

Designing a Predictive and Explainable Model for Stablecoin Prices by Integrating On-Chain Data and Deep Learning-Based Market Sentiment Analysis

1. Mohammadsaeid. Daliri¹: Department of IT Management, QA.C., Islamic Azad university , Qazvin, Iran
2. Mohammadreza. Sanaei^{2*}: Department of IT Management, QA.C., Islamic Azad university , Qazvin, Iran
3. Roozbeh. Balounejad Nouri³: Department of Economic Science, QA.C., Islamic Azad university, Qazvin, Iran

*corresponding author's email: mr.sanaei@iau.ac.ir

ABSTRACT

This study aimed to design and evaluate an explainable deep learning model for predicting stablecoin prices and peg deviations by integrating on-chain indicators, conventional market variables, and market sentiment data. This applied predictive-analytical study was conducted in Tehran, Iran, during 2025–2026. The behavioral validation sample consisted of 384 active cryptocurrency market participants. The longitudinal dataset included 2,190 synchronized daily observations for USDT, USDC, and DAI. Data comprised historical prices, trading characteristics, transaction volume, active addresses, exchange inflows and outflows, net exchange flows, large-holder transactions, supply changes, and deep learning-derived sentiment indicators. Sentiment was extracted from cryptocurrency-related textual data using transformer-based natural language processing. An integrated Attention-LSTM model was developed and compared with Support Vector Regression, Random Forest, Extreme Gradient Boosting, conventional LSTM, and GRU models. Chronological training, validation, and test partitions were used, and SHAP analysis assessed feature importance and direction. The integrated Attention-LSTM model demonstrated the strongest out-of-sample performance, with MAE = 0.00196, RMSE = 0.00324, MAPE = 0.195%, $R^2 = 0.939$, and directional accuracy = 83.23%. It outperformed the Attention-LSTM model without sentiment, which achieved RMSE = 0.00453 and $R^2 = 0.881$. Ablation analysis showed that the full integrated model reduced RMSE by 53.31% relative to the historical-price-only model. The model achieved R^2 values of 0.947, 0.942, and 0.921 for USDT, USDC, and DAI, respectively, and directional accuracy reached 90.32% during high-deviation periods. SHAP analysis identified negative sentiment intensity, net exchange inflow, previous-day peg deviation, on-chain transaction volume, sentiment dispersion, and large-holder transaction volume as the most influential predictors. Integrating on-chain, market, and sentiment information within an explainable Attention-LSTM framework substantially improves stablecoin price forecasting and provides interpretable early-warning signals of peg instability.

Keywords: Stablecoin; Price Prediction; Deep Learning; Attention-LSTM; On-Chain Data; Market Sentiment; Explainable Artificial Intelligence; Cryptocurrency

Introduction

The rapid expansion of cryptocurrencies has transformed the structure of contemporary financial markets by introducing decentralized, programmable, and globally accessible digital assets whose prices are shaped by a complex interaction of technological, behavioral, and market forces. Unlike conventional financial assets,



Article history:
Received 03 September 2026
Revised 07 September 2026
Accepted 10 September 2026
Initial Publish 16 October 2026
Published online 10 April 2027

How to cite this article:

Daliri, M., Sanaei, M., & Balounejad Nouri, R. (2027). Designing a Predictive and Explainable Model for Stablecoin Prices by Integrating On-Chain Data and Deep Learning-Based Market Sentiment Analysis. *Journal of Management and Business Solutions*, 5(4), 1-23. <https://doi.org/10.61838/jmbs.433>



cryptocurrencies trade continuously across decentralized and centralized venues, are exposed to heterogeneous investor populations, and respond rapidly to information disseminated through social media, news platforms, blockchain activity, and macro-financial developments. These characteristics have generated substantial interest in the development of predictive models capable of capturing nonlinear patterns and short-term dependencies in cryptocurrency prices. Deep learning has become particularly influential in this context because traditional linear time-series techniques are often unable to adequately represent the nonstationary, volatile, and highly interactive nature of crypto-asset markets (1, 2). The growing body of research on cryptocurrency forecasting therefore increasingly emphasizes architectures such as recurrent neural networks, Long Short-Term Memory networks, hybrid deep learning systems, and ensemble learning models that can process high-dimensional market data and identify latent temporal dependencies that are difficult to detect through conventional econometric approaches (3-5).

