Mixed-Frequency Modeling of Carbon Market Option Implied Volatility: A Multifactor Approach with GARCH-based and LSTM-based Frameworks

by

Ziyue Fang

An honors thesis submitted in partial fulfillment

of the requirements for the degree of

Bachelor of Science

Business and Economics Honors Program

NYU Shanghai

May 2025

Professor Marti G. Subrahmanyam
Professor Christina Wang
Professor Wendy Jin

Faculty Advisers

Professor Xin Zhou

Thesis Adviser

Contents

Acknowledgments

I would like to express my deepest gratitude to my thesis advisor, Dr. Xin Zhou, whose guidance introduced me to the carbon market and helped shape the direction of my research topic. His continuous support was instrumental in helping me complete this thesis step by step. I am also thankful to the Volatility Institute at NYU Shanghai, under Dr. Xin Zhou's leadership, for providing me with the opportunity to participate in a collaborative project on carbon markets. I appreciate the insightful suggestions I received from Mr. Xianzhen Wang and Mr. Haiping Wang, which helped refine the scope of my research during its early stages.

I would also like to thank the Business and Economics Honors Program and its dedicated faculty advisors — Professor Christina Wang, Professor Ye Jin, and Professor Marti G. Subrahmanyam — for organizing weekly seminars that enhanced my research and analytical skills.

Looking ahead, I will continue to deepen my understanding of financial markets and remain committed to applying rigorous research to address real-world challenges.

Abstract

This study investigates the forecasting of implied volatility (IV) for carbon options in the European Union Allowance (EUA) market using both econometric and machine learning methodologies. Given the carbon market's sensitivity to various macroeconomic and financial influences, we construct a series of volatility prediction models that integrate macroeconomic, commodity, and equity-related information. In addition to benchmark GARCH(1,1) and Long Short-Term Memory (LSTM) models, we extend the analysis by incorporating exogenous variables through the GARCH-MIDAS (Mixed Data Sampling) framework and further propose an LSTM-MIDAS architecture, which enables the integration of mixed-frequency data into a deep learning structure. Dynamic Principal Component Analysis (DPCA) is employed to reduce the dimensionality of factor groups while retaining their temporal dependencies.

Empirical results show that macroeconomic uncertainty—represented by the Economic Policy Uncertainty (EPU) index—significantly improves forecasting performance. Among machine learning models, the LSTM model incorporating commodity price factors (referred to as LSTM-COMO) achieves the highest accuracy, while GARCH-MIDAS-EPU outperforms other econometric alternatives. These findings highlight the importance of combining multifactor inputs and mixed-frequency techniques in modeling carbon option volatility.

Keywords: Carbon Market; Implied Volatility; GARCH-MIDAS; LSTM-MIDAS; Mixed-Frequency Data; Economic Policy Uncertainty; EUA Options

1. Introduction

1.1 Development of the Carbon Market and European Union Allowance (EUA)

The issue of global climate change has become increasingly critical in the past decades, prompting the international community to prioritize the reduction of greenhouse gas emissions. In response, various countries have been actively exploring the establishment of carbon market mechanisms as a means of achieving emission reduction targets. Carbon markets utilize market-based approaches to transform carbon emission allowances into tradable commodities, thereby incentivizing companies to voluntarily reduce their emissions. The fundamental principle behind carbon markets is to attribute economic value to emission allowances, enabling firms to make balanced decisions between environmental responsibility and economic cost.

The carbon market concept first gained traction as a policy instrument aimed at addressing climate change. One of the most significant milestones in the development of the carbon market was the adoption of the Kyoto Protocol in 1997, which set legally binding emission reduction targets for developed countries. The protocol laid the groundwork for the creation of emission trading systems (ETS), allowing countries with surplus emission allowances to trade them with countries exceeding their emission limits.

To accelerate the reduction of greenhouse gas emissions, the European Union Emissions Trading Scheme (EU ETS) was launched in 2005. As the world's most mature and largest carbon market, the EU ETS serves as a cornerstone of the EU's climate policy. The system is currently in its fourth phase (2021–present), during which the EU has implemented stricter carbon reduction targets, raising the 2030 emission reduction goal from 40% to 55%. Additionally, the annual reduction rate of the total volume of allowances has been increased, reflecting the EU's commitment to more aggressive climate action.

Within the EU ETS framework, European Union Allowances (EUA) function as the primary tradable unit, representing the right to emit one tone of carbon dioxide or its equivalent. EUA futures play a pivotal role not only in helping companies manage their emission risks but also in providing an effective financial instrument for speculative trading (Ye, 2021; Szolgayová et al., 2014). Consequently, EUA options are of considerable significance as they offer both risk management solutions for companies and trading opportunities for institutional and individual investors. Understanding the dynamics of EUA option pricing is essential for both market participants and policymakers, as it directly impacts the effectiveness of carbon market mechanisms.

1.2 Factors Influencing EUA Price Dynamics and Related Literature

Existing studies have identified numerous exogenous factors that may influence EUA price movements, including macroeconomic factors (Chevallier, 2009; Ding et al., 2024; Xiao, 2022), commodity market factors (Gronwald et al., 2011; Koenig, 2011), and equity market factors (Wen et al., 2020). However, many of these studies tend to focus on the effects of a single factor or a specific category of factors, and relatively few have attempted to comprehensively incorporate a wide range of potential variables to analyze the price behavior in the carbon market. This suggests that there may be room for further exploration by adopting a more holistic and multifactorial perspective.

In the context of carbon market options, various option pricing models have been developed to better capture the unique characteristics of these financial instruments. Traditional models, such as the Black-Scholes (B-S) model, have played a fundamental role in option pricing. However, recognizing the complexity of carbon market options, researchers have explored a range of enhanced methodologies. A significant portion of the literature has focused on combining the GARCH family of models with classical option pricing techniques, including the Black-Scholes model, fractional Brownian motion (FBM), and Monte Carlo simulations (Liu, 2021; Wu, 2022; Tang 2023). For example, Liu Z. applied a combination of the GARCH model and the B-S pricing formula to analyze the daily closing prices of EUA options, illustrating how volatility modeling can complement traditional pricing approaches.

