Hybrid Volatility Modeling with Calibration Under Structural Price Features for Option Pricing in Cryptocurrency Markets

Qiu Zhihao
Business School the Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong, China
ericqiu2323@gmail.com

ABSTRACT

The classical option pricing frameworks may be inefficient in the cryptocurrency market due to its characteristics of persistent volatility clustering, heavy-tailed return distributions, and structural price dynamics. This paper introduces a hybrid model that incorporates the econometric and deep learning models to estimate implied volatil- ity for option pricing in the cryptocurrency market.

In the framework of the hybrid model, the Student-t GARCH model and the Long Short-Term Memory (LSTM) model are used to capture the conditional volatility dy- namics and learn the nonlinear temporal dependencies in realized volatility, respec- tively. The expected realized volatility proxy generated by the hybrid model is cali- brated to market-implied volatility with the incorporation of structural features derived from Fibonacci retracement levels, which explain variations in volatility risk premiums of option prices.

Using cryptocurrency options data and corresponding market-implied volatility ob- servations, the empirical analysis shows that the calibrated volatility proxy is sig- nificantly related to market-implied volatility. Moreover, the inclusion of structural features improves the interpretation of the volatility gap between model-implied and market-implied measures. Additionally, the calibrated proxy, when embedded in the Black–Scholes framework, produces economically meaningful option prices. Overall, the findings support the view that a hybrid volatility-modeling approach, augmented by structural market features, provides a useful framework for cryptocurrency option valuation and volatility-premium analysis.

Keywords: Cryptocurrency options; Implied volatility; Hybrid model; GARCH–LSTM; Volatility risk premium; Fibonacci retracement; Option pricing.

1. INTRODUCTION

The significant development of cryptocurrency markets over the last decade has attracted an increasing number of institutional and retail investors participating in cryptocurrency trad- ing. However, an increasing number of market events and academic studies indicate that the cryptocurrency markets exhibit volatility clustering, heavy-tailed distributions, struc- tural breaks, and state-dependent dynamics [1,2,3]. highlighting the uncertainty of its price dynamics, which will boost the demand for its derivative trading, especially option trad- ing. The limitations of traditional option pricing models, such as the constant volatility assumption, are amplified in option pricing of cryptocurrency.

Implied volatility should not be explained as the predictor of future realized volatility, because it incorporates the expectation of realized volatility and the volatility risk premium [4,5,6]. Therefore, it is essential to estimate the expected realized volatility and identify the determinants of volatility risk premium.

The GARCH model measures volatility clustering effectively but is limited in capturing nonlinear dynamics [7,8]. In contrast, the LSTM model is able to learn complex nonlinear temporal dependencies but is limited in terms of economic interpretability [9,10]. The combination of these models offsets the individual limitations of each model.

This paper introduces the Garch-LSTM model for implied volatility modeling of cryp- tocurrency options and the further calibration between the expected realized volatility proxy generated by the hybrid model and market-implied volatility. The structural features derived from Fibonacci retracement levels are incorporated in the calibration as the determinants of volatility risk premium.

The remainder of the paper is organized as follows: Section 2 reviews the related lit- erature; Section 3 describes the data and preprocessing procedures; Section 4 presents the hybrid modeling framework and calibration methodology; Section 5 reports the empirical results; Section 6 demonstrates the limitation of this framework; Section 7 is the conclusion of this paper.

2. LITERATURE REVIEW

2.1 Cryptocurrency Volatility Dynamics

The differences between traditional financial assets and cryptocurrency are that the cryp- tocurrency markets display severe volatility dynamics in the aspect of persistent volatility clustering, heavy-tailed return distributions, and structural breaks [1,2,3]. The prior stud- ies indicate that cryptocurrency returns exhibit strong time-varying and conditionally het- eroskedastic volatility [3]. Such stylized facts motivate the application of econometric volatil- ity models to capture return variance dynamics in cryptocurrency markets, particularly the GARCH model.

