Modern Finance

Modern Finance Institute

Article

Predicting volatility of cryptocurrencies: Deep learning and GARCH family models

Abdul Moiz [1], **and Hassan Raza** [2, *]

1 MS Scholar, SZABIST University, Islamabad Campus
2 Associate Professor, SZABIST University, Islamabad Campus
* Correspondence: hassanrazaa@live.com

Abstract: This paper examines the application of econometric models, deep learning architectures, and hybrid combinations of both methods for volatility forecasting in cryptocurrency markets. Using daily data on 10 major cryptocurrencies from 2020–2025, this work benchmarks GARCH-family models, deep learning models, and hybrid frameworks against a uniform 30-day realized volatility target under a strict walk-forward evaluation protocol. The findings indicate that traditional GARCH models perform adequately for relatively stable assets, whereas deep learning models improve accuracy by capturing nonlinear dynamics. Hybrid models, particularly those combining GARCH models with TDNN and GRU components, achieve the most accurate and statistically significant volatility forecasts. These findings highlight the added value of hybrid methodologies in cryptocurrency risk modeling.

Keywords: cryptocurrency, GARCH, deep learning, hybrid modeling, volatility

1. Introduction

Cryptocurrency markets are very volatile. They regularly switch between regimes and react quickly to changes in sentiment. With these conditions, accurate volatility forecasting is important for investors, risk managers, and policymakers. Unlike traditional assets, cryptocurrencies show extreme price moves, asymmetric volatility, and nonlinear dynamics. These features challenge conventional models and motivate the development of more flexible frameworks.

This paper evaluates econometric, deep learning, and hybrid models for predicting cryptocurrency volatility. It is based on daily price data for ten leading cryptocurrencies from June 2020 to June 2025. The paper compares GARCH and GJR-GARCH models, deep learning approaches including LSTM, GRU, FFNN, and TDNN, and hybrid models that combine econometric and deep learning predictions. All models use a strict walk-forward framework and the same 30-day realized volatility target. This ensures fair and robust comparisons.

There are several major findings from the results. First, GARCH models perform well for stable cryptocurrencies such as Bitcoin and Ethereum, where volatility clustering persists. For volatile, sentiment-driven assets such as Dogecoin, Shiba Inu, and Toncoin, forecasting accuracy drops sharply. Second, deep learning models usually outperform GARCH models by capturing nonlinear volatility dynamics. Among these, GRU and TDNN achieve lower error rates than LSTM and FFNN, indicating better adaptability to complex patterns. However, standalone deep learning models are less robust during periods of extreme volatility. This shows they have limits in capturing asymmetric, shock-driven market behavior.

Hybrid models that combine GARCH-type forecasts with deep learning predictions show the strongest and most consistent performance. In particular, hybrids featuring GJR-GARCH improve forecasting accuracy by directly modeling asymmetric volatility. The

Check for updates

Citation: Moiz, A., & Raza, H. (2026). Predicting volatility of cryptocurrencies: Deep learning and GARCH family models. *Modern Finance*, *4*(1), 95-113.

Accepting Editor: Adam Zaremba

Received: 26 July 2025
Accepted: 12 March 2026
Published: 20 March 2026

Copyright: © 2026 by the authors. This article is an open-access article distributed under the terms and conditions of the Creative Commons Attribution 4.0 International (CC BY 4.0) license (https://creativecommons.org/licenses/by/4.0/).

TDNN–GJR and GRU–GJR models are the most reliable across many cryptocurrencies, as shown by low error metrics and strong QLIKE values. Diebold–Mariano tests confirm that improvements from hybrid models are statistically significant versus standalone models.

This paper advances the literature on cryptocurrency volatility forecasting in four direct ways. It provides a unified comparison of econometric, deep learning, and hybrid models under a shared forecasting period and evaluation process. It extends hybrid approaches by embedding asymmetric GJR-GARCH structures within deep learning architectures. It rigorously applies relevant evaluation metrics, such as the QLIKE loss and the Diebold-Mariano test, to ensure robust inference. Finally, it systematically analyzes a broad range of cryptocurrencies, both speculative and stable, to directly highlight model performance across diverse market conditions.

The rest of the paper is organized as follows: Section 2 reviews related literature. Section 3 describes the data and methodology. Section 4 presents the results. Section 5 discusses findings and implications. Section 6 concludes and suggests directions for future research.

2. Literature Review

Cryptocurrency markets are highly volatile, meaning their prices can change rapidly and unpredictably. This makes price movements difficult to predict. While volatility prediction in traditional markets using GARCH (Generalized Autoregressive Conditional Heteroskedasticity) models is well-studied, volatility prediction in cryptocurrency markets is less explored. Variants of GARCH models, such as GJR-GARCH (Glosten-Jagannathan-Runkle GARCH) and EGARCH (Exponential GARCH), have been applied to cryptocurrencies. However, these models do not fully capture key dynamics, such as price movements that do not follow a straight line (non-linear trends) and unusual, significant changes (extreme events).

Omari and Ngunyi (2021) proposed a heavy-tailed GARCH-EVT model for forecasting cryptocurrency volatility. They aimed to apply extreme-value theory to GARCH to improve volatility forecasts for Bitcoin. However, Micu and Dumitrescu (2022) found traditional GARCH models inadequate for capturing cryptocurrency volatility. Their research shows a need for more advanced modeling to match the unique characteristics of cryptocurrency movements. Catania and Grassi (2022) found that conventional GARCH models perform poorly for cryptocurrency returns, whereas score-driven models provide better estimates. This all points to ongoing research and the need for innovation in volatility modeling for this fast-changing market.

This gap has led researchers to test deep learning architectures to recognize nonlinear dynamics in cryptocurrency returns. Tripathy et al. (2025) compared econometric models such as GARCH and TGARCH with deep learning models, including LSTM, Bi-LSTM, and Multivariate Bi-LSTM. They used Bitcoin prices from 2014 to 2022. Their findings showed that Multivariate Bi-LSTM performed much better than the econometric models, giving the lowest MSE and MAE. This proves that multivariate deep learning methods excel at detecting complex volatility patterns. Shaimaa Alghamdi et al. (2022) used an LSTM-based sentiment analysis model to forecast cryptocurrency volatility. They found that adding investor sentiment helped capture sentiment-driven volatility. Seabe et al. (2023) primarily studied price forecasting rather than volatility. They compared LSTM, CNN, and GRU across cryptocurrencies and concluded that LSTM-based models performed best. This highlights LSTM's suitability for detecting nonlinear dynamics in crypto markets and suggests its use for volatility forecasting.

Hybrid techniques that combine GARCH (Generalized Autoregressive Conditional Heteroskedasticity) and Deep Learning models have proven to be strong forecasting tools. Hu et al. (2020) found that neural networks, inspired by the structure of the human brain, outperformed GARCH models in forecasting copper volatility. Using GARCH forecasts as explanatory variables further improved performance. Aras (2021) used a stacking ensemble—a method that combines multiple base learners (individual predictive models)

with an SVM (Support Vector Machine) meta-learner (a model that aggregates the base learners' outputs)—to forecast Bitcoin volatility. Hybrid models with feature selection, a technique for choosing the most relevant input variables, increased forecasting performance. Shaimaa Alghamdi et al. (2022) also used LSTM (Long Short-Term Memory) models for sentiment analysis. This process analyzes text for positive or negative sentiment to explain cryptocurrency volatility better, thereby widening the use of deep learning. Recent studies continue to build on these results. Garcia Medina et al. (2024) developed an LSTM-GARCH hybrid that combines GARCH-based volatility inputs with LSTM networks. They achieved better forecast accuracy. Tripathy et al. (2025) showed that a Multivariate Bi-LSTM (a Bidirectional LSTM that handles multiple input features) outperformed GARCH and TGARCH (Threshold GARCH) in predicting Bitcoin volatility, reinforcing the value of multivariate deep learning. Iqbal et al. (2025) combined Nonlinear GARCH (NGARCH) with machine learning classifiers, including SVM, KNN (K-Nearest Neighbors), and ANN (Artificial Neural Network). NGARCH-KNN stood out for its ability to predict directional volatility. This demonstrates the strengths of combining statistical and machine learning methods. In summary, hybrid approaches combining econometric and deep learning tools, such as CNNs (Convolutional Neural Networks), LSTMs, and ensembles, are versatile for capturing the complexity and nonlinearity of financial and cryptocurrency markets.

In short, the discussion of recent studies on cryptocurrency volatility forecasting reveals a wide variety of models and approaches, including GARCH-family models, deep learning models such as LSTM, and sentiment analysis techniques. In general, the above research helps us understand the dynamics of cryptocurrency markets and provides valuable insights for investors and researchers in this fast-evolving field. While GARCH models have been widely applied to volatility forecasting, they are not well-suited to modeling cryptocurrency markets. Hybrid models combining GARCH with deep learning approaches are identified as a better choice, offering enhanced forecasting accuracy and improved risk management in cryptocurrency markets.