In recent years, machine learning (ML) methods have gained significant traction in the field of option pricing due to their ability to model complex, non-linear relationships that traditional methods may not fully capture. Unlike classical approaches that rely on fixed assumptions, ML techniques offer a flexible and adaptive framework for analyzing dynamic market conditions. Several studies have demonstrated the effectiveness of ML models in improving option pricing accuracy and capturing intricate interactions among various influencing factors (Abrell et al., 2022; Zhang et al., 2025; Shang et al., 2025).

1.3 Predicting Carbon Option Prices Using Implied Volatility

In the field of financial research, implied volatility (IV) is widely recognized as an essential element for option pricing. Unlike historical volatility, which is derived from past price data, IV is forward-looking and reflects the market's expectations of future price fluctuations. This characteristic makes IV particularly valuable for analyzing options, as it inherently incorporates information from current market conditions and investor sentiment. Previous research has highlighted the predictive power of implied volatility over historical volatility (Beckers, 1981;

Chiras et al., 1978). Subsequent research has aimed to refine these models by incorporating larger sample sizes and accounting for additional variables, where researchers have increasingly sought to incorporate multifactorial perspectives, extending beyond traditional models to include sophisticated forecasting techniques based on machine learning and GARCH-type models. (Martens et al., 2004; Ahoniemi, 2009; Vrontos et al., 2021).

However, existing research on carbon market options seldom directly employs IV as the primary predictor for option pricing. Most studies focus on price dynamics or volatility modeling rather than utilizing IV itself for prediction. Considering the distinct nature of carbon market options—characterized by regulatory complexities, speculative interests, and their role in environmental compliance—leveraging IV as a central forecasting variable could possibly yield more accurate and insightful results. Hence, this study aims to bridge this gap by investigating the predictive potential of IV in the context of carbon option prices.

1.4 Overall Framework

To address the challenges in predicting carbon option prices, this study first explores the application of GARCH models and Long Short-Term Memory (LSTM) neural networks as benchmark models for IV prediction. Given the multifaceted nature of the carbon market, where IV is influenced by diverse exogenous factors, we integrate various types of factors, including commodity prices, equity indices, and macroeconomic indicators. Due to the large number of factors within the same category, we deployed Dynamic Principal Component Analysis (DPCA) to reduce dimensionality while preserving the temporal structure inherent in financial time series data.

However, some factors, such as Economic Policy Uncertainty (EPU) within macroeconomic indicators, do not have the same frequency as daily option data. Traditional GARCH models are not well-suited for handling mixed-frequency data. To address this issue, Engle et al. (2013) proposed the GARCH-MIDAS (Generalized Autoregressive Conditional Heteroskedasticity - Mixed Data Sampling) model. This model decomposes the volatility series into short-term and long-term components, with the short-term component captured by the GARCH model and the long-term component modeled through MIDAS to incorporate the influence of low-frequency variables. This approach effectively captures the impact of macroeconomic factors on daily volatility in financial markets.

In this study, we first adopt the GARCH-MIDAS framework to address the challenge of integrating mixed-frequency data in volatility forecasting. Additionally, we experiment with applying the MIDAS methodology within a neural network context by combining it with the LSTM model, thereby constructing the LSTM-MIDAS model. This attempt seeks to capture the influence of both high-frequency and low-frequency explanatory variables on carbon option volatility, utilizing the complementary strengths of both econometric and machine learning techniques.

By employing these advanced modeling techniques, we aim to investigate which factors significantly influence the implied volatility of carbon options and achieve more accurate predictions. This study contributes to the existing literature on carbon option pricing by integrating multifactor modeling and mixed-frequency analysis, providing a deeper understanding of how diverse market factors collectively shape volatility within the carbon trading environment. Ultimately, the findings from this research aim to support more informed decision-making and strategic planning for market participants, policymakers, and researchers engaged in the evolving landscape of carbon finance.

2. Data and Methodology

2.1 Data Selection

The European Union Allowance (EUA) is one of the most prominent carbon trading instruments under the EU Emissions Trading System (EU ETS), which represents the cornerstone of the region's carbon market. Considering data availability and market maturity, this study selects EUA option data spanning from October 12, 2022 to July 31, 2024. To ensure consistency and liquidity, the analysis focuses on at-the-money (ATM) call options, which are generally regarded as the most sensitive to implied volatility changes and are commonly used for volatility forecasting purposes. To improve data quality, observations with missing or zero transaction volume were removed. After applying these filters, a final sample of 463 daily observations was retained for modeling and empirical analysis.

In addition to option prices, several external variables are included to capture the broader market influences on implied volatility. The Economic Policy Uncertainty (EPU) index is retrieved from the official website of policyuncertainty.com, and the short-term risk-free rate is sourced from the European Central Bank (ECB). All the other data are obtained from Barchart.com.

The explanatory exogenous variables are grouped into three categories:

Macroeconomic Uncertainty Indicator:

- Economic Policy Uncertainty (EPU): Quantifies uncertainty in economic policymaking.

Commodity Price Factors:

- Brent Crude Oil Price: Global benchmark for energy costs.
- UK Natural Gas Futures (UKGAS): Industrial input affecting emissions.
- European Electricity Price Index (EEX Power): Proxy for energy demand and emissions.
- London Metal Exchange Index (LMEX): Reflects industrial activity and carbon usage.

- ICE Coal Price Index: Indicator of traditional fuel usage and carbon intensity.

Equity Market Factor:

- EURO STOXX Volatility Index (VSTOXX): Captures financial market uncertainty.

2.2 Main Methodology

2.2.1 Newton-Raphson Method for Implied Volatility Estimation

Implied volatility (IV) is a key unobservable parameter in option pricing models such as the Black-Scholes model. To estimate IV from observed market option prices, this study employs the Newton-Raphson iterative algorithm, a widely used numerical method that solves for the root of a nonlinear equation.

