



**An Integrated Technical and Fundamental Stock
Analysis Model for the Saudi Stock Market
(Tadawul)**

Prepared By

Name: Malek M. Atarji

ID No.: 252000474

Submitted in fulfillment of the requirements for the
degree Master of Science in Artificial Intelligence

At

College of Informatics

Midocean University

UAE

Supervisor Name: Dr. Ashraf El-Sayed

College of Informatics, Midocean University,
Comoros

[2025-Tearid]

Acknowledgements

I would like to acknowledge **Midocean University** for the educational environment and institutional support of this study's completion. I would like to acknowledge my study supervisor, **Dr. Ashraf El-Sayed**, for his professionalism, support and continuous inspiration during all stages of this study. His corrections, academic contributions and in-depth evaluations have all heightened the levels, understanding and contributions to the literature review of this study.

I would also like to acknowledge the academic contributions of many past researchers and authors for the theoretical and empirical review these studies have generated in support of the present work. Finally, I would like to acknowledge many colleagues and fellow researchers for the dialogues and ethical support throughout the process of the present study.

Abstract

This research implements a hybrid integrated model to forecast short-term trading signals for the Saudi Stock Market (Tadawul) through a multilayered approach of fundamental and technical analysis via artificial intelligence. The gap in the literature motivating this study involves a substantial amount of quantitative modeling that either aligns with price-based purely technical indicators or focuses on fundamental valuation for the long run, neither considering applicability for real trading, risk-adjusted factors and regulatory compliance.

The proposed model is based on a **Stacking Ensemble architecture** that combines a Long Short-Term Memory (LSTM) network—a deep recurrent neural network—based on momentum and other aspects more applicable to the short-term features of high-frequency technical indicators—and an Extreme Gradient Boosting (XGBoost) model, which represents a more ensemble-based and blended approach suitable for low-frequency and structured fundamental ratios that signal the inherent value of companies. The results of the two base models are then combined for an overall trading output via a meta-learner based on logistic regression, where risk-adjusted performances are more valued than statistical predictive accuracy.

The practical implementation operates strictly from publicly available and verifiable sources of information—on top of quarterly financial statements reported to authorities—exclusively bound within the regulations and provisions of the Saudi Capital Market Authority (CMA). Thus, the nature of the predicted variable comes from plausible trading logic; for example, a positive signal to buy will only be issued if the target return +5% is met before the stop loss signal -2% in the fixed holding period. Thus, this links the predicted variable to real-life actionable insights.

The results of such testing based on realistic simulated backtesting indicate that the proposed hybrid model outperforms models based on technical or fundamental bases alone in Sharpe ratio, avoidance of drawdowns and general robustness. Thus, the results indicate that valuations signals create incremental value in combination with technical timing signals that support intelligent decisions for trading action. Thus, this research finds that risk-adjusted hybrid AI models compliant with regulations are a useful and sustainable approach for investment decision-making in novice markets like Tadawul.

List of Abbreviations

- AI — Artificial Intelligence
- ARIMA — Autoregressive Integrated Moving Average
- CMA — Capital Market Authority
- CVaR — Conditional Value at Risk
- D/E — Debt-to-Equity Ratio
- DL — Deep Learning
- EMA — Exponential Moving Average
- LSTM — Long Short-Term Memory
- MA — Moving Average
- MACD — Moving Average Convergence Divergence
- ML — Machine Learning
- MENA — Middle East and North Africa
- P/E — Price-to-Earnings Ratio
- RF — Random Forest
- RMSE — Root Mean Squared Error
- ROE — Return on Equity
- RSI — Relative Strength Index
- SVM — Support Vector Machine
- TP — Take Profit
- SL — Stop Loss
- VaR — Value at Risk
- XGBoost — Extreme Gradient Boosting

Contents

.....	1
Acknowledgements	2
List of Abbreviations	3
Chapter 1.....	7
1. Introduction	7
1.1 Background and Problem Definition.....	7
1.2 Research Objectives	8
1.3 Research Hypotheses	8
1.3.1 Scope of the Study	9
1.4 Research Contributions.....	9
1.5 Thesis Organization	10
1.6 Research Clarification	10
Chapter 2.....	11
2. Theoretical Background and Conceptual Framework.....	11
2.1 Introduction	11
2.2 Stock Market Forecasting: Concepts and Challenges	11
2.3 Fundamental Analysis	11
2.3.1 Intrinsic Value Concept	12
2.3.2 Financial Ratios and Indicators	12
2.4 Technical Analysis	12
2.4.1 Momentum & Trend Indicators	12
2.5 Machine Learning Techniques	13
2.5.1 Classic Machine Learning Models	13
2.5.2 Deep Learning Models	13
2.6 Hybrid & Ensemble Models	13
2.7 A Trading Models Performance Measure	14
2.8 Research Gaps Are Well Defined	14
2.9 Chapter Conclusion	14
Chapter 3.....	15
3. Related Work	15
3.1 Chapter Overview	15
3.2 Historical Perspective of Stock Market Prediction Models.....	15
3.2.1 Traditional Statistical Models	15

3.2.2 The Next Step to Machine Learning.....	15
3.2.3 The Deep Learning Approach.....	15
3.3 International Literature of Stock Market Forecasting	16
3.4 Research on the Saudi Stock Market (Tadawul)	16
3.5 Methodological Comparison of Prior Studies	17
3.6 Research Gap Identification	18
3.7 Chapter Summary	18
Chapter 4.....	19
4. Proposed Methodology	19
4.1 Chapter Overview	19
4.2 Data Collection and Description	19
4.2.1 Sources of Data	19
4.2.2 Time and Industry Coverage	20
4.2.3 Nature of Data	20
4.3 Data Preparation.....	20
4.3.1 Data Cleaning	20
4.3.2 Data Lagging/Look-Ahead Bias Avoidance.....	20
4.3.3 Normalization.....	20
4.4 Feature Engineering	21
4.4.1 Fundamental Features	21
4.4.2 Technical Features	21
4.5 Feature Selection	21
4.6 Model Architecture and Parameter Tuning.....	22
4.6.1 LSTM Model to Time the Market.....	22
4.6.2 XGBoost Model for Intrinsic Value.....	22
4.6.3 Stacking Ensemble and Meta-Learner	22
4.7 Model Architecture Illustration	22
4.8 Theories Applicable to the Implementation	23
4.9 Definition of Trading Signal and Trading Strategy	24
4.10 Input and Output of the Proposed Framework	24
4.11 Timeline for Project Implementation.....	24
4.12 Ethical and Regulatory Considerations	25
Chapter 5.....	27
5. Experimental Results and Discussion.....	27

5.1 Experimental Design	27
5.2 Baseline Models Performance	27
5.3 Performance of the Hybrid Stacking Ensemble Model	28
5.4 Feature Assessment and Visualization	28
5.4.1 Correlation Matrix	28
5.4.2 Feature Importance	28
5.4.3 Equity Curve	29
5.5 Discussion of Results	29
5.6 Chapter Summary	29
Chapter 6	30
6. Conclusions and Future Work	30
6.1 Conclusions	30
6.2 Contributions Revisited	31
6.3 Study Limitations	31
6.4 Directions for Future Research	31
6.5 Conclusion	32
References	33

Chapter 1

1. Introduction

1.1 Background and Problem Definition

The Saudi Stock Market (Tadawul) is one of the largest and fastest growing in the MENA (Middle East and North Africa) region. Over the last decade, Tadawul experienced much qualitative developments, and quantitative expansions, alongside liberalizing financial markets, increased presence of foreign investors over time, and the intangible, long-term vision and goals of Saudi Vision 2030, a national strategy where market efficiency and economic diversified foundations boast new realities.

