Multi-Frequency Feature Integration for Commodity Price Prediction: A Temporal Resolution Perspective

FU-SHENG TSAI*, TZU-HSUAN HUANG* and PO-CHENG SU**

**Intelligent Solution Technology R&D Department*
***Inspection and Electromagnetic R&D Department*
China Steel Corporation

Accurate forecasting of coking coal and iron ore prices is essential for strategic procurement and cost management in the steel industry. Traditional approaches often rely on single-frequency time series, which fail to capture the complex and multi-scale nature of commodity markets. This study proposes a multi-frequency feature integration framework that incorporates monthly, weekly, and daily indicators to improve forecasting accuracy across short- and medium-term horizons (N+1 to N+3). Using a dual-stage learning architecture based on Support Vector Regression, we evaluate four temporal configurations and demonstrate that models combining all three frequencies consistently outperform their single-scale counterparts. Furthermore, explainable AI methods such as SHAP and permutation importance are applied to identify key drivers across different frequencies. Results reveal that coal prices are primarily influenced by slow-moving monthly indicators, while iron ore forecasts benefit more from high-frequency variables such as port inventory levels and daily price benchmarks. These findings not only validate the importance of temporal diversity but also offer interpretable insights into market dynamics, supporting more informed and responsive forecasting strategies for industrial practitioners.

Keywords: Coal and iron ore price forecasting, Multi-frequency features, Time series prediction, Support vector regression, Explainable AI

1. INTRODUCTION

Price volatility in raw material markets—especially coking coal and iron ore—remains a critical concern for the steel industry. These commodities constitute a major portion of total production costs, often accounting for over 60% of the variable expenditure in integrated steelmaking processes. Coking coal, in particular, plays a vital role in blast furnace operations as both a reducing agent and a source of heat, making its price and availability crucial to production planning and cost efficiency. Fluctuations in the prices of coking coal and iron ore are driven by a wide array of external forces, including geopolitical instability, environmental regulations, currency exchange fluctuations, global trade dynamics, and logistical bottlenecks. For instance, extreme rainfall events in Queensland, Australia—home to several key coal exporters—have historically disrupted global coking coal supply, causing price surges that ripple through steel production value chains(1).

Such volatility poses multidimensional risks. From an operational standpoint, unexpected price hikes can erode profit margins, strain procurement schedules, and complicate budgeting cycles. From a strategic perspective, accurate price forecasting enables firms to optimize raw material procurement timing, hedge cost exposure, and better align upstream supply with downstream demand. In globalized commodity markets, where demand and supply dynamics are often asynchronous and exogenous shocks can originate from non-market forces (e.g., export bans, port closures, or ESG policy shifts), building robust and responsive forecasting models becomes a foundational capability(2).

Despite technological advances in data collection and modeling, forecasting prices in raw material markets remains challenging. Traditional statistical methods such as ARIMA, SARIMA, and exponential smoothing have been widely used due to their simplicity and interpretability. However, these models typically assume linearity and stationarity, assumptions that are frequently violated in commodity markets(3). Furthermore, they tend to operate at a single temporal resolution—either daily or monthly—without capturing the latent patterns observable across other time scales. As a result, these models often perform poorly in the presence of abrupt structural changes or high-frequency noise.

To address nonlinearities, hybrid approaches combining statistical and machine learning techniques have gained popularity. Notable examples include ARIMA–SVM and ARIMA–ANN models, which seek to leverage the time series structure of ARIMA with the non-linear mapping capacity of machine learning[4]. Although these models show performance gains over classical baselines[5], they generally continue to operate within a single time scale, limiting their ability to detect transient signals that might only manifest across different temporal granularities.

To overcome these limitations, recent studies have explored decomposition-based multiscale approaches, such as Empirical Mode Decomposition (EMD), Variational Mode Decomposition (VMD), and Wavelet Transforms, which allow for the decomposition of price series into intrinsic mode functions or frequency components that are modeled separately and recombined[5]. These methods have achieved substantial improvements in short- and long-term prediction of coal, oil, and natural gas prices, particularly when coupled with deep learning architectures like LSTM, BiLSTM, and Transformers [6, 7].