2.2 Implied Volatility and Variance Risk Premium

The existing studies indicates that implied volatility incorporates not only expectation of realized volatility but also the volatility risk premium which are compensation to investors for volatility risk [4,5] such as demand-driven pricing pressure [6]. Therefore, IV cannot be interpreted as a pure statistical forecast of realized volatility (RV) and it is necessary to calibrate expected realized volatility proxy derived from historical return dynamics with market-implied volatility before being interpreted within an option-pricing framework.

2.3 Econometric Volatility Modeling

The GARCH model captures volatility clustering by modeling the time-varying conditional variance and allows for explicit specification of shock distributions [7,11]. Moreover, the Student-t GARCH model can be utilized to model the heavy-tailed distribution, such as cryptocurrency returns, which present excess kurtosis and extreme shocks [12].However, the GARCH model is a linear model and therefore limited in capturing nonlinear temporal dependencies [13].

2.4 Deep Learning in Volatility Forecasting

The LSTM model is a powerful type of current neural networks in financial time-series forecasting because of its capacity to capture complex nonlinear temporal dependencies and out-of-sample volatility prediction[9,14,15]. Nevertheless, the LSTM model is purely data- driven, which means it lacks structural interpretability and economic grounding, which are its severe limitations, affecting the application on derivative pricing and risk management, which highlights the economic interpretation importance [16].

2.5 Structural Market Features

Except for realized volatility dynamics based on statistics, the market structure and price- level behaviors also provide additional information for explaining variation in option-implied volatility[6]. Order-flow clustering and behavioral trading patterns may be reflected by the technical price structures such as support and resistance levels [17]. Those factors may affect how volatility risk premiums are embedded in implied volatility [5].

3. DATA

3.1 Underlying Market Data

The raw data source of this paper is the hourly price data of the Bitcoin perpetual futures contract traded on the Deribit exchange. The span of the sample is a rolling 180-day win- dow ending at the most recent observation, yielding nearly 4,000-5,000 observations after preprocessing.

There are two reasons for selecting the hourly frequency. First, short-term volatility dynamics and structural price movements of Bitcoin are sufficiently captured. Second, ex- cessive microstructure noise is mitigated, improving the stability of econometric volatility estimation.

Log returns are computed as

$$r_t = \log(P_t) - \log(P_{t-1}) \tag{1}$$

where Pt denotes the hourly closing price of the underlying asset.

The log transformation is consistent with the GARCH model and ensures approximate stationarity. Moreover, excess kurtosis and volatility clustering are found on preliminary inspection of the return series, which are consistent with the stylized facts mentioned by prior studies.

3.2 Implied Volatility Data

The Deribit Bitcoin Volatility Index (DVOL), which is derived from option prices traded on the Deribit exchange with a 30-day horizon, is utilized as an empirical benchmark for market-implied volatility to calibrate model-expected realized volatility proxies. Therefore, the DVOL index is transformed to hourly frequency to maintain consistency with the log return series and its values are converted into decimal form for estimation and calibration.

3.3 Sample Window Selection

The sample window is the most recent 180 trading days. The length of sample window is chosen for three reasons.

3.3.1 Parameter Stability

A large sample size is required for the GARCH model to estimate persistence parameters (α, β) reliably, so that a 180-day sample window provides sufficient observations for maximum likelihood estimation.

3.3.2 Structural Adaptability

Regime dependency is a characteristic of volatility behavior in cryptocurrency markets. Therefore, the model parameters based on a rolling 180-day sample window are adapted to evolving market conditions while maintaining a stable estimation environment.

3.3.3 Practical Implementation Considerations

Based on an empirical perspective, the rolling multi-month sample window is adopted in the institutional volatility model to maintain the relevance of model without excessive param- eter drift. Therefore, the 180-day maturity balances between responsiveness and statistical reliability.

Robustness checks in this paper utilize alternative window lengths (e.g., 90-day subsam- ples), which are presented in the empirical results section to evaluate the sensitivity of the model estimates.

3.4 Feature Construction and Alignment

All explanatory variables are derived from the raw dataset. The feature set includes:

• Log returns and rolling realized volatility measures;

• Conditional volatility and standardized residuals extracted from a Student-t GARCH (1,1) specification;

• Structural price features derived from rolling swing high–low ranges.