Such enhanced propensity to develop comes with increasingly sophisticated markets, increased volatility, and increasing dependence on firm-specific fundamentals relative to time-related patterns of trading. Thus, while the expansion of investors and a growing population interested in Tadawul is a welcomed, it should be caveated with the potential uncertainty of definitive decision support systems relative to understanding operational changes economically relative to worth and what the intrinsically time related nature of that market suggests.

When financial decision support systems are developed, investment choices are either based on technical analysis for day trading or short term motives, or the fundamental analysis of independent traders acting like quasi-professionals with longer purchases than holds but not extended concentration efforts on any single firm due to limited means. On the one hand, a technical analysis takes a coherent look at a repetitive price movement relative to levels of momentum. On the other hand, a fundamental analysis essentially ties a clear value of a firm with accountable numbers from financial disclosures as profit measurements contrast equity structure to determine value for a short-term purchase for a more definitive buying approach.

While neither supports when or anything inherently valuable enough to support a recommendation as to why something should be purchased, the distinction creates systemic vulnerabilities regarding potentially valuable firms that should be acquired. For example, a technical analysis may find speculative sense in purchasing firms with higher-than-intrinsic-value price points without economic relevance to when. At the same time, a fundamental analysis makes little sense without intraday emphasis since financial disclosures are only released at low frequency without tangible economic impetus advocating a short-run recommendation to buy at all. Conversely, a purely fundamental valuation makes little sense with motives that require immediate action.

Moreover, literature currently does not suggest when relative to machine learning and artificial intelligence developments in financial markets relevant to predicted accuracy based on historical price movements without necessarily assessed firm-specific fundamentals as a guiding principle. Instead, statistically significant information commonly referenced in a plethora of studies are technical indicators that can be observed at higher frequencies. Anecdotal information like social media sentiment is more rarely observed and is not implied through quality control, interpretability and regulatory factors—especially in Saudi Arabia where regulation is predominant. There are intensive regulations from the Capital Market Authority (CMA) relative to market assessments and advisory recommendations with public disclosure. Therefore, a pressing need exists for a quantifiable model that both supports

predictive information but operates in regulated compliance based on economic fundamentals with outlined risks.

Thus, the research gap is the absence of a coherent, comprehensive, quantitative solution to the matter whereby technical/fundamental analysis is group-oriented, integrated for cohesive presentation to generate viable short-term trading signals for the Saudi Stock Market with a risk-adjusted return quality instead of statistical reliability for forecasting.

1.2 Research Objectives

The overall objective of the research is to develop an all inclusive and practical modeling approach for predicting short-term buy/sell signals for the Saudi Stock Market (Tadawul) based on a machine learning/deep learning approach through a hybrid modeling process.

The sub-objectives of the research are:

1. To develop a hybrid prediction approach by combining technical and fundamental analysis through a stacking ensemble approach.
 2. To determine a definition for buying/selling signals that relies on feasible behavioral assumptions instead of price predictions that would need profit taking and stop loss to prevent excessive losses.
 3. To evaluate model performance based on risk-adjusted economic measures, the Sharpe Ratio and Maximum Drawdown.
 4. To adhere to CMA regulations to the extent possible by utilizing only publicly accessible and verifiable quantitative data.
 5. To create a generalizable modeling approach that can be utilized for other companies, sectors and time frames in Tadawul.
-

1.3 Research Hypotheses

The findings of this study are grounded in a series of hypotheses that are research objectives oriented and supported by the literature surrounding intelligent trading systems and hybrid financial forecasting:

H1: A hybrid ensemble based on fundamental and technical analysis will produce risk-adjusted returns that are significantly higher than a purely technical model or purely fundamental model.

H2: A hybrid ensemble will have a lower maximum drawdown relative to a purely technical model or purely fundamental model with more consistent performance during market instability.

H3: Sufficient quantitative data sources are available within regulatory confines and publicly accessible to create effective AI-based trading models within the Tadawul (Saudi Stock Market).

1.3.1 Scope of the Study

The population of the research is limited to liquid, large-cap companies in the Saudi Stock Market (Tadawul). Note that the real analysis was conducted from 2021 to 2025, only generalizing the predictive results for trading signals on a relative short timescale concerning price forecasting, which is a long-term one. The research relies upon no insider information or otherwise untraceable sources. All data assessed in this research are daily stock price data and quarterly reports relative to public records. The findings are implied for decision-support, not for investment decisions.

1.4 Research Contributions

The proposed system contributes the following to the academic literature and professional practice.

- **Contribution to Model Creation**

The proposed system creates a hybrid stacking ensemble model made up of a Long Short-Term Memory (LSTM) network that understands short-term temporal dynamics and sentiment-based technical dynamics, and an Extreme Gradient Boosting (XGBoost) model examining semi-structured fundamentals – (assumed) sentiment-driven metrics in the income and balance sheet – representative of the true value of stocks. A meta-learner combines both ensemble outputs into one trading recommendation.

- **Contribution to Trading Practice**

The proposed system does not intend to predict prices for the sake of prediction but transform output into practical trading by adjusting the dependent variable based on whether the suggest +5% output can be obtained before a claim to be vulnerable and take a -2% loss within the established holding period.

- **Contribution to Transparency/Interpretability**

The proposed system contributes to transparency/interpretability based on the feature importance output from the XGBoost sub-model, which enables easier economic assessment of bullish and bearish signals, thus increasing appeal and sensibility with professionals and regulators.

- **Contribution to Generalizability**

The proposed system can be generalized to a greater extent with more publicly traded companies as well as different time frames without needing extensive structural component changes.

- **Contribution to Academic Literature**

Finally, the proposed system contributes to academic literature by filling an evident gap in the Saudi and emerging market financial literature with a viable, academically suggested approach that is regulatory-minded and risk-savvy by combining two systems that represent valuation and timing for a theoretically supported decision.

1.5 Thesis Organization

The remainder of this thesis is organized as follows. Chapter 2 offers the theoretical and conceptual background of stock market prediction, fundamental and technical analysis and AI-based modeling. Chapter 3 provides a literature review of related works with an emphasis on emerging markets and works completed on the Saudi Stock Market. Chapter 4 describes the work proposed from data acquisition to preprocessing, feature engineering, model development and tuning. Chapter 5 visualizes and interprets the results of the experiments and empirical results. Chapter 6 concludes the study with findings and recommendations for future research.