Given the Black-Scholes formula for a European call option:

$$C_{\mathrm{BS}}(S, K, T, r, \sigma) = S \cdot N(d_1) - Ke^{-rT} \cdot N(d_2)$$

where:

$$d_1 = \frac{\ln(S/K) + \left(r + \frac{1}{2}\sigma^2\right)T}{\sigma\sqrt{T}},$$

$$d_2 = d_1 - \sigma\sqrt{T}$$

and N(·) denotes the cumulative distribution function of the standard normal distribution. The Newton-Raphson method finds the implied volatility σ such that the theoretical price C_{BS} matches the market-observed price C_{market}. The update formula is:

$$\sigma_{n+1} = \sigma_n - \frac{C_{\mathrm{BS}}(\sigma_n) - C_{\mathrm{market}}}{\nu(\sigma_n)}$$

Where $\nu(\sigma_n)$ is the Vega of the option, i.e., the derivative of the Black-Scholes price with respect to volatility: $\nu(\sigma_n) = S \cdot \sqrt{T} \cdot \phi(d_1)$

Here, $\phi(\cdot)$ is the standard normal probability density function. The iteration continues until convergence is achieved within a predefined tolerance. This method is implemented to recover the daily implied volatility (IV) series used throughout the empirical analysis in this study.

2.2.2 GARCH Model

To capture the time-varying nature of financial market volatility, this study employs the Generalized Autoregressive Conditional Heteroskedasticity (GARCH) model, originally introduced by Bollerslev (1986). GARCH models are designed to model the conditional variance of time series data, making them particularly suitable for financial return series, which often exhibit volatility clustering and heteroskedasticity.

Let r_t denote the return series, typically calculated as the difference or percentage change in implied volatility. The mean equation of the model is expressed as:

$$r_t = \mu + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma_t^2)$$

Here, μ is the constant mean return, and σ_t^2 represents the time-varying conditional variance. The GARCH(1,1) specification models the variance as a function of past squared innovations and past variances:

$$\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2$$

where:

- $\omega>0$ is a constant term,
- $\alpha\geq0$ measures the short-run impact of past shocks (ARCH effect),
- $\beta\geq0$ captures the persistence of volatility (GARCH effect).

The GARCH model assumes that large shocks to volatility can persist over time, and the condition $\alpha+\beta<1$ is required to ensure stationarity of the variance process.

In this study, the GARCH(1,1) model serves as the benchmark for volatility modeling due to its simplicity and proven empirical performance in capturing volatility clustering.

2.2.3 GARCH-MIDAS Model

To incorporate low-frequency explanatory variables into high-frequency volatility modeling, Engle et al. (2013) proposed the GARCH-MIDAS (Mixed Data Sampling) model. This approach decomposes the conditional variance into two multiplicative components: a short-term component driven by GARCH dynamics, and a long-term component driven by low-frequency macroeconomic or market variables.

The model is specified as follows:

$$r_t = \mu + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma_t^2)$$

$$\sigma_t^2 = \tau_t \cdot g_t$$

Where:

- r_t captures the long-term volatility component affected by low-frequency data,
- g_t is the short-term component following a standard GARCH process:

$$g_t = \omega + \alpha \left(\frac{\varepsilon_{t-1}}{\sqrt{\tau_{t-1}}} \right)^2 + \beta g_{t-1}$$

The long-term component r_t is modeled using a MIDAS polynomial:

$$\ln(\tau_t) = \theta_0 + \sum_{k=1}^{K} \theta_k B(k; \varphi_1, \varphi_2) \cdot Z_{t-k}$$

Where:

- Z_{t-k} represents lagged low-frequency explanatory variables
- $B(k; \varphi_1, \varphi_2)$ is a normalized Beta weighting function that ensures recent lags are more heavily weighted.

In this study, the GARCH-MIDAS framework is employed to capture the influence of macroeconomic and other low-frequency variables (Economic Policy Uncertainty) on the daily implied volatility of carbon options

2.2.4 LSTM Model

The Long Short-Term Memory (LSTM) model is a specialized recurrent neural network designed to capture long-term dependencies in sequential data. It uses a set of gates—forget, input, and output—to regulate information flow and maintain memory over time, making it well-suited for modeling complex dynamics in financial volatility forecasting.

The architecture of an LSTM cell is composed of a series of gating mechanisms that regulate the flow of information through the network:

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

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

$$\text{Candidate state:} \quad \widetilde{C}_t = \tanh(W_C \cdot [h_{t-1}, x_t] + b_C)$$

$$\text{Cell state update:} \quad C_t = f_t \odot C_{t-1} + i_t \odot \widetilde{C}_t$$

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

$$\text{Hidden state:} \quad h_t = o_t \odot \tanh(C_t)$$

Here, x_t represents the input at time t, h_t the hidden state, and C_t the cell state. σ denotes the sigmoid activation function, and $\odot$ represents element-wise multiplication.

In this study, the LSTM model is applied to forecast implied volatility (IV) based on sequences of past values and relevant exogenous variables. The model parameters—including the number of units, learning rate, batch size, and dropout rates—are tuned to ensure optimal predictive performance and to prevent overfitting.

2.2.5 Extended LSTM-MIDAS Model

To address the challenge of incorporating mixed-frequency data into deep learning models, this study adopts the extended LSTM-MIDAS framework inspired by Kamolthip (2021). Kamolthip (2021) proposes an elegant solution by transforming low-frequency variables into frequency-aligned high-frequency representations using the MIDAS (Mixed Data Sampling) technique. This process involves applying a weighted lag structure (such as the normalized exponential Almon lag) to compress historical low-frequency values into a single high-frequency input vector. These transformed features are then concatenated with other high-frequency predictors and passed into the LSTM architecture, enabling the network to simultaneously capture long-term macroeconomic effects and short-term dynamics.

The MIDAS transformation for a low-frequency predictor $x_t^{(m)}$ sampled at frequency m is defined as:

$$z_t = \sum_{j=0}^{J} \omega_j \cdot x_{t-j/m}^{(m)} \quad \text{where} \quad \omega_j = \frac{\exp(\theta_1 j + \theta_2 j^2)}{\sum_{j=0}^{J} \exp(\theta_1 j + \theta_2 j^2)}$$

The resulting aligned feature z_t is then fed into the LSTM along with other predictors. This approach enables the model to capture both low-frequency macroeconomic signals and high-frequency market patterns. This extension allows for a comprehensive investigation of how various classes of exogenous information contribute to the volatility dynamics in the carbon options market.