While these hybrid models offer promising results, they introduce new challenges. First, decomposition-based approaches often suffer from high computational complexity and limited scalability. Second, many studies focus on specific frequency combinations (e.g., daily + weekly) or are confined to particular commodities such as crude oil. Third, model interpretability remains a key barrier to adoption in industrial contexts, where business decisions depend not only on accuracy but also on explainability and feature traceability[8].

Moreover, current literature seldom compares the predictive utility of systematically combined multi-frequency features in a unified framework. Few studies have asked: Does integrating daily, weekly, and monthly features offer additive value? Or does the inclusion of high-frequency data introduce more noise than signal, particularly in forecasting horizons that exceed a few days? The answers to these questions have direct implications for practitioners seeking to deploy forecasting tools in volatile environments.

To address these gaps, the present study proposes a structured multi-frequency feature integration framework for coal and iron ore price forecasting. Specifically, we conduct a comprehensive evaluation of models built on four feature combinations: (a) monthly-only, (b) monthly + weekly, (c) monthly + daily, and (d) monthly + weekly + daily. Utilizing a consistent dual-stage learning architecture adapted from our previous work[9], we examine forecast accuracy across various horizons (N+1, N+2, N+3), assess model robustness, and analyze feature contributions using SHAP and permutation importance. This study aims to answer three research questions:

(1) Does multi-frequency integration consistently improve predictive accuracy over single-frequency models?
(2) Which temporal combination offers the best balance of accuracy and stability?
(3) How do feature contributions vary by frequency, and what does this reveal about underlying market dynamics?

Through this investigation, we aim not only to enhance forecasting accuracy but also to provide actionable insights into how temporal resolution shapes the value of input features. By exploring the balance between resolution granularity and signal robustness, our work contributes both methodologically and practically to the development of interpretable, adaptable price forecasting systems for the raw materials industry.

2. METHODOLOGY

2.1 Multi-Frequency Feature Collection

To effectively capture the heterogeneous and time-varying drivers of commodity price movements, a multi-frequency feature collection strategy was designed spanning three distinct temporal resolutions: monthly, weekly, and daily. This design reflects the premise that different classes of market signals emerge and evolve at different speeds. Macroeconomic shifts unfold gradually over months, while logistical dynamics or market sentiment may change weekly or even daily. By incorporating features from each frequency level, we aim to maximize the information diversity available to the forecasting model.

(1) *Monthly features* are designed to reflect slow-moving macroeconomic fundamentals and broad market trends. These include indices such as the OECD Composite Leading Indicators, the Baltic Dry Index (BDI), the Purchasing Managers' Index (PMI), and the S&P 500. Such indicators are generally updated on a monthly basis and serve as proxies for global demand, industrial activity, trade volume, and investment sentiment.
(2) *(eekly features* capture mid-frequency signals relevant to regional market conditions or supply chain dynamics. Key variables include port inventories at major exporting and importing hubs, as well as alternative pricing benchmarks for coking coal and iron ore, which are published on a weekly basis. These features are particularly useful for tracking stock levels and price spreads that indicate regional supply–demand imbalances.
(3) *Daily features* represent high-frequency fluctuations and are critical for capturing short-term shocks. These include spot prices, futures prices, and weather-related data such as rainfall levels in key mining regions or cyclone alerts near shipping ports. By integrating daily data, the model can better

respond to rapidly evolving market conditions and external disruptions.

This feature collection framework serves as the foundation for subsequent integration experiments, where different frequency combinations (e.g., monthly-only vs. monthly + weekly + daily) are compared to assess their impact on forecast performance.