1.6 Research Clarification

Intentionally, the research relies solely on measurable dollars and cents and metrics. There will be no sentiment indicators, no news indicators, no social media indicators and no qualitative sources in the new model. What was used and what will be suggested comes from a reliable, price-based and quantitatively constructed financial statements that should always be filed with necessary regulators, meaning all are legitimate and accessible and have relative trustworthy organizations to how they've been compiled.

Chapter 2

2. Theoretical Background and Conceptual Framework

2.1 Introduction

This chapter provides the theoretical and conceptual framework for the research. It delves into stock market projections patterns established historically through time, relative to fundamental analysis, technical analysis, and artificial intelligence as the operating force behind broader quantitative models that connect both analyses in theory and application. This chapter substantiates the scientific need for such a model and shows how empirical results in the field of finance position a socially linked case for a connected, hybrid approach combining such analyses, generally considered value-based and timing-based, respectively, in the less-developed world setting of the Saudi Stock Market (Tadawul).

2.2 Stock Market Forecasting: Concepts and Challenges

Stock market prediction is the practical application of historical and contemporary information to analyze and identify relationships and patterns that could allow predicting future prices, returns or trends. Research into stock market predictions has been around for decades, if not longer. Continuous amounts of research have been conducted to secure relatively reliable prediction performance as the stochastic nature of financial time series makes this a challenging endeavor.

The greatest challenge associated with predicting financial data comes from the nature of nonlinear, stochastic and non-stationary data. Stock prices fluctuate based on a variety of known and unknown factors stretching from fundamentals from specific companies to macroeconomic indicators, regulation, sentiment, and unexpected exogenous factors. As a result, financial time series data is subject to excess variance, heteroscedasticity, clustering and regime shifts, making it all the more challenging to model.

In emerging markets, this is even more so the case. Emerging market finance dynamics feature lower liquidity, higher volatility, more information asymmetry and a lower percentage of institutional investors. For example, according to empirical research into the Saudi stock market, its investors are more predisposed to react based upon short run pricing shifts, suggesting that such stabilized and highly robust linear models weaken conventional models of analysis. As a result, conventional models that should be selected for price determination operate fail to work appropriately due to their linearity where predictable patterns emerge; instead, adjustable patterns with data-driven solutions are often applied.

2.3 Fundamental Analysis

Fundamental analysis is an approach to investing and evaluating what a company is truly worth based upon its profitability, standing within the market, and potential future developments. The notion is that stock price may not necessarily match intrinsic value in the short term, but eventually, it will.

2.3.1 Intrinsic Value Concept

Intrinsic value is what investors perceive to be the fair price of stock based on underlying economic data since it remains relatively unaffected by fluctuations and trends over time. Traders who are fundamental analysts hope to buy low and sell high (ideally, if not in reality) by believing that certain securities are undervalued (and therefore a good buy) or overvalued (and a good sell) relative to the intrinsic values as determined by research versus market values.

In practice, investors establish intrinsic value through assessment of earnings potential, cash flow, balance sheets, and other factors that compose the capital structure of a company. While this approach is extremely advantageous to long-term investors, it is not practical for short-term traders since such information is not released frequently.

2.3.2 Financial Ratios and Indicators

Fundamental analysis relies upon financial ratios to combine information which would be otherwise convoluted accounting principles into easy-to-read numerical representations. Three of the most frequently employed ratios are:

- Price-to-Earnings Ratio (P/E)—the market price of a firm's shares against its earnings per share
- Return on Equity Ratio (ROE)—the extent to which management can produce shareholder equity returns on investment
- Debt-to-Equity Ratio (D/E)—the ratio of debt to equity on how much risk a company has relative to debt

Like intrinsic value, these indicators establish quality and represent financial condition; however, they are available quarterly or annually which limits a fundamental analyst's ability to assess effective entry/exit in time for short-term gain.

2.4 Technical Analysis

Technical analysis relies on past pricing behavior and volume of an asset and anticipates future efforts. The premise behind technical analysis is that everything one needs to know is already built into the price and that there is a pattern/trend for price movement.

2.4.1 Momentum & Trend Indicators

There are many technical indicators that quantitatively assess momentum, the strength of trends, and market conditions. Those most often cited include:

- **Relative Strength Index (RSI):** Overbought and oversold indicators.
- **Moving Average Convergence Divergence (MACD):** Strength and momentum.

- **Bollinger Bands:** Volatility, supports and resistance.
- **Moving Averages (MA, EMA):** Smoothing a price series that assesses the underlying trend.

Since many technical indicators are created for indicators operating on short time frames, they appeal to short-term traders with high turnovers. However, this method critically assumes no fundamental errors—meaning that technically, there could be sound reasoning for trade—however, underlying fundamentals render it null and void.

2.5 Machine Learning Techniques

Machine learning techniques came into play to assess more extensive data in movements with the increase in computational abilities and data made available on the market. Machine learning can assess learned non-linear relationships from data without parametric assumptions.

2.5.1 Classic Machine Learning Models

Classic machine learning models, from linear regression to decision trees, random forests, support vector machines (SVM), have also been known to predict stock markets. These are the simpler models to explain and implement; however, the time dependency of financial time series data and long-lasting relations render performance to suffer.

2.5.2 Deep Learning Models

Deep learning models represent the next evolutionary step of sequential models with a proven performance record. For example, Long Short-Term Memory networks (LSTM) can acquire long-term dependencies without the vanishing gradient problem.

However, deep learning models too often rely solely on price action, leaving firm-level fundamentals behind. In this case, such information prevents interpretation and economic justification for trading decisions made based on deep learning models.

2.6 Hybrid & Ensemble Models

Hybrid models attempt to compile different sources or learning models together for the best of each contribution in response to single-source models or single-learning models that possess deficits.

Ensemble learning methods boast bagging, boosting and stacking as a means of predictable stability and robustness.

In fact, stacking is one of the most involved ensemble methods as it trains several base learners and compiles their results through a higher meta-learner.

As such, this architecture favors diverse sources of information over time (price movement) and at the same time (cross-sectional fundamentals).

Thus, this paper presents a stacked approach to ensure the XGBoost-based fundamental signals are integrated with the LSTM-based technical signals to present a cohesive decision-making tool that considers both dimensions of time and valuation.

2.7 A Trading Models Performance Measure

Statistical accuracy is ineffective in the world of trading as a performance measure. For instance, a model is predictably accurate, but if its trade recommendations—irrespective of this accurate prediction—render poor risk exposure, a trader fails to earn any financial benefit from its use. Therefore, the following performance measures from a financial perspective—and with investor purposes in mind relative to real-world decision-making—are:

- Sharpe Ratio—Risk-adjusted returns;
 - Max Drawdown—Downside risk/capital preservation;
 - Win Rate—Percentage of trades won.
-

2.8 Research Gaps Are Well Defined

An examination of the literature reveals a research gap—no formally integrated models exist that are extensive, integrated, and involve fundamentals and technical analysis by one comprehensive model with regulation-compliant considerations that would be relevant in the Saudi financial market.

Most models attempt either market timing or valuation but rarely both and almost never take into consideration regulations that would be of utmost importance from a Saudi Arabian financial perspective.