3. Data and Methodology

This research utilizes historical daily cryptocurrency prices from Yahoo Finance with the yfinance package. The sample is for the period between June 1, 2020, and June 1, 2025, and includes ten of the largest cryptocurrencies: Bitcoin (BTC), Ethereum (ETH), Tether (USDT), Binance Coin (BNB), Solana (SOL), Ripple (XRP), Cardano (ADA), Dogecoin (DOGE), Shiba Inu (SHIB), and Toncoin (TON). Daily closing prices for each asset were collected, as this is the most commonly used variable in related research, such as that by Seo & Kim (2020), and converted to log returns to capture percentage changes and eliminate scale effects. Preprocessing involved (i) removing infinite and missing values, (ii) downscaling returns by a factor of 10 in the GARCH scripts to avoid numerical underflow, (iii) calculating realized volatility as the 30-day rolling standard deviation of returns, and (iv) normalizing inputs using MinMax Scaler into the range [0,1] for deep learning models. To ensure data homogeneity and compatibility across econometric and deep learning frameworks, we also applied the Augmented Dickey-Fuller (ADF) test to assess the stationarity of the return series, confirming that all series were stationary prior to model estimation. All these processes ensure reproducibility, comparability, and consistency for all modeling approaches.

In this paper, a two-part modeling strategy is considered that links traditional econometric volatility models with state-of-the-art deep learning architectures. The approach allows us to study both linear and nonlinear dependencies in cryptocurrency volatility while hybridizing the two paradigms to leverage their strengths.

3.1. Benchmark models

We have included two benchmark models for comparative purposes.

3.2. Naive Forecast

The naive model uses the previous day's realized volatility as the forecast for the next day. This simple benchmark is widely used in volatility research because of its robustness and interpretability, and allows the assessment of whether more sophisticated models yield meaningful improvements.

3.3. HAR-RV

The HAR-RV model incorporates lagged daily, weekly, and monthly realized volatilities to capture multiple volatility components. In this study, the specification used was:

- Lagged 1-day realized volatility.
- Lagged 5-day realized volatility.
- Lagged 22-day realized volatility

Ordinary least squares estimated coefficients. HAR-RV provides a well-established benchmark for daily volatility forecasts and serves as an intermediate benchmark between naive and GARCH-type models.

3.4. GARCH Models

Two conditional volatility models were estimated by using the arch library. The standard GARCH(1,1) model estimated by maximum likelihood captures volatility persistence and clustering. In cases where heavier-tailed distributions improved estimation stability, a Student-t distribution for residuals was allowed. The GJR-GARCH(1,1,1) model includes an asymmetric term, a leverage effect commonly found in financial time series, in which negative returns lead to higher volatility. Both specifications were estimated using a walk-forward approach in which, at each time point, the model is refitted on an expanding window and used to generate a one-day-ahead conditional variance forecast. This ensures that the predictions are based solely on information available at each point in time.

3.5. Deep learning Models

Four deep learning models were developed using TensorFlow (Keras) and LSTM, GRU, FFNN, and TDNN. They used a 30-day look-back of realized volatility to capture the short-term temporal structure of volatility dynamics on a day-to-day basis. Configurations included dropout 0.3, batch normalization, Adam optimizer, and mean squared error loss. Early stopping of patience 3–5 epochs prevented overfitting. Models followed a strictly chronological walk-forward approach: after each forecasting block, retrain the models and forecast only out-of-training periods, without allowing leakage, thereby enabling realistic forecasting. The random seeds were fixed for reproducibility of the results. By this design, the models can learn nonlinear, time-dependent features while remaining comparable to econometric and benchmark volatility models.

3.6. Hybrid Model

The hybrid modelling framework combines forecasts from deep learning architectures with outputs from GARCH-type volatility models. For each cryptocurrency, both the deep learning models (LSTM, GRU, FFNN, TDNN) and the econometric models (GARCH(1,1) and GJR-GARCH(1,1,1)) were estimated using a strictly walk-forward procedure, so that only out-of-sample predictions are generated at each point in time. These out-of-sample forecasts were then realigned to a common daily horizon so that the deep learning outputs and the conditional variance estimates refer to the same realized-volatility target. All series were standardized into a common numerical scale before hybridization.

Hybrid models were constructed by stacking the out-of-sample forecasts from each deep learning model with the corresponding GARCH or GJR-GARCH forecasts. A

Random Forest Regressor served as the meta learner. It was trained only on out-of-fold predictions, thereby precluding the use of any in-sample information and ensuring that all hybrid forecasts are based on completely unseen data. This procedure yielded eight hybrid specifications: LSTM–GARCH, GRU–GARCH, FFNN–GARCH, TDNN–GARCH, LSTM–GJR, GRU–GJR, FFNN–GJR, and TDNN–GJR. The hybrid forecasts thus combine the nonlinear sequential structure captured by deep learning models with the volatility-clustering dynamics captured by GARCH-type models.

3.7. Evaluation Metrics

Forecasting performance was evaluated using metrics, including MSE and MAE, which quantify predictive accuracy and are commonly used in volatility forecasting studies, such as Plevris et al. (2022). Moreover, the QLIKE loss function was used, as it is an appropriate method for evaluating volatility models due to its appealing statistical properties and heightened sensitivity to variance underestimation. Comparative model performance was further analyzed using DM tests to statistically evaluate differences in forecast accuracy among competing models, including benchmarks, econometric specifications, deep learning methods, and hybrid approaches.

4. Results

This study compares forecasting methods in cryptocurrency volatility between deep learning, GARCH-family models, and their hybrid combinations. Deep learning captures nonlinear, time-dependent patterns, while GARCH models capture volatility clustering. Hybrids incorporate out-of-sample predictions from both through a meta-learning framework. The models are benchmarked using mean squared error, mean absolute error, and the QLIKE loss to determine which model provides the most reliable volatility estimates across different cryptocurrencies.

4.1. Descriptive Statistics

There is considerable price variability among the selected cryptocurrencies. Bitcoin (BTC-USD) has the highest average price, 43,432.23 USD, with significant volatility (SD = 24,748.21) and a wide price range from 9,045.39 USD to 111,673.28 USD. The average price of Ethereum (ETH-USD) was 2,167.64 USD, with an SD of 1,029.30, and ranged from 222.95 USD to 4,812.08 USD. On the contrary, USDT-USD is seen as one of the most stable cryptocurrencies, with an average price of 1.00 USD and a very low SD of 0.01. However, it has an extremely high kurtosis of 99.41 and an exceptionally high skewness of 6.86, reflecting infrequent deviations from the peg but yielding large values. Binance Coin and Solana display middle-range price movements with average prices of 357.92 USD and 78.63 USD, respectively, and relatively low kurtosis and skewness. Other altcoins, such as XRP-USD (0.76 USD), ADA-USD (0.68 USD), and DOGE-USD (0.12 USD), have lower mean values but higher skewness, reflecting occasional rapid price surges. For example, SHIB-USD has a very small mean near zero, with a kurtosis of 4.60 and a skewness of 1.57, indicating that large upward price jumps may occur periodically. Finally, TON-USD had an average of 3.22 USD and an SD of 3.07 USD, indicating middle-range volatility with greater kurtosis (8.33) and higher skewness (2.44), suggesting greater speculative risk. Overall, Bitcoin and Ethereum lead in market value and volatility; stablecoins exhibit stability through minimal fluctuation; and low-cap altcoins exhibit more speculative behavior, as evidenced by high skewness and kurtosis.

4.2. Correlation Matrix

The correlation matrix identifies strong interdependencies between leading cryptocurrencies. Ethereum (ETH-USD) and Bitcoin (BTC-USD), for instance, have a high positive correlation (0.86), indicating that price changes in these assets tend to move in similar directions. Stablecoins such as Tether (USDT-USD) have little correlation with other cryptocurrencies, further establishing their position as relatively stable assets. The

