Informative Inputs Over Complex Features: Long Short-Term Memory Forecasting in the Saudi Stock Market

Munira AlBalla*, Arwa Alawajy, Seetah Alsalamah, Zahida Almuallem
College of Computer and Information Sciences-Department of Computer Science,
King Saud University, Riyadh 12371, Saudi Arabia

***Abstract*—This study investigates the effectiveness of deep learning, specifically Long Short-Term Memory (LSTM) networks, for forecasting stock closing prices in the Saudi Arabian market. Unlike prior research that focuses on narrow stock subsets or individual technical indicators, we present the first comprehensive evaluation across the top thirty companies by market capitalization on the Tadawul Exchange. We compare eight LSTM variants trained on different combinations of feature families, including technical indicators, calendar effects, and macroeconomic variables such as Brent oil prices and the Tadawul All Share Index (TASI). Despite the popularity of complex feature engineering, our results show that models using simpler macroeconomic inputs consistently outperform those based on technical indicators. In particular, combining the TASI with closing prices yielded the best results for 44.8% of stocks, improving median accuracy by 6.19% over the closing price-only baseline. Conversely, models incorporating extensive technical indicators or applying advanced feature selection techniques underperformed the baseline by 8 to 12%. These findings challenge the assumption that greater complexity leads to better performance in financial forecasting and highlight the value of focused, economically interpretable features in the Saudi market context.**

Keywords*—Deep learning; feature engineering; stock price forecasting; macroeconomic indicators; Brent crude oil; Saudi stock market; Long Short-Term Memory; Tadawul; Tadawul All Share Index (TASI)

I. Introduction

The financial market exhibits highly dynamic and complex behavior, shaped by multiple interrelated factors including the global economy, political events, and investor sentiment [1], [2], [3], [4]. This inherent volatility and non-linearity make it challenging for investors, risk managers, and automated trading systems to make well-informed decisions regarding asset management and profit maximization. Therefore, accurate stock price forecasting is still a complex problem, mainly due to the presence of noisy, non-stationary, and often incomplete financial data [5], [6], [7]. These challenges have motivated researchers to explore advanced predictive models capable of capturing hidden temporal dependencies and improving forecasting performance in such uncertain environments.

Various stock market prediction techniques have been developed to forecast future trends and returns. Traditional statistical methods, such as AutoRegressive Moving Average (ARMA) and Generalized AutoRegressive Conditional Heteroskedasticity (GARCH), suffer from key limitations when dealing with non-stationary financial data. These models rely on strong statistical assumptions such as linearity and noise-free data, which restrict their ability to capture the complex, nonlinear, and dynamic nature of financial markets. Furthermore, they often operate on univariate models and struggle to incorporate external factors such as macroeconomic variables that influence market behavior [8], [9], [6].

To address some of these limitations, Machine Learning (ML) models such as Support Vector Machines (SVM) and Random Forests (RF) have been introduced. These models leverage the availability of large, high-dimensional datasets and are capable of discovering hidden patterns and complex relationships in data. While ML-based models offer greater flexibility and improved predictive power over traditional statistical approaches, they still fall short in modelling temporal dependencies, which is an intrinsic characteristic of financial time series.

These challenges have led researchers to explore deep learning-based methods for stock market prediction. In particular, Recurrent Neural Networks (RNNs) and their advanced variant, LSTM networks, have gained popularity due to their ability to handle sequential data and model long-term dependencies. Standard RNNs, however, face challenges such as the vanishing gradient problem, which limit their effectiveness on long sequences [10], [11], [12], [6]. LSTM networks address this issue through specialized memory cells and gating mechanisms, enabling them to learn temporal patterns more effectively. As a result, LSTMs have been widely applied in financial forecasting, as shown by prior empirical studies that they are particularly well-suited for capturing the temporal dynamics and complex structure of financial markets [13], [14], [15], [17].

The main contributions of this study can be summarized as follows: First, this study addresses a notable gap in the literature by providing a comprehensive analysis of stock price forecasting across the top 30 market capitalization companies on the Saudi Tadawul Exchange from 2006 to 2025. Unlike much of the existing research that concentrates on individual stocks or narrow sectors, our work offers crucial, broad market coverage. This breadth allows us to identify patterns in how different input families perform that are stable across many large-cap stocks, rather than being tied to a single case study or sector representative. Second, we demonstrate that a simple

*Corresponding author.

univariate model relying solely on the lagged closing price serves as a surprisingly powerful baseline. This minimalistic approach provides effective forecasts and establishes a benchmark that more complex, feature engineered and multivariate models must clearly surpass to justify their added complexity, and it serves as a reference point for judging whether additional feature families genuinely add predictive value. Importantly, this evaluation is grounded in multiple performance metrics including a Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), Mean Absolute-Percentage Error (MAPE), and win rate, ensuring that the assessment of predictive accuracy is both rigorous and comprehensive. Third, we present a systematic comparison of different feature-engineering approaches by evaluating the predictive contribution of lag features including technical indicators, macroeconomic variables, and calendar effects through eight distinct experimental scenarios, under a shared LSTM architecture. Most significantly, our results demonstrate that incorporating the TASI as a feature substantially outperforms complex technical indicator approaches across all four evaluation metrics, challenging the conventional belief that more sophisticated features necessarily improve prediction accuracy. To our knowledge, this is the first study to systematically demonstrate the predictive value of the TASI in stock price forecasting and to turn that evidence into a simple feature selection rationale that prioritizes market-wide macro variables over large technical indicator sets.

The remainder of this study is organized as follows: Section II presents a review of recent literature pertaining to stock analysis and prediction. Section III presents the theoretical concepts underpinning the research. A comprehensive discussion of the experimental design, the data selection process, and the prediction methodology will be provided in Section IV. Section V presents the research results in detail. This is followed by Section VI, which provides an in-depth discussion of the key findings, interpreting their significance in the context of research objectives. Finally, Section VII summarizes the main insights and contributions of the study, as well as potential directions for future work.

II. Literature Review

LSTM networks have been extensively explored for stock-price forecasting, ranging from simple architectures to sophisticated hybrids. In [16], the authors conducted a study which demonstrated that a single-layer LSTM with nine input features including, macroeconomic indicators, like unemployment rate, technical signals like *Moving Average Convergence Divergence (MACD)*, and historical trading daily data can effectively predict next-day Standard & Poor's (S&P) 500 closing prices over a 2006–2020 large dataset with a RMSE of 12.45 and correlation above 0.92, highlighting the trade-off between depth and overfitting. In [17], the authors extended this to a two-layer LSTM to understand the patterns over time, followed by another two dense layers trained on only 60-day sliding windows of normalized daily closing prices and trading volumes gathered from Yahoo Finance for Apple, Google, Microsoft, and Amazon from 2012 to 2022. Their model achieved an RMSE of 18.89, but the narrow training set and the omission of potentially significant variables, such as news sentiment and macroeconomic indicators, may limit generalizability.

Similarly, in [5], the authors applied an LSTM model to predict closing prices of 31 Vietnamese stocks, incorporating historical price data and three technical indicators. The model achieved forecast accuracy above 93% for most of the stocks included in the study, highlighting the effectiveness of LSTM in time series prediction. However, the authors note that the high volatility of prices within Vietnam's emerging stock market could affect its predictive robustness in different contexts. In [15], the authors proposed a multi-value associated network model of LSTM to predict multiple prices of a stock simultaneously. The model predicts open, low, and high prices on the Shanghai Stock Exchange (SSE), reporting overall accuracy above 95% when compared to traditional LSTM and deep RNN models. However, this approach limited the model's ability to optimize each price prediction effectively, as the training process relied on a total loss function that did not account for the individual sub-losses associated with the three output branches (open, low, and high).

A study in [18], has pushed LSTM into hybrid and decision-making roles. They proposed an automated stock trading system using deep reinforcement learning with LSTM, trained on data from 120 stocks across four indices. LSTM extracts patterns from daily prices and technical indicators to predict when to buy, sell, or hold stocks to achieve higher returns. However, the model's effectiveness relies on a vast amount of training data to fully learn market characteristics. Similarly, in [19], the authors embedded an LSTM within a decision-support system, but for the Indian stock market. They used normalized *Open-High-Low-Close-Volume (OHLCV)* data along with technical indicators to predict the closing stock price for the next thirty days. Using an optimised model, the model yielded an RMSE of 4.13 and a MAE of 3.24 compared to prominent statistical and machine learning models. However, the high computational cost associated with processing historical stock data is the main limitation of the model. With similar limitation of heavy computation and strict data access requirements, reference[20] further enriched Bidirectional Long Short-Term Memory (BiLSTM) with features extracted by a Convolutional Neural Network (CNN) to predict the direction for the next day (up or down) and a trend prediction for several days ahead (up, down, no). In addition to typical market data such as stock prices and trading volume of 14 stocks of a Chinese investment bank, the model incorporates information about the relationships among various stocks based on a knowledge graph and includes detailed records of each investor's trading activities, organized into simplified clusters. The model achieved 74% accuracy in predicting the direction of the next day's closing price and 66% for 5-day trend prediction, which are higher than the accuracies obtained from Adaboost and RF models mentioned in their work. In a similar manner, the authors in [21], integrated investor sentiment data gathered from online stock forums with a set of 16 different technical indicators of thirty stocks in the SSE market, then fed that to an LSTM to predict the closing price one day ahead. The model achieved a MAPE of 0.05, which is superior to the other prediction models used in their work, such as CNN, SVM, and also to their proposed model without including sentiment analysis. However, the model is constrained by a single market-specific sentiment source.