The hourly frequency is benchmark of frequency across all datasetsand observations are matched using nearest-hour timestamp alignment, with those containing missing values re- moved within the tolerance window. Moreover, an initial warm-up period caused by rolling-window computation is excluded from the estimation sample to prevent forward-looking bias and ensure that all model inputs are fully observable at time t.

4. METHODOLOGY

4.1 Econometric Volatility Backbone

The Student-t GARCH (1,1) model specification is widely applied to model the time-varying conditional variance of cryptocurrency returns.

Let r_t denote the hourly log return of the Bitcoin perpetual futures contract. In the GARCH specification, the return equation is defined as

$$r_t = \sigma_t \epsilon_t \tag{2}$$

where σ_t^2 represents the conditional variance and ϵt denotes standardized innovations.

The conditional variance evolves according to

$$\sigma_t^2 = \omega + \alpha r_{t-1}^2 + \beta\sigma_{t-1}^2 \tag{3}$$

where $\omega > 0$, $\alpha > 0$, and $\beta > 0$. α measures the sensitivity of volatility to new shocks, while β captures volatility persistence. The sum $\alpha + \beta$ reflects the degree of volatility clustering.

Given the heavy-tailed nature of cryptocurrency returns, the innovation term ϵt is as- sumed to follow a Student-t distribution. The parameters of the GARCH model are es- timated by maximum likelihood over a rolling 180-day window to capture the volatility dynamics under evolving market conditions.

The Student-t GARCH model constructs two key outputs:

• Conditional volatility σ_t

• Standardized residuals ϵ_t

These outputs are subsequently utilized as the econometric volatility inputs in the hybrid modeling framework.

4.2 LSTM-Based Nonlinear Volatility Proxy

After the GARCH model accounts for the linear volatility dynamics, the nonlinear temporal dependencies in financial time series are modeled using the LSTM model to construct the forward-looking volatility proxy.

At each time t, the input feature vector consists of:

• Log returns r_t

• Conditional volatility σ_t from the GARCH model

• Standardized residuals ϵ_t from the GARCH model

A fixed lookback window of past observation, which allow the LSTM model to learn the nonlinear temporal dependencies in volatility dynamics, is utilized as the sequence of input for LSTM model.

The expected realized volatility proxy RVt generated by the LSTM model represents the model-implied expected realized volatility state of Bitcoin, incorporating both the lin- ear volatility dynamics and the nonlinear temporal dependencies, thereby enhancing both

economic interpretability and predictive flexibility.

4.3 Structural Price Features and Market Structure

Market structure is the one of main explanations for volatility risk premium affecting option pricing. Therefore, structural price features are extracted to characterize the relative position of the current price within rolling price ranges before being calibrated with market-implied volatility.

For each time t, let Ht and Lt denote the highest and lowest prices over a rolling lookback window. A normalized structural distance measure is defined as

$$d_t = \frac{P_t \quad L_t}{H_t - L_t} \tag{4}$$

where P_t denotes the current price. After normalizing the distance of the current price within the rolling price range, additional features are constructed to measure the proximity between the normalized price position and Fibonacci retracement levels derived from the range. The final structural features are incorporated into the implied-volatility calibration stage to explain structural variation in the volatility risk premium embedded in option prices.

4.4 Calibration to Market-Implied Volatility

The calibration between the expected realized volatility proxy and market-implied volatility yields the volatility level used in option pricing, as implied volatility reflects both expected realized volatility and volatility risk premium.

The Deribit Volatility Index (DVOL) of Bitcoin options over a 30-day horizon serves as the benchmark for market-implied volatility in calibration process.

The calibration utilizes a linear mapping and accounts for structural variation in the volatility risk premium:

$$IV_t^{model} = a + b \cdot RV_t + \gamma \cdot d_t \tag{5}$$

where RV_t denotes the expected realized volatility proxy produced by the hybrid model, and dt represents the structural position between the current price and Fibonacci retracement levels.