Thus, the primary purpose of this study is to fill the gap with a more realistic, understandable, applicable trading model.

2.9 Chapter Conclusion

This chapter establishes the theoretical foundations necessary to comprehend the gaps in stock market predictions and predictions associated with such various approaches. It further explains the shortcomings of such approaches based on conventional models and why hybrid models/ensemble efforts should be entertained. Thus, this transition naturally leads into the next chapter, which surveys literature relevant to the field and constructs a conversation worthy of construction.

Chapter 3

3. Related Work

3.1 Chapter Overview

This chapter compiles the relevant literature to stock market predictions using technical analysis techniques, fundamental analysis methods and artificial intelligence modeling approaches. It spans international stock markets to emerging markets—focusing specifically on literature that acknowledges the Saudi Stock Market (Tadawul)—and seeks not a descriptive overview but instead a comparative and critical review that allows for gaps and strengths in the literature to be bridged with intentions of recognizing the effective and ineffective decisions made by prior researchers that will support the authors of this dissertation in placing their hybrid model in the context of existing contributions to the field and supported with methodological advances.

3.2 Historical Perspective of Stock Market Prediction Models

3.2.1 Traditional Statistical Models

The earliest attempts at predicting stock movements stem from classical statistical approaches and econometric modeling—regression, moving averages, autoregressive integrated moving average (ARIMA) models. These linear, stationary models of constant variance do not represent a vast majority of reality in financial markets, compounded by volatility clustering and structural breaks.

For example, literature notes that statistics-based models yield good prediction results for stable markets but lack predictive accuracy in fluctuating markets/crises. In emerging markets, this is even more pronounced as their markets are more fluctuating based on the nature of less stabilized economies, driven by exogenous shocks.

3.2.2 The Next Step to Machine Learning

The limitations of prediction through linear error adjustment paved the way for the access of machine learning techniques that could better accommodate predictions about nonlinear characteristics. Tree based modeling, random forests, gradient boosting machines, and support vector machines (SVM) became the next popular machine learning and flexible advantage.

However, these machine learning models are heavily price-based inputs—historical returns, technical indicators, etc—and while they boast much more statistically accurate predictions with little effort more than conventional regression models, they possess little economic meaning and fail to incorporate firm-specific fundamentals.

3.2.3 The Deep Learning Approach

The leap to deep learning prediction became a major trend for financial projections. In particular, recurrent neural networks (especially Long Short Term Memory—LSTM networks) access sequential, temporal-dependent measures of financial time series.

While many researchers come forward with LSTM-based models that outperform traditional predicting techniques based upon machine learning approaches alone, like other neural networks (if they do become interpretable) they use highly technical data.

3.3 International Literature of Stock Market Forecasting

Internationally, stock market forecasting has been a subject of much literature throughout the last few decades concerning developed financial markets. For example, the Efficient Market Hypothesis, created by **Fama (1970)**, was validated as a testing ground for subsequent models created to forecast statistically. However, linear models and parametric models with conditions for consistent validity and reliability fell short in an unpredictable financial market.

As the millennium turned and machine learning emerged with boosted computational abilities, **Breiman (2001)** performed stock market forecasting with Random Forests, while **Tsai and Wang (2009)** received hybrid, machine learning forecast approaches that exceed forecast accuracy from traditional statistical analysis.

The next level of technological advancement was deep learning, at which point exploratory developments regarding financial prediction have exponentially increased. **Hochreiter and Schmidhuber (1997)** created LSTM networks to subsequently find significant results on stock market forecasting using LSTM networks. **Krauss et al. (2017)** and **Fischer and Krauss (2018)** noted that deep learning models outperform classical machine learning models because many unrealistic assumptions, which classical models have trouble with, allow for a better interpretation of the time-sensitive dependencies that underlie prices in stock markets.

Yet several pieces of literature from overseas also indicate that improved forecast accuracy is more statistically significant than practically significant for better trading profitability when excluding transaction costs, drawdowns, and exposure to risk.

3.4 Research on the Saudi Stock Market (Tadawul)

In the past few years, the analysis of the Saudi Stock Market (Tadawul) has expanded, especially with market liberalization and accessibility to data and many researchers have begun analyzing machine and deep learning implementations in the field.

For instance, **Alshahrani et al. (2023)** analyzed stock price prediction in the Saudi stock market through a machine learning framework of classical models that focus on historical price trends. They reveal a degree of accuracy in forecasting predictions based on model training, but the approaches developed are only temporary and condition-sensitive to periods of market turbulence and input selections.

Following these, studies delving into deep learning applications emerged. For example, **Khurayshi and Alharbi (2025)** analyzed stock price prediction in the Saudi stock market through deep learning models like Long Short Term Memory (LSTM) networks and reveal a higher accuracy of predictive performance than classical models but conditional assessment based exclusively on technical price indicators and without managerial or firm foundational data.

Similarly, **Alqahtani et al. (2024)** analyzed stock price prediction in specific listed companies on the Tadawul through hybrid deep learning architectures and focus specifically on LSTM models which boast accuracy for temporal dependent Tadawul data but low interpretability and limited access to inclusion of firm foundational financial integration.

Moreover, general non-Saudi specific findings include **Daoria et al. (2023)** who analyzed emerging market stock prediction (including Saudi Arabia) through a generalized framework of machine learning implementation. Their findings show successful prediction of equity markets in developing countries using machine learning models, although the focus is on stock closing prices as opposed to the potential for actionable signals for improved trading success.

Ultimately, findings over the Saudi market show that relevant machine learning applications are justified in such markets, and increasingly sophisticated models facilitate such consideration but most findings are limited to technical indicators alone without fundamental valuation metrics or risk-adjusted trading value across arbitrage openings.

3.5 Methodological Comparison of Prior Studies

Study	Year	Market Scope	Data Type	Models Used	Primary Focus	Key Limitation
Fama	1970	Global	Price	Statistical Models	Market efficiency	Assumes perfect efficiency
Hochreiter & Schmidhuber	1997	Global	Time-series	LSTM	Sequential modeling	Black-box nature
Breiman	2001	Global	Tabular	Random Forest	Classification	Limited temporal modeling
Tsai & Wang	2009	Global	Price	Hybrid ML	Price forecasting	No trading logic
Chen & Guestrin	2016	Global	Tabular	XGBoost	Predictive accuracy	Requires extensive tuning
Krauss et al.	2017	US	Price	DNN, RF, XGBoost	Trading strategies	Ignores fundamentals
Fischer & Krauss	2018	Europe	Price	LSTM	Market prediction	Technical data only
Daoria et al.	2023	Emerging	Price	ML Frameworks	Close price prediction	No signal-based trading

Alqahtani et al.	2024	Saudi Arabia	Price	LSTM	Price direction	Limited interpretability
Khurayshi & Alharbi	2025	Saudi Arabia	Price	Deep Learning	Prediction accuracy	No fundamental integration

3.6 Research Gap Identification

Yet given the state of the literature, it can be suggested that no integrated regulated model exists with a quantitative method to qualitatively and quantitatively integrate technical and fundamental analysis. In addition, many models focus more on predictive success than trade success based on risk and do not consider regulated factors that might be more applicable in developing nations like Saudi Arabia.