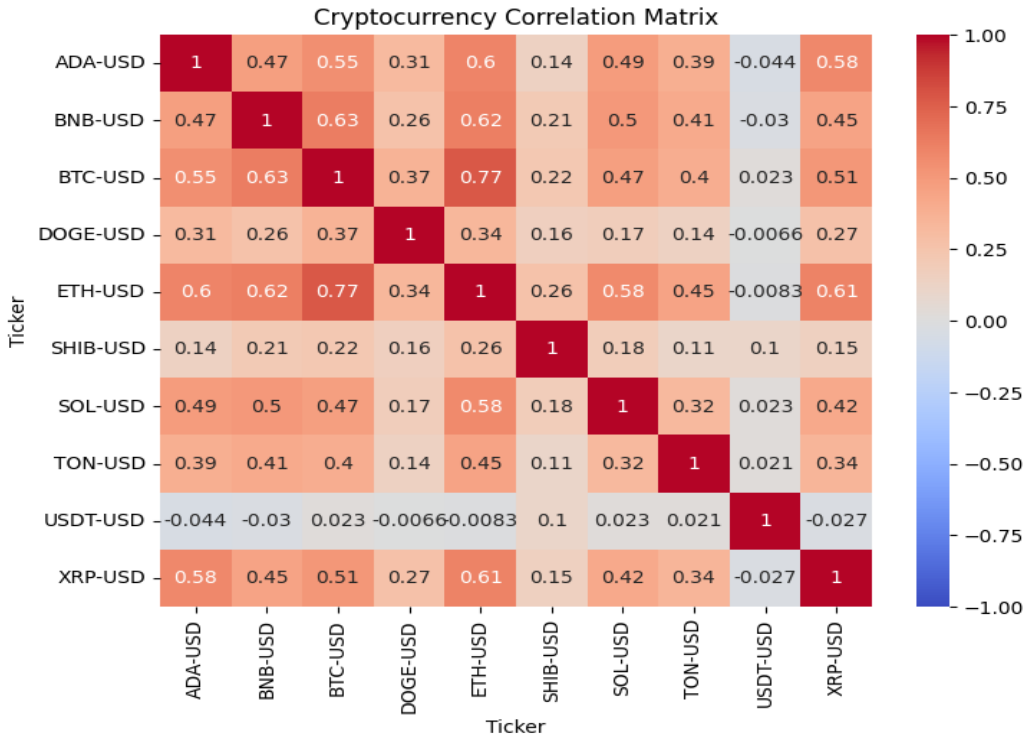
matrix offers insights into market behavior, assisting investors and analysts with risk management for diversified cryptocurrency portfolios.

Table 1. Descriptive Statistics of Cryptocurrency

Crypto	Mean	SD	Min	Max	Kurtosis	Skewness
BTC-USD	43,432.23	24,748.21	9,045.39	111,673.28	0.13	0.77
ETH-USD	2,167.64	1,029.30	222.95	4,812.08	(0.55)	0.62
USDT-USD	1.00	0.01	0.99	1.02	99.41	6.86
BNB-USD	357.92	194.67	15.19	750.27	(0.85)	0.06
SOL-USD	78.63	71.09	0.56	261.87	(0.92)	0.64
XRP-USD	0.76	0.61	0.18	3.29	3.44	2.04
ADA-USD	0.68	0.55	0.08	2.96	2.38	1.59
DOGE-USD	0.12	0.10	0.00	0.68	2.28	1.37
SHIB-USD	0.00	0.00	0.00	0.00	4.60	1.57
TON-USD	3.22	3.07	0.02	24.71	8.33	2.44

Note. Summary statistics of daily log returns for the period June 2020–June 2025. All return series are stationary.

Figure 1. Correlation Matrix



4.3. Econometric Volatility Forecasting Results

The Econometric models help to predict the volatility of the various cryptocurrencies, as illustrated in Table 2. It presents the volatility-forecasting performance of the GARCH, GJR-GARCH, Naive, and HAR-RV models across 10 cryptocurrencies. Taken as a whole, econometric models capture volatility features with varying degrees of accuracy, and their relative performance varies significantly across stable and highly speculative assets.

Table 2. Econometric Models

Crypto	GARCH			GJR-GARCH			Naïve			HAR-RV		
	MSE	MAE	QLIKE	MSE	MAE	QLIKE	MSE	MAE	QLIKE	MSE	MAE	QLIKE
BTC-USD	0.000621	0.020842	-2.87	0.000626	0.020884	-2.78	0.000906	0.020808	2.00	0.000492	0.015915	-2.86
ETH-USD	0.000988	0.025263	-2.55	0.001006	0.025389	-2.54	0.001507	0.02628	0.22	0.000836	0.020316	-2.58
USDT-USD	0.000031	0.000386	-7.06	0.000001	0.00034	-7.13	0.000051	0.000344	2.50	0.000001	0.000292	-7.17
BNB-USD	0.001379	0.0253	-2.66	0.001538	0.025874	-2.65	0.001696	0.024689	2.29	0.001058	0.019413	-2.69
SOL-USD	0.002222	0.038244	-2.12	0.002228	0.038249	-2.12	0.003497	0.040198	0.85	0.002064	0.031155	-2.13
XRP-USD	0.002561	0.036748	-2.36	0.002865	0.036954	-2.37	0.003267	0.031089	1.95	0.001973	0.024618	-2.48
ADA-USD	0.001802	0.031037	-2.35	0.001793	0.030976	-2.35	0.002427	0.031059	2.35	0.001451	0.024182	-2.39
DOGE-USD	0.006127	0.054269	-2.14	0.004165	0.045555	-2.20	0.016047	0.038145	7.89	0.009615	0.031317	-2.31
SHIB-USD	0.004332	0.056167	-2.24	0.004684	0.057259	-2.24	0.032394	0.051774	1.84	0.004019	0.041976	-2.38
TON-USD	7.481054	0.239104	-1.80	7.479499	0.239759	-1.79	12.95085	0.33144	5.69	6.454518	0.289428	8.92

Note. This table reports the volatility forecasting performance of the GARCH, GJR-GARCH, Naive, and HAR-RV models. The Naive forecasts are obtained from lagged realized volatility. HAR-RV refers to the Heterogeneous Autoregressive model of realized volatility. Forecast accuracy is evaluated using MSE, MAE, and QLIKE; lower values indicate better performance.

Table 3. Deep Learning Models

Crypto	LSTM-MSE	LSTM-MAE	LSTM-QLIKE	GRU-MSE	GRU-MAE	GRU-QLIKE	FFNN-MSE	FFNN-MAE	FFNN-QLIKE	TDNN-MSE	TDNN-MAE	TDNN-QLIKE
BTC-USD	5.32E-06	0.001387	-2.65	2.72E-06	0.000969	-2.65	3.85E-06	0.001232	-2.65	2.53E-05	0.003849	-2.64
ETH-USD	7.98E-06	0.001835	-2.42	5.56E-06	0.001416	-2.42	7.84E-06	0.001774	-2.42	4.10E-05	0.004866	-2.40
USDT-USD	9.59E-09	4.60E-05	-7.19	4.07E-09	3.40E-05	-7.21	3.55E-09	3.47E-05	-7.18	3.06E-08	0.000106	-7.10
BNB-USD	1.02E-05	0.002084	-2.60	5.59E-06	0.001412	-2.60	7.50E-06	0.001791	-2.60	3.97E-05	0.004703	-2.58
SOL-USD	1.99E-05	0.002406	-2.04	1.31E-05	0.001882	-2.05	2.64E-05	0.003016	-2.05	0.000106	0.007297	-2.03
XRP-USD	3.16E-05	0.002164	-2.31	2.93E-05	0.001889	-2.31	4.90E-05	0.003041	-2.30	0.000168	0.008367	-2.27
ADA-USD	3.52E-05	0.002617	-2.24	1.85E-05	0.001606	-2.24	2.63E-05	0.002263	-2.24	0.000127	0.00681	-2.21
DOGE-USD	5.94E-05	0.005044	-2.16	3.33E-05	0.003918	-2.16	3.78E-05	0.004458	-2.16	0.000196	0.009857	-2.12
SHIB-USD	8.01E-05	0.00395	-1.99	3.72E-05	0.002979	-1.97	1.89E-05	0.002435	-1.97	0.000164	0.007199	-1.95
TON-USD	5.738408	0.559505	-1.20	5.918656	0.57268	-1.10	0.14572	0.057565	-1.54	0.667178	0.188792	-1.50

Note. This table summarizes the forecasting performance of the deep learning models LSTM, GRU, FFNN, and TDNN. All models were estimated under a rolling walk-forward framework using a 30-day look-back window. Models evaluated using MSE, MAE, and QLIKE computed over strictly out-of-sample forecasts.

Both GARCH and GJR-GARCH perform relatively well for major assets such as BTC-USD, ETH-USD, and BNB-USD, where volatility clustering appears more structured. The GARCH and GJR-GARCH models for Bitcoin have low error values (GARCH MSE: 0.000621; GJR-GARCH MSE: 0.000626) and negative QLIKE scores, indicating well-calibrated volatility forecasts. Conversely, highly speculative cryptocurrencies such as DOGE-USD and SHIB-USD exhibit significantly higher error values and QLIKE scores, underscoring the challenge of capturing underlying sentiment-driven volatility jumps with variance-based frameworks. Likewise, TON-USD exhibits similarly high errors, further underscoring the limitations of GARCH-type modeling for assets susceptible to extreme market shocks.

The Naive benchmark performs unexpectedly well in several cases, especially for highly stable assets where volatility tends to evolve smoothly over successive periods, such as BTC-USD and ETH-USD. Conversely, the Naive forecasts deteriorate markedly for highly volatile assets like TON-USD, confirming that simple persistence assumptions are insufficient in unstable markets. The HAR-RV model proves competitive across all instances, with the lowest error on the majority of assets, such as BTC-USD (HAR-RV MSE: 0.000492), and the most favorable QLIKE scores overall. This result indicates that including a multi-horizon realized volatility component yields a more robust description of volatility dynamics in markets where short-run movements interact with medium- and long-run volatility trends.

