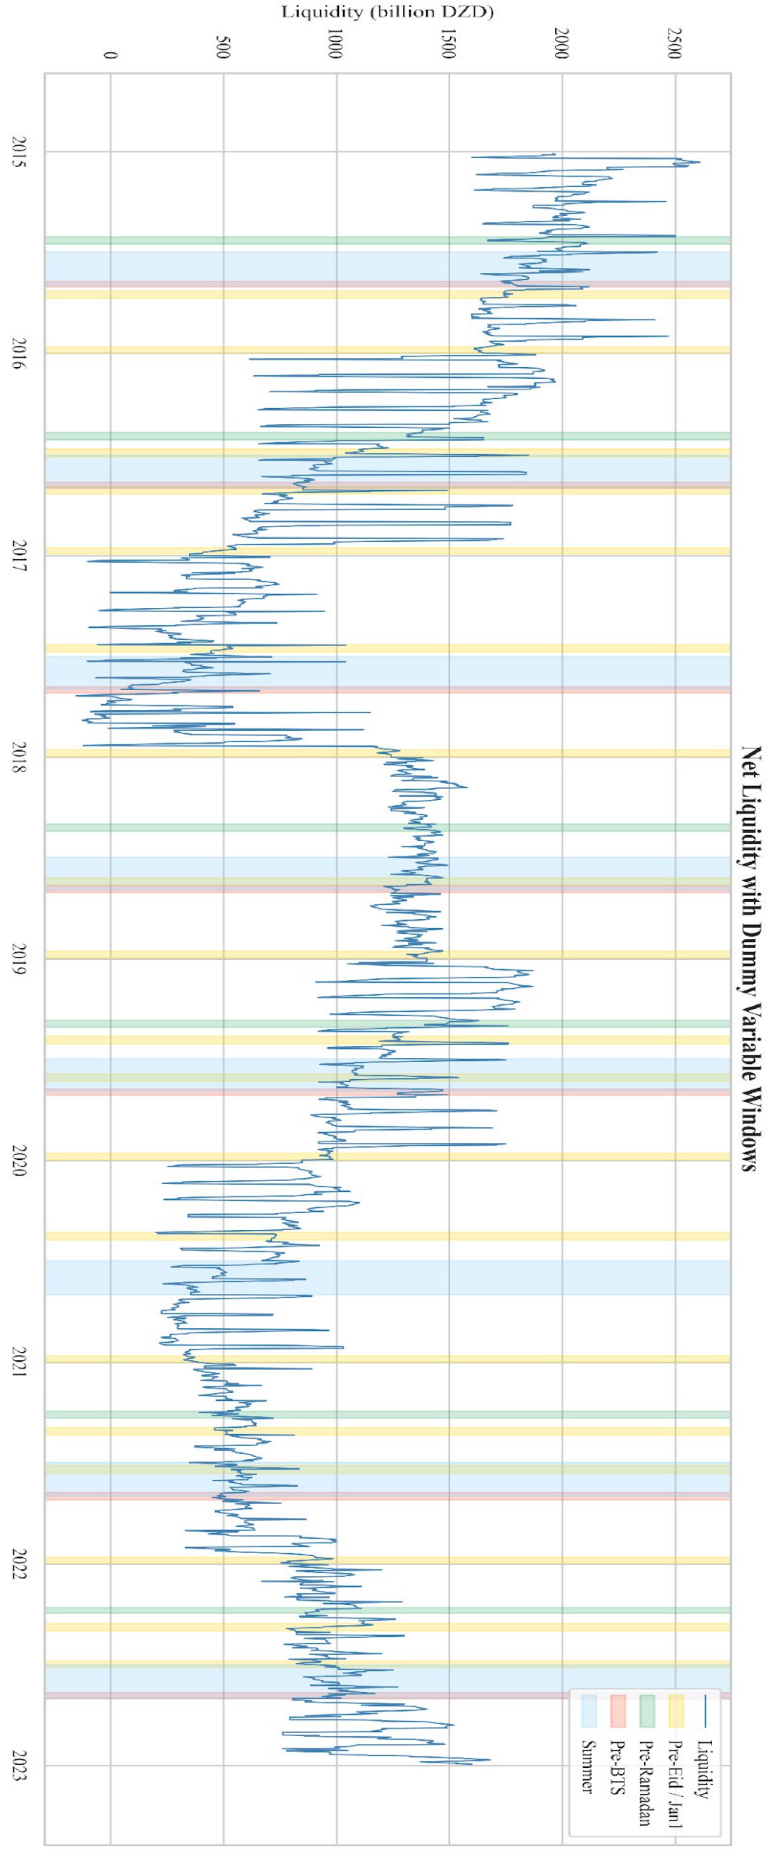


Table of content

ACKNOWLEDGMENT **1**
DEDICATION **1**
CONTENTS ..... **I**
LIST OF TABLES ..... **II**
LIST OF FIGURES ..... **III**
LIST OF APPENDICES ..... **IV**
LIST OF ABBREVIATIONS ..... **V**
ABSTRACT ..... **VIII**
GENERAL INTRODUCTION **1**
CHAPTER I: THEORETICAL FOUNDATIONS OF BANKING LIQUIDITY AND THE ALGERIAN CONTEXT **6**
Section 1: The Concept of Banking Liquidity 8
1. Definitions and Nature of Liquidity 8
2. The Typology of Liquidity 9
3. The Interconnection: Virtuous and Vicious Cycles 11
4. The Importance of Liquidity for Financial Stability 11
Section 2: Liquidity Risk and its Management 13
1. The Role of the Central Bank as the Ultimate Liquidity Provider 13
2. Understanding Liquidity Risk 14
3. Liquidity Management Framework 16
3.2. Unconventional Monetary Policy Instruments 17
3.3. Prudential and Regulatory Management Tools 17
3.4. The Role of Forecasting in Effective Liquidity Management 18
Section 3: Banking Liquidity in Algeria and the Challenges of Forecasting 21
1. The Algerian Banking Liquidity Context 21
2. The Impact of Oil Shocks and COVID-19 on Algerian Liquidity 23
3. Liquidity Management Practices by the Bank of Algeria 24
CHAPTER II: THEORETICAL FRAMEWORK AND LITERATURE REVIEW: FROM TRADITIONAL FORECASTING TO DEEP LEARNING FOR CENTRAL BANK LIQUIDITY PREDICTION **26**
Section 1: Traditional Forecasting Approaches for Central Bank Liquidity and Their Limitations 28
1. ARIMA Foundations 28
2. SARIMA: The Seasonal Baseline 30