The parameters of the calibration equation is estimated by ordinary least squares (OLS) over the same rolling window, providing an interpretable mapping between the expected realized volatility proxy and the volatility level used in option pricing.

4.5 Option Pricing Interpretation

Given spot price S_t, strike price K, and time to maturity τ, the calibrated implied volatility IV_t^{model} is applied in the Black–Scholes option pricing framework to compute the theoretical option prices and evaluate the economic relevance of the volatility estimates.

The implied volatility gap between market-implied volatility and model-implied volatility:

$$\Delta IV_t = IV_t^{mkt} - IV_t^{model} \tag{6}$$

measure the volatility risk premium embedded option prices in an economically inter- pretable way.

This framework decomposes market-implied volatility into two components: a realized- risk component captured by the hybrid model and a structural premium component related to market structure.

5. EMPIRICAL RESULTS

5.1 Volatility Proxy and Calibration

The expected realized volatility proxy generated by the GARCH–LSTM framework exhibits a strong positive correlation and statistically meaningful mapping in calibration with the market-implied volatility using DVOL, indicating that the expected realized volatility proxy can explain a substantial portion of the variation in market-implied volatility.

The slope coefficient of calibration regression is positive, which is consistent with fi- nancial intuition, suggesting that higher value of the expected realized volatility proxy is associated with higher market-implied volatility. RMSE and MAE, which are moderate cali- bration errors, demonstrate that the expected realized volatility proxy provides a reasonable approximation to market-implied volatility dynamics.

Overall, the GARCH-LSTM model succeeds in extracting an expected volatility factor economically consistent with observed option market pricing.

5.2 Option Pricing Consistency

In option pricing processing, the calibrated implied volatility is embedded into the Black–Scholes option pricing framework to assess the economic relevance.

Because the illustrative pricing exercise considers a 7-day option maturity, the volatility inputs are transformed to the 7-day horizon using the square-root-of-time rule, aligning the volatility measure with the maturity of the option used in the Black–Scholes pricing example.

There are three regimes of calibrations:

- The calibrated implied volatility without structural feature: $IV_{7d}^{model-baseline}$
- The calibrated implied volatility whose distance to the nearest Fibonacci retracement level: $IV_{7d}^{near\text{-}fib}$
- The calibrated implied volatility under the binary indicator representing proximity to Fibonacci levels fib min

At the final observation (March 4, 2026), the following values are observed:

$$S = 72092.5, \qquad K = 72000, \qquad \tau = \frac{7}{365} \tag{7}$$

The market-implied volatility derived from DVOL is

$$IV_{7d}^{mkt} = 0.2621 \tag{8}$$

while the calibrated model implied volatility in three regimes is:

$$IV_{7d}^{model-baseline} = 0.2327 \tag{9}$$

$$IV_{7d}^{near-fib} = 0.2324 \tag{10}$$

$$IV_{7d}^{fib\_min} = 0.2296 \tag{11}$$

Using these volatility inputs, the corresponding Black–Scholes option prices are: Market price:

$$C^{mkt} = 1090.10 \tag{12}$$

Prices in three regimes:

$$C^{model-baseline} = 973.04 \tag{13}$$

$$C^{near-fib} = 971.86 \tag{14}$$

$$C^{fib\text{-}min} = 961.13 \tag{15}$$

The results in a pricing difference of

$$\Delta C^{model-baseline} = 117.06 \tag{16}$$

$$\Delta C_{near\_fib} = 118.25 \tag{17}$$

$$\Delta C_{fib\_min} = 128.94 \tag{18}$$

Importantly, the differences between market-implied volatility and the calibrated implied volatility closely match the vega approximations in three regimes:

$$Vega \times (IV_{7d}^{mkt} - IV_{7d}^{model-baseline}) \approx 117.07 \tag{19}$$

$$Vega \times (IV_{7d}^{mkt} - IV_{7d}^{near\text{-}fib}) \approx 118.25 \tag{20}$$

$$Vega \times (IV_{7d}^{mkt} - IV_{7d}^{fib\text{-}min}) \approx 128.98 \tag{21}$$

There are three contributions from the above empirical results

First the near equivalence between the actual pricing difference and the vega-based ap- proximation confirms internal consistency between the volatility calibration and the optionpricing framework.