Figure 2 presents a visual comparison of BTC-USD and TON-USD. In the case of BTC-USD, econometric models track observed volatility well, exhibiting well-behaved volatility clustering patterns. In contrast, for TON-USD, all models, especially GARCH and GJR-GARCH, fail to anticipate volatility jumps induced by speculative trading. These results collectively suggest that while GARCH-type models remain effective for large, relatively stable cryptocurrencies, HAR-RV has proved the most reliable of the econometric frameworks, especially when market conditions feature multiscale volatility dynamics.

4.4. Deep Learning Models

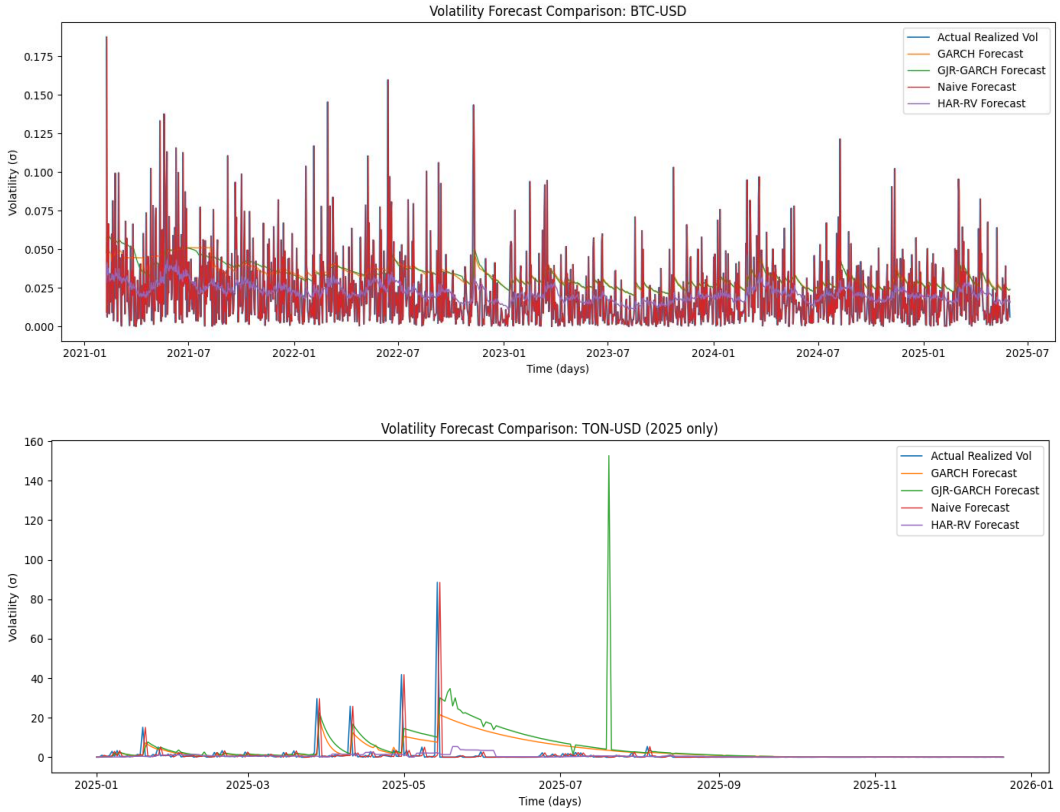
In the deep learning model, four main models, that are LSTM, FFNN, GRU, and TDNN, were evaluated using the parameters of mean square error and mean absolute error, which are illustrated in Table 3.

The overall results suggest that the deep learning models significantly outperform the econometric benchmark for most cryptocurrencies, offering a higher degree of nonlinearity and temporal resolution in volatility dynamics. Recurrent models yield the lowest consistent forecasting errors across the investigated architectures.

As a result, LSTM and GRU are the strongest models for most assets. Most notably, GRU has very low error metrics for BTC-USD, ETH-USD, ADA-USD, and SOL-USD. This indicates its strengths in modeling long-lasting temporal dependencies with greater computational efficiency. LSTM also shows competitive performance, especially for large-cap cryptocurrencies, which indicates its ability to capture longer-term volatility structures.

TDNN yields results that are robust across a variety of assets; for DOGE-USD and SHIB-USD, the performance is particularly pronounced, further supporting the idea that convolution-based temporal modeling is efficient at capturing local volatility patterns and short-term fluctuations. In contrast, the FFNN model consistently underperforms recurrent and convolutional architectures, particularly for assets such as BTC-USD and BNB-USD, which exhibit strong temporal dependence. This corroborates the earlier conclusion that feedforward structures are not well-suited for volatility forecasting because they lack explicit modeling of sequential dependencies.

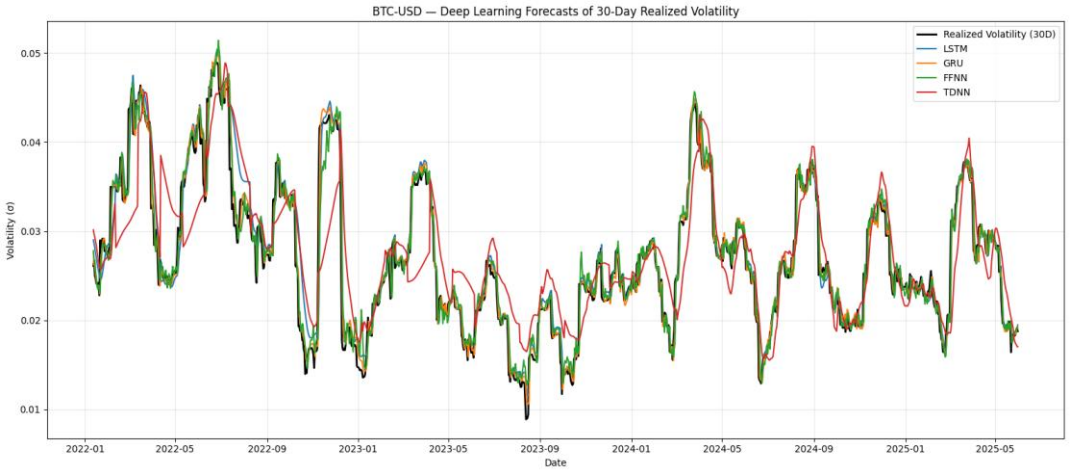
Figure 2. GARCH Actual vs Predicted

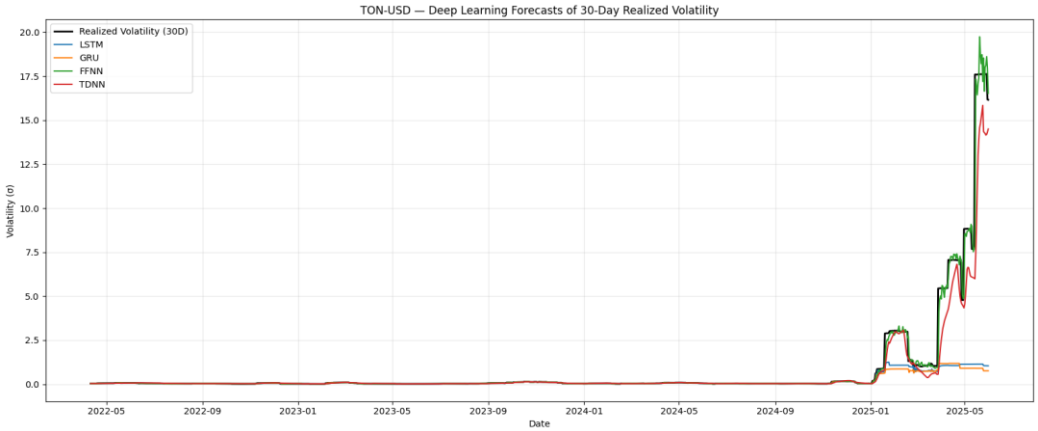


Note. Comparison of realized volatility and out-of-sample model forecasts. Volatility is measured as 30-day realized volatility.

Figure 3 shows Actual vs. predicted volatility, BTC-USD (strong) and TON-USD (weak). In the case of BTC-USD, the GRU and TDNN track realized volatility well, both in magnitude and direction. The FFNN shows larger deviations, most of which occur during high-activity phases. For TON-USD, all models are less precise because of the very low, stable volatility. Still, the recurrent models have lower error rates.

Figure 3. Deep learning Actual vs Predicted Volatility





Note. Comparison of realized volatility and out-of-sample model forecasts. Volatility is measured as 30-day realized volatility.