To enhance the accuracy of stock market predictions, a recent study in [22] optimized LSTM hyperparameters via a

metaheuristic, Artificial Rabbits Optimization (ARO), on five years of Dow Jones Industrial Average (DJIA) data with a time window sized 20 previous days. The model demonstrates the lowest Mean squared error (MSE) and MAE for most of the ticker symbols of the market, compared to the other four LSTM models. Still, this improvement in predictive accuracy came at the cost of increased computational overhead during the search phase, as the metaheuristic nature of the ARO algorithm requires extensive exploration of the search space. In [23], the authors showed that preprocessing data using Complementary Ensemble Empirical Mode Decomposition (CEEMD) for decomposing price series to remove noise, Principal Component Analysis (PCA) for reducing data dimensions, and LSTM for predicting closing price of the next trading day for various equity indices reduces RMSE by up to 7% over models that use raw prices series. The comprehensive survey in [24] reviewed the developments in machine learning techniques, including LSTM and CNN, and highlights the advantages of hybrid models and the integration of sentiment analysis to enhance forecasting capabilities. The authors also acknowledged several challenges associated with model building, such as the need for large datasets containing non-stationary data and the complexity of training advanced models.

The Saudi stock market (Tadawul) closely aligns with global markets in its suitability for applying deep learning, particularly LSTM networks for stock price prediction, offering a valuable context for advancing internationally established models. In [25], the authors developed a binary classification model to predict next-day stock price direction (up or down), utilizing a comprehensive set of 31 features, including *OHLCV* data from four companies over ten years, supplemented by various technical indicators. Among the models evaluated, the RF algorithm demonstrated the highest predictive accuracy at 73.5%, outperforming both LSTM and Artificial Neural Network (ANN). The predictive model was subsequently integrated into an automated investment system incorporating risk management strategies, which yielded substantial returns ranging from 65% to 86% across ten companies listed on the Tadawul; however, the study was limited by its focus on a mostly bullish market condition. In [26], the authors found that multivariate LSTM models utilizing *Open-High-Low-Close (OHLC)* inputs outperformed univariate models, achieving a prediction accuracy of 97.49% when forecasting closing prices over the following seven days. Their study was based on a decade of historical data up to November 2022 for three major communication firms (Mobily, Saudi STC, and Zain KSA). Still, the limited scope of input features that exclude other influential factors, such as sentiment and macroeconomic indicators, restricts the generalizability. It limits the model's effectiveness in capturing broader market dynamics. Extending this, in [12], the authors recently compared the performance of the statistical analysis AutoRegressive Integrated Moving Average (ARIMA) and LSTM algorithms for stock price forecasting, following the application of Discrete Wavelet Transform (DWT) to denoise historical closing price data from four companies listed on Tadawul, and showed that the proposed LSTM model achieved a prediction accuracy of 97.54%, outperforming the ARIMA model in terms of forecasting performance.

Similarly, in [27], the authors explored LSTM and a multiLayer perceptron (MLP) models across four key sectors of the Saudi market using data spanning three years from 2018 to 2020, finding that while the LSTM model achieved higher prediction accuracy compared to MLP, the generalizability of both models is limited as they do not integrate broader feature sets. In a more applied setting, in [10], the authors incorporated oil closing prices into an LSTM model to predict next-day trading signals (buy or sell) for three stocks in the Saudi market, training on historical stock data including closing prices, trading volume, technical indicators, and oil price data from June 2018 to August 2019. Similar to [18], they fed the predicted signals to a trading agent to decide whether to buy, sell, or hold based on current share value, available balance, and shareholdings. The proposed LSTM model generated next-day trading signals with 55% accuracy, resulting in a 23% investment return that outperformed the traditional buy-and-hold strategy, which achieved 1.2%. However, its simplicity of feature set and absence of a robust feature selection mechanism, which relies on trial-and-error, remain key limitations. Table I and Table II provide a structured comparison of the literature, organized by Saudi and international markets for better readability. Collectively, these studies illustrate LSTM's flexibility for capturing temporal dependencies in financial series, while also revealing recurring challenges in model complexity, evaluation rigour, dataset diversity, and real-world applicability.

In the Saudi market, the evidence on input design remains fragmented. Some studies rely on closing prices alone and compare LSTM with traditional or shallow models for a small set of stocks or sector indices [12], [27]. While others use multivariate OHLC inputs and report gains over univariate specifications but only for three large communication stocks [26]. Further work combines OHLCV data with long lists of technical indicators [25], or mixes oil prices with technical indicators for three stocks [10], yet none of these studies justify the incremental value of technical indicators over a close-only LSTM baseline or tests whether their findings generalize across the main Tadawul listings. As a result, we are left with scattered evidence on close-only, OHLC, technical indicators, and oil-augmented designs. The present study combines these strands within a single, unified LSTM framework by evaluating close-only, OHLCV, technical, calendar, and macroeconomic configurations across the top market-cap Saudi stocks, providing a cross-sectional view of when close alone suffices and when richer inputs add value.

III. Theoretical Basis

This section outlines the key concepts underlying our approach to predicting stock prices in the Saudi market. First, we look at the Efficient Market Hypothesis (EMH), which suggests that stock prices reflect all available information and are hard to predict. Then, we introduce LSTM networks, a type of machine learning model that can learn patterns in time-based data, such as stock prices. Finally, we discuss how investor behaviour and psychological factors can create patterns that LSTM models may be able to detect. These ideas provide the foundation for our research and help explain why combining finance theory with machine learning can be helpful.

TABLE I. RELATED WORK ON INTERNATIONAL MARKETS PREDICTION USING LSTM ORDERED BY PUBLICATION YEAR, MOST RECENT FIRST, 2024–2020

Reference	Data	Period	Input	Technique	Output	Performance Measure
[5]	31 Vietnamese stocks	Start listed- Apr2021	*OHLCV*, three technical indicators	LSTM	1-day-ahead closing price	Accuracy
[18]	Dow Jones, Chinese, India, UK Stock Markets	Jan2009-Aug2020	Closing price, share, balance, Four technical indicators	Cascaded LSTM	128 temporal features	Cumulative Return, Max Earning Rate, Sharpe Ratio
[22]	Dow Jones	Jan2018-Jan2023	Closing price	LSTM ARO	1-day-ahead closing price	MAPE, MAE, MSE, R^2
[17]	Four US tech stocks	2012-2022	Closing price	LSTM	1-day-ahead closing price	RMSE
[16]	S&P 500	2006-2020	OC, three technical indicators, five macroeconomic factors	LSTM	1-day-ahead closing price	RMSE, MAPE
[19]	India Stock Market	Jan2020-July2020	*OHLCV*, closing price of industry average	LSTM	*OHLCV*, closing price of industry average	RMSE, MAE, MAPE
[21]	Chinese Stock Market	Jan2017-July2019	16 technical indicators, sentiment factor	LSTM	1-day-ahead closing price	MAPE
[15]	Chinese Stock Market	-	*OHLCV* and two technical parameters	Associated Net, LSTM	Next-day OLH	MAE
[20]	Chinese Stock Market	Mar2012-June2018	29 basic stock data & stock correlation, investor clusters	BiLSTM	Next-day direction, trend over *n* days	Accuracy, AUC
[23]	Hong Kong, Dow Jones, S&P 500, Chinese Stock Markets	Jan 2010-Apr2018	Four abstract features out of 12	CEEMD PCA LSTM	1-day-ahead closing price, direction of change	RMSE, MAE, NMSE, DS

TABLE II. RELATED WORK ON THE SAUDI MARKETS PREDICTION USING LSTM ORDERED BY PUBLICATION YEAR, MOST RECENT FIRST, 2024–2020

Reference	Data	Period	Input	Technique	Output	Performance Measure
[12]	Four Saudi Stocks: Saudi Arabian Mining Co, Yanbu Cement Co., Sabic, Saudi Indian	Jan2019-July 2023	Closing price	LSTM	7-day-ahead closing price	RMSE, MSE
[26]	Three Saudi Stocks: STC, Mobily, Zain	Jan2017-Nov2022	*OHLC*	LSTM	7-day-ahead closing price	RMSE, MAE, MAPE, MSE
[27]	Four Saudi Sectors: Energy, Industrial, Financial, Communication	2018-2020	Closing price	LSTM	1-day-ahead closing price	RMSE, MSE, R^2
[25]	Four Saudi Stocks: Sabic, Jarir, Alrajhi, STC	Jan2010-Aug2021	*OHLCV*, 26 technical indicators	LSTM	Next-day price direction	Accuracy, Precision, Recall
[10]	Three Saudi Stocks: Alinma, Alrajhi, SABIC	Mar2012-April2020	*OHLCV*, Oil price, month number, five technical indicators	LSTM	Next-day buy/sell signals	Accuracy

A. Market Efficiency versus Predictability

The EMH, first proposed by the authors in [28], suggests that financial markets are efficient, meaning that asset prices quickly and fully reflect all available information. If this is true, it becomes challenging to consistently make profits by predicting price movements or employing trading strategies. However, many studies have shown that fundamental markets do not always behave this way. Factors such as slow information flow, investor psychology, market structure, and regulation can cause inefficiencies [29], [30]. These inefficiencies may allow room for prediction. In such cases, machine learning methods, especially those that can handle patterns over time, can provide a valuable alternative to traditional financial models by challenging some of the assumptions of EMH, particularly.

B. LSTM Mechanics in Financial Time Series Prediction

LSTM networks, a type of RNN, have been shown to work very well with sequential data [31], including financial time series forecasting [32], [33].

Traditional RNNs often struggle with problems such as vanishing gradients. They cannot handle long-term patterns, but LSTMs overcome this by using a gated memory structure that stores and updates information over time [34], [35]. This structure consists of three main parts: a forget gate that removes unnecessary information, an input gate that adds new relevant data, and an output gate that determines which information to pass on. These mechanisms enable LSTMs to capture complex time-based patterns, making them highly suitable for forecasting stock prices [36]. In Tadawul, stock movements are influenced by various factors, including changes in oil prices, political events, and seasonal effects such as those observed during Ramadan [10], [37]. LSTMs can learn these nonlinear patterns without heavy manual feature design. While the EMH suggests that prices cannot be predicted from past information, studies in behavioural finance show that inefficiencies do exist [29], [38], [39], especially

in emerging markets [30], [37]. Therefore, combining LSTMs with technical and macroeconomic indicators can help uncover proper predictive signals in Tadawul's market environment [40].