Within the broader cryptocurrency ecosystem, stablecoins occupy a distinctive position because they are designed to maintain a relatively stable value, usually by linking their market price to a reference asset such as the U.S. dollar. Their intended stability has made them integral to cryptocurrency trading, decentralized finance, liquidity management, settlement, collateralization, and the movement of capital across digital-asset markets. However, the assumption that stablecoins remain continuously anchored to their target value is not always supported by actual market behavior. Stablecoins can deviate from their peg in response to liquidity stress, changes in reserve confidence, sudden redemption pressure, collateral instability, exchange disruptions, market panic, regulatory uncertainty, or broader crypto-market shocks. These deviations may be temporary and small, but they can also become economically significant, particularly during periods of systemic stress. Research on the stability of crypto-assets has therefore emphasized that stablecoins should not be treated merely as fixed-value instruments, but rather as dynamically priced assets whose apparent stability may conceal meaningful short-term fluctuations and structural vulnerabilities (6, 7). High-frequency evidence further indicates that stablecoins can display distinct microstructural and price-adjustment behaviors, reinforcing the need for models that are specifically designed for their unique market dynamics rather than simply adapting conventional cryptocurrency forecasting methods (7).

The challenge of predicting stablecoin prices is fundamentally different from predicting highly volatile assets such as Bitcoin or Ethereum. In conventional cryptocurrencies, predictive models are often evaluated primarily on their ability to anticipate direction, returns, or large price movements. In stablecoin markets, however, the economically relevant variation is often concentrated in small deviations around the target peg. This makes the forecasting problem more demanding because the model must distinguish routine micro-fluctuations from meaningful early signals of de-pegging. Small numerical errors may have substantial economic implications when the underlying asset is expected to trade near one dollar. Moreover, stablecoins differ in their stabilization mechanisms, reserve structures, collateral designs, redemption processes, and degrees of decentralization. Accordingly, their price behavior can reflect not only historical market data but also token issuance and redemption, blockchain transfers, exchange flows, changes in wallet concentration, collateral conditions, and the expectations of market participants. This multidimensional structure suggests that prediction systems based only on historical prices or technical indicators are unlikely to capture the full informational environment relevant to stablecoin stability.

The broader cryptocurrency forecasting literature has demonstrated that deep learning methods offer clear advantages in extracting temporal and nonlinear relationships from financial data. Early and subsequent studies have shown that recurrent architectures can outperform simpler models when price formation depends on complex

historical interactions and nonlinear dynamics (3, 5). Time-series deep learning methods have also been shown to effectively model cryptocurrency prices by capturing patterns across sequential observations and adapting to rapidly changing market conditions (2). Ensemble deep learning models extend this capability by combining complementary learners or representations, thereby reducing model-specific error and improving predictive robustness (4). More recent comparative evidence likewise suggests that advanced deep learning and ensemble learning approaches can achieve strong cryptocurrency forecasting performance, although the relative success of different architectures depends heavily on the type of input information, time horizon, market regime, and evaluation strategy (8). These findings establish an important methodological foundation for stablecoin prediction, but they also reveal a persistent limitation: many existing forecasting studies rely heavily on historical price and market variables while insufficiently integrating blockchain-native and sentiment-related information.

This limitation is important because blockchain markets generate a level of transactional transparency that is rarely available in conventional financial systems. On-chain data provide direct measures of economic activity within blockchain networks and can reveal changes in token circulation, exchange flows, active addresses, transfer volumes, large-holder transactions, minting and burning behavior, and concentration of holdings. In the case of stablecoins, such indicators are particularly relevant because price stability may depend on the capacity of the market to absorb redemption pressure, maintain liquidity, and preserve confidence in the issuing or collateralization mechanism. An increase in exchange inflows, for example, may indicate a shift in selling pressure or liquidity demand, while sudden changes in large-holder transfers may reflect repositioning by influential market participants. Similarly, changes in token supply can reveal minting or redemption activity that precedes price adjustment. These indicators are fundamentally different from historical price variables because they capture behavioral and transactional activity at the infrastructure level rather than merely the market outcome generated by that activity. A model that combines these data with conventional prices may therefore detect early signs of instability before they become fully reflected in observed market prices.