To assess whether differences in forecast accuracy among deep learning models are statistically significant, we conducted DM tests using GRU as the reference model, as it provides strong, consistent performance across assets. Table 4 shows that GRU significantly outperforms LSTM and FFNN across most cryptocurrencies, as indicated by large, positive DM statistics in BTC-USD, ETH-USD, USDT-USD, BNB-USD, ADA-USD, DOGE-USD, and SHIB-USD. This implies that such reductions in error are not due to luck. However, large negative DM statistics for GRU–TDNN across several assets, such as BTC-USD, ETH-USD, BNB-USD, and SOL-USD, indicate that TDNN can compete with and outperform GRU in capturing short-term volatility. No clear dominance of any architecture was evident in the TON-USD results. Overall, DM statistical tests confirm a significant difference among models, with recurrent or temporal architectures, GRU and TDNN, offering better volatility forecasts than feedforward architectures.

4.5. Hybrid Models

In this study, four deep learning algorithms, LSTM, GRU, FFNN, and TDNN, were experimented with individually and collectively combined with GARCH and GJR-GARCH. The GARCH-based hybrid in Table 5 and the GJR-GARCH-based hybrid in Table 7 were evaluated for efficacy using MSE, MAE, and QLIKE scores to assess their volatility prediction performance.

Table 4. Diebold–Mariano (DM) test

Crypto	GRU vs LSTM	GRU vs FFNN	GRU vs TDNN
BTC-USD	7.576806	4.48043	-18.46522
ETH-USD	6.307682	0.347928	-17.91745
USDT-USD	5.171367	5.036355	-10.771874
BNB-USD	9.796416	5.015393	-18.219974
SOL-USD	3.302697	-3.21191	-16.72956
XRP-USD	2.791456	-3.9499	-8.60426
ADA-USD	5.406791	4.235004	-8.5286
DOGE-USD	6.718025	4.337965	-11.14163
SHIB-USD	4.5186224	4.4175826	-10.2536
TON-USD	-5.285654	5.697014	5.57596

Note. Diebold-Mariano test statistics are reported to compare forecast accuracy. Positive values indicate that the first model performs better than the second.

Table 5 Hybrid GARCH and DL Models

Crypto	Hybrid (LSTM-GARCH)			Hybrid (GRU-GARCH)			hybrid (FFNN-GARCH)			hybrid (TDNN-GARCH)		
	MSE	MAE	QLIKE	MSE	MAE	QLIKE	MSE	MAE	QLIKE	MSE	MAE	QLIKE
BTC	1.97E-09	2.51E-05	-7.15	1.77E-09	2.41E-05	-7.15	2.57E-09	2.96E-05	-7.15	1.61E-09	2.24E-05	-7.15
ETH	5.15E-06	1.59E-03	-2.56	5.07E-06	1.57E-03	-2.56	5.77E-06	1.74E-03	-2.56	4.53E-06	1.47E-03	-2.56
USDT	1.97E-09	2.51E-05	-7.15	1.77E-09	2.41E-05	-7.15	2.57E-09	2.96E-05	-7.15	1.61E-09	2.24E-05	-7.15
BNB	5.15E-06	1.59E-03	-2.56	5.07E-06	1.57E-03	-2.56	5.77E-06	1.74E-03	-2.56	4.53E-06	1.47E-03	-2.56
SOL	1.91E-05	2.89E-03	-2.02	1.92E-05	2.93E-03	-2.02	2.07E-05	3.08E-03	-2.02	1.76E-05	2.74E-03	-2.02
XRP	4.87E-05	4.00E-03	-2.27	4.65E-05	4.06E-03	-2.27	5.16E-05	4.41E-03	-2.27	3.76E-05	3.86E-03	-2.27
ADA	2.10E-05	2.77E-03	-2.22	2.06E-05	2.72E-03	-2.22	1.97E-05	2.72E-03	-2.22	1.64E-05	2.49E-03	-2.22
DOGE	1.70E-05	2.84E-03	-2.14	1.67E-05	2.81E-03	-2.14	2.20E-05	3.26E-03	-2.14	1.34E-05	2.53E-03	-2.14
SHIB	1.92E-05	2.87E-03	-1.95	1.55E-05	2.53E-03	-1.95	1.89E-05	2.82E-03	-1.95	2.93E-05	3.55E-03	-1.95
TON	5.10E-02	3.89E-02	-1.63	5.70E-02	3.90E-02	-1.63	6.56E-02	5.58E-02	-1.63	8.63E-02	5.20E-02	-1.63

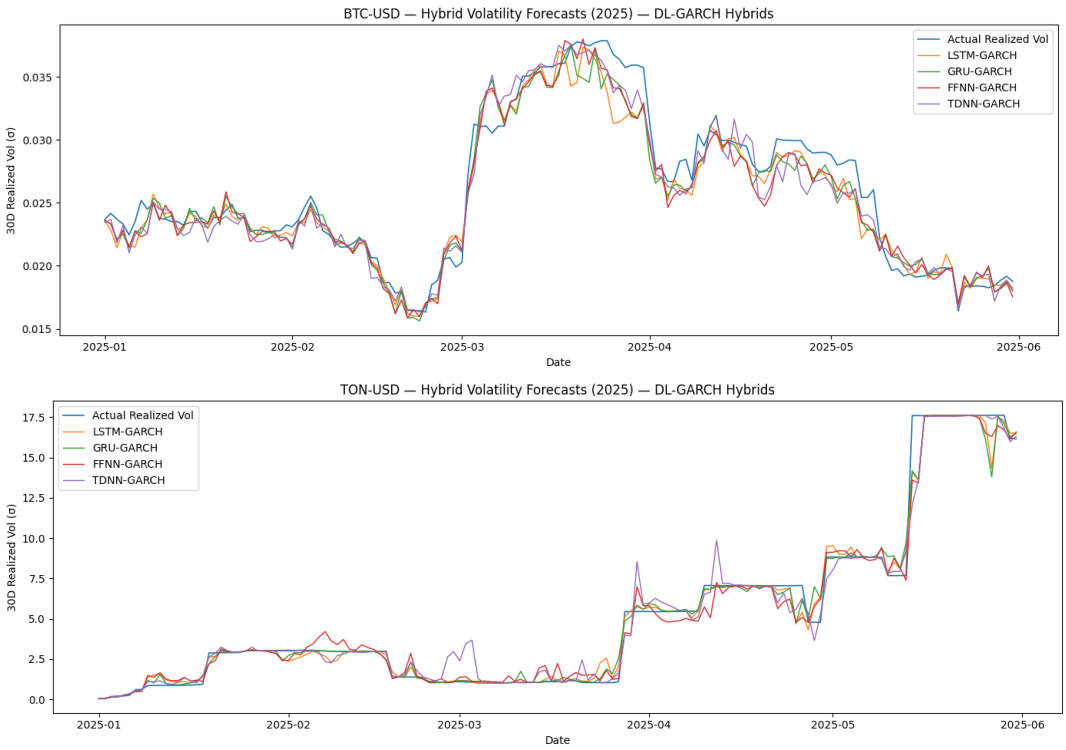
Note. The table reports results for hybrid models that combine GARCH-type forecasts with deep learning forecasts. The hybrid forecasts are generated using the Random Forest meta-learner, trained on out-of-sample predictions. The forecast performance is evaluated using mean squared error (MSE), mean absolute error (MAE), and QLIKE. Lower values denote superior forecasting accuracy.

Table 5 reports the hybrid volatility forecasts that combine deep learning outputs with GARCH/GJR-GARCH via a Random Forest meta-learner. The hybrid setups include LSTM–GARCH, GRU–GARCH, FFNN–GARCH, and TDNN–GARCH. Hybrid combinations uniformly outperform their pure econometric and DL counterparts for all coins in terms of MSE, MAE, and QLIKE.

The biggest gains are from volatile assets like DOGE-USD, SHIB-USD, and TON-USD, for which hybrids substantially reduce forecast errors by blending nonlinear pattern recognition with volatility clustering. GRU–GARCH and TDNN–GARCH are the top-performing methods most of the time, indicating that recurrent or convolutional neural networks with conditional variance forecasts can yield higher accuracy. Notably, the FFNN-GARCH outperforms the standalone FFNN, while hybridization appears to offset feedforward limits by adding econometric insights into volatility. Similarly, QLIKE favors hybrids, indicating better alignment of densities with respect to realized volatility.

Figure 4, Actual vs. predicted volatility for the BTC-USD and TON-USD. Hybrids better capture the persistence in regimes and shocks. Generally, hybrid GARCH-DL frameworks provide the most robust volatility forecasts across conditions of the crypto market.

Figure 4. Hybrid Model Actual vs Predicted Volatility



Note. Comparison of realized volatility and out-of-sample model forecasts. Volatility is measured as 30-day realized volatility.

DM tests were run to check whether the improvements of the GRU–GARCH hybrid over other hybrids are statistically significant. DM stats compare the forecast accuracy of the GRU–GARCH against those of LSTM–GARCH, FFNN–GARCH, and TDNN–GARCH. Large negative DM values imply that GRU–GARCH outperforms its rivals. Indeed, for the major coins such as BTC-USD, ETH-USD, USDT-USD, and BNB-USD, the GRU–GARCH outclasses LSTM–GARCH and FFNN–GARCH with strongly negative DM values (for example, BTC-USD: −4.226 vs FFNN-GARCH; ETH-USD: −9.824 vs LSTM-GARCH), therefore meaning that the combination of nonlinear learning by GRU with GARCH volatility yields statistically significant gains in forecast accuracy. However, compared to TDNN–GARCH, some assets show positive DM values, such as BTC-USD

and ADA-USD, suggesting that TDNN–GARCH can be competitive under market conditions characterized by strong local patterns. The DM results confirm that the proposed GARCH–DL hybrid offers meaningful improvements over other hybrids, with GRU–GARCH proving to be the most consistently superior across many cryptocurrencies.