C. Behavioral Tadawul Market Efficiency

While the EMH assumes that investors act rationally and that markets always reflect all available information, behavioral finance shows that emotions and mental shortcuts often lead to mistakes in judgment. These behaviors can create patterns in the market that the EMH cannot explain [29], [38]. In the Tadawul, several examples of such behavior affect how prices move. One typical pattern is *herding*, where individual investors follow the crowd instead of making their own decisions. This often causes trends to grow stronger, even when they are not supported by fundamentals [39]. Another example is *overreacting to news*, such as sudden changes in oil prices or updates related to major plans like Vision 2030 [41], [42], [43]. These reactions can lead to bigger price changes than necessary [30]. Additionally, some *seasonal effects* have been observed, particularly during Ramadan, where studies such as [37] have found unusually high returns, likely due to shifts in mood and trading habits. These behaviours create inefficiencies that are hard to capture with traditional models. That's why advanced machine learning techniques, such as LSTM networks, are helpful. They can detect complex and hidden patterns in market data, making them well-suited for forecasting in behavior-driven markets like Tadawul.

IV. Data and Methodology

Fig. 1 summarizes the end-to-end workflow, executed independently for each stock in our top-30 Tadawul universe (29 after cleaning). Details of every step appear in the subsections below:

A. Dataset and Market Universe

This study utilizes a comprehensive dataset comprising historical financial and macroeconomic data relevant to the Saudi Arabian equity market. Historical daily OHLCV data were collected from Yahoo Finance using ticker symbols with the ".SR" suffix, which marks shares listed on the Saudi Exchange (Tadawul). The primary focus is on the thirty largest publicly traded companies on the Tadawul Exchange, selected based on free float market capitalization to represent the most liquid and actively traded equities. The data collection period spans from each company's first available trading day to May 8, 2025. Following data cleaning procedures, 29 companies remained for analysis. Table III lists the top 30 stocks listed on Tadawul based on free float market capitalization as of 2025. Fig. 2 illustrates the proportional allocation of stocks by sector.

B. Inputs and Feature Set

To capture the factors influencing Saudi stock price movements, we use **37 predictive variables**. The predictors are grouped by their source and role into four families: *Raw OHLCV*, *Technical indicators*, *Calendar features*, and *Macroeconomic features*, as shown in Table IV.

1) Raw OHLCV: Our feature set begins with the raw daily *Open*, *High*, *Low*, *Close*, and *Volume* series downloaded from Yahoo Finance (".SR" tickers). These variables capture each session's price range and trading activity. These stock-level series serve as the basis from which all technical indicators are computed.

2) Technical indicators: All technical indicators are computed directly from each stock's OHLCV.

a) Moving averages: SMA and EMA over rolling windows of 5, 10, 20, 50, 100, and 200 days to capture multi-horizon trend dynamics. Short windows (5–20 days) reflect immediate moves; long windows (50–200 days) reflect underlying trends.

b) Momentum indicators: MACD (main MACD line, signal line MACDs, and histogram MACDh) to detect momentum shifts. RSI to identify overbought/oversold conditions and potential turning points.

c) Volatility: We record the lower Bollinger Band (BBL_20_2.0), middle Bollinger Band (BBM_20_2.0), and upper Bollinger Band (BBU_20_2.0). From these, we compute Bollinger Bandwidth (BBB_20_2.0), which measures volatility through the distance between the bands, and Bollinger Band Percentile (BBP_20_2.0), which shows where the current price falls within the envelope.

d) Volume-based On-Balance Volume (OBV): which measures sustained buying or selling pressure.

3) Calendar features: To capture temporal patterns and specific events that may influence trading behaviour, we incorporated several calendar-based features. Day of the week and day of the year are represented through sinusoidal functions (*dow_sin*, *dow_cos*, *doy_sin*, *doy_cos*) to preserve temporal continuity and capture any recurring seasonal patterns in market behavior. Additionally, we created binary indicators for specific calendar events known to affect Saudi market behavior. The *Ramadan* indicator identifies trading days during the holy month, which has been documented to influence market volatility and returns in Islamic markets. The quarter-end proximity flag (*qtr_end_5d*) marks the five days preceding quarterly financial reporting periods, typically associated with increased trading activity and volatility. All variables in this group are computed from the trading date only and do not use price or volume information.

4) Macroeconomic features: Given Saudi Arabia's position as a major oil exporter and the importance of global economic conditions, we incorporated key macroeconomic indicators. *Brent* crude oil prices serve as a critical external factor, supplemented by its daily return (*Brent_ret*) and 60-day rolling volatility (*Brent_vol60*) to capture both level and risk dimensions of oil market dynamics. TASI represents overall market sentiment and provides context for individual stock performance within the broader Saudi equity market. These inputs are external series merged to the stock-level table by date and are not derived from the stock's OHLCV.

C. Data Preprocessing

1) Data merging: We align and join, by date, each stock's OHLCV with the engineered technical and calendar features

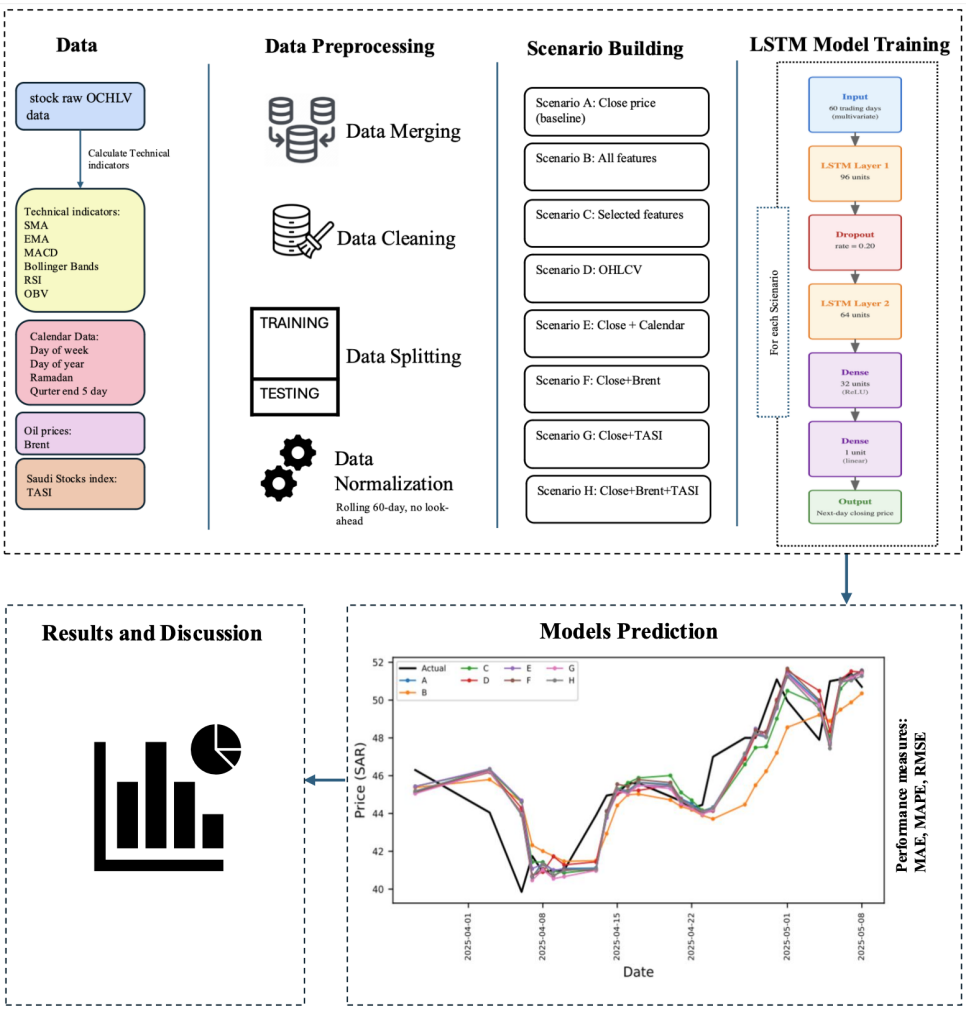

Fig. 1. Methodology overview: Daily OHLCV per stock is combined with engineered technical features, calendar features, and external macro features (TASI, Brent). After data merging, cleaning, a fixed train/test split, we build eight input scenarios (A–H) and train the same LSTM per stock and per scenario. This pipeline is executed for each stock among the top market-cap Tadawul companies (29 after cleaning). The bottom band illustrates test predictions, performance metric computation, and cross-stock/sector comparison.

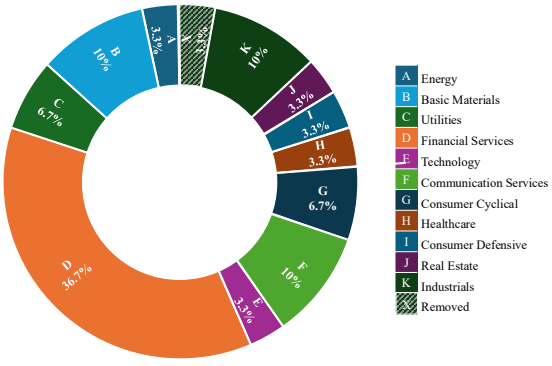

Fig. 2. Percentage breakdown of the top 30 Tadawul stocks by sector.

and the external macro series (TASI, Brent). The merge is restricted to common trading days, with one record per stock per date after alignment.

2) Data cleaning: The data cleaning process involved removing rows containing NaN values resulting from rolling-window calculations in technical indicators, where the first several observations could not be computed due to insufficient lookback data. Any residual missing values after alignment were dropped. Stock 4325 (Umm Al Qura Development & Construction Co.) had no data remaining after cleaning, yielding a final sample of 29 stocks.

