EAI Endorsed Transactions
on Energy Web

Research Article EAI.EU

AI-Driven LSTM-Copula Hybrid Model for Joint Risk Dependence Modeling in Carbon–Electricity Portfolio Management: Implications for Grid Cost-Effectiveness and Stability

Runxin Hua[1,*]

[1] Northeast Agricultural University, Harbin, China

Abstract

INTRODUCTION: The increasing coupling between carbon emission trading and electricity markets creates significant joint risk challenging grid cost-effectiveness and stability. Existing approaches apply LSTM and Copula models separately, lacking a unified framework capturing both non-linear temporal dynamics and asymmetric tail dependence.

OBJECTIVES: This paper proposes an end-to-end LSTM-Copula hybrid model integrating deep learning-based marginal modeling with time-varying Copula dependence estimation for joint risk measurement and optimal portfolio allocation.

METHODS: The framework employs LSTM-GARCH for conditional mean and volatility modeling, EVT-GPD for tail fitting, and probability integral transform to obtain uniform variates. A time-varying t-Copula with DCC-type evolution captures dynamic joint dependence. Monte Carlo simulation estimates VaR and CVaR, followed by Min-CVaR portfolio

RESULTS: The LSTM-GARCH model achieves RMSE reductions of 42.2% and 38.0% for carbon and electricity price prediction versus standalone GARCH. The integrated model attains a VaR failure rate of 5.2% at the 95% confidence level, outperforming GARCH-Copula, GARCH-Normal, and Historical Simulation in Kupiec and Christoffersen backtesting.

CONCLUSION: The proposed model provides a unified framework for carbon–electricity joint risk modeling, insights for AI-driven energy portfolio optimization and grid stability enhancement.

Keywords: carbon–electricity portfolio, Copula, CVaR, energy storage, grid stability, LSTM, risk dependence, VaR

Received on 12 April 2026, accepted on 29 June 2026, published on 20 July 2026

doi: 10.4108/ ew.12583

1. Introduction

The global transition toward low-carbon energy systems has intensified the interdependence between carbon emission trading schemes (ETS) and electricity markets. As carbon costs are directly embedded in power generation expenses, the price dynamics of these two markets exhibit significant co-movement, risk transmission, and tail dependence. Wang et al. [1] demonstrated that volatility spillover among Chinese carbon, energy, and electricity markets is significant but asymmetric, with hedging effectiveness varying across market regimes. Fang [2] confirmed that electricity markets tend to act as net risk receivers while carbon markets function as net risk

*Corresponding author. Email: 18595229365@163.com

transmitters. Yuan et al. [3] further revealed that electricity-carbon market coupling leads to substantial fluctuations in marginal clearing prices, increasing power supply costs and necessitating compensation mechanisms for flexible generation resources. Cao et al. [4] provided empirical evidence from China's electricity sector showing that carbon emission trading significantly alters plant-level generation decisions and marginal abatement costs. Understanding and quantifying the joint risk dependence between carbon and electricity assets is therefore essential for both regulatory oversight and investment portfolio management aimed at enhancing grid cost-effectiveness and operational stability [5].

Meanwhile, the rapid deployment of artificial intelligence (AI) in energy systems has opened new avenues for optimizing hybrid energy storage and grid operations. Zhang and Strbac [6] provided a comprehensive review of AI applications for energy storage, highlighting that deep learning and reinforcement learning techniques substantially improve storage scheduling, cost reduction, and grid stability. Wang et al. [7] demonstrated that deep reinforcement learning enables multi-objective energy dispatch in wind-solar-thermal-storage hybrid systems, achieving simultaneous reductions in fuel costs and carbon emissions. Sepehrzad et al. [8] applied deep reinforcement learning to control active networked microgrids with multi-level battery energy storage participation, significantly enhancing computational efficiency and frequency regulation. These advances underscore the potential of AI-driven approaches in managing the growing complexity of energy portfolios, yet a critical gap remains in integrating AI-based predictive modeling with rigorous statistical frameworks for joint risk quantification across carbon and electricity markets.

Existing literature has established that the risk spillover between carbon and electricity markets is bidirectional but asymmetric, with dependence structures that are highly time-varying. Luo et al. [5] analyzed dynamic risk spillover effects using TVP-VAR-DY models, revealing that spillover intensity amplifies significantly during crisis periods. Naeem and Arfaoui [9] explored downside risk dependence across energy markets and found that tail dependence strengthens substantially during episodes of market crises such as the COVID-19 pandemic and the Russia-Ukraine conflict. Wu et al. [10] further demonstrated that high-frequency spillovers dominate short-term dynamics while structural linkages persist at lower frequencies. These stylized facts necessitate modeling frameworks that can simultaneously capture non-linear temporal dynamics, heavy-tailed marginal distributions, and time-varying asymmetric dependence.