The importance of richer information sets is also reflected in the growing literature on high-frequency cryptocurrency forecasting. Akouaouch and Bouayad introduced a deep learning framework for high-frequency short-term cryptocurrency price prediction, highlighting the ability of deep models to extract meaningful patterns under fast-changing conditions (9). High-frequency environments are particularly relevant for stablecoins because deviations from the peg can develop rapidly and may resolve before they are visible in low-frequency summaries. This raises the methodological importance of feature synchronization, lag construction, and temporal modeling. Similarly, Bakhit and Rabiei emphasized the potential of deep machine learning techniques for forecasting virtual cryptocurrency prices, reinforcing the broader shift from purely statistical forecasting toward data-driven learning systems (10). Collectively, these studies suggest that predictive performance depends not only on the choice of algorithm but also on whether the model receives sufficiently informative variables that represent the underlying mechanisms of price formation.

Market sentiment constitutes another critical dimension of cryptocurrency pricing. Digital-asset markets are unusually sensitive to rapidly changing narratives disseminated through social media, online communities, specialized financial media, and investor discussions. Unlike traditional financial markets, where institutional research and regulated disclosures often dominate information flow, cryptocurrency prices can respond almost immediately to shifts in collective expectations expressed through decentralized communication channels. Sentiment analysis therefore offers a way to quantify behavioral information that may not yet be visible in price or

volume data. Early research integrating sentiment and cryptocurrency prediction demonstrated that deep learning-based sentiment representations can contribute useful information for forecasting market movements (11). Later studies examining Twitter-based cryptocurrency sentiment further showed that social-media discussions contain detectable patterns of positive, negative, and neutral attitudes that can reflect changing perceptions of digital assets (12). These findings are particularly relevant to stablecoins because concerns about reserves, liquidity, redemptions, exchange solvency, or de-pegging can spread rapidly through social platforms and trigger behavioral responses before the full market effect is visible in transactional data.

Sentiment may also behave nonlinearly across market regimes. The informational effect of sentiment is unlikely to be constant under stable and stressed conditions. A small increase in negative sentiment during calm periods may have little effect, whereas a similar increase during a period of uncertainty may amplify selling pressure and liquidity withdrawal. Recent evidence suggests that extreme sentiment regimes can be associated with adverse selection and market inefficiency in cryptocurrency markets, indicating that the predictive relevance of sentiment may be particularly strong when emotions become unusually concentrated or polarized (13). This is consistent with findings showing that global uncertainty and sentiment factors can play an important role in cryptocurrency market crashes and stress episodes (14). For stablecoins, this implies that sentiment should not be treated merely as a secondary contextual variable. Instead, it may constitute an early-warning component that interacts with exchange flows, transaction intensity, liquidity conditions, and prior peg deviations.

The behavioral role of sentiment is also supported by evidence from broader financial and cryptocurrency settings. Nourollahi, Jafari, and Pakmaram examined the effect of investor sentiment on portfolio optimization across the Tehran Stock Exchange and cryptocurrency markets, emphasizing that investor psychology has practical implications for asset allocation and decision-making in both traditional and digital markets (15). Their findings are important for the present context because they demonstrate that sentiment is not limited to online textual expression; it also reflects the expectations and risk perceptions of actual market participants. In stablecoin markets, confidence in the peg, perceived redemption safety, and expectations regarding liquidity may influence whether investors hold, exchange, or redeem stablecoins. These behavioral responses can subsequently become visible in blockchain and market variables. The combination of computational sentiment measures with participant-based perception data may therefore enhance the external validity of sentiment indicators by confirming that text-derived signals correspond to actual market expectations rather than merely reflecting linguistic noise.

The interaction between sentiment and broader market conditions is especially important during crisis periods. Soltani, Taleb, and Boujelbene Abbes documented spillover effects among stock markets, cryptocurrencies, and investor sentiment during the COVID-19 period, demonstrating that behavioral and financial shocks can transmit across markets rather than remaining isolated within a single asset class (16). Such spillovers are highly relevant to stablecoins because their price stability depends partly on the broader liquidity and risk environment in which they operate. During episodes of systemic uncertainty, investors may rapidly rotate between volatile cryptocurrencies, stablecoins, fiat currencies, and other assets. Consequently, a stablecoin may temporarily experience unusually strong demand as a safe-haven-like instrument or, conversely, intense redemption pressure if confidence in its backing or exchangeability deteriorates. A forecasting system that incorporates both broader market variables and sentiment information is therefore better positioned to distinguish general crypto-market stress from stablecoin-specific instability.