Table 6 Diebold–Mariano (DM) test

Crypto	GRU-GARCH Vs LSTM-GARCH	GRU-GARCH Vs FFNN-GARCH	GRU-GARCH Vs TDNN-GARCH
BTC-USD	-1.077	-4.226	9.673
ETH-USD	-9.824	-3.591	3.972
USDT-USD	-1.077	-4.226	9.673
BNB-USD	-0.386	-3.368	2.135
SOL-USD	0.057	-1.774	1.340
XRP-USD	-1.104	-2.833	2.148
ADA-USD	-4.270	1.152	2.217
DOGE-USD	-3.494	-5.079	3.270
SHIB-USD	-3.200	-2.200	-5.141
TON-USD	1.134	-0.678	-1.293

Note. Diebold-Mariano test statistics are reported to compare forecast accuracy. Positive values indicate that the first model performs better than the second.

DM tests as reported in Table 6, were run to check whether the improvements of the GRU–GARCH hybrid over other hybrids are statistically significant. DM stats compare the forecast accuracy of the GRU–GARCH against those of LSTM–GARCH, FFNN–GARCH, and TDNN–GARCH. Large negative DM values imply that GRU–GARCH outperforms its rivals. Indeed, for the major coins such as BTC-USD, ETH-USD, USDT-USD, and BNB-USD, the GRU–GARCH outclasses LSTM–GARCH and FFNN–GARCH with strongly negative DM values (for example, BTC-USD: −4.226 vs FFNN-GARCH; ETH-USD: −9.824 vs LSTM-GARCH), therefore meaning that the combination of nonlinear learning by GRU with GARCH volatility yields statistically significant gains in forecast accuracy. However, compared to TDNN–GARCH, some assets show positive DM values, such as BTC-USD and ADA-USD, suggesting that TDNN–GARCH can be competitive under market conditions characterized by strong local patterns. The DM results confirm that the proposed GARCH–DL hybrid offers meaningful improvements over other hybrids, with GRU–GARCH proving to be the most consistently superior across many cryptocurrencies.

According to the results in Table 7, hybrid models combining GJR–GARCH with deep learning techniques-LSTM–GJR, GRU–GJR, FFNN–GJR, and TDNN further enhance the crypto volatility forecasts. Such hybrid models, which combine asymmetric volatility with nonlinear learning, produce lower error rates than traditional econometric or pure deep learning models.

TDNN–GJR has the best overall performance, especially for the highly volatile assets TON-USD (MSE ≈ 0.00000, MAE ≈ 0.00061) and SHIB-USD, demonstrating better capture of short-term shocks and longer-term dynamics. Another great performance is seen with GRU-GJR in DOGE-USD and USDT-USD, showing its strength across both speculative and stablecoins. It is competitive with large-cap assets like BTC-USD and ETH-USD, which may suggest its effectiveness in periods of persistent volatility. FFNN–GJR improves upon FFNN but lags behind the recurrent and convolution-based hybrids.

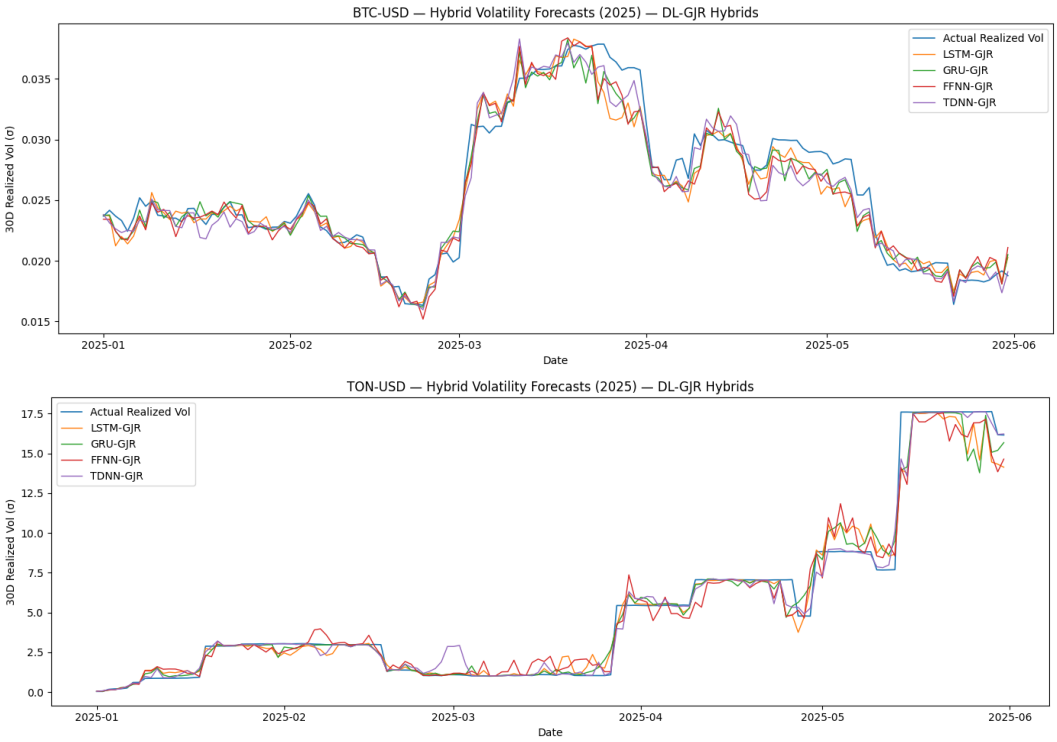
Overall, the conclusion of asymmetrical volatility via GJR–GARCH improves hybrid forecasting performance, especially in markets that exhibit leverage effects and sudden regime changes. TDNN–GJR and GRU–GJR are the most robust hybrids under crypto conditions.

Table 7 Hybrid GJR-GARCH and DL Models

Crypto	Hybrid (LSTM-GJR)			Hybrid (GRU-GJR)			Hybrid (FFNN-GJR)			Hybrid (TDNN-GJR)		
	MSE	MAE	QLIKE	MSE	MAE	QLIKE	MSE	MAE	QLIKE	MSE	MAE	QLIKE
BTC	2.74E-09	3.11E-05	-7.15	2.38E-09	2.89E-05	-7.15	3.39E-09	3.51E-05	-7.15	2.41E-09	2.75E-05	-7.15
ETH	5.33E-06	1.61E-03	-2.56	5.44E-06	1.62E-03	-2.56	5.82E-06	1.73E-03	-2.56	4.80E-06	1.51E-03	-2.56
USDT	2.74E-09	3.11E-05	-7.15	2.38E-09	2.89E-05	-7.15	3.39E-09	3.51E-05	-7.15	2.41E-09	2.75E-05	-7.15
BNB	5.33E-06	1.61E-03	-2.56	5.44E-06	1.62E-03	-2.56	5.82E-06	1.73E-03	-2.56	4.80E-06	1.51E-03	-2.56
SOL	1.84E-05	2.86E-03	-2.02	1.91E-05	2.91E-03	-2.02	2.07E-05	3.06E-03	-2.02	1.75E-05	2.72E-03	-2.02
XRP	5.15E-05	4.39E-03	-2.27	5.23E-05	4.44E-03	-2.27	5.68E-05	4.73E-03	-2.27	3.76E-05	4.08E-03	-2.27
ADA	2.20E-05	2.84E-03	-2.22	2.08E-05	2.73E-03	-2.22	2.11E-05	2.76E-03	-2.22	1.86E-05	2.62E-03	-2.22
DOGE	2.89E-05	3.48E-03	-2.14	3.12E-05	3.65E-03	-2.14	3.13E-05	3.84E-03	-2.14	2.53E-05	3.27E-03	-2.14
SHIB	1.90E-05	2.86E-03	-1.95	1.50E-05	2.50E-03	-1.95	1.81E-05	2.78E-03	-1.95	2.91E-05	3.54E-03	-1.95
TON	8.71E-02	5.97E-02	-1.63	8.32E-02	5.41E-02	-1.63	1.07E-01	7.53E-02	-1.63	5.07E-02	4.11E-02	-1.63

Note. This table highlights the performance of various hybrids when GJR-GARCH is integrated into deep learning architectures. The inclusion of GJR-GARCH can model asymmetric volatility effects. The model performance is evaluated based on three metrics: mean squared error (MSE), mean absolute error, and QLIKE, within a strict walk-forward evaluation framework.