In the prediction literature, LSTM-based hybrid models have shown superior performance in carbon and electricity price forecasting by capturing complex non-linear patterns and long-range dependencies. Huang et al. [11] proposed a hybrid GARCH-LSTM model that combines volatility modeling with deep learning, achieving notable improvements over standalone methods. Cheng et al. [12] developed a decomposition-enhanced LSTM framework that substantially reduces forecasting errors. Pentsos et al. [13] introduced a hybrid LSTM-Transformer architecture for power load forecasting that leverages both sequential memory and self-attention mechanisms, achieving state-of-the-art accuracy across multiple time horizons. Nauman et al. [14] proposed a genetic algorithm-optimized LSTM model for concurrent electricity price and load forecasting in smart grids, demonstrating enhanced robustness in real-time prediction scenarios. Meanwhile, Copula-based risk models have been widely adopted for multi-asset dependence modeling and VaR/CVaR estimation in energy and carbon markets [15], [16]. Manner et al. [17] applied dynamic vine Copula models to forecast the joint distribution of Australian electricity prices, showing that time-varying dependence structures are essential for accurate risk assessment. However, most existing studies apply LSTM and Copula separately — LSTM for time-series prediction and Copula for cross-sectional dependence — without integrating them into a unified framework [11], [15]. The marginal distribution in conventional Copula-GARCH models relies solely on GARCH-type specifications, which are limited in capturing the non-linear dynamics inherent in carbon and electricity returns.

To bridge this gap, this paper proposes an LSTM-Copula hybrid model that unifies deep learning-based marginal modeling with Copula-based dependence estimation for joint risk measurement in carbon–electricity portfolios. By providing accurate joint risk quantification, the proposed framework supports AI-driven energy portfolio optimization that can enhance grid cost-effectiveness and improve power system stability through informed investment and hedging decisions. The main contributions are threefold:

This paper constructs an end-to-end LSTM-GARCH-EVT framework for marginal distribution modeling, where LSTM captures the conditional mean and GARCH models the conditional volatility, followed by EVT-GPD tail fitting.

This approach overcomes the non-linearity limitations of pure GARCH specifications, drawing on the established effectiveness of EVT for heavy-tailed electricity price modeling. This paper employs a time-varying t-Copula to characterize the dynamic, asymmetric, and tail-dependent joint distribution between carbon and electricity returns, reflecting the empirical evidence of intensified coupling during market stress periods.

This paper integrates the joint distribution model into a Min-CVaR portfolio optimization framework, providing a more robust risk management tool compared to traditional VaR-based or mean-variance approaches. This framework aligns with recent CVaR-based energy storage and renewable energy investment optimization methods, offering actionable implications for grid cost-effectiveness.

2. Related work

2.1. Risk Dependence Between Carbon and Electricity Markets

The risk relationship between carbon and electricity markets has received growing attention in recent years. Wang et al. [1] examined volatility spillover and hedging strategies among Chinese carbon, energy, and electricity markets, finding significant but asymmetric bidirectional risk transmission with hedging effectiveness that varies across market regimes. Fang [2] studied risk spillovers between electricity and carbon markets using frequency-domain methods, confirming that electricity markets tend to be net risk receivers in the medium to long term. Luo et al. [5] analyzed dynamic risk spillover effects using TVP-VAR-DY models, revealing that spillover intensity is time-varying and amplified during crisis periods such as the COVID-19 pandemic. Naeem and Arfaoui [9] explored downside risk dependence across energy markets including electricity, conventional energy, carbon, and clean energy, finding that tail dependence strengthens substantially during episodes of market crises. Wu et al. [10] further demonstrated that return and volatility connectedness among carbon and energy markets exhibits distinct patterns across time and frequency domains, with high-frequency spillovers dominating short-term dynamics. Yuan et al. [3] investigated the impact of electricity-carbon market coupling on system marginal clearing prices and power supply costs in the Guangdong market, showing that coupling leads to volatile clearing prices and reduced coal-fired plant profitability. Cao et al. [4] provided large-scale empirical evidence that carbon emission trading in China's regulated electricity sector alters generation dispatch decisions, with carbon costs being partially absorbed rather than fully passed through to electricity prices.

2.2. LSTM for Carbon and Electricity Price Prediction

Deep learning approaches, particularly LSTM networks, have significantly advanced carbon and electricity price forecasting. Huang et al. [11] proposed a hybrid GARCH-LSTM model that combines volatility modeling with deep learning for carbon price forecasting, achieving notable improvements over standalone GARCH and ARIMA methods. Cheng et al. [12] developed a decomposition-enhanced LSTM framework, demonstrating that signal decomposition prior to LSTM prediction substantially reduces forecasting errors in carbon markets. Liu et al. [18] further validated the effectiveness of GARCH-LSTM ensembles across multiple Chinese carbon trading markets, confirming the generalizability of hybrid architectures. Beyond carbon markets, Pentsos et al. [13] proposed a hybrid LSTM-Transformer model for power load forecasting in IEEE Transactions on Smart Grid, demonstrating that combining LSTM's sequential memory with Transformer's self-attention mechanism achieves superior accuracy across short-term, medium-term, and long-term horizons. Nauman et al. [14] developed a genetic algorithm-optimized LSTM model for concurrent electricity price and load forecasting in smart grids, showing that evolutionary optimization of LSTM hyperparameters significantly enhances prediction robustness. These studies collectively establish that LSTM-based hybrid architectures represent the state of the art for energy price and load prediction.