Another important challenge in the literature concerns model interpretability. Deep learning models are often criticized as black-box systems because they can achieve strong predictive performance without providing a clear explanation of why particular predictions were made. This problem is especially important for stablecoins, where users, financial institutions, regulators, exchanges, and risk managers may need to understand whether a predicted de-pegging event is being driven by sentiment deterioration, exchange inflows, supply changes, whale activity, market volatility, or some combination of these factors. A model that generates accurate forecasts but cannot identify the variables contributing to the forecast has limited practical value for decision support. The growing emphasis on explainable artificial intelligence therefore creates an opportunity to combine high-performing deep learning architectures with methods that attribute model outputs to individual predictors. In this context, feature-attribution techniques such as SHapley Additive exPlanations can quantify the contribution of market, on-chain, and sentiment variables to specific predictions, thereby converting a purely predictive model into an interpretable risk-monitoring framework.

This emphasis on explainability also addresses an important gap in cryptocurrency forecasting research. Much of the existing literature focuses on whether one algorithm performs better than another, often comparing LSTM, GRU, ensemble learning, support vector methods, or tree-based techniques (4, 8). Although such comparisons are methodologically useful, they do not necessarily explain the economic mechanisms underlying prediction. For stablecoins, the central research question is not only whether future price deviations can be forecast, but also which market processes are most responsible for those deviations. Historical prices may capture persistence, but they do not identify whether the underlying cause is redemption pressure, negative sentiment, exchange concentration, broader market volatility, or sudden on-chain transfer activity. Thus, an explainable multimodal model is conceptually superior to a purely price-based black-box forecast because it links predictive accuracy with economically meaningful interpretation.

The literature also indicates that no single data source is likely to provide a complete representation of cryptocurrency price formation. Historical prices summarize previous market outcomes, on-chain variables capture transactional and network behavior, and sentiment data approximate behavioral expectations. Each source therefore represents a different layer of information. Deep learning is particularly suitable for integrating these heterogeneous inputs because it can learn complex nonlinear representations from high-dimensional data and combine multiple feature domains within a unified architecture (1, 3). The use of attention mechanisms can further improve this process by allowing the model to assign greater weight to temporally or informationally salient observations. This is important for stablecoins because the most relevant signals may occur only during short episodes of instability. For example, a sudden increase in negative sentiment may have limited predictive importance during a stable period but become highly informative when combined with rising exchange inflows and prior deviations from the peg.

The need for integrated modeling is further reinforced by the distinctive risk structure of stablecoins. Bullmann, Klemm, and Pinna raised fundamental questions regarding whether stablecoins can actually provide the stability they are intended to deliver, emphasizing that stability mechanisms themselves introduce specific risks and constraints (6). Barucci, Moncayo, and Marazzina similarly showed that high-frequency stablecoin behavior warrants detailed empirical analysis because stablecoins are not perfectly invariant around their target values (7). These observations imply that stablecoin prediction should focus not only on average price accuracy but also on abnormal deviations from the peg. A model may appear highly accurate if it simply predicts values close to one

dollar in a dataset dominated by normal observations. However, such a model may fail precisely when prediction is most economically valuable—during rare but consequential de-pegging events. Accordingly, an appropriate modeling framework must assess performance under both normal and stressed regimes and should separately examine its capacity to identify moderate and severe peg deviations.

This consideration also has implications for evaluation methodology. Random train-test splitting, commonly used in general machine learning research, can generate unrealistic results in financial time series because future observations may indirectly influence model development through temporal leakage. Chronological partitioning is therefore essential for evaluating whether the model can generalize to genuinely unseen future periods. Similarly, benchmark comparison should include not only alternative deep learning models but also traditional machine learning approaches in order to determine whether the additional complexity of deep learning produces meaningful performance gains. Bouteska et al. emphasized the importance of comparative evaluation between ensemble and deep learning methods in cryptocurrency forecasting (8), while Pintelas et al. highlighted the need to carefully evaluate deep learning architectures within the specific characteristics of cryptocurrency prediction problems (5). These considerations support a study design that combines chronological validation, benchmark comparison, ablation analysis, and explainability assessment.