3) Data splitting: The dataset was divided into training and testing periods to ensure robust out-of-sample evaluation. Training data spans from the earliest available date (after data cleaning) to December 31, 2024, while the testing period covers January 1, 2025, to May 8, 2025.

4) Data normalization: All numeric features are standardized using a 60-day rolling window, computed strictly from prior observations at time t. The same transformation is applied during validation and testing, preserving temporal causality and avoiding look-ahead bias.

TABLE III. TOP 30 STOCK COMPANIES IN TADAWUL BY SECTOR AND TICKER

No.	Ticker	Company	Sector
1	2222	Saudi Arabian Oil Co. (Saudi Aramco)	Energy
2	1120	Al Rajhi Bank	Financial Services
3	2082	ACWA Power Co.**	Utilities
4	7010	Saudi Telecom Co. (STC)	Communication Services
5	1180	The Saudi National Bank	Financial Services
6	2010	Saudi Basic Industries Corp. (SABIC)	Basic Materials
7	1211	Saudi Arabian Mining Co. (Ma'aden)	Basic Materials
8	4013	Dr. Sulaiman Al Habib Medical Svcs. Grp.	Healthcare
9	1010	Riyad Bank	Financial Services
10	7203	Elm Co.**	Technology
11	1150	Alinma Bank	Financial Services
12	1060	Saudi Awwal Bank	Financial Services
13	5110	Saudi Electricity Co.	Utilities
14	2280	Almarai Co.	Consumer Defensive
15	2020	SABIC Agri-Nutrients Co.	Basic Materials
16	7020	Etihad Etisalat Co. (Mobily)	Communication Services
17	1080	Arab National Bank	Financial Services
18	1050	Banque Saudi Fransi	Financial Services
19	1140	Bank Albilad	Financial Services
20	7202	Arabian Internet & Comm. Svcs. Co.**	Communication Services
21	4280	Kingdom Holding Co.	Industrials
22	4325	Umm Al Qura Dev. & Construction Co.*	Industrials
23	4250	Jabal Omar Development Co.	Consumer Cyclical
24	8210	Bupa Arabia for Cooperative Insurance Co.	Financial Services
25	4300	Dar Al Arkan Real Estate Dev. Co.	Real Estate
26	1111	Saudi Tadawul Group Holding Co.**	Financial Services
27	4030	National Shipping Co. of Saudi Arabia (Bahri)	Industrials
28	4100	Makkah Construction & Development Co.	Consumer Cyclical
29	8010	The Co. for Cooperative Insurance (Tawuniya)	Financial Services
30	4142	Riyadh Cables Group Co.**	Industrials

*Umm Al Qura Dev. & Construction Co. (4325) was subsequently excluded during data cleaning due to no data left after cleaning.

**Recently listed companies such as ACWA Power (2082), Elm Company (7203), Riyadh Cables Group (4142), Saudi Tadawul Group Holdings (1111), and Arabian Internet & Communications Services (7202) have limited trading histories ($\leq$650 days).

D. Model Architecture and Evaluation Procedure

To isolate the effect of inputs, all scenarios share a fixed network and training setup; the full specification appears in Table V. Briefly, the model ingests a multivariate 60-day window, stacks two LSTM layers (96 units with sequence return and 0.20 dropout; then 64 units), applies a 32-unit ReLU dense layer, and outputs the next-day closing price with a linear head. Models are trained with mean-squared error using Adam (learning rate $\eta = 0.001$), mini-batches of 32 for up to 80 epochs with early stopping after 6 consecutive validation epochs without improvement.

To evaluate predictive accuracy, we use three standard metrics: MAE, RMSE, and MAPE. MAE is defined by Eq. (1), and RMSE and MAPE by Eq. (2) and Eq. (3).

$$MAE = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i| \tag{1}$$

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \tag{2}$$

$$MAPE = \frac{1}{n}\sum_{i=1}^{n}\left|\frac{y_i - \hat{y}_i}{y_i}\right| \times 100\% \tag{3}$$

where, n is the number of observations in the test set, y_i is the actual closing price for observation i, and $\hat{y}_i$ is the corresponding predicted closing price.

E. Scenario Descriptions

To assess the marginal contribution of different feature groups to forecasting performance, we define eight LSTM-based experimental scenarios:

1) Scenario A (Close-only baseline): Uses only the Close price as input. This baseline tests whether price memory alone is sufficient for short-horizon forecasting.

2) Scenario B (Full feature set): Uses all 37 features, including prices, technical indicators, calendar patterns, and macroeconomic data. This is the most complete model, assuming access to all possible signals at the end of the day.

3) Scenario C (Selected features): Uses only a carefully chosen subset of features based on importance scores (correlation and SHAP). This tests if fewer, more relevant signals can perform as well as the complete set.

4) Scenario D (OHLCV): Open, High, Low, Close prices, and Volume only, focusing on raw trading activity without engineered indicators or external inputs.

5) Scenario E (Close + calendar): Combines Close with calendar features (day-of-week and day-of-year sine/cosine, Ramadan, qtr_end_5d) to test seasonal and temporal effects.

6) Scenario F (Close + Brent): Uses Close with Brent oil inputs. Since oil plays a key role in Saudi Arabia's economy, this model tests whether oil price movements help improve the forecasting of stock prices.

7) Scenario G (Close + TASI): Uses Close and TASI. TASI reflects the overall market mood and helps capture broader trends that may influence individual stocks.

8) Scenario H (Close + Brent + TASI): Combines Close with both Brent and TASI to test whether combining global and local market signals improves predictions.

Each scenario is designed to test the model's ability to predict stock prices and to understand the significance of specific inputs. By comparing them, we can determine which types of information enhance predictions in the Saudi stock market. Table VI summarizes the inputs and intent of each scenario.

TABLE IV. PREDICTOR INPUTS GROUPED BY SOURCE AND ROLE, WITH REPRESENTATIVE EXAMPLES

Group	Examples
Raw OHLCV (stock-level)	*Open*, *High*, *Low*, *Close*, *Volume*
Technical indicators (engineered from OHLCV)	Moving averages: *SMA_n*, *EMA_n* ($n = 5, 10, 20, 50, 100, 200$); Momentum: *MACD*, *MACDs*, *MACDh*, *RSI*; Volatility (Bollinger): *BBL_20_2.0*, *BBM_20_2.0*, *BBU_20_2.0*, *BBB_20_2.0*, *BBP_20_2.0*; Volume-based: *OBV*
Calendar features (engineered from date)	*dow_sin*, *dow_cos*, *doy_sin*, *doy_cos*; flags: *Ramadan*, *qtr_end_5d*
Macroeconomic features (external series)	*TASI*; *Brent*, *Brent_ret*, *Brent_vol60*

TABLE V. UNIFIED LSTM ARCHITECTURE AND TRAINING SETTINGS USED FOR ALL SCENARIOS (A–H)

Component	Setting
Input window	60 trading days (multivariate sequence)
Target	Next-day closing price
Layer 1	LSTM (96 units), `return_sequences`=True
Regularization	Dropout = 0.20 (after Layer 1)
Layer 2	LSTM (64 units)
Head	Dense(32, ReLU) $\rightarrow$ Dense(1, linear)
Loss	Mean Squared Error (MSE)
Optimizer	Adam, learning rate = 0.001
Batch size	32
Epoch cap	80
Early stopping	Patience = 6; restore_best_weights=True
Validation split	15% of training data
Evaluation metrics	MAE, RMSE, MAPE (%)

V. RESULTS

Eight model variants were evaluated, ranging from a simple baseline, Variant A (Close price only) to models applying advanced features (Variants B–H). Performance was assessed using MAE, RMSE, and MAPE. The analysis focuses on MAE as the primary evaluation metric due to its direct interpretability and robustness to outliers, though results are validated across all three measures. Full per-ticker MAE/RMSE/MAPE values are provided in the Supplementary Material (Table IX).

Table VII shows that market-aware models achieved superior predictive performance. Variant G (Close + TASI) achieved the highest performance for 44.8% of stocks with a median 6.19% improvement over baseline. Variant F (Close + Brent oil) ranked second, optimal for 24.1% of stocks with a 3.17% median improvement. Combined, these macroeconomic variants were top-performing for nearly 70% of evaluated stocks.

Complex feature engineering approaches proved counterproductive. In contrast, **Variant B** (*Close* + all features) achieved optimal results for only 10.3% of stocks and underperformed the baseline by 11.66%. **Variant C** (*Close* + selected features), despite employing advanced techniques including mutual information, correlation filtering, and SHAP-based ranking, was optimal for just 3.4% of stocks with an 8.32% performance decline. **Variant D** (*OHLCV*) failed to achieve optimal performance for any stock, while **Variant E** (*Close* + calendar features) was optimal for only 3.4% of stocks, despite incorporating seasonal effects specific to the Saudi market. Even combining individually successful features failed to yield benefits, **Variant H** (*Close* + *Brent* + *TASI*) was optimal for only 6.9% of stocks despite incorporating the two most effective individual inputs.

To assess variation in model performance across market segments, we initially considered a sector-level analysis. However, several sectors included fewer than three stocks, making statistical comparisons infeasible. To address this, we aggregated sectors into three broader, economically meaningful groups: **Oil-Sensitive** (Energy, Basic Materials, Utilities), **Financial** (Financial Services), and **Demand & Services** (Technology, Communication Services, Consumer Cyclical, Healthcare, Industrials, Consumer Defensive, Real Estate). These groupings reflect shared economic drivers: resource price exposure, interest rate sensitivity, and consumer demand patterns, respectively. Fig. 3 shows the three groups and the percentage of stocks in each group.