2.2.6 Dynamic Principal Component Analysis (DPCA)

To address the high dimensionality and multicollinearity among input variables within each category (e.g., multiple commodity prices), this study adopts the Dynamic Principal Component Analysis (DPCA) method to extract the most informative components for forecasting. Unlike traditional Principal Component Analysis (PCA), which ignores the temporal dependence inherent in time series data, DPCA captures both cross-sectional and dynamic (time-lagged) correlations among variables, making it more suitable for financial and economic time series. DPCA assumes that the multivariate time series $\mathrm{X_t} = (X_{1t}, X_{2t}, \dots, X_{Nt})^\top$ can be decomposed into a lower-dimensional set of dynamic factors that preserve the structure of lagged covariances. The goal is to estimate the dynamic principal components (DPCs) that explain the maximum proportion of variance not only at time t, but also across a range of lags.

Let $\Gamma\,(k)$ be the autocovariance matrix of $\mathrm{X_t}$ at lag k. Then the spectral density matrix of the process is:

$$f(\omega) = \frac{1}{2\pi} \sum_{k=-\infty}^{\infty} \Gamma\,(k) e^{-i\omega k}$$

DPCA estimates the eigenvectors of this spectral density matrix, rather than the sample covariance matrix used in PCA. These eigenvectors are used to construct filters that extract dynamic principal components from $\mathrm{X_t}$, which are then used as inputs for forecasting models. In this study, DPCA is applied separately to the commodity price variables to derive the leading dynamic component, which captures the most relevant temporal information across multiple commodity series. The extracted DPC is subsequently used as a consolidated input in both GARCH-X and LSTM-based models to enhance predictive efficiency and mitigate overfitting due to variable redundancy.

3. Results

To evaluate the forecasting performance of various models for implied volatility (IV) in the carbon options market, this section presents the empirical results derived from both traditional econometric methods and machine learning techniques. The models include the benchmark GARCH(1,1), several GARCH-MIDAS extensions that incorporate exogenous information from mixed-frequency macroeconomic and market indicators, as well as LSTM-based models designed to capture nonlinear patterns in the data. Furthermore, an extended LSTM-MIDAS model is constructed to integrate low-frequency variables into the LSTM framework, addressing the challenges of mixed-frequency data.

The analysis proceeds in the following order: we begin with preliminary statistical diagnostics on the IV series, including tests for stationarity, normality, and the presence of ARCH effects. Next, we implement and evaluate the GARCH(1,1) model as a baseline. This is followed by the GARCH-MIDAS variants, which progressively introduce macroeconomic, commodity, and equity-related factors. Finally, we apply the LSTM-based models, including their MIDAS-enhanced versions, to assess whether deep learning models can offer improved volatility forecasting performance. All models are evaluated based on their ability to predict future implied volatility and reconstruct option prices, allowing for a comprehensive comparative analysis across different modeling approaches.

3.1 Preliminary Diagnostics for GARCH Modeling

3.1.1 Descriptive Statistics

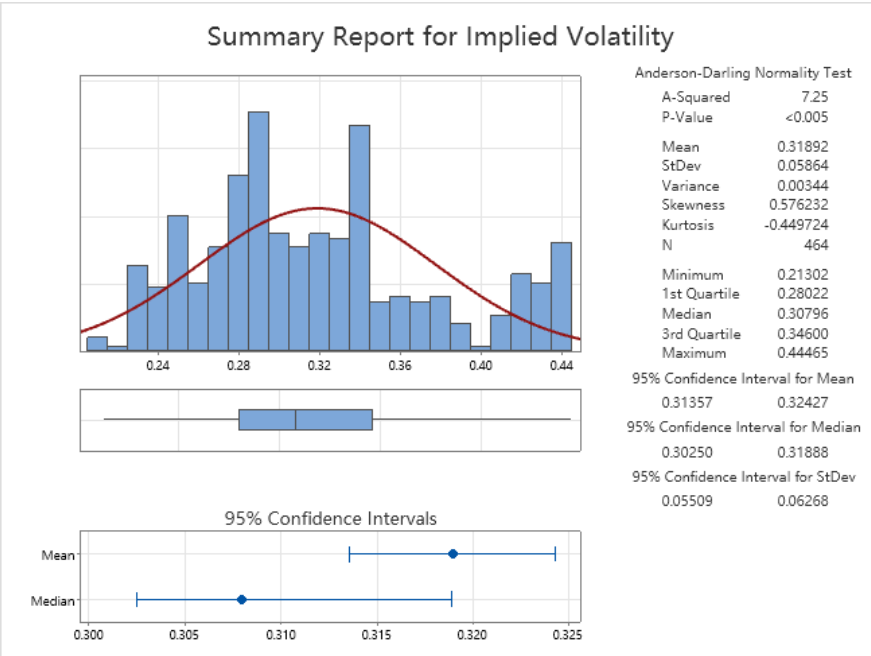

Figure 1: Summary Statistics

Figure 3.1 provides a comprehensive summary of the implied volatility (IV) series. The histogram suggests a moderately right-skewed distribution, which is further confirmed by the skewness value of 0.576 and kurtosis of -0.449, indicating the presence of a long right tail and a relatively flatter peak compared to the normal distribution. The IV values range between 0.213 and 0.445, with a mean of 0.3198 and a median of 0.3079. This slight discrepancy between the mean and median reflects the underlying asymmetry in the distribution.

3.1.2 Stationarity and White Noise Check

Table 1: Augmented Dickey-Fuller Test Results for IV and differenced IV

Series	ADF Statistic	P-value	Stationarity
IV	-1.79401	0.383	No
Diff(IV)	-17.2705	0.000	Yes

As shown in Table 1, the original implied volatility series fails the ADF test at the 5% significance level ($p = 0.383 > 0.05$), indicating non-stationarity. To address this, we take the first difference of the IV series. The differenced series passes the ADF test ($p = 0.000 < 0.05$), confirming stationarity.