An additional gap concerns the limited integration of market-participant perceptions into computational cryptocurrency forecasting. While textual sentiment analysis can process large volumes of online information, it may also be affected by bots, duplicated content, promotional messages, sarcasm, and platform-specific linguistic patterns. Behavioral validation using actual market participants can therefore improve confidence in computational sentiment measures. In the Iranian context, this issue is especially relevant because cryptocurrency participation has expanded despite institutional, monetary, and regulatory constraints, and market users may interact with digital assets for trading, hedging, settlement, and value preservation. Incorporating data from active cryptocurrency participants in Tehran provides an additional behavioral layer that can be used to test whether deep learning-derived sentiment scores align with human perceptions of stablecoin safety, de-pegging risk, and market confidence. This does not replace large-scale digital sentiment analysis, but it provides a useful external validity check.

Taken together, the existing literature demonstrates four important points. First, cryptocurrency prices exhibit nonlinear and temporally dependent behavior that can be effectively modeled using deep learning (2, 9). Second, ensemble and hybrid learning approaches can improve forecasting performance when they capture complementary representations of market information (4, 8). Third, investor sentiment and uncertainty are important determinants of cryptocurrency market behavior, especially during periods of stress and regime change (11-14). Fourth, stablecoins require distinct analytical treatment because their economic relevance lies in maintaining a peg and because deviations from that peg may be driven by both market and structural mechanisms (6, 7). Despite these advances, few studies have simultaneously integrated historical market variables, on-chain indicators, deep learning-based sentiment measures, participant-level behavioral validation, and explainable artificial intelligence within a unified stablecoin forecasting framework.

Addressing this gap is important both theoretically and practically. From a theoretical perspective, an integrated model can clarify whether stablecoin instability is better understood as a purely market-based phenomenon or as the result of interacting technological, behavioral, and liquidity mechanisms. From a practical perspective, improved prediction and explanation of stablecoin prices can support traders, exchanges, portfolio managers, decentralized finance participants, and risk-monitoring institutions by identifying not only when peg instability may occur, but also

which variables are contributing to the emerging risk. Deep learning can provide the predictive capacity needed for complex nonlinear relationships, while on-chain data can capture blockchain-native behavior and sentiment analysis can quantify collective market expectations. When combined with explainability techniques, these elements can provide a more transparent and decision-relevant framework than conventional black-box forecasting models.

Accordingly, the aim of the present study was to design and evaluate an explainable deep learning model for predicting stablecoin prices and peg deviations by integrating on-chain data, conventional market indicators, and market sentiment information, while also examining the incremental predictive contribution and interpretive importance of each information domain.

Methods and Materials

The present study employed an applied, quantitative, predictive-analytical research design aimed at developing and evaluating an explainable deep learning model for predicting stablecoin prices through the integration of on-chain indicators and market sentiment information. The study was conducted in Tehran, Iran, during 2025–2026. Because stablecoin price dynamics are influenced simultaneously by blockchain-level activity, market microstructure, liquidity conditions, and behavioral reactions of market participants, a multimodal time-series modeling framework was adopted. The analytical framework combined structured numerical data derived from blockchain networks and cryptocurrency markets with unstructured textual data reflecting investor sentiment. To supplement computational sentiment indicators and provide a Tehran-based behavioral component for model development and validation, the study population included active cryptocurrency market participants residing in Tehran who had practical experience in buying, selling, or holding digital assets and stablecoins. Using purposive eligibility screening, 384 participants were included in the final sample. Eligibility criteria consisted of residence in Tehran, age of at least 18 years, a minimum of one year of cryptocurrency trading or investment experience, familiarity with at least one major stablecoin, such as Tether (USDT), USD Coin (USDC), or DAI, and active participation in cryptocurrency markets during the six months preceding data collection. Individuals who provided incomplete responses or reported no recent market activity were excluded. The participant-based data were used primarily to validate the direction and intensity of computationally extracted market sentiment rather than as the sole source of predictive variables. The principal unit of analysis for the prediction model was therefore the daily observation of stablecoin market behavior. Daily observations were synchronized across on-chain, market, and sentiment datasets to construct a unified longitudinal dataset. The principal dependent variable was the daily market price of the selected stablecoins relative to the U.S. dollar, with particular emphasis on deviations from the intended one-dollar peg. The modeling architecture was designed to estimate both the subsequent stablecoin price and the probability and magnitude of short-term de-pegging. To prevent information leakage, all observations were ordered chronologically, and future information was not incorporated into the construction of predictors for earlier periods.