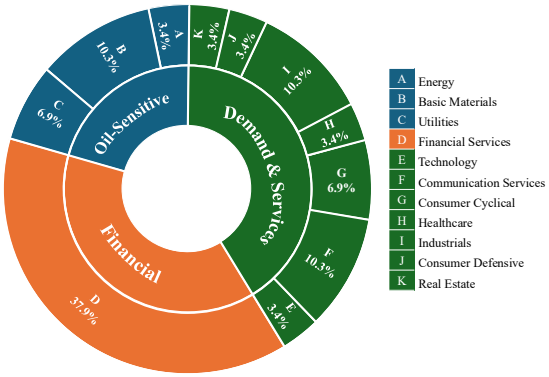

Fig. 3. Percentage breakdown of the stocks in three groups.

Table VIII demonstrates Variant G (Close + TASI) consistent effectiveness across groups. Under random chance, any variant would win roughly 12.5% of the time (1/8). Variant G (Close + TASI) significantly exceeds this benchmark, capturing 54.55% of Financial stocks (4.4$\times$ chance), 50% of Oil-Sensitive stocks (4$\times$ chance), and 33.33% of Demand & Services stocks (2.7$\times$ chance).

Fig. 4, Fig. 5, and Fig. 6 provide individual stock visualizations for each sector group. Variant G (Close + TASI) appears alone when optimal; otherwise, the best-performing alternative appears in orange. While no single variant proves optimal for all stocks, Variant G (Close + TASI) emerges as the dominant general-purpose solution across diverse economic sectors. The orange lines visible in over half the stocks demonstrate opportunities for stock-specific optimization, yet Variant G (Close + TASI) broad effectiveness confirms its value as the preferred baseline approach. Notably, recently listed companies, marked

TABLE VI. SIMPLIFIED LSTM INPUT SCENARIOS AND THEIR PURPOSE

Scenario	What It Uses	Why It Matters
A	*Close* price only	Tests if recent closing prices alone can predict future prices (price memory).
B	All 37 features	Uses all available signals, including prices, indicators, calendar data, and economic information, to test maximum prediction ability.
C	Selected top features	Checks if a smaller set of the most useful features can perform as well as the full list.
D	*OHLCV*	Focuses on basic trading data without additional indicators or external inputs.
E	*Close* + calendar patterns	Tests whether time-based patterns (e.g., weekdays, *Ramadan*) contribute to price prediction.
F	*Close* + *Brent* oil price	Evaluates how oil prices, crucial to the Saudi economy, influence stock movements.
G	*Close* + *TASI*	Adds the general market trend (*TASI*) to assess the impact of market sentiment on individual stocks.
H	*Close* + *Brent* + *TASI*	Combines both oil and market trend information to determine if their joint use enhances forecast accuracy.

TABLE VII. PERFORMANCE SUMMARY OF MODEL VARIANTS

Variant	Description	Win Rate (%)	Median Improvement over Baseline (%)
G	*Close* + *TASI*	44.8	6.19
F	*Close* + *Brent* oil	24.1	3.17
B	*Close* + all features	10.3	-11.66
H	*Close* + *Brent* + *TASI*	6.9	1.55
A	*Close* price only	6.9	baseline
E	*Close* + calendar features	3.4	0.48
C	*Close* + selected features	3.4	-8.32
D	*OHLCV*	0	0.45

TABLE VIII. SECTOR GROUP PERFORMANCE ANALYSIS

Sector Group	Number of Stocks	Best Variant	Win Rate (%)
Oil-Sensitive	6	G (Close + TASI)	50.00
Financial	11	G (Close + TASI)	54.55
Demand & Services	12	G (Close + TASI)	33.33

with double asterisks (**) in Table III, show substantially worse forecasting performance, highlighting the critical importance of sufficient trading history for reliable predictions.

VI. DISCUSSION

The results challenge the commonly held belief that increased feature complexity leads to superior predictive accuracy in financial markets. All models employed identical LSTM architectures, isolating the effect of feature selection and revealing that simpler, macroeconomic-based inputs consistently outperform sophisticated technical indicators.

Viewed against recent LSTM-based forecasting studies, our results help clarify what different feature choices actually contribute. Many works report strong performance, but treat the input design as fixed, offering limited insight into how much each additional feature family improves on a simple price-based baseline. Some studies argue that multivariate OHLC or OHLCV inputs yield better predictions than univariate closing prices [15], [19], [26], while others rely on closing prices by default [22], [27] or augment them with large sets of technical indicators and hybrid components [5], [16], [18], [21], [23], [25]. Our eight-scenario framework brings these variations together in a single empirical setting: with the LSTM architecture held constant, we systematically compare close-only, OHLCV, technical indicator, and macroeconomic configurations across the largest Tadawul-listed companies. The evidence shows that closing prices form a strong baseline, that OHLCV data do not systematically improve upon this baseline, and that extensive technical-indicator sets rarely deliver gains. By contrast, a small set of macroeconomic variables (TASI and Brent) provides consistent, albeit moderate, improvements over the closing price baseline for a large share of stocks. In this way, the study synthesizes scattered findings in the literature into a coherent picture of what each class of inputs adds beyond the closing price itself.

The systematic failure of complex feature engineering approaches reveals a fundamental insight about information processing in LSTM networks. Variant C (Close + selected features), despite employing advanced techniques including mutual information, correlation filtering, and SHAP-based ranking that are customized for each stock, degraded performance by 8.32% compared to the baseline. Variant B (Close + all features) underperformed by 11.66%. Similarly, Variant D (OHLCV) achieved zero optimal results, demonstrating that even basic microstructural data cannot enhance LSTM predictions beyond what is captured by closing price alone. Variant E (Close + calendar features) proved equally ineffective, with only 3.4% optimal performance despite incorporating known calendar effects such as Ramadan. This suggests that LSTM architectures efficiently capture price-based patterns from closing prices alone, making technical indicators largely redundant. Technical indicators such as moving averages, RSI, and MACD are essentially mathematical transformations of the same underlying price information that LSTM networks can learn implicitly through their temporal modeling capabilities. The addition of these derived features introduces noise rather than signal, as they provide no genuinely new information beyond what the network can extract from the raw price sequence.

The superior performance of macroeconomic variants provides a contrast to the redundancy of this technical indicator. Variant F (Close + Brent oil) achieved optimal performance for 24.1% of stocks, reflecting the Kingdom's central role in global energy markets. Variant G (Close + TASI) performed

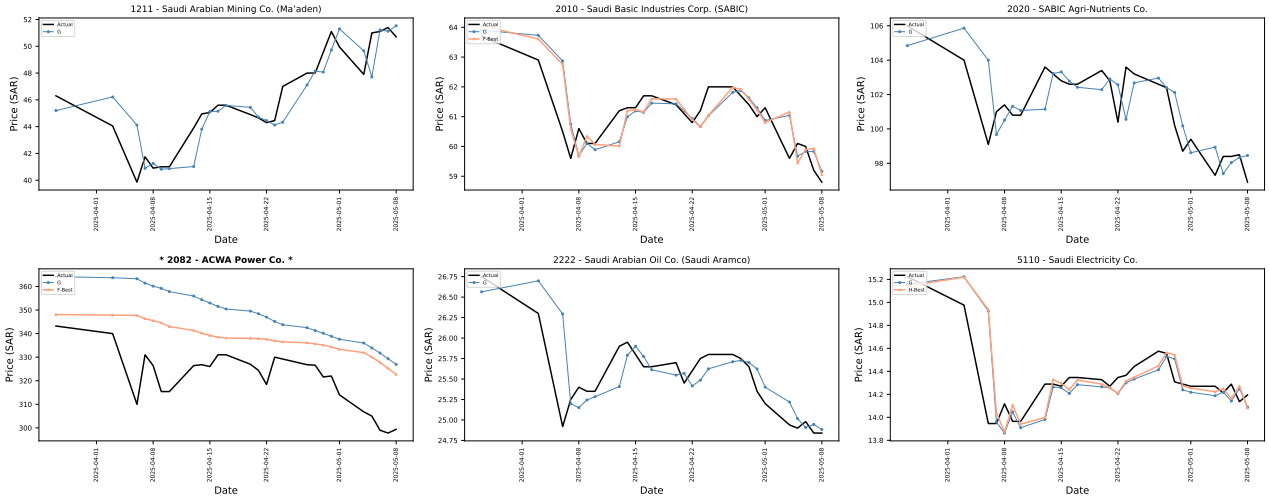

Fig. 4. Oil-Sensitive group (Energy, Basic Materials, and Utilities). Variant G (Close + TASI) predictions (blue) vs. Actual prices (black). If Variant G (Close + TASI) is not the top model for a stock, the best-performing alternative is overlaid in orange.