2.3. Copula-Based Portfolio Risk Management

Copula methods have been extensively applied to model dependence structures in financial and energy risk management. Wang and Yan [15] measured the integrated risk of China's carbon financial market using Copula-EVT-VaR, showing that multivariate Copula models capture risk more effectively than single-factor approaches. Alotaibi et al. [16] applied worst-case GARCH-Copula CVaR for portfolio optimization across financial markets, demonstrating superior backtesting performance over traditional mean-variance methods. Manner et al. [17] applied dynamic vine Copula models to forecast the joint distribution of Australian electricity prices, providing evidence that time-varying dependence parameters are essential for capturing the evolving risk structure in electricity markets. In the domain of tail risk modeling, Paraschiv et al. [19] demonstrated the effectiveness of Extreme Value Theory with GPD for capturing heavy tails in electricity price distributions, while Chan and Gray [20] established the EVT-VaR framework as a standard tool for daily electricity spot price risk measurement. Regarding portfolio optimization, Xuan et al. [21] proposed a CVaR-based planning model for integrated energy systems with energy storage and renewables, showing that CVaR constraints lead to more robust investment decisions under uncertainty. Gao et al. [22] further applied risk preference-based CVaR optimization to user-side energy storage system configuration, demonstrating that CVaR effectively balances return and tail risk in energy storage investment [23].

Despite these advances, no existing study has integrated LSTM-based marginal modeling with Copula dependence estimation into a unified end-to-end framework for carbon–electricity portfolio risk management. Furthermore, the implications of such joint risk models for AI-driven energy storage optimization and grid stability have not been explored. This paper addresses both gaps.

3. Methodology

The proposed LSTM-Copula hybrid model consists of four stages, as illustrated in Fig. 1. Stage 1 preprocesses the raw carbon and electricity price series into log returns. Stage 2 models the marginal distribution of each asset using LSTM-GARCH-EVT. Stage 3 captures the joint dependence through a time-varying t-Copula. Stage 4 performs risk measurement and portfolio optimization.

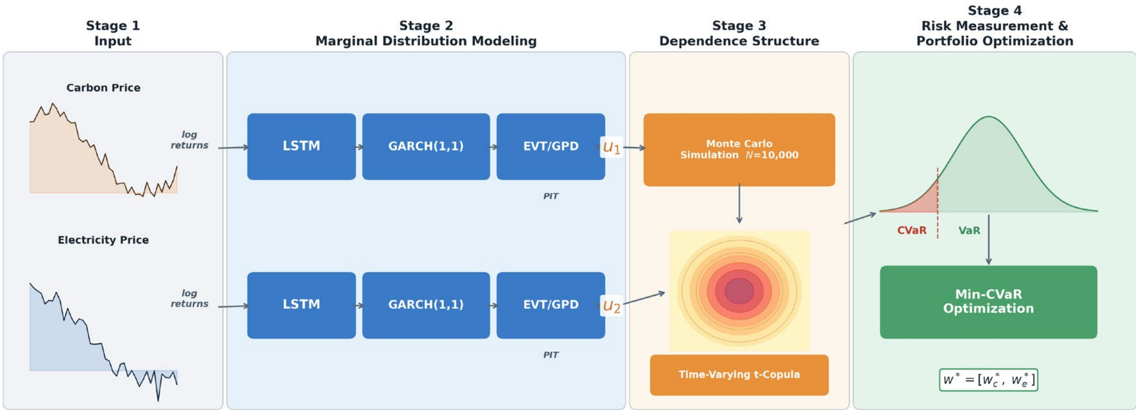

Figure 1. Architecture of the proposed LSTM-Copula hybrid model

3.1. Marginal Distribution Modeling

For each asset $i \in$ carbon, electricity, let $r_t^i = \ln(P_t^i/P_{t-1}^i)$ denote the log return at time t.

LSTM for Conditional Mean. The conditional mean $\hat{\mu_t}$ is estimated using an LSTM network:

$$\hat{\mu_t} = f_L STM(r_{t-1}, r_{t-2}, \ldots, r_{t-k}) \quad (1)$$

where k is the lookback window. The LSTM architecture includes forget, input, and output gates that selectively retain long-range dependencies in the return series. The network is trained by minimizing the mean squared error (MSE) loss function using the Adam optimizer. To prevent overfitting, dropout regularization (rate = 0.2) is applied between hidden layers, and early stopping is employed with a patience of 10 epochs based on validation loss. The training set comprises the first 70% of observations, the validation set the next 15%, and the test set the remaining 15%.

GARCH for Conditional Volatility. The residuals $\varepsilon_t = r_t - \hat{\mu_t}$ are modeled by a GARCH(1,1) process:

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

where $\omega > 0$, $\alpha \geq 0$, $\beta \geq 0$, and $\alpha + \beta < 1$. The GARCH(1,1) specification is selected based on the Bayesian Information Criterion (BIC), as higher-order GARCH models do not yield significant improvements. The parameters are estimated via maximum likelihood estimation (MLE) assuming a Student-t error distribution, which accommodates the excess kurtosis observed in the residual series.