Second, the structural feature influences the magnitude of the volatility gap between model-implied and market-implied volatility.

Third, the volatility gap relative to the market tends to increase when the underlying price is further away from structural Fibonacci levels. This contribution indicates that option markets may embed additional volatility risk premium when the price deviates from structural price regions, and structural market features contain incremental explanatory power for implied volatility modeling beyond the statistical realized volatility dynamics.

5.3 Robustness Checks

The robustness checks are organized along two dimensions:

(1) the assumed innovation distribution in the GARCH backbone (Student-t versus normal)

(2) the estimated sample window length (180 days versus 90 days)

Both a baseline calibration without structural feature and an augmented calibration that includes a Fibonacci-based structural premium term will be reported for each specification.

Table 1. Robustness checks across sample windows and GARCH distributions

Setting	Model	RMSE	MAE	R^2	γ	$p(\gamma)$	Call Diff
180d-normal	Baseline	0.06508	0.05406	0.30518	–	–	100.57
180d-normal	+Fib Premium	0.06521	0.05412	0.30240	-0.04295	7.08×10^{-5}	111.88
180d-t	Baseline	0.06535	0.05436	0.29935	–	–	98.79
180d-t	+Fib Premium	0.06548	0.05442	0.29663	-0.04247	9.08×10^{-5}	109.98
90d-t	Baseline	0.08121	0.07804	-9.79	_	_	130.29
90d-t	+Fib Premium	0.08213	0.07882	-10.03	0.11131	6.26×10^{-14}	99.91

According to Table 1 the overall calibration error metrics maintain similar magnitude across the Student-t and normal GARCH assumptions under the same estimated sample window length, suggesting that a particular distributional choice does not affect the quali- tative results.

Additionally, the shorter calendar span substantially reduces the effective information available for calibration in the cryptocurrency market characterized by volatility clustering and heavy-tailed dynamics, making the estimates more sensitive to local extreme observa- tions and regime-specific dynamics. Both RMSE and MAE of the 90-day setting increase, while R2 of the 90-day setting turns negative, suggesting that the specification performs worse than a naive mean benchmark over that local sample. The reversal of γ from astable negative coefficient in the 180-day specifications to a positive coefficient in the 90-day specification further indicates that the structural premium term is not sign-robust under a short estima- tion window. Therefore, the 90-day outcome supports that a multi-month rolling window is a practical trade-off between adaptability and statistical reliability.

The pricing example presented in the previous subsection represents a single-point il- lustration of the implied-volatility gap at the final observation, rather than the robustness analysis which evaluates the statistical relationship between the volatility proxy and market- implied volatility across the entire estimation sample. As a result, the magnitude of the pricing gap observed in the prior example may differ from the average effects observed in the regression-based robustness analysis. This discrepancy reflects the time-varying nature of the volatility risk premium in cryptocurrency option markets.

6. DISCUSSION

The empirical results of this paper demonstrate that the hybrid model decomposes the implied volatility in economically meaningful terms.

The expected realized volatility proxy captures the realized risk component embedded in cryptocurrency price dynamics, while the calibration step maps this proxy into a market- consistent implied volatility measure.

Furthermore, the analysis of Fibonacci-based structural features indicates that market structure may influence the magnitude of this volatility premium. When prices deviate from structural price regions, option markets tend to increase implied volatility relative to the model-implied realized volatility proxy, consistent with the presence of addition risk compensation under structurally uncertain market conditions.

6.1 Limitations

There are several limitations in the current framework.

First, the sample window length is a 180-day rolling window. The longer historical windows may provide additional information about long-term volatility dynamics, which may be beneficial to parameter stability and structural adaptability.

Second, the research on volatility dynamics are restricted to price data, while the market microstructure variables, such as order-book liquidity, funding rates, and liquidation flows, may contain additional predictive signals for option-implied volatility.