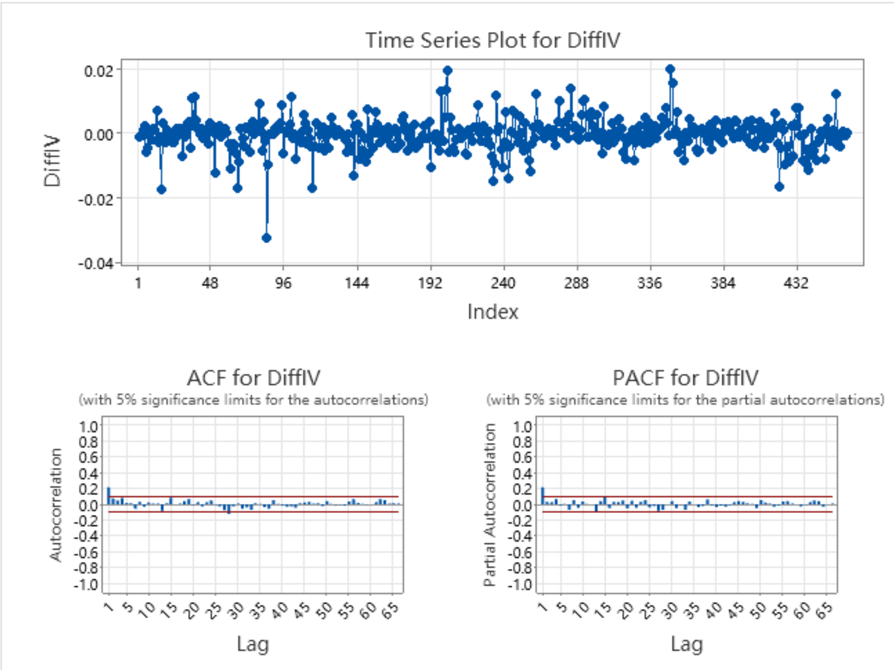

Figure 2: Time Series Plot, ACF and PACF for Diff IV

Figure 2 illustrates the differenced IV series along with its ACF and PACF plots. The time series plot of Diff(IV) fluctuates around a constant mean, showing no obvious trend or seasonality. Additionally, the ACF and PACF suggest that autocorrelation is weak and most values lie within the 95% confidence bands, which is consistent with a white noise process.

Table 2: Ljung-Box Statistics

Lag	Chi-Square	Degree of Freedom	P-value
12	8.55	10	0.575
24	22.33	22	0.441
36	34.91	34	0.425
48	42.20	46	0.632

To further evaluate whether the differenced IV series exhibits white noise behavior, we conduct the Modified Box-Pierce (Ljung-Box) Chi-Square test at multiple lags. The results in Table 2 show that all p-values are greater than 0.05, indicating that we fail to reject the null hypothesis.

This suggests that the differenced IV series does not exhibit significant autocorrelation and is approximately white noise.

3.1.3 Conditional Heteroskedasticity and ARCH Effect Test

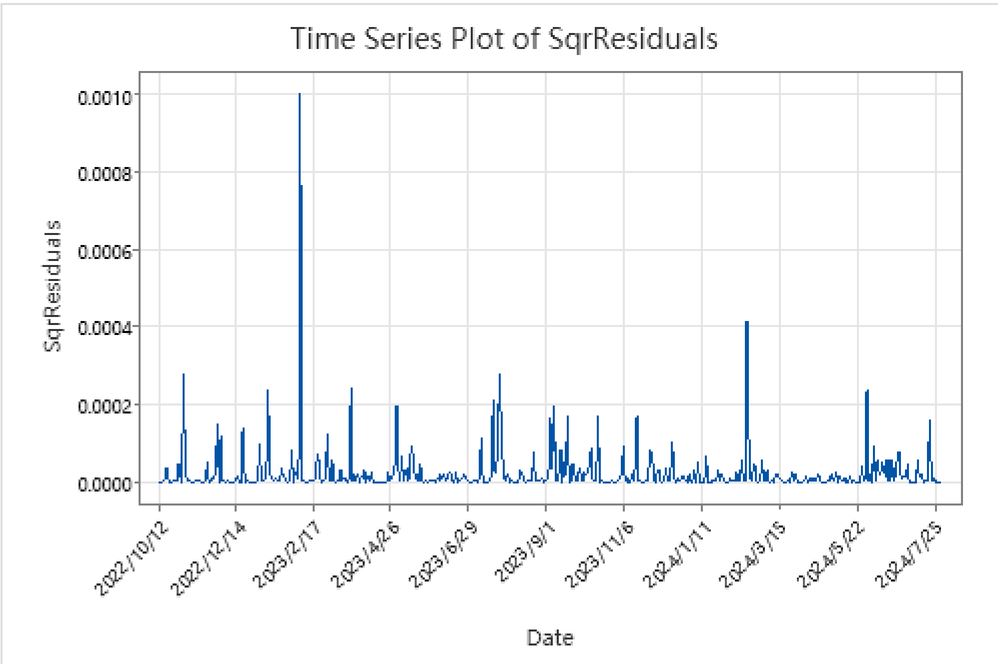

Figure 3: Time Series Plot of Residuals' Square

Table 3: ARCH-LM Test Results

Lag	LM Statistics	P-value
1	52.31	0.000
2	58.19	0.000
3	60.24	0.000
5	70.15	0.000
10	98.21	0.000

To examine the presence of heteroskedasticity in the differenced implied volatility (IV) series, a time series plot of the squared residuals is presented in Figure 3. The plot displays visible clustering of volatility over time, suggesting potential conditional heteroskedasticity.

To formally test for the presence of ARCH effects, the Lagrange Multiplier (LM) test was performed across various lag orders. From Table 3, the LM statistics are statistically significant at all selected lags, with corresponding p-values below the 1% significance level. These results confirm that the error variance in the differenced IV series is time-dependent, thereby validating the use of GARCH-type models in subsequent volatility modeling and forecasting.

3.1.4 GARCH(1,1) Model Estimation

Following confirmation of significant ARCH effects, a GARCH(1,1) model was fitted to the differenced implied volatility (DiffIV) series.

Table 4: GARCH(1,1) model results

Coefficient	Estimate	Std. Error	t value	Pr(>|t|)	Significance
α_0	5.524e-05	3.621e-05	1.525	0.127	
α_1	8.019e-02	1.440e-02	5.568	2.57e-08	***
β_1	9.155e-01	1.518e-02	60.293	< 2e-16	***

The GARCH(1,1) model estimated for the differenced implied volatility series is expressed as follows:

$$\sigma_t^2 = 5.524 \times 10^{-5} + 0.08019\, \varepsilon_{t-1}^2 + 0.9155\, \sigma_{t-1}^2$$

where $\alpha_0 = 5.524\times10{-}5 > 0$, α_1, $\beta_1 > 0$, and $\alpha_1 + \beta_1 = 0.08019 + 0.9155 = 0.99569 < 1$, indicating that the volatility process is stationary.