Figure 5. Hybrid Model (GJR-GARCH & Deep Learning) Actual vs Predicted Volatility



Note. Comparison of realized volatility and out-of-sample model forecasts. Volatility is measured as 30-day realized volatility.

The results of the Diebold–Mariano test for the hybrid GJR-GARCH–DL models are in Table 8, with GRU–GJR as the reference. Also, GRU–GJR outperforms LSTM–GJR and FFNN–GJR significantly for most cryptocurrencies, as indicated by mostly negative DM statistics, especially for BTC-USD, USDT-USD, and SHIB-USD. Comparisons between GRU–GJR and TDNN–GJR reveal mixed results: positive DM statistics for ETH-USD, BNB-USD, and DOGE-USD, suggesting that TDNN–GJR may also excel in markets with high short-run volatility and asymmetric responses. These results confirm that incorporating asymmetric volatility through GJR-GARCH enhances hybrid performance, and GRU- and TDNN-based hybrids are the most statistically competitive across diverse crypto markets.

Table 8. Diebold–Mariano (DM) test

Crypto	GRU-GJR vs LSTM-GJR	GRU-GJR vs FFNN-GJR	GRU-GJR vs TDNN-GJR
BTC-USD	-2.063	-4.705	-0.135
ETH-USD	1.421	-1.810	3.231
USDT-USD	-2.063	-4.705	-0.135
BNB-USD	0.463	-1.739	2.354
SOL-USD	0.899	-1.969	1.411
XRP-USD	0.357	-1.972	3.614
ADA-USD	-1.152	-0.429	1.419
DOGE-USD	1.505	-0.103	3.141
SHIB-USD	-3.130	-1.975	-5.354
TON-USD	-0.423	-1.321	1.846

Note. Diebold-Mariano test statistics are reported to compare forecast accuracy. Positive values indicate that the first model performs better than the second.

5. Discussion

The dynamics of crypto market volatility demand forecasts that account for persistence, asymmetry, and nonlinearity. This paper compared GARCH-family models-GARCH and GJR-GARCH-with deep learning-LSTM, GRU, FFNN, and TDNN-and their hybrids that merge both, in search of robust volatility predictors across diverse assets.

The results demonstrate that GARCH-family models perform reasonably well for relatively stablecoins such as BTC-USD, ETH-USD, and USDT-USD, thereby affirming their ability to capture volatility clustering. This supports Queiroz et al. (2023), who report that GARCH models perform well in stable or moderately volatile markets. However, GARCH and GJR-GARCH fail to capture the dynamics of highly speculative assets such as DOGE-USD, SHIB-USD, and TON-USD, which exhibit regime shifts and sentiment shocks. Including HAR-RV and naïve benchmarks, we observe that HAR-RV often outperforms GARCH in MSE and QLIKE, suggesting the value of multi-horizon realized-volatility benchmarks for longer horizons.

DL models consistently outperformed stand-alone GARCH models across most coins, capturing nonlinear temporal dependencies more effectively. Among the DLs, GRU and TDNN generally yielded lower errors than LSTM and FFNN, echoing Ter-Avanesov & Beigi (2024) on gated recurrent models in finance. TDNN's strong showing supports the findings of Khan et al. (2019) and Vancsura et al. (2025) that time-delay mechanisms capture local temporal patterns. However, pure DL models show reduced robustness under extreme volatility, echoing Lalbakhsh et al. (2022) on data-driven models during abrupt shocks.

Hybrid models provided the best and most consistent performance throughout all assets. Merging GARCH forecasts with DL predictions leverages statistical clustering and nonlinear learning. Among the GARCH hybrids, GRU-GARCH and TDNN-GARCH achieved significant reductions in MSE, MAE, and QLIKE for highly volatile assets such as DOGE-USD and SHIB-USD. This extends Kristjanpoller and Minutolo (2018) by showing that hybrids boost predictive accuracy by combining strengths.

GJR-GARCH hybrids further improved the results by modeling leverage effects. TDNN-GJR and GRU-GJR were the most reliable, with TDNN-GJR almost completely eliminating errors on TON-USD and GRU-GJR performing well on both stable and volatile coins, such as USDT-USD and DOGE-USD. These results support both Vancsura et al. (2025) on asymmetric volatility and Tian et al. (2024) on improved stability under regime shifts when combining asymmetric GARCH with DL.

Diebold–Mariano tests support the significance of the performance differences. Frequently, the specifications for GRU and TDNN architectures outperformed alternatives in both standalone deep learning models and hybrid approaches, thereby enabling robustness inference for the proposed hybrids beyond average performance metrics. Though econometric and deep learning models have distinct merits, their integration, particularly within a GJR-GARCH framework with either a GRU or a TDNN, offers significant benefits.

First, the Diebold-Mariano test confirms the earlier results: the observed differences are not coincidental. Among both pure deep learning models and hybrid configurations, many GRU- and TDNN-based setups perform among the best, confirming that the proposed hybrids are indeed robust, beyond what simple averaging over error metrics might suggest.

In general, hybridizing conventional econometric models with state-of-the-art deep learning is the most efficient, especially in an extended framework such as GJR-GARCH with either GRU or TDNN. This combination provides the best volatility forecast in terms of robustness and accuracy. Altogether, these results imply that modeling the market risk of cryptocurrencies has advantages from adaptive modeling strategies that can jointly capture volatility clustering, leverage effects, and nonlinear relationships.

6. Conclusion

This paper examines the relative performance of econometric, deep learning, and hybrid approaches to forecasting cryptocurrency volatility using daily data for 10 significant digital assets over 2020–2025. The systematic comparison between traditional GARCH-family models and several deep learning architectures, as well as their hybrid combinations, provides comprehensive evidence of the relative efficiency of different modeling paradigms across regimes of varying volatility.

The results suggest that the traditional GARCH and GJR-GARCH models remain suitable for cryptocurrencies with fairly low volatility, whose primary characteristic is volatility clustering. However, these methods have limited ability to predict highly volatile and speculative assets, indicating limited capacity to model nonlinear dynamics and rapid regime switches. Deep learning models improve forecasts by modeling nonlinear temporal dependencies: Recurrent and time-delay models systematically outperform feedforward topologies. However, individual deep learning models tend to fail during periods of extreme volatility, indicating limitations in modeling asymmetric and heavy-tailed risks using pure data-driven strategies.

Hybrids that combine deep learning forecasts with GARCH-type volatility estimates show the best and most consistent performance. The hybrids with GJR-GARCH achieve the highest accuracy because they explicitly model asymmetric volatility responses. Of those, the TDNN-GJR and GRU-GJR are the most reliable across a wide range of cryptocurrencies, with lower forecast errors and more robust performance in both stable and highly volatile market conditions. The results also show that Diebold-Mariano tests indicate that hybrid models have superior predictive validity relative to their respective base models in most comparisons.

In general, the findings imply that no single modeling approach is universally optimal for forecasting cryptocurrency volatility. Instead, hybrid frameworks-fusing the structural advantages of econometric models with nonlinear learning capabilities of deep neural networks-offer an effective and practical approach. With implications for risk management, portfolio allocation, and volatility-sensitive trading strategies, these results underscore the need for adaptive modeling techniques that reflect the heterogeneous, rapidly changing nature of cryptocurrency markets.

6.1. Implications for Investors

This research adds to the body of knowledge on cryptocurrency volatility prediction, providing useful insights for policymakers, traders, and investors. The combination of GARCH and deep learning models offers improved predictive power, enabling market participants to make better-informed decisions.

For investors, the volatility characteristics of various cryptocurrencies are important to understand to manage risk. This study demonstrates that GARCH models are highly effective for forecasting the volatility of major cryptocurrencies such as Bitcoin and Ethereum, enabling investors to forecast price volatility more accurately. Hybrid models such as GRU-GARCH and TDNN-GJR are even more accurate and thus serve as valuable tools for portfolio management and hedging strategies.

Policymakers can also learn from these results by gaining a better grasp of the volatility dynamics that drive cryptocurrency markets. Forecast models can help regulators anticipate how markets will respond to regulatory adjustments, economic changes, or shifts in market sentiment, making it easier to create effective policies that promote investment while safeguarding investors from potential risks.

6.2. Limitations

Although this research has gone further than the general GARCH model by including GJR-GARCH, it is still within the GARCH family of models; these models have the capacity to treat asymmetry in terms of leverage effects but carry forward the

structural inflexibility of the GARCH architecture, such as rigid distributional assumptions and modest capacity for regime change or nonlinear effects outside of conditional variance. Which might be covered by other volatility modeling methods, such as stochastic volatility models, realized volatility models based on high-frequency data, or regime-switching models.

In addition, the study considered only the top 10 cryptocurrencies by market cap. The universe of cryptocurrencies is vast and constantly evolving, with many new assets exhibiting abnormal volatility. A larger dataset that encompasses smaller or newer cryptocurrencies might provide better insights into market trends and improve model stability.