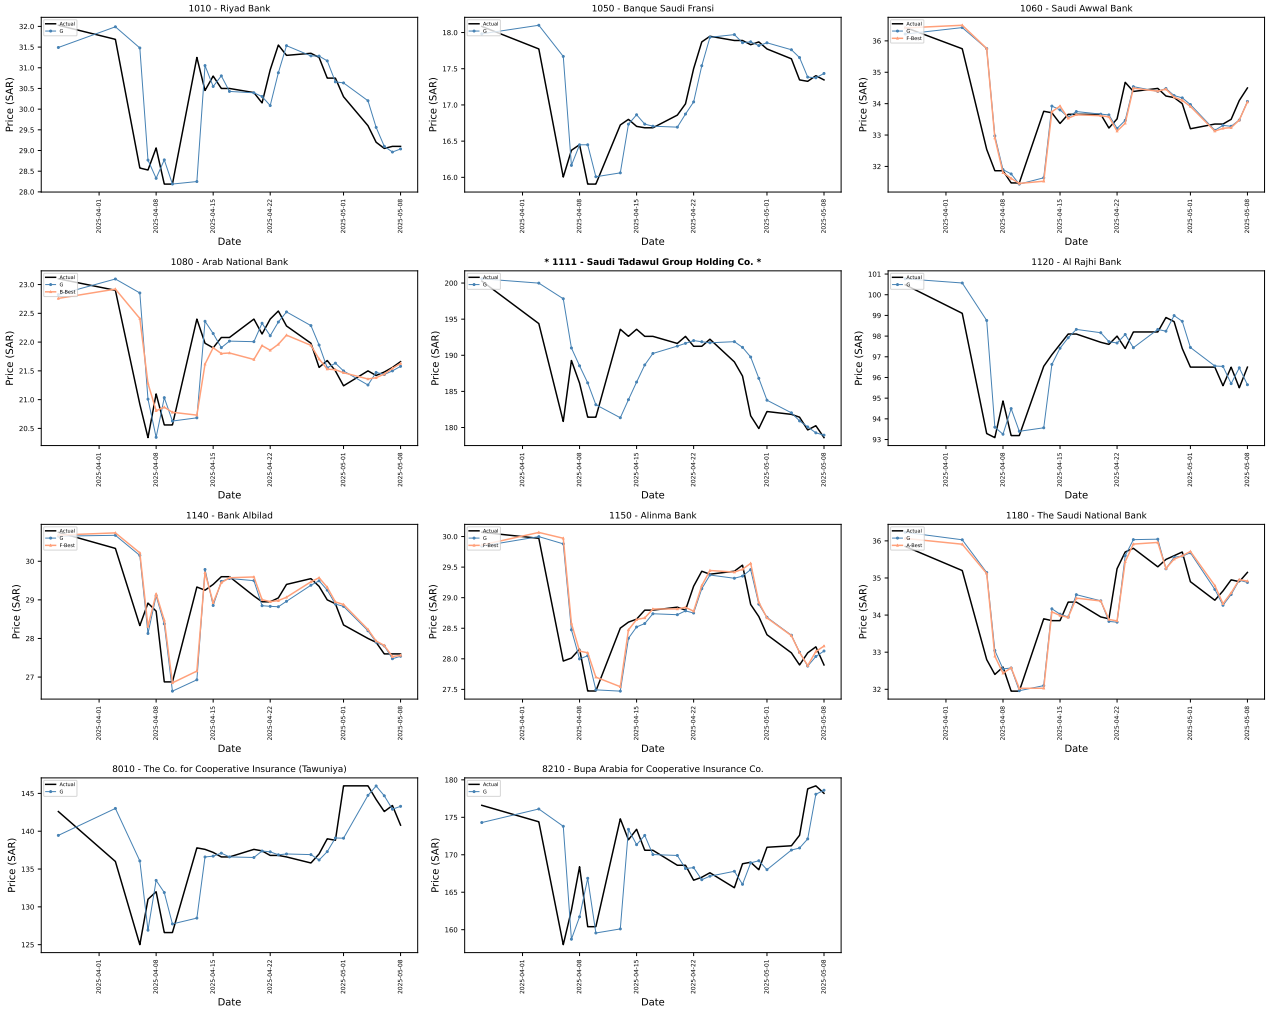

Fig. 5. Financial group (financial services). Variant G (Close + TASI) predictions (blue) vs. Actual prices (black). If Variant G (Close + TASI) is not the top model for a stock, the best-performing alternative is overlaid in orange.

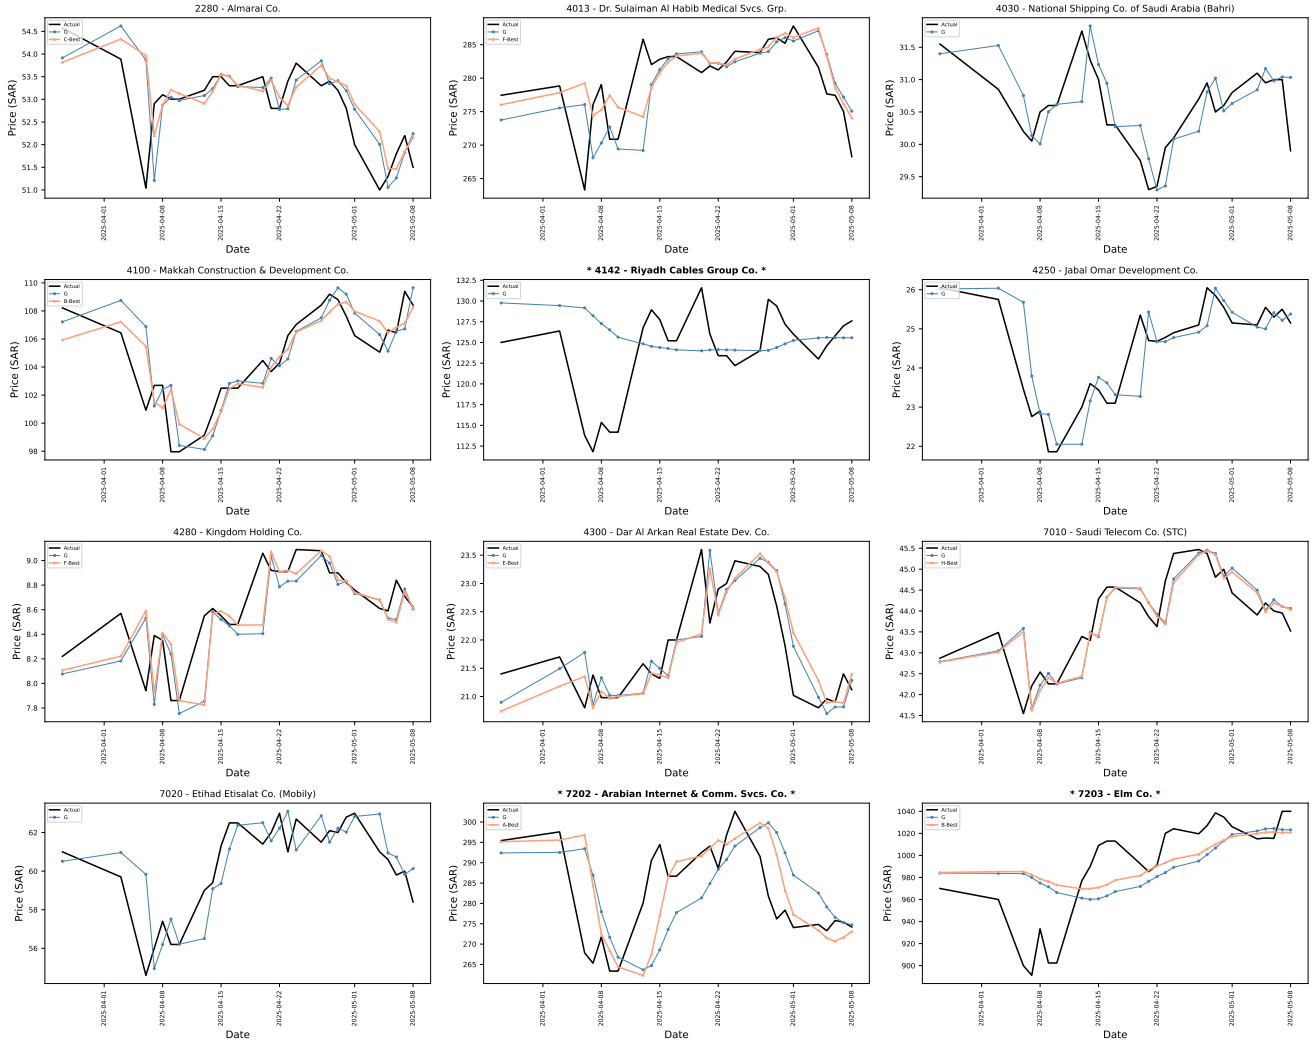

Fig. 6. Demand & services group (technology, communication services, consumer cyclical, healthcare, industrials, consumer defensive, and real estate). Variant G (Close + TASI) predictions (blue) vs. Actual prices (black). If Variant G (Close + TASI) is not the top model for a stock, the best-performing alternative is overlaid in orange.

best for 44.8% of stocks with a median 6.19% improvement over baseline. Unlike technical indicators that derive from the same price data, TASI provides independent information not captured in individual stock prices. TASI captures market-wide sentiment, liquidity conditions, and systematic risk factors that affect individual stock behavior. This represents a fundamentally different information source rather than a more complex processing of the same data.

The failure of technical indicators and success of individual macroeconomic variants reflects a fundamental principle: practical feature engineering requires genuinely distinct information sources rather than mathematical transformations of the same data. However, Variant H (Close + Brent + TASI) achieved optimal performance for only 6.9% of stocks despite combining individually successful features, demonstrating that even orthogonal information sources can create harmful interactions when combined, making simpler approaches more reliable.

A critical constraint emerges from data availability that transcends model sophistication. Recently listed companies such as ACWA Power (2082), Elm Company (7203), Riyadh Cables Group (4142), Saudi Tadawul Group Holdings (1111), and Arabian Internet & Communications Services (7202) have limited trading histories (≤ 650 days) and show dramatically worse performance (average MAE: 16.936) compared to established stocks ($\geq 3,000$ days, average MAE: 0.844). This twenty-fold performance deterioration underscores that data availability is more fundamental than feature complexity.

These findings have immediate practical implications for portfolio management and algorithmic trading systems. Variant G (Close + TASI) emerges as the most effective general-purpose approach, showing consistent performance across diverse sectors even though it may not be optimal for every individual stock. The results provide clear guidance: models based on fundamental economic signals outperform complex feature engineering in Saudi Arabia's market. Portfolio implementations may achieve further gains through selective model customization when sufficient trading history is available.

It is important to acknowledge several assumptions and limitations of this study. First, only a predetermined set of daily predictors (OHLCV, TASI, Brent, and calendar effects) and the biggest Tadawul-listed companies are included in the analysis. Therefore, the results might not apply to small-cap stocks, intraday timeframes, or situations where firm fundamentals or other information sources are accessible. Second, to isolate the contribution of different feature families, we use the same LSTM architecture and hyperparameters for all stocks and scenarios. While this design ensures a fair comparison across inputs, it may underestimate the best achievable performance for particular stocks or alternative architectures. Third, the testing evaluation window is relatively short (January to May 2025), longer or multi-regime evaluation periods could lead to different relative rankings of the scenarios. Finally, we assume that daily closing prices and a 60-day lookback window are sufficient for one-day-ahead forecasting, and other horizons and window lengths are not explored here. Together with the earlier observation that performance sharply deteriorates for newly listed stocks with short trading histories, these limitations define the scope within which our conclusions should be interpreted.