3. SARIMAX: Incorporating Exogenous Drivers 31
4. Limitations of SARIMA and SARIMAX for Central Bank Liquidity Forecasting 33
Section 2: Deep Learning Architectures for Liquidity Time Series Forecasting 35

1. Artificial Neural Networks: Foundations 35
Section 3: Literature Review on Deep Learning Applications to Liquidity Forecasting and Comparative Studies 42
1. Deep Learning in Financial Time Series Forecasting 42
2. Comparative Studies: Traditional vs. Deep Learning Models 44
3. The Role of Exogenous Variables in Liquidity Forecasting 47
4. Research Gaps and Positioning 49

CHAPTER III: EMPIRICAL STUDY OF CENTRAL BANK LIQUIDITY FORECASTING USING STATISTICAL AND DEEP LEARNING MODELS 52
Section 1: Data Construction and Descriptive Analysis 55
1. Methodological Approach 55
2. Descriptive characteristics of the central bank Liquidity 55
Section 2: Forecasting using Classical Models 65
1. Univariate Time-Series Forecasting: The SARIMA Framework 65
2. Multivariate Integration: The SARIMAX Framework 74
3. Comparative Assessment: SARIMA vs SARIMAX 81
Section 3: Forecasting using Deep Learning Models 83
1. Data preparation and experimental design 83
2. Model configuration 84
3. Synthesis: Comparison of Deep Learning Models 94
Section 4: Final Benchmarking Across Daily, Weekly, and Monthly Horizons 96
1. Models Retained from Sections 2 and 3 96
2. Final Synthesis and Recommended Model by Horizon 99

GENERAL CONCLUSION 101
REFERENCES 105
APPENDICES 108
TABLE OF CONTENT 120