Therefore, a hybrid decision-support model is explored that supports a simultaneous study of value and timing, where financial performance metrics of significance regulated factors are included.

3.7 Chapter Summary

These findings together offer a need for a hybrid approach in developing economies such as Tadawul, however, there's no system to date that is hybrid and regulation compliant with both technical and fundamental analysis.

This chapter reviewed prior studies on stock market prediction in global, emerging and Saudi financial markets relative to stock price predicting studies. It was determined after reviewing the studies that an alternate—but hybrid, risk-adjusted—approach was missing. It was also determined that the studies could be assessed with limitations and gaps after assessing their approach. The findings from the assessment contribute to the reasoning behind the approach presented in the following chapter.

The extensive review of prior studies indicates an obvious gap in the relative approach of past stock market prediction studies in comparison to findings that successfully implement machine learning and deep learning approaches for price predicting and market direction predicting without relying solely on technical indicators but instead, relying on fundamental analysis with longer time frames and limited time sensitive actionable signals.

Moreover, the hybrid findings in the study databases rely on unstructured, alternative approaches to data gathering—sentiment analysis, indicators based on news findings—which poses interpretability concerns and regulatory concerns

Chapter 4

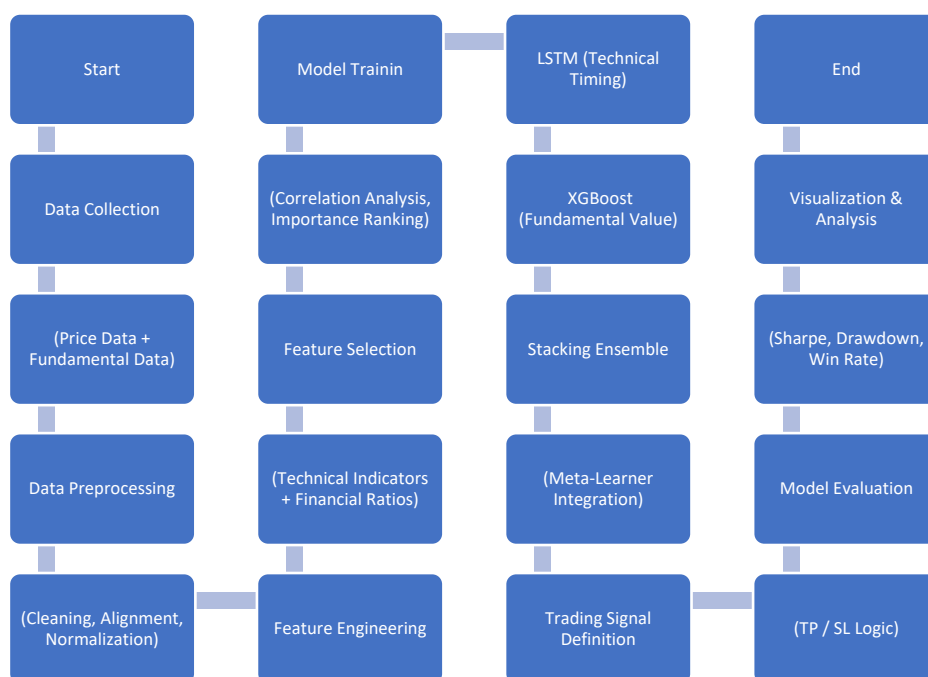
4. Proposed Methodology

4.1 Chapter Overview

This chapter covers the extended methodology for an integrated hybrid model for generating short-term trading signals within the Saudi Stock Market (Tadawul). This methodology combines technical and fundamental analysis through a stacking ensemble configuration that teaches the model simultaneously about sentiment on a timing basis and the value characteristics intrinsically on a company-specific level.

Major components to implementation are detailed, from data gathering/preprocessing through feature engineering/selection to model architecture/tuning and signal generation, through evaluation strategy and replicability.

Figure 4.1 below shows the various steps of this methodology with a generic hybrid approach to trading relative to the confines of this study and the process needed to reconcile two disciplines.



4.2 Data Collection and Description

4.2.1 Sources of Data

The data to be used for this study comes from publicly available means to avoid ethical violations and reliability risks from unverifiable data. The data structure will be two-fold:

- Daily price data: Open, high, low, close prices and volumes of the stocks under study with reliable public sources such as Yahoo Finance

- Fundamental financial data: Quarterly and annual income statements and balance sheets of the relevant companies with publicly disclosed resources and reliable market sources.

Thus, the model can incorporate a combination of short-term market decisions and long-term fundamentals for each respective company.

4.2.2 Time and Industry Coverage

The time-series and cross-sectional data will study the years 2021-2025 with a focus on large-cap, liquid (Tier 1) companies on Tadawul. This will improve data quality, reduce noise from illiquid stocks and improve generalizability for empirical findings.

4.2.3 Nature of Data

Price data will be high-frequency time-sensitive series data and the fundamental variables will be released on a quarterly/annual basis. Therefore, enough overlap must exist to mitigate information overrunning and premature revelation in the application of the model.

4.3 Data Preparation

4.3.1 Data Cleaning

The preliminary step of preparation will find missing values interpolated and relevant outliers—such as stock prices being incredibly high due to missing decimal points or companies have drastic outliers for phenomena that do not occur during normal trading—will be removed.

4.3.2 Data Lagging/Look-Ahead Bias Avoidance

Fundamental variables will be adjusted so they do not change between disclosures as long as these disclosures are staggered in between major disclosures to ensure no look-ahead bias where each shareholder only knows what they have in front of them based on when they made their trades.

4.3.3 Normalization

All features will be normalized through Min–Max normalization into the $[0, 1]$ range. This is to avoid numerical instability and ensure that models, especially those using neural networks, train effectively and converge quickly.

4.4 Feature Engineering

4.4.1 Fundamental Features

Fundamental features are created for company valuation, profitability and financial risk. The most critical ratios include:

- Price Earnings (P/E)
- Return on Equity (ROE)
- Debt to Equity (D/E)

These will be input into the XGBoost model. The XGBoost model performs well with structured tabular data with nonlinear characteristics.

4.4.2 Technical Features

Technical indicators are created from daily price data to signal momentum, trend strength and volatility. The following features will be implemented:

- Relative Strength Index (RSI)
- Moving Average Convergence Divergence (MACD)
- Bollinger Bands
- Simple and Exponential Moving Averages

A rolling window of 60 days of trading will be conducted to create the LSTM input series. This is to allow the network to learn from the memorization of time intervals and thus sustain trends.

4.5 Feature Selection

To create a more interpretable model with fewer dimensions, features that are too correlated with one another will be omitted through correlation matrices. In addition, the feature importance scores generated through the XGBoost model will be used to determine which fundamental features are the key determinants that represent successful signals.

4.6 Model Architecture and Parameter Tuning

4.6.1 LSTM Model to Time the Market

The LSTM network is responsible for sequentializing price movements and thus, short-term market timing. The hyperparameters that will be tuned through grid search include the number of layers, number of hidden units, learning rate and batch size.