Third, the framework focuses primarily on Fibonacci-based retracement levels derived from rolling price ranges. However, more sophisticated structural features may further im- prove the performance of modeling framework.

7. CONCLUSION

This study proposes a hybrid model incorporating a Student-t GARCH model and a LSTM model, with calibration under structural price features, providing an economically inter- pretable decomposition of option-implied volatility and a new approach to option pricing.

The empirical analysis demonstrates that the expected realized volatility proxy generated by the hybrid model exhibits a strong positive relationship with market-implied volatility, producing economically meaningful option prices within a Black–Scholes framework after calibration. The close equivalence between pricing differences and the vega-based approxi- mation further enhances the internal consistency of the model.

Furthermore, the structural features derived from Fibonacci retracement levels are shown to explain the volatility risk premium component in implied volatility modeling by influencing the gap between the model-implied and market-implied volatility. When the underlying asset prices deviate from structural price regions, option markets appear to embed larger volatility premium.

Nevertheless, the window length, limited raw data sources and limited structural features remain significant limitations and require further investigation in further research.

The proposed framework provides a transparent bridge between realized volatility mod- eling and option pricing interpretation, offering a practical approach for volatility analysis and option pricing analysis in cryptocurrency derivatives markets.

REFERENCES

[1] Elie Bouri, Peter Moln´ar, Georges Azzi, David Roubaud, and Lars Hagfors. On the hedge and safe haven properties of bitcoin: Is it really more than a diversifier? Finance Research Letters, 20:192–198, 2017.
[2] Andrew Phillip, Joshua Chan, and Shelton Peiris. A new look at cryptocurrencies. Economics Letters, 163:6–9, 2018.
[3] David Ardia, Keven Bluteau, and Maxime R¨uede. Regime changes in bitcoin garch volatility dynamics. Finance Research Letters, 29:266–271, 2019.
[4] Bent J. Christensen and Narasimhan R. Prabhala. The relation between implied and realized volatility. Journal of Financial Economics, 50(2):125–150, 1998.
[5] Tim Bollerslev, George Tauchen, and Hao Zhou. Expected stock returns and variance risk premia. Review of Financial Studies, 22(11):4463–4492, 2009.
[6] Nicolas P. B. Bollen and Robert E. Whaley. Does net buying pressure affect the shape of implied volatility functions? Journal of Finance, 59(2):711–753, 2004.
[7] Tim Bollerslev. Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics, 31(3):307–327, 1986.
[8] Torben G. Andersen, Tim Bollerslev, Francis X. Diebold, and Paul Labys. Modeling and forecasting realized volatility. Econometrica, 71(2):579–625, 2003.
[9] Sepp Hochreiter and J¨urgen Schmidhuber. Long short-term memory. Neural Computa- tion, 9(8):1735–1780, 1997.
[10] Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nature Machine Intelligence, 1(5):206– 215, 2019.
[11] Robert F. Engle. Autoregressive conditional heteroskedasticity with estimates of the variance of united kingdom inflation. Econometrica, 50(4):987–1007, 1982.

[12] Tim Bollerslev. A conditional heteroskedastic time series model for speculative prices and rates of return. Review of Economics and Statistics, 69(3):542–547, 1987.
[13] Ruey S. Tsay. Analysis of Financial Time Series. Wiley, 2005.
[14] Thomas Fischer and Christopher Krauss. Deep learning with long short-term memory networks for financial market predictions. European Journal of Operational Research, 270(2):654–669, 2018.
[15] David M.Q. Nelson, Adriano C.M. Pereira, and Renato A. de Oliveira. Stock market's price movement prediction with lstm neural networks. International Joint Conference on Neural Networks, pages 1419–1426, 2017.
[16] John C. Hull. Options, Futures, and Other Derivatives. Pearson, 10edition, 2018.
[17] Jean-Philippe Bouchaud, J. Doyne Farmer, and Fabrizio Lillo. How markets slowly digest changes in supply and demand. In Handbook of Financial Markets: Dynamics and Evolution, pages 57–160. Elsevier, 2009.