According to the estimated GARCH(1,1) model, the volatility for the next 14 trading days is forecasted and will be used to reconstruct future implied volatility values. These results serve as a benchmark for evaluating alternative models introduced in the following sections.

3.2 Forecast results with GARCH models

To capture the volatility dynamics of carbon option markets, this section applies a series of GARCH-based models to forecast implied volatility (IV). Starting from a benchmark GARCH(1,1) framework, we extend the modeling approach by incorporating various external variables to assess their contribution to forecasting performance. Depending on the frequency of the incorporated factors, we distinguish between standard GARCH-X models (for high-frequency variables) and the GARCH-MIDAS framework (for low-frequency macroeconomic indicators). The MIDAS (Mixed Data Sampling) structure enables the integration of lower-frequency information—such as monthly policy indices—into daily volatility modeling. Specifically, the models implemented include:

- A benchmark GARCH(1,1) model without external regressors;
- GARCH models that incorporate single exogenous information such as macroeconomic uncertainty;
- A multifactor GARCH-MIDAS model that integrates multiple factors derived from macroeconomic, commodity, and equity indicators.

Each model's forecasted IV series is then input into the Black-Scholes pricing formula to generate predicted option prices, allowing for a comparative assessment of model accuracy under different informational structures.

3.2.1 Benchmark GARCH Model

Table 5: IV Forecast with GARCH

Day1	Day2	Day3	Day4	Day5	Day6	Day7
0.2432	0.2522	0.2482	0.2453	0.2489	0.2492	0.2392
Day8	Day9	Day10	Day11	Day12	Day13	Day14
0.2424	0.2412	0.2403	0.2395	0.2358	0.2385	0.2389

The benchmark model employs a standard GARCH(1,1) specification without external regressors to forecast implied volatility. As shown in Table 5, the model produces 14-day ahead IV forecasts based solely on historical volatility patterns. These predicted IV values are then used as inputs in the Black-Scholes pricing formula to reconstruct option prices over the forecast horizon.

Table 6: Forecasted Price with B-S model (GARCH)

Date	1	2	3	4	5	6	7
Forecast	2.95	3.25	3.08	2.73	2.98	2.93	3.23
Actual	3.88	3.58	2.70	3.78	3.50	3.82	3.21
Date	8	9	10	11	12	13	14

Forecast	3.18	4.02	3.67	3.88	3.87	4.02	4.09
Actual	4.65	3.78	4.18	4.21	4.47	4.75	4.61

Table 7: Evaluation Metrics for GARCH Model

MSE	MAE	MAPE
4.58	2.52	11.23%

Table 6 presents a comparison between the predicted and actual option prices during the same 14-day forecast window. While the model successfully captures the overall direction of price movement, deviations from actual values are observed. Evaluation metrics shown at Table 7 indicate a mean squared error (MSE) of 4.58, a mean absolute error (MAE) of 2.52, and a mean absolute percentage error (MAPE) of 11.23%, reflecting a reasonable level of forecasting accuracy for a baseline model.

3.2.2 GARCH-MIDAS and GARCH-X Models

To enhance the predictive power of volatility modeling, this section explores a suite of extended GARCH frameworks that incorporate macroeconomic, commodity, and equity-based explanatory variables. In particular, we distinguish between two modeling approaches based on the frequency characteristics of the external variables: GARCH-MIDAS and GARCH-X.

The macroeconomic component is represented by the Economic Policy Uncertainty (EPU) index, a low-frequency monthly variable. To properly integrate this mixed-frequency data, we apply the GARCH-MIDAS (Mixed Data Sampling) model, which decomposes volatility into short- and long-term components and allows for low-frequency drivers of high-frequency volatility.

In contrast, both commodity and equity variables are available at the daily level. Commodity factors include Brent Crude Price, UK Gas, EEX Power, LMEX, and ICE Coal Price. To address multicollinearity and capture the underlying structure across multiple series, Dynamic Principal

Component Analysis (DPCA) is used to extract the leading factor. The equity market is represented by the EURO STOXX Volatility Index (VSTOXX). Given their high frequency, both commodity and equity variables are incorporated using the GARCH-X model, which allows daily exogenous regressors to influence conditional variance dynamics directly.

The models implemented include:

- GARCH-MIDAS with EPU only,
- GARCH-X with commodity factors (COMO),
- GARCH-X with equity volatility (EQU),
- A combined model incorporating all three factors using GARCH-MIDAS-X (EPU-COMO-EQU), where EPU is treated through MIDAS and the others through standard exogenous terms.

Table 8: Evaluation Metrics for Extended GARCH Models

Model Type	MSE	MAE	MAPE
GARCH-MIDAS-EPU	3.52	1.87	8.28%
GARCH-COMO	4.25	2.24	10.40%
GARCH-EQU	5.92	3.25	14.72%
GARCH-MIDAS-EPU-COMO-EQU	3.98	2.04	9.42%

As shown in Table 8, the inclusion of EPU significantly improves the model’s accuracy across all evaluation metrics, achieving the lowest MSE (3.52), MAE (1.87), and MAPE (8.28%). In contrast, the EQU-based model performs the worst, suggesting that equity market volatility alone may not be a reliable predictor of implied volatility in the carbon market. The model based on commodity prices performs moderately well, and the all-factor GARCH-MIDAS variant achieves a balanced improvement across metrics.

These results highlight the pronounced influence of macroeconomic uncertainty on carbon option volatility, likely due to the carbon market's sensitivity to policy shifts and regulatory announcements. This finding is consistent with previous studies (Li et al., 2022; Yan et al., 2021; Adediran et al., 2023), which underscore the role of macro-level shocks in shaping price dynamics in emission trading systems.