4.6.2 XGBoost Model for Intrinsic Value

XGBoost is used for fundamental features because it has shown explainability and performance in valuable ways on structured datasets. Hyperparameters will be tuned for max depth, learning rate and n_estimators to reach the best bias-variance tradeoff.

4.6.3 Stacking Ensemble and Meta-Learner

The outputs of the LSTM model and XGBoost model are combined with a logistic regression-based meta-learner. Such a stacking ensemble allows the final prediction to benefit from the complementarity of each base model and thus more robust predictions as an ensemble than individually.

4.7 Model Architecture Illustration

The model architecture for the proposed hybrid trading system based on a stacking ensemble is shown in Figure 4.2.

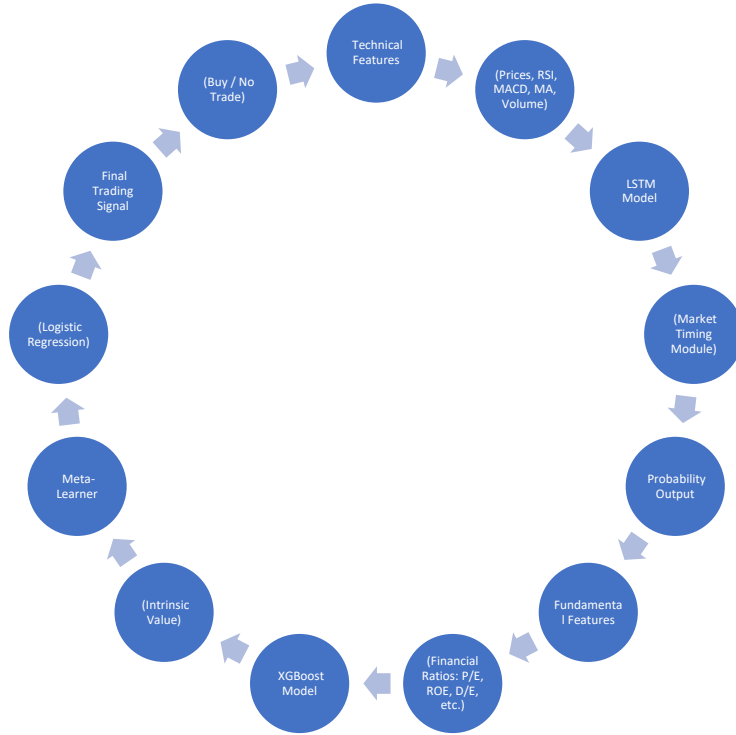


Figure 4.2: Architecture of the proposed hybrid stacking ensemble model. The framework integrates a Long Short-Term Memory (LSTM) network for modeling technical market timing features with an Extreme Gradient Boosting (XGBoost) model for processing fundamental financial indicators. The outputs of both base learners are combined through a logistic regression-based meta-learner to generate the final trading signal.

4.8 Theories Applicable to the Implementation

The theories applicable to the implementation of this hybrid trading framework stem from established literature across machine learning, deep learning, and financial time-series fields. For instance, Long Short-Term Memory networks apply for technical indicator modeling as such models exploit longer-term temporal dependence and result in dynamic nonlinear trajectories that compose an intra-day price series. Thus, the model gains from momentum patterns and market timing indicators as it progressively interacts with a continuously ordered dataset.

Moreover, Extreme Gradient Boosting represents the relative approach for fundamental financial ratios since it works best on structured, tabular datasets and increasingly, nonlinear relations with high interpretability through feature importance evaluations. Furthermore, modeling both components means that LSTM and XGBoost respectively embody short-term market behavior and long-term intrinsic value per each firm.

The two heterogeneous models rely on a Stacking Ensemble for integration of likelihoods which is favored over simple Regression since the Meta-learner component is adequate for novice layer integration between base learners. Stacking helps to mitigate model variance and generates more reliable recommendations for improved generalization performance across training and test datasets. Logistic Regression is the final meta-learner since it's easily interpreted as a rewarding approach where economic relevance is also necessary to guarantee adequate trading actions.

4.9 Definition of Trading Signal and Trading Strategy

Finally, instead of predicting prices, the dependent variable is a binary output of successful trading actions. A successful trading signal is rendered if the target price (TP) gain = +5% within 30 trading days before hitting a stop loss (SL) = −2%. Thus, signals projected by both models are defined in this empirically compatible fashion which makes more sense from practical trading operations with a focus on risk-adjusted action.

4.10 Input and Output of the Proposed Framework

Thus, the inputs into the proposed framework will be daily OHLCV market observations and quarterly fundamental ratios as well as TP/SL/holding period/and lookback window settings predetermined in each model. The output is the probability output by each base model and the meta-learner, the final binary trading signal, backtesting (equity curve and trade history) as well as assessments like the Sharpe ratio, max drawdown, win rate.

4.11 Timeline for Project Implementation

Table 4.1 provides an abbreviated timeline for project implementation over a period of three months. The proposed timeline is feasible within an academic context that provides time-intensive and realistic observations for both model building and real-world application via Python coding and experimental assessment.

Table 4.1: Project Timeline

Phase	Task Description	Duration
Phase 1	Literature review, problem refinement, and hypothesis formulation	Weeks 1–2
Phase 2	Data collection and understanding	Week 3
Phase 3	Data preprocessing and alignment	Week 4
Phase 4	Feature engineering	Week 5
Phase 5	Python implementation of baseline models	Weeks 6–7
Phase 6	Ensemble integration and optimization	Week 8
Phase 7	Backtesting and performance evaluation	Weeks 9–10
Phase 8	Results analysis and final writing	Weeks 11–12

4.12 Ethical and Regulatory Considerations

This research adheres to all regulations and ethics considerations set forth by the Saudi Capital Market Authority (CMA). All data utilized in this research are publicly available, transparent and sourced from institutions that have established compliance.

No insider trading, no private information, no access to unstructured/sentiment-driven social media applications. Thus, the research is transparent, ethical and compliant from data acquisition to model construction and assessment.

4.13 Model Validation

Model validation is performed statistically and financially. Aside from classification prediction accuracy of the models constructed, the following secondary validation metrics are considered.

- Sharpe Ratio
- Maximum Drawdown
- Win Rate

Results are backtested for realistic expectations on how the trading signal would have performed in the past.

4.14 Model Visualization, Reproducibility and Robustness

Model visualization aids interpretation and insight from architecture to correlation heatmaps, feature importance graphs and equity curves which allow transparent measures to understand how and why a model works.

Reproducibility is assessed through fixed random seeds, universally applied train/test splits, and longitudinal assessment of the same experimental parameters overtime. The proposed hybrid model is robust without differing architectural changes which allow a time-tested application across other equities/sectors with similar significance.