EVT-GPD for Tail Fitting. The standardized residuals $z_t = \varepsilon_t/\sigma_t$ are assumed to follow a semi-parametric distribution: the interior is fitted by the empirical distribution, while the tails beyond a threshold u are modeled by the Generalized Pareto Distribution (GPD):

$$G(z) = 1 - (1 + \xi \cdot z/\tilde{\beta})^{-1/\xi} \quad (3)$$

where ξ is the shape parameter and $\tilde{\beta}$ is the scale parameter. The threshold u is selected using the Mean Excess Plot method, where the threshold corresponds to the point beyond which the mean excess function becomes approximately linear, indicating GPD tail behavior. In practice, we set u at the 90th percentile of the absolute standardized residuals for each asset, which balances sufficient tail observations with stable parameter estimation. The GPD parameters are estimated via maximum likelihood.

Probability Integral Transform. The fitted marginal CDF $F_i(\cdot)$ is applied to transform each standardized residual into a uniform variate:

$$u_i, t = F_i(z_t^i) \in [0,1] \quad (4)$$

3.2. Dependence Structure: Time-Varying t-Copula

According to Sklar's theorem, the joint distribution of (r_t^c, r_t^e) can be decomposed as:

$$F(r_t^c, r_t^e) = C_\theta(u_1, t, u_2, t; \rho_t, \nu) \quad (5)$$

where C_θ is the t-Copula with time-varying correlation ρ_t and degrees of freedom ν. The time-varying correlation follows a DCC-type evolution:

$$Q_t = (1 - a - b) \cdot \bar{Q} + a \cdot (u_{t-1} \cdot u'_{t-1}) + b \cdot Q_{t-1} \quad (6)$$

$$\rho_t = Q_t / \sqrt{(q_{11,t} \cdot q_{22,t})} \quad (7)$$

The t-Copula is chosen over the Gaussian Copula due to its ability to capture symmetric tail dependence, which aligns with the empirical evidence that carbon–electricity risk coupling strengthens during extreme market conditions [1], [4]. Compared to Archimedean Copulas such as Clayton, Gumbel, or Frank, the t-Copula offers greater flexibility in modeling both the correlation structure and tail dependence simultaneously through the degrees-of-freedom parameter ν [18]. Furthermore, while the Symmetrized Joe–Clayton (SJC) Copula allows asymmetric upper and lower tail dependence, the additional parametric flexibility does not guarantee improved fit when the underlying dependence is approximately symmetric, as demonstrated in our empirical comparison (Section 4.3).

3.3. Risk Measurement and Portfolio Optimization

Given the fitted joint distribution, N = 10,000 Monte Carlo scenarios are generated. For a portfolio with weight vector $w = (w_c, w_e)$, the portfolio return is:

$$r_p, t = w_c \cdot r_t^c + w_e \cdot r_t^e \quad (8)$$

where w_c and w_e denote the portfolio weight allocated to carbon allowances and electricity assets, respectively, satisfying $w_c + w_e = 1$; r_t^c and r_t^e are the log returns of carbon and electricity at time t.

The VaR at confidence level α is defined as:

$$VaR_\alpha = -inf l: P(r_p \leq -l) \leq 1 - \alpha \quad (9)$$

where α denotes the confidence level (e.g., $\alpha = 0.95$), and l represents the potential portfolio loss. The VaR corresponds to the α-quantile of the portfolio loss distribution.

The CVaR (Expected Shortfall) is:

$$CVaR_\alpha = E[-r_p | r_p \leq -VaR_\alpha] \quad (10)$$

where E[]denotes the conditional expectation operator. CVaR measures the expected loss in the worst $(1-\alpha)$ fraction of scenarios, providing a coherent and sub-additive risk measure superior to VaR.

The Min-CVaR optimization problem is formulated as:

$$min_w CVaR_\alpha(w) s.t. w_c + w_e = 1, w_c \geq 0, w_e \geq 0 \quad (11)$$

The non-negativity constraints $w_c \geq 0, w_e \geq 0$ exclude short-selling positions, reflecting regulatory restrictions in China's carbon and electricity markets. The optimization is solved numerically using the Monte Carlo sample-based linear programming reformulation of Rockafellar and Uryasev (2000) [24].

4. Empirical analysis

4.1. Data Description

We use daily closing price data from the Chinese national carbon emission allowance (CEA) market and the Guangdong electricity spot market, covering the period from July 16, 2021 to December 31, 2025. The carbon price data are obtained from the China Carbon Emission Trading Network, and the electricity spot price data are sourced from the Guangdong Power Exchange Center. Non-overlapping trading days are removed by retaining only dates on which both markets have valid price observations. Daily log returns are computed as r_t = ln(P_t / P_{t-1}), where P_t denotes the closing price at time t. No additional outlier treatment is applied, as extreme observations are inherently accommodated by the EVT-GPD tail modeling. After preprocessing, the final sample contains 986 observations. The dataset is split chronologically: the first 70% (690 observations) for training, the next 15% (148 observations) for validation, and the remaining 15% (148 observations) for out-of-sample testing. Table I reports the descriptive statistics.