Data collection was conducted using three complementary sources consisting of on-chain blockchain indicators, cryptocurrency market variables, and market sentiment data. On-chain data were collected for the selected stablecoins and their underlying blockchain ecosystems and included daily transaction volume, transaction count, number of active addresses, number of new addresses, exchange inflows, exchange outflows, net exchange flow, circulating supply, token issuance and redemption, large-holder or whale transactions, transfer value, network activity, concentration of holdings, and changes in major wallet balances. Where applicable, smart-contract information related to minting, burning, collateralization, and token redemption was also incorporated. These

variables were selected because they represent liquidity movements, changes in token supply, investor activity, and potential pressure on the stability mechanism of stablecoins. For decentralized or collateralized stablecoins, additional indicators relating to collateral conditions and changes in backing assets were considered whenever consistent data were available. All blockchain variables were transformed into a uniform daily frequency, and variables with substantially different scales were normalized before their inclusion in the deep learning models.

Market-level data consisted of the daily opening price, highest price, lowest price, closing price, trading volume, market capitalization, liquidity measures, bid-ask-related indicators where available, and the return and realized volatility of each stablecoin. Because deviations from the reference value are particularly important in stablecoin markets, additional variables were constructed to represent absolute price deviation from USD 1.00, percentage deviation from the peg, duration of deviation, maximum daily deviation, and the frequency of deviations exceeding predetermined thresholds. Broader cryptocurrency market indicators were also collected, including Bitcoin and Ethereum returns and volatility, overall cryptocurrency market capitalization, major market trading volume, and selected indicators of systemic market conditions. Lagged variables and rolling statistics were generated to allow the models to capture temporal dependencies. These included one-day, three-day, seven-day, fourteen-day, and thirty-day lags as well as rolling means, rolling standard deviations, momentum indicators, and rates of change. All timestamps were standardized prior to merging the datasets so that information belonging to the same analytical period was accurately aligned.

Market sentiment constituted the second major explanatory domain of the study. Textual information related to USDT, USDC, DAI, stablecoins, de-pegging, reserves, liquidity, cryptocurrency risk, exchange conditions, and related market events was collected from publicly available cryptocurrency-related social media posts, discussion channels, news headlines, and online financial content. After removal of duplicate records, advertisements, hyperlinks, spam-like messages, irrelevant content, and records with insufficient textual information, the remaining texts were subjected to preprocessing. Preprocessing included normalization of characters, removal of unnecessary symbols, tokenization, handling of emojis and hashtags when they contained sentiment-relevant information, and linguistic normalization. Persian and English content were processed separately during the initial preprocessing stage to reduce distortions caused by language-specific structures and were subsequently converted into comparable sentiment representations. Deep learning-based natural language processing was used to classify the textual observations into positive, negative, and neutral market sentiment and to estimate a continuous sentiment intensity score. Transformer-based language representations were employed because they permit contextual interpretation of financial statements and can distinguish between superficially similar words used in different market contexts. Daily sentiment indices were then constructed on the basis of the volume, polarity, and intensity of relevant textual records. Measures of sentiment dispersion, negative-sentiment concentration, changes in sentiment, and the difference between positive and negative sentiment were also calculated. The questionnaire data obtained from the 384 Tehran-based cryptocurrency participants included perceived market confidence, perceived stablecoin safety, expected price stability, perceived de-pegging risk, liquidity concerns, and short-term market expectations. These responses served as an external behavioral validation source for the machine-generated sentiment indicators. Data from all sources were subsequently integrated by date and stablecoin identifier to construct the final modeling dataset.

Data analysis was conducted in several sequential stages consisting of preprocessing, feature engineering, sentiment modeling, predictive modeling, explainability analysis, and model evaluation. Initially, all datasets were

examined for missing values, duplicated observations, inconsistent timestamps, extreme values, and anomalous observations. Missing values were treated according to the nature and temporal structure of each variable; short discontinuities in continuous time-series variables were handled using appropriate time-series interpolation procedures, whereas observations with substantial missing information were excluded. Extreme observations were not automatically removed because abrupt changes in transaction activity, exchange flows, liquidity, sentiment, and prices may represent genuine de-pegging or market-stress events. Instead, suspicious observations were examined in relation to adjacent time points and corresponding market conditions. Numerical variables were standardized or normalized to prevent variables measured on larger scales from disproportionately affecting model training. Stationarity characteristics and temporal patterns of the principal variables were also assessed, while first differences, logarithmic transformations, percentage changes, and rolling measures were generated where necessary to obtain informative representations of changes in market conditions.