Overall, with the LSTM architecture held fixed, our results show that supplying inputs that introduce new information is more important than augmenting the dataset with technical indicators and other complex transformations derived from the same price data.

VII. Conclusion

This study examined how different input features affect the performance of LSTM models in forecasting stock prices within the Saudi market. Through a systematic evaluation of eight structured model variants across the 30 largest Tadawul-listed companies (29 after cleaning), we show that simpler, macroeconomic-based inputs consistently outperform complex technical indicators and advanced feature engineering techniques.

Feature engineering was most effective when it introduced genuinely distinct information. Models incorporating TASI and Brent oil prices, which reflect broader market conditions, delivered the strongest performance. The Close + TASI model (Variant G), in particular, reduced median error by 6.19% compared to a closing-price baseline. In contrast, models built on extensive technical indicators or automated selection pipelines degraded accuracy by 8 to 12%.

This performance gap likely reflects the LSTM's ability to capture temporal patterns directly from raw price sequences. When technical indicators such as moving averages, RSI, and MACD replicate patterns already present in historical data, they contribute more noise than a proper signal.

Data availability emerged as an even more critical determinant of accuracy. Recently listed firms with limited trading histories ($\leq$650 days) experienced MAE that rose dramatically, reaching levels twenty times higher than established stocks. This finding highlights the importance of historical depth and cautions against applying deep learning models in data-scarce settings.

Overall, our results challenge the notion that large and advanced feature sets lead to better forecasts, demonstrating that a parsimonious set of economically meaningful variables enhances both performance and interpretability in LSTM-based stock forecasting. To transition this single-market (Saudi) finding into a generalizable framework, our future research will be built on three pillars. First, we will systematically test indicator combinations across diverse markets. Second, we will challenge the sufficiency of closing prices by incorporating richer, complementary data sources, such as intraday transactions and sentiment measures. Third, we will benchmark these findings against advanced architectures, including hybrid LSTM-Transformer models. This grand unified study will be essential for validation, enabling us to isolate a core set of universal features and quantify the predictive power of market specific indicators.

References

[1] X. Li, P. Wu, and W. Wang, "Incorporating stock prices and news sentiments for stock market prediction: A case of Hong Kong," *Information Processing & Management*, vol. 57, p. 102212, 2020.

[2] J. Bialkowski, H.D. Dang, and X. Wei, "High policy uncertainty and low implied market volatility: An academic puzzle?" *Journal of Financial Economics*, vol. 143, pp. 1185–1208, 2022.

[3] N. Sapkota, "News-based sentiment and bitcoin volatility," *International Review of Financial Analysis*, vol. 82, p. 102183, 2022.

[4] W. Alim, N.U. Khan, V.W. Zhang, H.H. Cai, A. Mikhaylov, and Q. Yuan, "Influence of political stability on the stock market returns and volatility: GARCH and EGARCH approach," *Financial Innovation*, vol. 10, p. 139, 2024.

[5] T. Phuoc, P.T.K. Anh, P.H. Tam, and C.V. Nguyen, "Applying machine learning algorithms to predict the stock price trend in the stock market—The case of Vietnam," *Humanities and Social Sciences Communications*, vol. 11, pp. 1–18, 2024.

[6] C. Zhang, Y. Zhang, M. Cucuringu, and Z. Qian, "Volatility forecasting with machine learning and intraday commonality," *Journal of Financial Econometrics*, vol. 22, pp. 492–530, 2024.

[7] D.P. Gandhmal and K. Kumar, "Systematic analysis and review of stock market prediction techniques," *Computer Science Review*, vol. 34, p. 100190, 2019.

[8] H. Han, Z. Liu, M. Barrios Barrios, J. Li, Z. Zeng, N. Sarhan, and E.M. Awwad, "Time series forecasting model for non-stationary series pattern extraction using deep learning and GARCH modeling," *Journal of Cloud Computing*, vol. 13, p. 2, 2024.

[9] M. Adhikari, "Forecasting stock index closing points using ARIMA-GARCH with a rolling data window," *International Journal of Scientific Research in Archive*, vol. 13, 2024.

[10] F.A. Alturki and A.M. Aldughaiyem, "Trading Saudi stock market shares using a multivariate recurrent neural network with a long short-term memory layer," *International Journal of Advanced Computer Science and Applications*, vol. 11, 2020.

[11] J. Sen and S. Mehtab, "Long-and-Short-Term Memory (LSTM) Networks Architectures and Applications in Stock Price Prediction," in *Emerging Computing Paradigms: Principles, Advances and Applications*, Wiley, 2022, pp. 143–160.

[12] M. Jarrah, "Long Short-Term Memory and Discrete Wavelet Transform-based Univariate Stock Market Prediction Model," *Journal of Information and Organizational Sciences*, vol. 48, pp. 263–277, 2024.

[13] K. Olorunnimbe and H. Viktor, "Deep learning in the stock market—a systematic survey of practice, backtesting, and applications," *Artificial Intelligence Review*, vol. 56, pp. 2057–2109, 2023.

[14] S. Sang and L. Li, "A novel variant of the LSTM stock prediction method incorporating an attention mechanism," *Mathematics*, vol. 12, p. 945, 2024.

[15] G. Ding and L. Qin, "Study on the prediction of stock price based on the associated network model of LSTM," *International Journal of Machine Learning and Cybernetics*, vol. 11, pp. 1307–1317, 2020.

[16] H.N. Bhandari, B. Rimal, N.R. Pokhrel, R. Rimal, K.R. Dahal, and R.K.C. Khatri, "Predicting stock market index using LSTM," *Machine Learning with Applications*, vol. 9, p. 100320, 2022.

[17] Z. Li, H. Yu, J. Xu, J. Liu, and Y. Mo, "Stock market analysis and prediction using LSTM: A case study on technology stocks," *Innovations in Applied Engineering and Technology*, pp. 1–6, 2023.

[18] J. Zou, J. Lou, B. Wang, and S. Liu, "A novel deep reinforcement learning based automated stock trading system using cascaded LSTM networks," *Expert Systems with Applications*, vol. 242, p. 122801, 2024.

[19] S. Banik, N. Sharma, M. Mangla, and S.N. Mohanty, "LSTM-based decision support system for swing trading in the stock market," *Knowledge-Based Systems*, vol. 239, p. 107994, 2022.

[20] J. Long, Z. Chen, W. He, T. Wu, and J. Ren, "An integrated framework of deep learning and knowledge graph for prediction of stock price trend: An application in the Chinese stock exchange market," *Applied Soft Computing*, vol. 91, p. 106205, 2020.

[21] N. Jing, Z. Wu, and H. Wang, "A hybrid model integrating deep learning with investor sentiment analysis for stock price prediction," *Expert Systems with Applications*, vol. 178, p. 115019, 2021.

[22] B. Gülmez, "Stock price prediction with optimized deep LSTM network with artificial rabbits optimization algorithm," *Expert Systems with Applications*, vol. 227, p. 120346, 2023.

[23] Y. Zhang, B. Yan, and M. Aasma, "A novel deep learning framework: Prediction and analysis of financial time series using CEEMD and LSTM," *Expert Systems with Applications*, vol. 159, p. 113609, 2020.

[24] J. Shah, D. Vaidya, and M. Shah, "A comprehensive review on multiple hybrid deep learning approaches for stock prediction," *Intelligent Systems with Applications*, vol. 16, p. 200111, 2022.

[25] M. Alsulmi and N. Al-Shahrani, "Machine Learning-Based Decision-Making for Stock Trading: Case Study for Automated Trading in Saudi Stock Exchange," *Scientific Programming*, vol. 2022, p. 6542862, 2022.

[26] M. Jarrah and M. Derbali, "Predicting Saudi stock market index by using multivariate time series based on deep learning," *Applied Sciences*, vol. 13, p. 8356, 2023.

[27] A.H. Al-Nefaie and T.H.H. Aldhyani, "Predicting close price in the emerging Saudi Stock Exchange: time series models," *Electronics*, vol. 11, p. 3443, 2022.

[28] E.F. Fama, "Efficient capital markets: A review of theory and empirical work," *Journal of Finance*, vol. 25, pp. 383–417, 1970.

[29] R.J. Shiller, "From efficient markets theory to behavioral finance," *Journal of Economic Perspectives*, vol. 17, pp. 83–104, 2003.

[30] C.R. Harvey, "The evolving risks of emerging markets," *Journal of Financial Economics*, vol. 132, pp. 145–168, 2018.

[31] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, pp. 1735–1780, 1997.

[32] T. Fischer and C. Krauss, "Deep learning with long short-term memory networks for financial market predictions," *European Journal of Operational Research*, vol. 270, pp. 654–669, 2018.

[33] D.M.Q. Nelson, A.C.M. Pereira, and R.A. de Oliveira, "Stock market's price movement prediction with LSTM neural networks," in *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*, IEEE, 2017, pp. 1419–1426.

[34] F.A. Gers, J. Schmidhuber, and F. Cummins, "Learning to forget: Continual prediction with LSTM," *Neural Computation*, vol. 12, pp. 2451–2471, 2000.

[35] K. Greff, R.K. Srivastava, J. Koutník, B.R. Steunebrink, and J. Schmidhuber, "LSTM: A search space odyssey," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 28, pp. 2222–2232, 2017.

[36] K. Chen, Y. Zhou, and F. Dai, "Stock price prediction using long short-term memory and convolutional neural network," *Expert Systems with Applications*, vol. 185, p. 115716, 2021.