2.2 Multi-Frequency Feature Integration

To systematically represent the temporal context preceding each prediction target, a unified feature integration scheme is developed to combine multi-frequency signals into a single aligned input vector. As illustrated in Figure X, for any given forecasting target month *t*, the corresponding feature vector X_t is constructed by aggregating past information from three temporal resolutions: monthly, weekly, and daily.

(1) *Monthly features*: For the target month *t*, the five preceding months *t-1* to *t-5* are used. This includes macroeconomic indicators that evolve gradually over time, capturing long-term structural trends.
(2) *Weekly features:* The five weeks immediately preceding the first day of month t are collected. These mid-frequency signals, such as port inventory levels and regional pricing benchmarks, provide a medium-term snapshot of market conditions.
(3) *Daily features:* To incorporate short-term market dynamics, we include the 14 trading days immediately prior to the start of month *t*, i.e., from *t-14* to *t-1*. These high-frequency indicators include spot prices, futures market signals, and weather-related disruptions.

All features from the three frequency levels are concatenated into a single vector X_t, which serves as the input for predicting the target value Y_t—the monthly average price of coking coal or iron ore for month *t*. This construction ensures that no future information is leaked into the model and that all predictions rely solely on historical data, thereby closely simulating a real-world forecasting environment. The feature integration strategy enables consistent comparison across different configurations (e.g., monthly-only, monthly + weekly, full integration), allowing for an empirical assessment of the added value of incorporating multi-frequency temporal structures.

2.3 Feature Preprocessing

Prior to model training, all raw features undergo standardized preprocessing procedures to ensure data quality and comparability across different temporal resolutions. First, missing values are imputed using frequency-specific strategies. For monthly and weekly data, linear interpolation and forward/backward filling methods are applied, depending on the consistency of reporting cycles. For more volatile daily data, statistical techniques such as rolling mean imputation or nearest-neighbor filling are employed to preserve short-term dynamics. Next, feature scaling is conducted to normalize the distribution of all inputs. Z-score normalization is adopted, transforming each feature into a zero mean and unit variance. This ensures that the features contribute proportionately during model training and reduces potential bias introduced by magnitude differences across variables from different frequencies.

2.4 Feature Selection and Validation

To reduce redundancy, prevent overfitting, and enhance generalization, a backward elimination feature selection strategy is implemented. Starting from the full candidate feature set, features are iteratively removed based on their marginal contribution to model performance.

The primary evaluation metric during elimination is the root mean square error (RMSE) on validation data. Given the temporal nature of the prediction task, we adopted a time-aware validation scheme based on an expanding window cross-validation protocol. In this approach, the training set is incrementally enlarged while keeping the test window fixed, mimicking real-world forecasting settings where only past information is available for future prediction. This methodology helps capture temporal drift, evaluate model robustness across time, and mitigate look-ahead bias.

2.5 Forecasting Model and Explainability

The final forecasting model is designed to predict the continuous value of coal and iron ore prices across a 3-month horizon. To capture complex, nonlinear patterns inherent in commodity markets, we employ a Support Vector Machine (SVM) regression model. The SVM is trained separately for each target commodity and forecast horizon (e.g., N+1, N+2, N+3), using the selected features from different temporal combinations. To enhance model transparency and usability, we incorporate explainable artificial intelligence (XAI) techniques to interpret model behavior. Specifically, we apply SHapley Additive exPlanations (SHAP) and permutation importance analysis to quantify the marginal contribution of each input feature. These explainability tools help identify which variables—across monthly, weekly, and daily frequencies—are most influential for each prediction horizon, providing actionable insights for practitioners and validating the design of the multi-frequency feature architecture.

3. EXPERIMENTAL SETUP AND RESULT

3.1 Experimental Setup

To evaluate the effectiveness of multi-frequency feature integration, we conduct separate forecasting experiments for coking coal and iron ore. The dataset spans from January 2016 to December 2024, and is split into a training set (2016/01–2022/12), a validation set (2023/01–2024/06), and a testing set (2024/07–2024/12). Forecasting models are trained to predict monthly average prices at three horizons: one, two, and three months ahead (N+1 to N+3). Model performance is evaluated using Mean Absolute Percentage Error (MAPE), computed on both validation and testing sets.