Feature engineering was performed to transform raw blockchain and market observations into variables capable of representing short-term and medium-term dynamics. In addition to contemporaneous indicators, lagged observations and rolling-window statistics were constructed for the principal on-chain, price, volume, liquidity, and sentiment variables. Correlation analysis and feature-importance screening were initially used to identify highly redundant variables, while final feature selection was determined within the predictive modeling process to avoid eliminating variables with nonlinear explanatory value. The complete feature set was organized into three conceptual channels consisting of on-chain features, conventional market features, and sentiment features. These channels were first analyzed separately and subsequently integrated to determine whether multimodal data fusion improved predictive accuracy relative to models relying exclusively on historical price information or a single category of predictors.

The principal predictive framework was based on deep learning for multivariate time-series analysis. Long Short-Term Memory networks were used as the primary sequential architecture because of their ability to model long-range temporal dependencies and nonlinear relationships in financial time series. Gated Recurrent Unit and temporal attention-based models were additionally estimated as comparative deep learning architectures. The sentiment representation generated from transformer-based textual analysis was synchronized with corresponding daily on-chain and market variables before being entered into the integrated model. In the final architecture, separate feature representations were initially learned for blockchain-market variables and sentiment variables, after which the latent representations were concatenated and passed through fully connected layers to generate the final price prediction. An attention mechanism was incorporated to allow the model to assign different weights to temporally and informationally important observations. The principal model output consisted of the predicted closing price for the subsequent trading period. A complementary output estimated the expected absolute deviation from the USD 1.00 peg so that the performance of the model could be evaluated not only in terms of general price prediction but also in relation to economically meaningful de-pegging behavior.

To obtain a reliable out-of-sample assessment, the dataset was divided chronologically rather than randomly. Seventy percent of the earliest observations were allocated to the training set, 15% to the validation set, and the final 15% to the test set. This procedure ensured that observations used for testing occurred chronologically after those used for model estimation and reduced the risk of look-ahead bias. Hyperparameters, including the number of hidden layers, number of neurons, sequence length, learning rate, batch size, dropout rate, and number of training epochs, were optimized using the validation dataset. Early stopping was employed to terminate training

when validation performance ceased to improve, and dropout and regularization procedures were used to reduce overfitting. The Adam optimization algorithm was employed for network training. Model performance was assessed using Mean Absolute Error, Mean Squared Error, Root Mean Squared Error, and Mean Absolute Percentage Error. The coefficient of determination was additionally calculated to evaluate the proportion of variation in observed stablecoin prices explained by the predictions. Directional accuracy was examined to determine the ability of the model to correctly identify the direction of price movements and changes in deviation from the peg.

To determine the incremental contribution of each information source, several benchmark and ablation models were estimated. The performance of a historical-price-only model was first established as the baseline. Models containing price and conventional market variables, on-chain variables, and sentiment variables were then evaluated separately. Subsequently, an integrated model incorporating historical prices, market indicators, on-chain variables, and deep learning-based sentiment measures was estimated. Comparison of the prediction errors across these specifications made it possible to determine whether on-chain information and sentiment analysis produced additional predictive value beyond conventional price history. Traditional machine learning models, including Random Forest, Extreme Gradient Boosting, and Support Vector Regression, were also employed as benchmark algorithms to evaluate whether the proposed deep learning framework achieved superior out-of-sample predictive performance.

Because the purpose of the study was not limited to prediction but also included explanation of the factors contributing to stablecoin price behavior, explainable artificial intelligence techniques were incorporated into the analytical procedure. SHapley Additive exPlanations were applied to the final predictive model to quantify the contribution of individual variables to each prediction and to rank the overall importance of on-chain, market, and sentiment indicators. SHAP summary and dependence analyses were used to investigate both the direction and magnitude of the effects of major predictors. Particular attention was given to whether exchange inflows and outflows, transaction volume, liquidity conditions, token issuance and redemption, large-holder activity, negative sentiment, sentiment volatility, and broader cryptocurrency market instability increased the predicted probability or magnitude of stablecoin price deviation. The attention weights generated by the deep learning architecture were additionally examined as complementary evidence regarding the temporal periods that contributed most strongly to predictions. This explainability framework enabled the proposed model to move beyond a black-box forecasting approach and provide interpretable evidence concerning the mechanisms associated with stablecoin price fluctuations.