[37] P. Andrikopoulos, Y.Cui, S.Gad, and V. Kallinterakis, "Feedback trading and the Ramadan effect in frontier markets," *Research in International Business and Finance*, vol. 51, p. 101085, 2020.

[38] S. Bikhchandani, D. Hirshleifer, and I. Welch, "A theory of fads, fashion, custom, and cultural change as informational cascades," *Journal of Political Economy*, vol. 100, pp. 992–1026, 1992.

[39] D.Gabbori, N.Virk, N.Aftab, and B.Awartani, "The impact of Islamic events on herding behaviour in Saudi Arabian equities market," *International Journal of Finance & Economics*, vol. 29, no. 1, pp. 119, 2024.

[40] W. Zhang, Y. Wang, and Y. Liu, "A machine learning approach to predict stock returns: Evidence from the Chinese stock market," *Economic Modelling*, vol. 89, pp. 307–317, 2020.

[41] D.W.H. Alrabadi and S.Y. Al-Abdallah, "Oil prices and stock markets in GCC countries: Empirical evidence from panel analysis," *International Journal of Energy Economics and Policy*, vol. 8, pp. 185–192, 2018.

[42] F. Alqahtani and D.G. Mayes, "Financial markets' reaction to major national policy announcements: Evidence from Saudi Vision 2030," *Research in International Business and Finance*, vol. 45, pp. 529–545, 2018.

[43] M.K. Khan, S. Bibi, and Y. Lyu, "Vision 2030 and Saudi stock market performance: A sectoral analysis," *Cogent Economics & Finance*, vol. 8, p. 1727624, 2020.

SUPPLEMENTARY MATERIAL

TABLE IX. PERFORMANCE METRICS ACROSS ALL TICKER VARIANTS

	MAE								RMSE								MAPE (%)							
Ticker	A	B	C	D	E	F	G	H	A	B	C	D	E	F	G	H	A	B	C	D	E	F	G	H
1010	0.561	0.533	0.532	0.540	0.519	0.514	**0.504**	0.535	0.944	**0.855**	0.873	0.887	0.912	0.903	0.889	0.915	1.9%	1.8%	1.8%	1.8%	1.7%	1.7%	**1.7%**	1.8%
1050	0.233	0.239	0.252	0.235	0.235	0.245	**0.219**	0.230	0.419	0.411	0.427	0.412	0.422	0.436	**0.395**	0.408	1.4%	1.4%	1.5%	1.4%	1.4%	1.5%	**1.3%**	1.4%
1060	0.492	0.740	0.680	0.535	0.503	**0.490**	0.499	0.495	**0.836**	1.059	1.056	0.885	0.868	0.873	0.860	0.856	1.5%	2.2%	2.0%	1.6%	1.5%	**1.5%**	1.5%	1.5%
1080	0.376	**0.342**	0.361	0.393	0.361	0.371	0.369	0.374	0.596	**0.538**	0.591	0.600	0.605	0.597	0.580	0.595	1.7%	**1.6%**	1.7%	1.8%	1.7%	1.7%	1.7%	1.7%
1111	4.200	9.105	10.642	4.646	5.912	5.741	**3.605**	6.324	**5.395**	10.611	11.951	5.756	7.058	6.861	5.445	7.596	2.3%	4.9%	5.8%	2.5%	3.2%	3.1%	**1.9%**	3.4%
1120	0.989	1.105	1.072	0.999	0.970	0.939	**0.896**	0.976	1.556	1.682	1.759	1.634	1.501	1.474	**1.409**	1.499	1.0%	1.2%	1.1%	1.0%	1.0%	1.0%	**0.9%**	1.0%
1140	0.458	0.630	0.449	0.473	0.441	**0.410**	0.436	0.416	0.735	0.832	0.769	0.803	0.757	**0.693**	0.714	0.681	1.6%	2.2%	1.6%	1.7%	1.5%	**1.4%**	1.5%	1.5%
1150	0.311	0.505	0.445	0.336	0.342	**0.301**	0.312	0.306	0.513	0.626	0.593	0.507	0.540	0.505	**0.492**	0.496	1.1%	1.8%	1.6%	1.2%	1.2%	**1.1%**	1.1%	1.1%
1180	**0.478**	0.670	0.567	0.510	0.482	0.485	0.495	0.479	**0.726**	0.924	0.752	0.728	0.732	0.743	0.739	0.738	**1.4%**	2.0%	1.7%	1.5%	1.4%	1.4%	1.5%	1.4%
1211	1.100	1.573	1.159	1.096	1.117	1.102	**1.071**	1.085	1.541	2.009	1.634	1.544	1.630	1.569	**1.538**	1.545	2.4%	3.4%	2.5%	2.4%	2.4%	2.4%	**2.3%**	2.4%
2010	0.576	0.602	0.723	0.588	0.576	**0.514**	0.520	0.576	0.791	0.857	1.029	0.783	0.803	**0.731**	0.735	0.777	0.9%	1.0%	1.2%	1.0%	0.9%	**0.8%**	0.9%	0.9%
2020	1.246	1.179	1.234	1.188	1.174	1.217	**1.129**	1.262	1.693	1.675	1.632	1.578	1.576	1.664	**1.568**	1.666	1.2%	1.2%	1.2%	1.2%	1.2%	1.2%	**1.1%**	1.3%
2082	30.639	27.682	20.465	16.954	23.890	**16.612**	26.532	26.150	31.701	29.162	22.304	19.388	25.747	**18.684**	28.070	27.431	9.6%	8.7%	6.5%	5.4%	7.5%	**5.3%**	8.3%	8.2%
2222	0.254	0.389	0.256	0.232	0.230	0.218	**0.204**	0.252	0.410	0.496	0.386	0.383	0.335	0.365	**0.326**	0.399	1.0%	1.5%	1.0%	0.9%	0.9%	0.9%	**0.8%**	1.0%
2280	0.533	0.595	**0.497**	0.574	0.519	0.505	0.507	0.514	0.798	0.802	**0.744**	0.820	0.784	0.769	0.775	0.785	1.0%	1.1%	**0.9%**	1.1%	1.0%	1.0%	1.0%	1.0%
4013	3.498	3.906	4.820	3.635	3.637	**3.097**	3.553	4.088	5.223	5.000	5.949	5.451	5.280	**4.761**	5.307	5.547	1.3%	1.4%	1.7%	1.3%	1.3%	**1.1%**	1.3%	1.5%
4030	0.382	0.468	0.447	0.379	0.372	0.370	**0.347**	0.364	0.495	0.576	0.552	0.485	0.491	0.485	**0.465**	0.471	1.3%	1.5%	1.5%	1.2%	1.2%	1.2%	**1.1%**	1.2%
4100	1.430	**1.301**	1.456	1.381	1.532	1.423	1.422	1.371	1.933	**1.732**	1.932	2.059	2.103	1.989	1.918	1.880	1.4%	**1.3%**	1.4%	1.3%	1.5%	1.4%	1.4%	1.3%
4142	5.993	4.608	7.202	5.873	5.570	9.513	**4.499**	16.513	7.341	**6.083**	8.166	7.217	6.381	11.337	6.524	17.429	5.0%	3.8%	5.9%	4.8%	4.5%	7.9%	**3.8%**	13.5%
4250	0.524	0.867	0.609	0.513	0.546	0.509	**0.492**	0.497	0.784	1.085	0.881	0.749	0.802	0.781	**0.745**	0.751	2.2%	3.6%	2.6%	2.1%	2.3%	2.1%	**2.0%**	2.1%
4280	0.209	0.192	0.192	0.226	0.193	**0.173**	0.188	0.175	0.290	**0.254**	0.279	0.292	0.274	0.277	0.282	0.276	2.5%	2.3%	2.3%	2.7%	2.3%	**2.0%**	2.2%	2.1%
4300	0.477	0.574	0.528	0.442	**0.425**	0.437	0.437	0.427	0.601	0.700	0.658	0.581	**0.563**	0.559	0.577	0.561	2.2%	2.6%	2.4%	2.0%	**1.9%**	2.0%	2.0%	1.9%
5110	0.145	0.258	0.213	0.142	0.146	0.130	0.133	**0.129**	0.238	0.312	0.278	0.248	0.254	0.233	0.226	**0.227**	1.0%	1.8%	1.5%	1.0%	1.0%	0.9%	0.9%	**0.9%**
7010	0.456	0.503	0.451	0.447	0.465	0.456	0.437	**0.432**	0.628	0.685	0.641	0.623	0.650	0.628	0.604	**0.593**	1.0%	1.2%	1.0%	1.0%	1.1%	1.0%	1.0%	**1.0%**
7020	1.174	1.157	1.167	1.182	1.169	1.179	**1.151**	1.197	1.543	**1.455**	1.566	1.528	1.601	1.579	1.553	1.546	2.0%	1.9%	2.0%	2.0%	2.0%	2.0%	**1.9%**	2.0%
7202	**7.346**	11.765	12.087	8.921	11.138	8.895	10.624	15.104	**10.837**	14.462	13.861	11.977	14.057	12.322	13.244	18.579	**2.6%**	4.2%	4.3%	3.2%	4.0%	3.2%	3.8%	5.3%
7203	34.362	**28.388**	38.477	33.888	43.120	32.036	31.619	50.806	42.167	**37.639**	45.277	38.987	54.061	38.695	38.929	54.429	3.5%	**3.0%**	3.9%	3.5%	4.3%	3.3%	3.3%	5.1%
8010	2.613	3.186	2.850	2.577	2.599	2.448	**2.427**	2.489	3.911	4.831	4.064	3.862	4.067	3.803	**3.789**	3.796	1.9%	2.4%	2.1%	1.9%	1.9%	1.8%	**1.8%**	1.8%
8210	3.248	3.746	3.308	3.216	3.138	3.291	**3.047**	3.113	5.198	5.654	5.243	5.132	5.117	5.122	**4.957**	5.009	1.9%	2.2%	2.0%	1.9%	1.9%	2.0%	**1.8%**	1.9%