To investigate the impact of temporal resolution, we test four types of feature combinations: (1) using only monthly features, (2) combining monthly and weekly features, (3) combining monthly and daily features, and (4) integrating all three frequencies—monthly, weekly, and daily—into a single input vector. This design allows us to isolate the incremental value of each additional frequency level.

3.2 Results and Analysis

As shown in Table 1, models trained using only monthly features yield the highest error rates on the validation set for both coal and iron ore. For coal, the two-month-ahead forecast (N+2) reaches a MAPE of 2.98% using monthly features alone, whereas the full-frequency model—combining monthly, weekly, and daily data—reduces the error to 2.02% and achieves the lowest error at N+1 (1.11%). In the case of iron ore, adding daily features significantly improves short-term predictions, with the monthly-plus-daily configuration achieving the best performance at N+1 (0.99%). For N+3, the full-frequency model yields the lowest errors, suggesting that richer temporal information becomes increasingly valuable for longer-range forecasts.

The testing results, also summarized in Table 1, provide further validation of these patterns. For coal, although the full-frequency model demonstrates superior performance in more stable months (e.g., October and November 2024), it still struggles during highly volatile periods such as September, where error spikes to 29.3%. In contrast, the monthly-plus-daily model for iron ore delivers consistently strong performance across all six testing months, achieving the lowest MAPE in multiple periods (e.g., August: 12.9%, September: 14.2%, and October: 7.3%). This suggests that integrating cross-scale features not only improves average accuracy but also enhances resilience to market fluctuations.

The experimental results highlight the value of integrating features from multiple temporal resolutions. While monthly-only models serve as a basic baseline, augmenting them with weekly and daily signals leads to substantial improvements in accuracy. The full-frequency model, which combines inputs from all three timescales, consistently delivers the most balanced and robust performance across validation and testing sets, particularly for iron ore. These findings validate the hypothesis that capturing multi-scale temporal dynamics enhances the predictability of commodity prices.

Table 1 It summarizes the mean absolute percentage error (MAPE) results for different combinations of temporal features used in forecasting coking coal and iron ore prices.

Commodity	Freq.	Validation			Test					
		N+1	N+2	N+3	202407	202408	202409	202410	202411	202412
Coal	M	1.98	2.98	2.19	14.4	35.6	78.1	2.8	3.1	48.7
	M+W	1.54	2.28	2.45	7.0	29.6	57.6	10.5	10.8	38.4
	M+D	1.27	2.71	2.78	9.1	11.9	13.4	2.1	1.8	38.5
	M+W+D	1.11	2.02	2.63	19.0	10.8	29.3	6.5	6.2	29.3
Iron Ore	M	1.45	1.67	2.15	8.2	13.9	16.4	11.4	3.9	1.3
	M+W	1.74	1.32	1.69	25.3	19.4	14.6	14.0	2.9	0.2
	M+D	0.99	1.48	1.26	10.1	12.9	14.2	7.3	17.0	8.9
	M+W+D	1.05	1.47	1.23	14.5	24.4	18.8	9.8	9.0	1.3

The abbreviations M, W, and D represent monthly, weekly, and daily features, respectively. Models labeled as “M” use only monthly indicators, while “M+W” and “M+D” incorporate additional weekly or daily variables. The “M+W+D” configuration integrates all three temporal frequencies. Results are reported for three validation horizons (N+1 to N+3) and six test months from July to December 2024.