In summary, the empirical results demonstrate that incorporating appropriate exogenous information—especially low-frequency macroeconomic indicators such as EPU—substantially enhances the accuracy of volatility forecasts in the carbon options market. The GARCH-MIDAS framework proves especially effective in capturing the long-term influence of macroeconomic uncertainty. These findings reinforce the importance of macroeconomic and policy-driven signals in shaping volatility in emission trading systems.

3.3 Forecast results with LSTM models

To establish a benchmark for machine learning-based volatility forecasting, a univariate Long Short-Term Memory (LSTM) neural network model is implemented using only the historical values of implied volatility (IV) as input. The model is designed to capture the temporal dependencies within the IV series and generate forward-looking volatility estimates.

3.3.1 Benchmark LSTM Model

The benchmark LSTM model is constructed with one hidden LSTM layer comprising 50 units, followed by a dense output layer with a single neuron. The model is trained over 50 epochs with a batch size of 16, using the Adam optimizer and mean squared error (MSE) as the loss function. Input data is scaled using MinMax normalization to enhance learning efficiency.

Table 9: IV Forecast with LSTM

Day1	Day2	Day3	Day4	Day5	Day6	Day7
0.2498	0.2525	0.2502	0.2476	0.2467	0.2470	0.2501
Day8	Day9	Day10	Day11	Day12	Day13	Day14
0.2502	0.2513	0.2500	0.2487	0.2490	0.2497	0.2499

Table 10: Forecasted Price with B-S model (LSTM)

Date	1	2	3	4	5	6	7
Forecast	3.46	4.02	3.56	3.34	3.28	3.45	3.87
Actual	3.88	3.58	2.70	3.78	3.50	3.82	3.21
Date	8	9	10	11	12	13	14
Forecast	4.05	4.20	3.97	3.66	3.78	4.02	4.09
Actual	4.65	3.78	4.18	4.21	4.47	4.75	4.61

Table 11: Evaluation Metrics for LSTM Model

MSE	MAE	MAPE
3.24	1.48	7.48%

As shown in Table 9, the model yields a relatively smooth IV forecast trajectory. These predicted IV values are then substituted into the Black-Scholes formula to compute option prices (Table 10). The resulting prediction performance metrics, summarized in Table 11, include an MSE of 3.24, MAE of 1.48, and MAPE of 7.48%, indicating a solid baseline performance.

3.2.2 Extended LSTM Models

Beyond the benchmark model, this section investigates several extended LSTM architectures that incorporate additional explanatory factors. In the LSTM-COMO and LSTM-EQU models, commodity and equity-related variables are introduced as direct inputs to the LSTM network, aiming to evaluate their potential to enhance volatility forecasting.

To incorporate lower-frequency macroeconomic information, this study adopts the LSTM-MIDAS framework. While MIDAS structures have been widely explored in econometric volatility modeling, their integration with deep learning models remains rare in the existing literature. Following the approach proposed by Kamolthip (2021), this study applies a transformation that converts low-frequency sequences into high-frequency vectors using MIDAS weighting functions, enabling the LSTM network to effectively capture the influence of macro-level indicators.

Two MIDAS-based variants are constructed: LSTM-MIDAS-EPU and LSTM-MIDAS-EPU-COMO-EQU, with the latter combining macroeconomic, commodity, and equity information.

Table 12: Evaluation Metrics for Extended LSTM Models

Model Type	MSE	MAE	MAPE
LSTM-MIDAS-EPU	3.89	1.89	9.28%
LSTM-COMO	3.02	1.45	6.52%
LSTM-EQU	3.59	1.99	9.33%
LSTM-MIDAS-EPU-COMO-EQU	3.44	1.53	8.12%

As presented in Table 12, the LSTM-COMO model exhibits the best performance across all metrics, achieving the lowest MSE (3.02), MAE (1.45), and MAPE (6.52%). This indicates that commodity prices provide informative short-term signals, which LSTM models can effectively exploit. Although the MIDAS-based models also perform reasonably well, the additional benefit over direct high-frequency inputs appears limited in this context.

3.4 Comparative Performance Analysis

Table 13: Evaluation Metrics for All Models

Model Type	MSE	MAE	MAPE
GARCH	4.58	2.52	11.23%
GARCH-MIDAS-EPU	3.52	1.87	8.28%

GARCH-COMO	4.25	2.24	10.40%
GARCH-EQU	5.92	3.25	14.72%
GARCH-MIDAS-EPU-COMO-EQU	3.98	2.04	9.42%
LSTM	3.24	1.48	7.48%
LSTM-MIDAS-EPU	3.89	1.89	9.28%
LSTM-COMO	3.02	1.45	6.52%
LSTM-EQU	3.59	1.99	9.33%
LSTM-MIDAS-EPU-COMO-EQU	3.44	1.53	8.12%

To provide a holistic evaluation of forecasting performance, Table 13 presents the Mean Squared Error (MSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE) across all GARCH-based and LSTM-based models. Overall, models that incorporate additional explanatory factors tend to outperform their benchmark counterparts, highlighting the importance of integrating relevant market information into volatility modeling.

Among all specifications, the LSTM-COMO model achieves the lowest prediction errors, with an MSE of 3.02, MAE of 1.45, and MAPE of 6.52%, suggesting that commodity price movements are particularly informative for short-term volatility in the carbon options market. The baseline LSTM model also performs strongly, benefiting from its ability to capture nonlinear temporal dynamics.

In the GARCH family, GARCH-MIDAS-EPU demonstrates superior performance, significantly reducing errors compared to the benchmark GARCH model. This again underscores the relevance of macroeconomic uncertainty—especially policy-related shocks—in shaping volatility patterns. On the contrary, models that rely solely on equity market information (GARCH-EQU and LSTM-EQU) exhibit relatively poor accuracy across all three metrics.

Interestingly, while combining all factor categories yields moderately strong results (e.g., GARCH-MIDAS-EPU-COMO-EQU and LSTM-MIDAS-EPU-COMO-EQU), the marginal improvements over single-factor models suggest potential redundancy or noise when aggregating disparate information sources.

In summary, models that carefully integrate targeted exogenous factors—notably commodity prices and macroeconomic uncertainty—demonstrate clear advantages in improving the accuracy of implied volatility and option price forecasts. These findings emphasize the value of hybrid approaches that combine financial domain knowledge with flexible machine learning techniques.