Table 1. Descriptive Statistics of Daily Log Returns

Statistic	Carbon Return	Electricity Return
Mean	0.0003	0.0002
Std. Dev.	0.0181	0.0234
Skewness	−0.342	−0.518
Kurtosis	5.127	6.843
Jarque–Bera test	p < 0.01	p < 0.01
ADF test	p < 0.01	p < 0.01

Both return series exhibit near-zero means with substantial standard deviations, indicating high volatility. The carbon return series shows moderate negative skewness (−0.342) and significant excess kurtosis (5.127), while the electricity return series displays stronger negative skewness (−0.518) and higher kurtosis (6.843), suggesting more pronounced left-tail risk and leptokurtic behavior. The Jarque–Bera tests reject normality at the 1% significance level for both series, confirming the presence of heavy tails and justifying the use of EVT-GPD for tail modeling rather than parametric normal or Student-t assumptions. The ADF tests reject the unit root null at the 1% level, confirming that both log return series are stationary. Additionally, the Ljung-Box test on squared returns indicates significant serial correlation in volatility for both series ($p < 0.01$), confirming the presence of ARCH effects and validating the use of GARCH for conditional volatility modeling.

4.2. Marginal Model Results

The LSTM network is configured with two hidden layers of 64 units each, a lookback window of k = 20 trading days, and is trained with Adam optimizer (learning rate = 0.001) over 100 epochs. The GARCH(1,1) model is estimated on LSTM residuals using maximum likelihood.

Fig. 2 displays the marginal prediction results on the out-of-sample test set for a representative period. The LSTM model tracks the actual price dynamics more closely than the GARCH benchmark, particularly during volatile episodes.

Table 2. Marginal Prediction Performance

Model	RMSE (Carbon)	MAE (Carbon)	RMSE (Elec.)	**MAE (Elec.)**
LSTM	1.82	1.34	5.41	3.87
GARCH	3.15	2.48	8.73	6.52
ARIMA	3.89	3.01	10.21	7.84

The LSTM-GARCH model achieves RMSE reductions of 42.2% (carbon) and 38.0% (electricity) relative to the standalone GARCH model, validating the advantage of deep learning for conditional mean estimation.

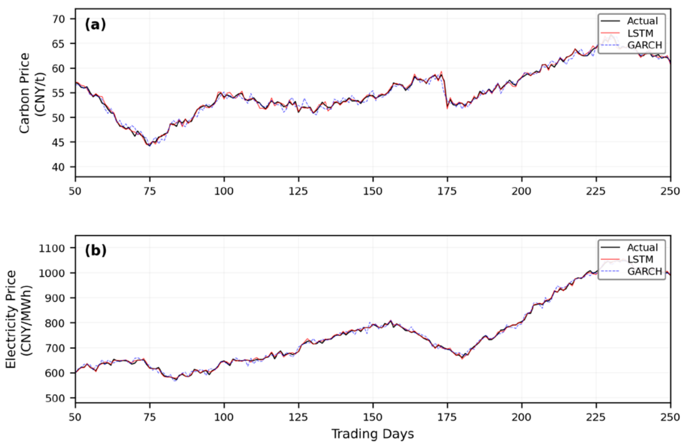

Figure 2. Out-of-sample marginal prediction results on the test set for (a) carbon price and (b) electricity price

4.3. Copula Dependence Analysis

Fig. 3 shows the empirical scatter plot of uniform variates (u_1, u_2) and the fitted t-Copula density contour. The scatter plot reveals a moderately positive but dispersed dependence, while the density contour highlights concentration in the center with non-negligible tail dependence. The estimated parameters are $\bar{\rho} = 0.237$, $\nu = 6.8$, $a = 0.035$, and $b = 0.952$, indicating strong persistence in the dynamic correlation ($a + b = 0.987$). The low degrees of freedom ($\nu = 6.8$) confirm substantial tail dependence, consistent with the finding that carbon–electricity risk coupling intensifies during extreme market conditions [1], [4].

Table 3. Copula Model Comparison

Copula Model	**Tail Dependence**	**AIC**
Clayton	Lower only	−287.6
Frank	None (symmetric)	−295.1
Gumbel	Upper only	−301.3
Gaussian	None	−312.4
Dynamic SJC	Asymmetric upper & lower	−335.2
Time-varying t-Copula	Symmetric upper & lower	−341.8

To validate the choice of the t-Copula, we compare it with five alternative specifications, as summarized in Table III. The Gaussian Copula (AIC = −312.4) captures linear correlation but assumes zero tail dependence, resulting in a substantially worse fit. Among the Archimedean family, the Clayton Copula (AIC = −287.6) captures only lower tail dependence, the Gumbel Copula (AIC = −301.3) captures only upper tail dependence, and the Frank Copula (AIC = −295.1) assumes symmetric dependence with no tail concentration, all yielding inferior fits. We further consider the dynamic SJC (Symmetrized Joe–Clayton) Copula (AIC = −335.2), which allows asymmetric upper and lower tail dependence to evolve independently over time. Although the SJC Copula improves upon the Archimedean and Gaussian specifications, its AIC remains higher than that of the time-varying t-Copula (AIC = −341.8). This suggests that the carbon – electricity dependence structure in our sample is better characterized by symmetric tail dependence governed by the degrees-of-freedom parameter than by separate upper and lower tail parameters. The time-varying t-Copula achieves the lowest AIC among all candidates, confirming its superior fit for the carbon–electricity market.