3.3 Model Explainability Analysis

To enhance model transparency and deepen understanding of the forecasting mechanism, we apply explainable AI (XAI) techniques to identify the most influential features across different temporal resolutions. The analysis is based on SHapley Additive exPlanations (SHAP) and permutation importance, focusing on the top 20 features extracted from each model configuration. We categorize these features by their originating data frequency—monthly (M), weekly (W), or daily (D)—to summarize their relative contributions and evaluate their role in improving forecast performance. This analysis enables us to answer not only which features are important, but also how temporal granularity affects the informativeness of each input. It offers practical guidance on which frequency combinations are more effective for modeling commodity prices under different conditions.

3.3.1 Important Factors in Coal Forecast

The XAI results for coking coal models reveal a clear dominance of monthly features, especially in models built on single-frequency (M) or monthly-plus-weekly (M+W) input sets. These models are primarily driven by slow-moving macroeconomic indicators such as steel product indices, industrial electricity consumption, and international hot-rolled coil prices—factors that reflect aggregated industrial demand and supply cycles. However, the introduction of daily features in the M+D and M+W+D configurations brings noticeable diversification in the sources of feature importance. While monthly indicators still contribute significantly, daily features like high-frequency market signals or short-term futures spreads begin to appear more frequently among the top predictors. This shift underscores the importance of integrating real-time market activity to enhance responsiveness to short-term price fluctuations.

Interestingly, in the fully integrated model (M+W+D), the share of total importance attributed to monthly features decreases. This suggests that some explanatory power previously captured by coarse indicators is now being redistributed to more volatile, event-driven variables. The redundancy or overlap among time scales implies that market sensitivity to near-term shocks is increasingly captured through daily and weekly signals—aligning with the operational need for real-time adjustments in coal procurement strategies. Moreover, the diversity of feature origins in the integrated models supports the hypothesis that coal prices, while largely macro-driven, can benefit from multi-scale inputs to better reflect temporal shifts in policy, demand, and energy mix transitions.

3.3.2 Important Factors in Iron Ore Forecast

In contrast to coal, iron ore models display a more balanced distribution of feature importance across all three time scales. Monthly features dominate in the baseline M and M+W models but gradually lose prominence when daily variables are incorporated. In the M+W+D configuration, daily and weekly features constitute a larger share of the top 20 contributors, emphasizing the relevance of short-term supply chain movements and trading activity in iron ore price formation. Specifically, daily features such as port inventory volumes, shipping delays, iron content benchmarks (e.g., 62% Fe or 58% Fe), and daily spot prices appear frequently in the top-ranking inputs. These features capture real-time logistics constraints and market volatility that are characteristic of the iron ore market, where seaborne trade disruptions or rapid inventory shifts can immediately influence price signals.

The more even contribution from daily, weekly, and monthly frequencies in iron ore forecasts suggests that this market responds to a broader range of temporal signals than coal. From policy shifts and macroeconomic indicators to inventory cycles and futures spreads, price dynamics are shaped by both long-term trends and short-term noise. This complexity reinforces the value of multi-frequency integration in capturing such layered interactions. In summary, the XAI results affirm that coking coal and iron ore exhibit distinct temporal sensitivities. Coal price forecasting benefits from macro-level indicators supplemented by event-driven signals, while iron ore forecasting requires more granular, real-time data reflecting the operational pulse of the global supply chain. These insights not only validate the model structure but also provide domain-relevant guidance for feature selection in future industrial applications.

4. CONCLUSION

This study presents a structured and interpretable framework for commodity price forecasting through the integration of multi-frequency features—namely, monthly, weekly, and daily data. By focusing on coking coal and iron ore, two of the most cost-sensitive inputs in steel production, we demonstrate that combining features across temporal scales enhances predictive performance and reveals meaningful insights into the structure of commodity markets. Empirical results from both validation and testing sets consistently show that models incorporating higher-frequency inputs outperform those using only monthly data. This performance gap is especially pronounced in short-term forecasts (N+1), where daily signals capture market volatility and short-lived events that macro-level indicators tend to miss. For longer-term horizons (N+2, N+3), the inclusion of weekly features contributes to improved stability and robustness, suggesting that medium-frequency data serve as a valuable bridge between long-term trends and real-time shocks. Notably, the full-frequency model (M+W+D) demonstrates the most balanced and resilient performance across a wide range of market conditions.