4. Conclusion & Discussion

This paper presents an in-depth investigation into the forecasting of implied volatility (IV) for European Union Allowance (EUA) options by integrating both econometric models and machine learning methodologies. Through the construction of benchmark and extended versions of GARCH and LSTM models, the study systematically examines how diverse sources of market information—including macroeconomic uncertainty, commodity prices, and equity market volatility—contribute to the modeling of IV in the carbon financial market.

The results reveal that incorporating carefully selected exogenous variables significantly enhances forecasting performance relative to benchmark specifications. In particular, macroeconomic uncertainty, represented by the EPU index, proves highly effective when modeled within the GARCH-MIDAS framework. This highlights the carbon market's sensitivity to policy-driven signals and regulatory expectations. From the machine learning perspective, commodity-related factors emerge as the most informative, with the LSTM-COMO model delivering the most accurate predictions across all evaluation metrics. This suggests that non-linear relationships and short-term fluctuations in energy-related commodities play a crucial role in shaping carbon option volatility.

Notably, the combination of all factor categories does not always result in superior performance. While multifactor models achieve reasonable accuracy, the marginal gain over single-factor models is often limited. This may point to overlapping signals or the introduction of noise when aggregating heterogeneous variables without further selection or refinement.

These findings offer several contributions to the field. First, they confirm the utility of implied volatility as a predictive target in carbon markets, an area where research remains relatively sparse.

Second, they illustrate how combining economic theory with data-driven modeling—especially through the integration of MIDAS structures into both statistical and deep learning models—can provide flexible yet interpretable tools for volatility forecasting. Lastly, the study underscores the importance of factor relevance over factor quantity: careful curation of input variables, especially those capturing macroeconomic and commodity-specific shocks, is more valuable than broad, undifferentiated inclusion.

Looking ahead, future research could refine factor selection using more advanced techniques such as attention mechanisms, explore the implications of policy regime shifts, or apply similar modeling strategies to other emerging environmental derivatives. As the carbon finance ecosystem expands, the ability to forecast volatility with precision will become increasingly vital for traders, regulators, and institutional participants alike.

Reference

Abrell, J., Kosch, M., & Rausch, S. (2022). How effective is carbon pricing?—A machine learning approach to policy evaluation. Journal of Environmental Economics and Management, 112, 102589.

Adediran, I. A., & Swaray, R. (2023). Carbon trading amidst global uncertainty: The role of policy and geopolitical uncertainty. Economic Modelling, 123, 106279.

Ahoniemi, K. (2009). Modeling and forecasting implied volatility. Helsinki School of Economics.

Beckers, S. (1981). Standard deviations implied in option prices as predictors of future stock price variability. Journal of Banking & Finance, 5(3), 363-381.

Chevallier, J. (2009). Carbon futures and macroeconomic risk factors: A view from the EU ETS. Energy Economics, 31(4), 614-625.

Chiras, D. P., & Manaster, S. (1978). The information content of option prices and a test of market efficiency. Journal of Financial Economics, 6(2-3), 213-234.

Ding, X., Ma, G., & Cao, J. (2024). The Emission-Reduction Effect of Green Demand Preference in Carbon Market and Macro-Environmental Policy: A DSGE Approach. Sustainability, 16(16), 6741.

Gronwald, M., Ketterer, J., & Trück, S. (2011). The relationship between carbon, commodity and financial markets: A copula analysis. Economic Record, 87, 105-124.

Kamolthip, S. (2021). Macroeconomic forecasting with LSTM and mixed frequency time series data. arXiv preprint arXiv:2109.13777.

Koenig, P. (2011). Modelling correlation in carbon and energy markets.

Li, X., Li, Z., Su, C. W., Umar, M., & Shao, X. (2022). Exploring the asymmetric impact of economic policy uncertainty on China's carbon emissions trading market price: do different types of uncertainty matter?. Technological Forecasting and Social Change, 178, 121601.

Liu, Z., & Huang, S. (2021). Carbon option price forecasting based on modified fractional Brownian motion optimized by GARCH model in carbon emission trading. The North American Journal of Economics and Finance, 55, 101307.

Martens, M., & Zein, J. (2004). Predicting financial volatility: High‐frequency time‐series forecasts vis‐à‐vis implied volatility. Journal of Futures Markets: Futures, Options, and Other Derivative Products, 24(11), 1005-1028.

Shang, D., Pang, Y., & Wang, H. (2025). Carbon price fluctuation prediction using a novel hybrid statistics and machine learning approach. Energy, 135581.

Szolgayová, J., Golub, A., & Fuss, S. (2014). Innovation and risk-averse firms: Options on carbon allowances as a hedging tool. Energy Policy, 70, 227-235.

Tang, J. (2023). Option Pricing Practices Based on the BSM Model in Carbon Markets. Highlights Bus. Econ. Manag., 21, 21-26.

Vrontos, S. D., Galakis, J., & Vrontos, I. D. (2021). Implied volatility directional forecasting: a machine learning approach. Quantitative Finance, 21(10), 1687-1706.

Wen, F., Zhao, L., He, S., & Yang, G. (2020). Asymmetric relationship between carbon emission trading market and stock market: Evidences from China. Energy Economics, 91, 104850.

Wu, H., Hua, H., Chen, X., & Jiang, Y. (2022, December). Carbon emission option pricing via Monte Carlo method. In Frontier Academic Forum of Electrical Engineering (pp. 1015-1025). Singapore: Springer Nature Singapore.

Xiao, B., Fan, Y., Guo, X., Voigt, S., & Cui, L. (2022). Effects of linking national carbon markets on international macroeconomics: An open-economy E-DSGE model. Computers & Industrial Engineering, 169, 108166.

Ye, J., & Xue, M. (2021). Influences of sentiment from news articles on EU carbon prices. Energy Economics, 101, 105393.

Yu, J., Shi, X., Guo, D., & Yang, L. (2021). Economic policy uncertainty (EPU) and firm carbon emissions: evidence using a China provincial EPU index. Energy Economics, 94, 105071.

Zhang, X., Li, Z., Zhao, Y., & Wang, L. (2025). Carbon trading and COVID-19: a hybrid machine learning approach for international carbon price forecasting. Annals of Operations Research, 345(2), 1267-1295.