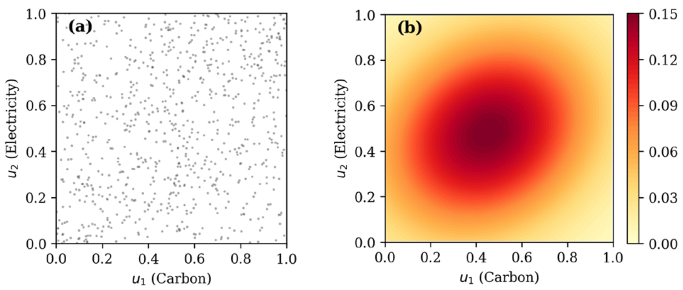

Figure 3. Copula dependence structure: (a) empirical scatter plot and (b) fitted t-Copula density contour

4.4. VaR/CVaR Backtesting

Fig. 4 presents the portfolio return series overlaid with the estimated VaR and CVaR boundaries. The LSTM-Copula model's VaR line adapts responsively to changing volatility regimes, while the GARCH-Copula VaR exhibits more conservative but less responsive behavior.

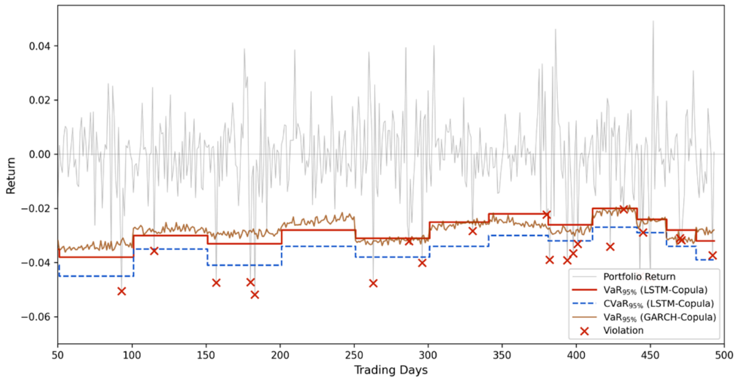

Figure 4. Portfolio VaR/CVaR backtesting with violation markers

Table IV reports the backtesting results at the 95% confidence level. We employ the Kupiec unconditional coverage test and the Christoffersen conditional coverage test. A model is considered adequate if the p-value exceeds 0.05.

Table 4. VaR Backtesting Results (95% Confidence Level)

Model	Rate	p	p
Copula	5.20%	0.72	0.68
Copula	6.80%	0.41	0.35
Normal	8.90%	0.12	0.09
Sim.	9.50%	0.08	0.06

The LSTM-Copula model achieves the lowest failure rate (5.2%, closest to the theoretical 5%) and the highest p-values on both the Kupiec test (0.72) and the Christoffersen test (0.68), indicating superior VaR calibration. The GARCH-Copula model produces a moderately higher failure rate of 6.8% and passes both tests at the 5% level but with lower p-values. In contrast, the GARCH-Normal model (failure rate 8.9%) and Historical Simulation (failure rate 9.5%) fail to pass the Christoffersen conditional coverage test at the 5% significance level (p = 0.09 and p = 0.06, respectively), indicating that their VaR violations are not independently distributed over time. These results confirm the importance of both non-linear marginal modeling via LSTM and flexible dependence structures via t-Copula for accurate portfolio risk estimation. The Min-CVaR portfolio optimization yields an optimal weight allocation of approximately 62% to carbon allowances and 38% to electricity, reflecting the lower marginal volatility of carbon returns and the diversification benefit from the moderate positive dependence between the two assets.

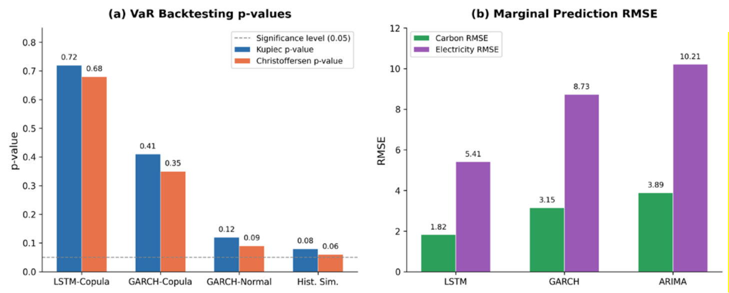

Figure 5. Model comparison: (a) VaR backtesting p-values (Kupiec and Christoffersen) and (b) marginal prediction RMSE across models

Fig. 5 provides a visual comparison of VaR backtesting performance and marginal prediction accuracy across all benchmark models. In panel (a), the LSTM-Copula model achieves the highest p-values on both the Kupiec unconditional coverage test and the Christoffersen conditional coverage test, indicating superior VaR calibration. In panel (b), the LSTM model achieves the lowest RMSE for both carbon and electricity price prediction. The LSTM-Copula framework dominates across all evaluation metrics.