Our feature explainability analysis further reinforces these findings. SHAP and permutation importance rankings indicate that while monthly features remain influential, their relative importance diminishes as more granular data are introduced. In coal forecasting, macroeconomic indicators continue to play a dominant role, but daily features—such as short-term futures spreads—emerge as critical for identifying turning points and local volatility. In contrast, iron ore forecasting reveals a more even distribution of influential features across all frequencies, highlighting the market's sensitivity to logistical dynamics, port-level inventory shifts, and near-real-time pricing benchmarks.

Beyond numerical improvements, the proposed framework contributes methodologically by offering a unified approach to integrating and evaluating multi-scale inputs in commodity price forecasting. It addresses a key limitation in prior research—namely, the lack of systematic comparisons across different frequency combinations—and introduces a scalable architecture that balances accuracy with interpretability. The application of time-aware validation (via expanding window) and model-agnostic XAI methods ensures that the insights derived from the model are not only statistically sound but also actionable for industry practitioners.

From a practical standpoint, this study provides valuable implications for steelmakers, commodity traders, and procurement managers. In volatile markets where timing and information asymmetry significantly impact cost structures, access to interpretable and frequency-aware forecasts can support more informed sourcing decisions, contract structuring, and risk mitigation strategies. Furthermore, by identifying which temporal signals offer the greatest marginal utility at different horizons, this framework enables decision-makers to tailor their data acquisition and monitoring strategies accordingly.

Looking forward, future research could extend this framework in several directions. First, more advanced ensemble architectures or temporal attention mechanisms could be explored to better learn hierarchical dependencies across frequencies. Second, incorporating sentiment-based or unstructured textual signals (e.g., from market reports or news) could further enhance short-term responsiveness. Lastly, domain adaptation across regions or commodities could generalize this approach beyond the iron and steel sector. In summary, our work underscores the importance of temporal resolution in commodity price forecasting. By moving beyond single-frequency models and embracing the complementary nature of multi-scale data, we advance both the theoretical and applied dimensions of predictive analytics in raw material markets.

REFERENCES

1. Caldara, D., Cavallo, M., & Iacoviello, M. (2016). "Oil price elasticities and oil price fluctuations." *Journal of Monetary Economics*, 103, 1-18.
2. Pindyck, R. S. (2001). "The dynamics of commodity spot and futures markets: A primer." *The Energy Journal*, 22(3), 1-29.
3. Geman, H., & Ohana, S. (2009). "Time to build and the real options approach to capital budgeting." *Journal of Applied Corporate Finance*, 21(2), 98-107.
4. Matyjaszek, M., Fernández, P. R., Krzemień, A., Wodarski, K., & Valverde, G. F. (2019). "Forecasting coking coal prices by means of ARIMA models and neural networks, considering the transgenic time series theory." *Resources Policy*, 61, 283-292.
5. Zhang, K., Cao, H., Thé, J., & Yu, H. (2022). "A hybrid model for multi-step coal price forecasting using the decomposition technique and deep learning algorithms." *Applied Energy*, 306, 118011.
6. Guan, K., & Gong, X. (2023). "A new hybrid deep learning model for monthly oil prices forecasting." *Energy Economics*, 128, 107136.
7. Huang, L., et al. (2024). "Crude Oil Futures Price Forecasting Based on Variational and Empirical Mode Decompositions and Transformer Model." *Mathematics*, 12(24), 4034.
8. Molnar, C. (2022). *Interpretable Machine Learning* (2nd ed.).
9. Tsai, F.-S., Lo, H.-Y., Feng, C.-H., & Chang, K.-F. (2025). A dynamic dual-stage learning framework for interpretable coal and iron ore price trend prediction. *Paper presented at the SEAISI Conference & Exhibition* 2025.