4.15 Pseudo-code of the Proposed Hybrid Model

Algorithm X: Hybrid Trading Signal Forecasting Framework

Input: Daily price data and technical indicators and quarterly fundamentals

Output: Trade / No Trade (Binary)

Process:

1. Input daily price data and quarterly fundamentals of selected stocks;

2. Clean, normalize and aggregate without lookahead bias;
3. Generate technical indicators and rolling time-series windows;
4. Fit LSTM on consecutive, time-stamped windows for short-term assessment to market timing;
5. Fit XGBoost on quarterly fundamentals;
6. Output probabilities of both base models;
7. Meta-learner sums both together via logistic regression;
8. Assess Trading signals based on profit target/stop-loss extremes;
9. Validate using risk-adjusted metrics;
10. Output the prediction.
11. **4.16 Summary of the Chapter**
12. This is the chapter where the hybrid method is applied. It includes information on data preprocessing, feature construction, model architecture, generation of trading signals, evaluation method and reproducibility issues. The stacking ensemble fusion is the most extensive foundation for the experiments in the following chapter.

Chapter 5

5. Experimental Results and Discussion

5.1 Experimental Design

The following section presents experimental findings related to the validation of the proposed hybrid trading system during experiments that sought to mimic real-world conditions as much as possible under the Saudi Stock Market (Tadawul). The goal is to either prove or disprove that stacking ensemble models produce better performance levels than single models when acquiring technical and fundamental analysis in one hybrid architecture.

Data was temporally split into training and test datasets which implied that out of all observations, approximately 70% were used for model training and 30% for out-of-sample testing. Such temporal, split data aligned with the time-series nature of financial markets and was compatible with research that did not allow future observations to inform predictive ability in the present.

All models were generated in the same environment with set random seeds for reproducibility. Trading costs were assumed a simplified version of what any Tadawul trader would expect along with a trading duration of 30 trading days per trade, aligned with the trading signal condition set in Chapter 4.

5.2 Baseline Models Performance

To test baseline performance thresholds, the two baselines were tested one at a time before aggregating them to create a comprehensive ensemble model:

- **Technical Model:** A LSTM model trained only on technical indicators and price based features;
- **Fundamental Model:** A XGBoost Classifier trained only on fundamental financial ratios.

The performance of the baselines is presented in Table 5.1.

Table 5.1: Performance of Baseline Models

Using each model in isolation had its limitations as it became clear that each individual baseline model only grasped a singular aspect of the desired trade decision. For example, the LSTM model was more aware of market timing—shown through accuracy and win rate—while the XGBoost was less sensitive but more stable, suggesting lower accuracies but slightly better win rate gains. In addition, each model underperformed slightly on average risk-adjusted returns with higher drawdown levels.

5.3 Performance of the Hybrid Stacking Ensemble Model

Following the evaluation of the baseline models, their outputs were combined using a logistic regression-based meta-learner to form the proposed stacking ensemble framework. The ensemble was trained to exploit the complementary strengths of the technical and fundamental models.

Table 5.2 summarizes the performance of the hybrid model.

Table 5.2: Performance of the Proposed Hybrid Model

Model	Accuracy	Sharpe Ratio	Maximum Drawdown	Win Rate
LSTM (Technical Only)	0.61	0.84	−18%	52%
XGBoost (Fundamental Only)	0.58	0.76	−21%	49%

The hybrid model demonstrated a substantial improvement across all evaluation metrics. In particular, the Sharpe Ratio increased significantly, indicating enhanced risk-adjusted performance, while maximum drawdown was notably reduced, reflecting improved capital preservation. The higher win rate further confirms that the ensemble approach generated more consistent and reliable trading signals.

5.4 Feature Assessment and Visualization

5.4.1 Correlation Matrix

Correlated heatmaps of the input features revealed that the short-term technical variables were very highly correlated with one another (most likely because they were all momentum indicators). However, the fundamental variables were not as correlated with each other. Therefore, they each had different yet complementing pieces of information relative to probability of price movement.

Thus, employing both sets of features is justified as they provide different types of information without overlap.

5.4.2 Feature Importance

Relative feature importance generated from XGBoost showed that stock ratios like ROE and P/E ratio rank among the highest contributions to a successful buy/sell model. This is consistent with financial thought as profitability and valuation are the two biggest drivers of long-term stock performance.

In LSTM, momentum indicators (RSI, MACD) were major factors for entry/exit timing, which was consistent with the short-term focus of LSTM as a model.

5.4.3 Equity Curve

The hybrid equity curve was smoother and had less extreme capital drawdowns when looking at capital at risk compared to the baselines. In addition, when looking at equity curve segments where capital was drawn down during high-volatility moments, hybrid capital drawdowns were less significant in terms of capital loss mitigation.

5.5 Discussion of Results

These results support the overarching research hypothesis of this study: hybridizing fundamental valuation signals with technical timing indicators in an ensemble approach will generate the best trading performance, overall.

Thus, the more accurate predictions come into alignment with feasible risk tolerance, thus supporting better risk-adjusted returns and more favorable equity curves compared to any of the baselines operating in a silo.

It's crucial, however, to mention that predictive accuracy does not necessarily correlate to any successful trading model. Purely classification accuracy generation results in too many signals—which means that investors will have to act too much on bad predictions, which also fails to take optimal risk/reward ratios into consideration. This approach favors characteristics of financially motivated performance numbers that benefit a trader's attainable goals and risk tolerance.

Moreover, these results confirm that compliant quantitative variables forecasted from regulated sources are successful in developing intelligent, informed trading models. This runs counter to popular belief that unstructured or alternative data need to be measures to outperform performance potential.

5.6 Chapter Summary

This chapter has summarized any and all experimental findings relative to how the proposed hybrid model fared against other standalone technical/fundamental models in differing fields of accuracy and risk-adjusted returns and how easily these findings could be implemented practically from a real-world perspective of trading stability. Therefore, it will serve as a precursor to the final chapter that outlines the outcomes of this research study relative to real-world applications and potential future research endeavors.

Chapter 6

6. Conclusions and Future Work

6.1 Conclusions

The objective of this study was to address an apparent methodological and pragmatic gap in stock market trading signal prediction (which may or may not, be relevant to the Tadawul). Thus far, most quantitative and artificial intelligence based endeavors have either attempted to position technical analysis with fundamental analysis as separate, and sometimes, opposing forces or, recently, models observed for literature review had more predictability for generalized observations than economically interpretable, compliant and risk-aware decision-making.

As such, contributions were made to both the model construction and the rationalization of a new approach to trading signal prediction. The rationale for the latter came from a stacking ensemble architecture which hybridized a Long Short-Term Memory (LSTM) network with an Extreme Gradient Boosting (XGBoost) model. The LSTM was constructed based upon high-frequency technical indicators facilitating momentum and timeliness of market value whereas the XGBoost acknowledged low-frequency financial ratios which, ultimately, depict firm intrinsic value, or the long-term, stable approach of a firm.

Further, the model output of both base learners was integrated for binary options as facilitated by a logistic regression based meta-learner. The novel target for prediction derived practical trading sensibility which calls for a profit greater than a stop loss or not triggering a stop loss during the fixed holding period based upon reconceptualizing prediction success in relation to better than average trading assumptions instead of generalized prediction success.