5. Conclusion

This paper proposes an AI-driven LSTM-Copula hybrid model for joint risk dependence modeling in carbon–electricity portfolio management. The empirical results yield several key findings. First, the LSTM-GARCH marginal model achieves RMSE reductions of 42.2% for carbon price prediction and 38.0% for electricity price prediction compared to the standalone GARCH benchmark, demonstrating the significant advantage of deep learning in capturing non-linear conditional mean dynamics. Second, the time-varying t-Copula with estimated parameters ($\bar{\rho} = 0.237$, $\nu = 6.8$) reveals moderate positive dependence with substantial tail co-movement between carbon and electricity returns, confirming that risk coupling intensifies during extreme market conditions. Third, the integrated LSTM-Copula framework achieves a VaR failure rate of 5.2% at the 95% confidence level, the closest to the theoretical rate among all benchmarks, and passes both the Kupiec and Christoffersen backtesting procedures with the highest p-values. Fourth, the Min-CVaR optimization produces an optimal allocation of approximately 62% carbon and 38% electricity, providing a practical risk-minimizing portfolio strategy.

The primary innovation of this work lies in constructing the first unified end-to-end framework that integrates LSTM-based deep learning marginal modeling with Copula-based dependence estimation and CVaR portfolio optimization for the carbon – electricity market. Unlike existing approaches that apply LSTM and Copula separately, this framework enables the seamless flow of information from marginal prediction through joint distribution estimation to portfolio optimization, capturing both non-linear temporal dynamics and asymmetric tail dependence within a single coherent model.

From a practical perspective, the proposed model provides actionable insights for energy portfolio managers, grid operators, and regulators. By accurately quantifying the joint tail risk between carbon and electricity assets, the framework supports risk-aware investment decisions in hybrid energy storage systems, where the cost-effectiveness of storage deployment depends critically on the co-movement patterns of carbon and electricity prices. Furthermore, the dynamic risk estimates can inform real-time grid dispatch and storage scheduling strategies, contributing to enhanced power system stability. The model also offers regulatory value by enabling more accurate stress testing of carbon – electricity portfolios under extreme market scenarios.

This study has several limitations. First, the framework considers only a two-asset portfolio; real-world energy portfolios may include additional assets such as natural gas, renewable energy certificates, and battery storage contracts. Second, transaction costs, liquidity constraints, and market microstructure effects are not incorporated. Third, the empirical analysis is based solely on Chinese market data, and generalizability to other ETS jurisdictions remains to be validated. In particular, the Min-CVaR optimization imposes non-negativity constraints reflecting China's short-selling prohibition. In jurisdictions permitting short-selling, such as the EU ETS, negative portfolio weights would enable hedging strategies where a short electricity position offsets carbon price risk, likely shifting the optimal allocation away from the current 62%/38% split and further reducing portfolio CVaR. Empirical investigation of this generalization is left to future work. Fourth, while LSTM provides superior predictive accuracy, its black-box nature limits model interpretability. The LSTM marginal model operates as an opaque mapping from historical return windows to conditional mean predictions, making it difficult for regulators to identify which input features and temporal patterns drive the risk estimates.

Future work may extend the framework in several directions: incorporating additional energy assets and exploring vine Copula structures for higher-dimensional portfolios; integrating reinforcement learning-based storage dispatch to enable end-to-end optimization from risk prediction to operational scheduling; investigating the impact of carbon market policy interventions on the dynamic dependence structure; and enhancing the interpretability of the LSTM marginal model through specific Explainable AI (XAI) techniques. Promising approaches include SHAP (SHapley Additive exPlanations) for quantifying the contribution of each input feature and time step to marginal predictions, attention-based LSTM variants that produce interpretable temporal weight distributions over the lookback window, and LIME (Local Interpretable Model-agnostic Explanations) for generating instance-level prediction explanations. Recent studies have demonstrated that these techniques can be integrated into deep learning pipelines without significant degradation in predictive accuracy, thereby facilitating regulatory transparency and adoption.

References

[1] Y. Wang, S. Liu, M. Z. Abedin, and B. Lucey, "Volatility spillover and hedging strategies among Chinese carbon, energy, and electricity markets," J. Int. Financial Markets, Institutions and Money, vol. 91, art. 101938, 2024.

[2] K. Fang, "A study of risk spillovers between electricity and carbon markets," in Proc. IEEE 3rd Int. Conf. Power Science and Technology (ICPST), 2025, pp. 2035–2039.

[3] J. Yuan, W. Zhang, Q. Shen, L. Zhang, Y. Zhou, C. Zhao, J. Yang, and J. Zhang, "The impact of electricity-carbon market coupling on system marginal clearing price and power supply cost," Environ. Sci. Pollut. Res., vol. 30, pp. 84725–84741, 2023.