Lastly, the study excluded potential external influences, such as policy announcements, macroeconomic conditions, and market sentiment, which significantly affect cryptocurrency volatility. Future studies can incorporate these variables into models to improve accuracy and provide a more complete picture of the crypto market.

6.3. Future Research Directions

Future research needs to investigate a variety of promising avenues. First, the application of EVT may provide insight into tail risks that matter most to cryptocurrency markets. Besides the GARCH family, other econometric methods, such as MSM, HAR-RV, or copula-based dependence structures, could depict more accurate volatility patterns. On the machine learning side, ensemble methods can be enhanced with reinforcement learning and attention-driven architectures, such as Transformers, thereby improving forecast accuracy while adapting to evolving market conditions. Other studies should, where relevant, consider the impact of DeFi protocols and NFTs on volatility, as well as broader macroeconomic and regulatory developments. Integrating sentiment measures from social networks and news sources could help in identifying behavioral drivers of valuation and further enhance predictive performance.

Author Contributions: Conceptualization, H.R.; methodology, H.R.; software, A.M.; validation, A.M. and H.R.; formal analysis, A.M.; investigation, A.M.; resources, H.R.; data curation, A.M.; writing—original draft preparation, A.M.; writing—review and editing, H.R.; visualization, A.M.; supervision, H.R.; project administration, H.R.; funding acquisition, H.R. All authors have read and agreed to the published version of the manuscript..

Funding: This research received no external funding.

Data Availability Statement: The data used in this study are publicly available from online financial data platforms, including Yahoo Finance. All datasets used in the analysis are available from publicly accessible sources, and no proprietary or restricted data were used in this research.

Conflicts of Interest: The authors declare no conflict of interest.

AI Use Statement: The authors used AI-assisted tools, including ChatGPT (OpenAI) and Grammarly, for grammar checking and language refinement during manuscript preparation. All scientific content, model design, data analysis, results, and interpretations were developed and verified by the authors. The authors remain fully responsible for the accuracy, integrity, and originality of the work.

References

Alghamdi, S., Alqethami, S., Alsubait, T., & Alhakami, H. (2022). Cryptocurrency price prediction using forecasting and sentiment analysis. *International Journal of Advanced Computer Science and Applications*, 13(10). DOI:10.14569/IJACSA.2022.01310105

Aras, S. (2021). Stacking hybrid GARCH models for forecasting Bitcoin volatility. *Expert Systems with Applications*, 174, 114747. https://doi.org/10.1016/j.eswa.2021.114747

Catania, Leopoldo, and Stefano Grassi (2022). Forecasting cryptocurrency volatility. *International Journal of Forecasting* 38, no. 3 (2022): 878-894. https://doi.org/10.1016/j.ijforecast.2021.06.005

Ge, W., Lalbakhsh, P., Isai, L., Lenskiy, A., & Suominen, H. (2022). Neural network–based financial volatility forecasting: A systematic review. *ACM Computing Surveys* (CSUR), 55(1), 1-30. https://doi.org/10.1145/3483596

García-Medina, A., & Aguayo-Moreno, E. (2024). LSTM–GARCH hybrid model for the prediction of volatility in cryptocurrency portfolios. *Computational Economics*, 63(4), 1511-1542. https://doi.org/10.1007/s10614-023-10373-8

Hu, Y., Ni, J., & Wen, L. (2020). A hybrid deep learning approach by integrating LSTM-ANN networks with the GARCH model for copper price volatility prediction. *Physica A: Statistical Mechanics and its Applications*, 557, 124907. https://doi.org/10.1016/j.physa.2020.124907

Iqbal, R., Riaz, M., Sorwar, G., & Qadir, J. (2024). Cryptocurrency Market Volatility and Forecasting: A Comparative Analysis of Modern Machine Learning Models for Cryptocurrencies Predicting Accuracy. *Review of Pacific Basin Financial Markets and* Policies, 27(04), 2450028. https://doi.org/10.1142/S0219091524500280

Jimenez, G. A., & Garcia, J. J. Analyzing cryptocurrency volatility: an EGARCH model. *Panorama Económico*, 20(41). https://doi.org/10.29201/pe-ipn.v20i41.183- ISSN 1870-2171

Khan, A. S., & Augustine, P. (2019). Predictive analytics in cryptocurrency using neural networks: A comparative study. International *Journal of Recent Technology and Engineering*, 7(6), 425-429.

Kristjanpoller, W., & Hernández, E. (2017). Volatility of main metals forecasted by a hybrid ANN-GARCH model with regressors. *Expert Systems with Applications*, 84, 290-300. https://doi.org/10.1016/j.eswa.2017.05.024

Kristjanpoller, W., & Minutolo, M. C. (2018). A hybrid volatility forecasting framework integrating GARCH, artificial neural network, technical analysis, and principal components analysis. *Expert Systems with Applications*, 109, 1-11. https://doi.org/10.1016/j.eswa.2018.05.011

Micu, R., & Dumitrescu, D. (2022). Study regarding the volatility of main cryptocurrencies. In *Proceedings of the International Conference on Business Excellence* (Vol. 16, No. 1, pp. 179-187). https://doi.org/10.2478/picbe-2022-0018

Omari, C., & Ngunyi, A. (2021). The predictive performance of extreme value analysis-based models in forecasting cryptocurrency volatility. https://doi.org/10.4236/jmf.2021.113025

Plevris, V., Solorzano, G., Bakas, N. P., & Ben Seghier, M. E. A. (2022, November). Investigation of performance metrics in regression analysis and machine learning-based prediction models. In the 8th European Congress on Computational Methods in Applied Sciences and Engineering (ECCOMAS Congress 2022). *European Community on Computational Methods in Applied Sciences*. https://doi.org/10.23967/eccomas.2022.155

Queiroz, R. G., & David, S. A. (2023). Performance of the Realized-GARCH Model against Other GARCH Types in Predicting Cryptocurrency Volatility. *Risks*, 11(12), 211. https://doi.org/10.3390/risks11120211

Rane, N., Choudhary, S. P., & Rane, J. (2024). Ensemble deep learning and machine learning: applications, opportunities, challenges, and future directions. *Studies in Medical and Health Sciences*, 1(2), 18-41.
https://doi.org/10.48185/smhs.v1i2.1225

Samson, T.K., Onwukwe, C.E., Lawal, A.I. (2023). An examination of cryptocurrency volatility: Insights from skewed error innovation distributions within GARCH model frameworks. *Mathematical Modelling of Engineering Problems*, Vol. 10, No. 4, pp. 1297-1306. https://doi.org/10.18280/mmep.100423

Seabe, P. L., Moutsinga, C. R. B., & Pindza, E. (2023). Forecasting cryptocurrency prices using LSTM, GRU, and bi-directional LSTM: a deep learning approach. *Fractal and Fractional*, 7(2), 203. https://doi.org/10.3390/fractalfract7020203

Seo, M., & Kim, G. (2020). Hybrid forecasting models based on neural networks for the volatility of Bitcoin. *Applied Sciences*, 10(14), 00. http://dx.doi.org/10.3390/app10144768

Song, Y., Tang, X., Wang, H., & Ma, Z. (2023). Volatility forecasting for stock market incorporating macroeconomic variables based on GARCH-MIDAS and deep learning models. *Journal of Forecasting*, 42(1), 51-59. https://doi.org/10.1002/for.2899

Ter-Avanesov, B., & Beigi, H. (2024). MLP, XGBoost, KAN, TDNN, and LSTM-GRU Hybrid RNN with Attention for SPX and NDX European Call Option Pricing. arXiv preprint arXiv:2409.06724. https://doi.org/10.48550/arXiv.2409.06724

Tian, B., Yan, T., & Yin, H. (2024). Forecasting the Volatility of the CSI 300 Index with a Hybrid Model of LSTM and Multiple GARCH Models. *Computational Economics*, 1-31. https://doi.org/10.1007/s10614-024-10785-0

Tripathy, N., Mishra, D., Hota, S., Mishra, S., Das, G. C., Dalai, S. S., & Nayak, S. K. (2025). Bitcoin volatility forecasting: a comparative analysis of conventional econometric models with deep learning models. *International Journal of Electrical & Computer Engineering* (2088-8708), 15(1). DOI: 10.11591/Ijece.v15i1.pp614-623

Vancsura, L., Tatay, T., & Bareith, T. (2025). Navigating AI-Driven Financial Forecasting: A Systematic Review of Current Status and Critical Research Gaps. *Forecasting*, 7(3), 36. https://doi.org/10.3390/forecast7030036

Vidal, A., & Kristjanpoller, W. (2020). Gold volatility prediction using a CNN-LSTM approach. *Expert Systems with Applications*, 157, 113481. https://doi.org/10.1016/j.eswa.2020.113481

Disclaimer: All statements, viewpoints, and data featured in the publications are exclusively those of the individual author(s) and contributor(s), not of MFI and/or its editor(s). MFI and/or the editor(s) absolve themselves of any liability for harm to individuals or property that might arise from any concepts, methods, instructions, or products mentioned in the content.