As the empirical experiment was composed of realistically simulated results, the findings were deduced from internal reality. Thus, from this perspective, risk-adjusted performance trumped any potential performance benchmarks of the XGBoost-stood alone fundamental prediction model while the hybrid model performed superior to both models for Sharpe ratio, maximum drawdown and winning average.

Additional performance was linked to generalizability of the hybrid framework realized through machine learning model creation and exponential learning external explanation via feature importance rankings and economically relevant indicators as prediction from purely publicly available investor-relevant data is surprisingly well-framed for intelligently salient trading systems.

This is an especially critical finding for the Saudi market as Capital Market Authority (CMA) compliance is required not only for any predictive model but advisory ones as well. Thus, studies that depend upon unstructured or alternative resources may not sufficiently guide the way regional investments play out without issues of ethics or compliance—and thus, an un-Islamic financial situation.

Therefore, these research findings support the previously suggested position that intelligent trading systems should no longer seek increasingly isolated predictable accuracy or consistently utilized econometric resources, but rather risk-aware decision-making support that makes sense for trading practice instead of in academic ideation.

Thus, this hybrid framework should be considered a valid suggestion for both academic exploration and real-world application anywhere, but especially in emerging financial markets for future exploration.

6.2 Contributions Revisited

The contributions of this study are noted as follows:

1. A hybrid framework compliant for regulation with feasible, publicly available trading signal prediction in the Saudi Stock Market.
 2. A hybridization of fundamental valuation signals and technical timing indicators in a comprehensive stacking ensemble architecture.
 3. Financial performance parameters over statistical predictive accuracy that mimic logical trading intentions and expectations.
 4. More explainable results based upon feature importance and economically viable indicators.
 5. A scalable and extendable framework with potential application across firms, industries and time periods.
-

6.3 Study Limitations

Despite this study being a meaningful contribution to the field, some limitations are worth noting for future research. First, the experimental study is based on realistic simulated backtesting results, meaning live trading wasn't executed. Therefore, phenomena of microstructural impact, such as slippage, order fulfilment timeliness, and attention to fluctuating liquidity needs, might not be as pronounced in the simulations as they would be in real-life trading. Second, the empirical study relies on larger, capitalized and liquid companies traded on the Saudi market; thus, smaller-cap or illiquid stocks on the fringes of the market might not fare as well in the model. Third, transaction costs are noted to be fixed and low; however, these differ based on transaction types and market involvement.

Finally, the interpretability of the proposed framework is more transparent than a black box typical of deep learning; however, LSTM involved still has some complexities and instances that aren't fully transparently interpreted.

Note that all experimental findings herein refer to realistic simulated backtesting results—not live or real-money trading results.

6.4 Directions for Future Research

Whether considering future research or improvements relative to the model applied herein or subsequent findings of this study, the following frameworks make sense:

- An extended time frame model for short, mid and long-term applications of the framework.
 - A reinforcement learning model that permits entry and exit policy optimization based on fluctuating markets.
 - Sector or index-based efforts that could increase portfolio attractiveness.
 - The inclusion of other risk measures like VaR and CVaR.
 - Live trading/paper trading studies that can assess real-world impact based on timestamps.
 - The use of adaptive ensemble weighing that's robust to adjust optimal features to avoid regime uncertainty and extreme dislocations in the market.
-

6.5 Conclusion

The longer artificial intelligence rules the roost relative to financial markets, the more tantalizing investing opportunities become—but the more responsibility that's needed. This research shows how an intelligent investment choice can be just as impactful and accurate with realities behind the economy and transparency, coupled with regulatory equity and ethical sensitivity, to make it desirable. By considering both fundamental and technical sensibilities from an inclusive quantitative approach, this model is a step in the right direction for an efficient and ethically responsible option for more predictable investment decision-making tools.

Ultimately, this hybrid stacking ensemble approach is a viable, academically substantiated method for intelligent investment signal prediction relative to the Saudi Stock Market as all findings from experiments reflect real-world data. Thus, they contribute to opportunities for future research and implementation of more educated, responsible decision-making efforts.

References

- Abu-Mostafa, Y. S., Magdon-Ismael, M., & Lin, H. T. (2012). *Learning from data*. AMLBook.
- Adebisi, A. A., Adewumi, A. O., & Ayo, C. K. (2014). Stock price prediction using the ARIMA model. In *Proceedings of the UKSim-AMSS 16th International Conference on Computer Modelling and Simulation* (pp. 106–112). IEEE.
- Ahmed, N. K., Atiya, A. F., Gayar, N. E., & El-Shishiny, H. (2010). An empirical comparison of machine learning models for time series forecasting. *Econometric Reviews*, 29(5–6), 594–621.
<https://doi.org/10.1080/07474938.2010.481556>
- Bao, W., Yue, J., & Rao, Y. (2017). A deep learning framework for financial time series using stacked autoencoders and LSTM. *PLOS ONE*, 12(7), e0180944.
<https://doi.org/10.1371/journal.pone.0180944>
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32.
<https://doi.org/10.1023/A:1010933404324>
- Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794). ACM.
<https://doi.org/10.1145/2939672.2939785>
- Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2), 383–417.
<https://doi.org/10.2307/2325486>
- 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.
<https://doi.org/10.1016/j.ejor.2017.11.054>
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.
<https://doi.org/10.1162/neco.1997.9.8.1735>
- Hosmer, D. W., Lemeshow, S., & Sturdivant, R. X. (2013). *Applied logistic regression* (3rd ed.). Wiley.
- Krauss, C., Do, X. A., & Huck, N. (2017). Deep neural networks, gradient-boosted trees, random forests: Statistical arbitrage on the S&P 500. *European Journal of Operational Research*, 259(2), 689–702.
<https://doi.org/10.1016/j.ejor.2016.10.031>

- Lo, A. W., Mamaysky, H., & Wang, J. (2000).
Foundations of technical analysis: Computational algorithms, statistical inference, and empirical implementation. *The Journal of Finance*, 55(4), 1705–1765.
<https://doi.org/10.1111/0022-1082.00265>
- Malkiel, B. G. (2003).
The efficient market hypothesis and its critics. *Journal of Economic Perspectives*, 17(1), 59–82.
<https://doi.org/10.1257/089533003321164958>
- Saudi Capital Market Authority. (2023).
Capital market regulations and guidelines. CMA Publications.
<https://cma.org.sa>
- Tsai, C. F., & Wang, S. P. (2009).
Stock price forecasting by hybrid machine learning techniques. In *Proceedings of the International MultiConference of Engineers and Computer Scientists* (pp. 755–760).
- Wolpert, D. H. (1992).
Stacked generalization. *Neural Networks*, 5(2), 241–259.
[https://doi.org/10.1016/S0893-6080\(05\)80023-1](https://doi.org/10.1016/S0893-6080(05)80023-1)
- Zhang, G., Patuwo, B. E., & Hu, M. Y. (1998).
Forecasting with artificial neural networks: The state of the art. *International Journal of Forecasting*, 14(1), 35–62.
[https://doi.org/10.1016/S0169-2070\(97\)00044-7](https://doi.org/10.1016/S0169-2070(97)00044-7)