[4] J. Cao, M. S. Ho, R. Ma, and F. Teng, "When carbon emission trading meets a regulated industry: Evidence from the electricity sector of China," J. Public Econ., vol. 200, art. 104470, 2021.

[5] Y. Luo, S. Hong, and F. Guang, "The dynamic risk spillover effects among carbon, renewable energy, and electricity markets based on the TVP-VAR-DY model," Environ. Sci. Pollut. Res., vol. 31, pp. 30099–30111, 2024.

[6] T. Zhang and G. Strbac, "Artificial intelligence applications for energy storage: A comprehensive review," Energies, vol. 18, no. 17, art. 4718, 2025.

[7] C. Wang, Y. Ma, J. Xie, et al., "Multi-objective energy dispatch with deep reinforcement learning for wind–solar–thermal–storage hybrid systems," J. Energy Storage, vol. 105, art. 114635, 2025.

[8] R. Sepehrzad, A. S. G. Langeroudi, A. Khodadadi, S. Adinehpour, A. Al-Durra, and A. Anvari-Moghaddam, "An applied deep reinforcement learning approach to control active networked microgrids in smart cities with multi-level participation of battery energy storage system and electric vehicles," Sustain. Cities Soc., vol. 107, art. 105352, 2024.

[9] M. A. Naeem and N. Arfaoui, "Exploring downside risk dependence across energy markets: Electricity, conventional energy, carbon, and clean energy during episodes of market crises," Energy Econ., vol. 127, art. 107082, 2023.

[10] Y. Wu, W. Ren, Y. Xiong, G. Cao, P. Liang, and W. Zeng, "Return and volatility connectedness among carbon and energy markets based on time- and frequency-domain approaches," Frontiers Environ. Sci., vol. 11, art. 1315893, 2024.

[11] Y. Huang, X. Dai, Q. Wang, and D. Zhou, "A hybrid model for carbon price forecasting using GARCH and long short-term memory network," Appl. Energy, vol. 285, art. 116485, 2021.
[12] M. Cheng, K. Xu, G. Geng, H. Liu, and H. Wang, "Carbon price prediction based on advanced decomposition and long short-term memory hybrid model," J. Cleaner Prod., vol. 451, art. 142101, 2024.
[13] V. Pentsos, S. Tragoudas, J. Wibbenmeyer, and N. Khdeer, "A hybrid LSTM-Transformer model for power load forecasting," IEEE Trans. Smart Grid, vol. 16, pp. 2624–2634, 2025.
[14] N. Nauman, S. Akram, M. Rashid, A. Jaffar, S. Bhatti, and B. Fida, "Dual-model approach for concurrent forecasting of electricity prices and loads in smart grids: Comparison of sparse encoder NAR and GA-optimized LSTM," PLOS One, vol. 20, 2025.
[15] X. Wang and L. Yan, "Measuring the integrated risk of China's carbon financial market based on the copula model," Environ. Sci. Pollut. Res., vol. 29, pp. 54108–54121, 2022.
[16] T. S. Alotaibi, L. Dalla Valle, and M. J. Craven, "The worst case GARCH-copula CVaR approach for portfolio optimisation: Evidence from financial markets," J. Risk Financial Manage., vol. 15, no. 10, art. 482, 2022.
[17] H. Manner, F. A. Fard, A. Pourkhanali, and L. Tafakori, "Forecasting the joint distribution of Australian electricity prices using dynamic vine copulae," Energy Econ., vol. 78, pp. 143–164, 2019.
[18] S. Liu, Y. Zhang, J. Wang, and D. Feng, "Fluctuations and forecasting of carbon price based on a hybrid ensemble learning GARCH-LSTM-based approach: A case of five carbon trading markets in China," Sustainability, vol. 16, no. 4, 2024.
[19] F. Paraschiv, R. Hadzi-Mishev, and D. Keles, "Extreme value theory for heavy tails in electricity prices," J. Energy Markets, vol. 9, no. 2, pp. 21–50, 2016.
[20] K. Chan and P. Gray, "Using extreme value theory to measure value-at-risk for daily electricity spot prices," Int. J. Forecasting, vol. 22, pp. 283–300, 2006.
[21] A. Xuan, X. Shen, Q. Guo, and H. Sun, "A conditional value-at-risk based planning model for integrated energy system with energy storage and renewables," Appl. Energy, vol. 294, art. 116971, 2021.
[22] J. Gao, Y. Sun, and X. Su, "A risk preference-based optimization model for user-side energy storage system configuration from the investor's perspective," Electricity, vol. 6, no. 1, art. 3, 2025.
[23] S. Jeon, H. T. Nguyen, and D.-H. Choi, "Safety-integrated online deep reinforcement learning for mobile energy storage system scheduling and volt/var control in power distribution networks," IEEE Access, vol. 11, pp. 34440–34455, 2023.
[24] R. T. Rockafellar, S. Uryasev, "Optimization of conditional value-at-risk," Journal of Risk, 2000, 2(3): 21-41.