Robustness analysis was performed by evaluating the model under normal market periods and periods characterized by relatively large deviations from the stablecoin peg. Predictive performance was also examined separately for the selected stablecoins to determine whether the contribution of on-chain and sentiment variables differed across stablecoin designs. The statistical association between participant-based sentiment measures and computational sentiment scores was assessed using Pearson correlation coefficients, and the consistency of the direction of sentiment between the two sources was examined as evidence of external validity. All statistical tests were interpreted at a significance level of .05. Data preprocessing, natural language processing, deep learning modeling, conventional machine learning estimation, and explainability analyses were conducted using Python and relevant scientific computing, machine learning, deep learning, and natural language processing libraries. The final analytical procedure was designed to maximize temporal validity, minimize information leakage, compare the

incremental predictive contribution of alternative information sources, and provide an interpretable account of the factors underlying stablecoin price stability and short-term de-pegging.

Findings and Results

The participant-based component of the study included 384 active cryptocurrency market participants residing in Tehran. Of the participants, 268 (69.79%) were men and 116 (30.21%) were women. The mean age of the participants was 32.74 years ($SD = 7.41$), with ages ranging from 19 to 56 years. Regarding educational attainment, 54 participants (14.06%) had a high school diploma or lower level of education, 68 (17.71%) held an associate degree, 162 (42.19%) held a bachelor's degree, 79 (20.57%) held a master's degree, and 21 (5.47%) held a doctoral degree. The average duration of cryptocurrency market experience was 3.86 years ($SD = 2.17$), while the median duration was approximately three years. A total of 103 participants (26.82%) reported between one and two years of cryptocurrency market experience, 148 (38.54%) reported between two and four years, 86 (22.40%) reported between four and six years, and 47 (12.24%) had more than six years of active market experience. Regarding trading frequency, 139 participants (36.20%) reported daily cryptocurrency market activity, 157 (40.89%) reported several transactions per week, and 88 (22.92%) reported less frequent but continued participation in cryptocurrency investment activities. USDT was the most frequently used stablecoin, reported by 354 participants (92.19%), followed by USDC among 217 participants (56.51%) and DAI among 108 participants (28.13%); because participants could use more than one stablecoin, these percentages were not mutually exclusive. Overall, the demographic and market-experience characteristics indicated that the sample consisted predominantly of relatively young and actively engaged cryptocurrency market participants with sufficient familiarity with stablecoin trading to provide a meaningful behavioral validation of the computational sentiment indicators.

The final time-series dataset consisted of 2,190 synchronized coin-day observations, including 730 daily observations for each of the three stablecoins, USDT, USDC, and DAI. After synchronization of market, on-chain, and sentiment variables, less than 1.8% of the initial observations contained isolated missing values, which were treated during preprocessing according to the procedures described in the Methods section. No observations associated with major market fluctuations were deleted solely on the basis of extreme values, because unusually large movements in transaction volume, exchange flows, sentiment, and price deviation represented substantively meaningful market events. The chronological allocation of observations resulted in 1,533 observations in the training set, 329 observations in the validation set, and 328 observations in the final test set. Table 1 presents the descriptive characteristics of the principal market, on-chain, and sentiment indicators included in the integrated predictive framework.

Table 1. Descriptive Statistics of the Principal Market, On-Chain, and Sentiment Variables

Variable	Mean	SD	Minimum	Maximum
Stablecoin closing price (USD)	1.0007	0.0094	0.9276	1.0583
Absolute deviation from USD 1.00	0.0049	0.0081	0.0000	0.0724
Daily return (%)	0.01	0.72	-6.48	5.91
Daily trading volume (USD billions)	18.64	15.27	0.41	87.93
Market capitalization (USD billions)	44.82	33.61	4.73	86.57
On-chain transaction volume (USD billions)	12.73	10.48	0.29	62.84
Daily transaction count (thousands)	617.42	418.76	73.19	2,486.73
Active addresses (thousands)	184.36	127.41	21.58	794.62
Exchange inflows (USD billions)	2.71	2.48	0.08	17.83
Exchange outflows (USD billions)	2.64	2.37	0.06	16.74

Net exchange flow (USD billions)	0.07	1.13	-5.82	6.47
Large-holder transaction volume (USD billions)	4.83	4.12	0.09	25.61
Daily supply change (%)	0.04	0.31	-2.16	3.74
Overall sentiment score	0.086	0.327	-0.842	0.791
Negative sentiment intensity	0.294	0.183	0.031	0.914
Positive sentiment intensity	0.371	0.194	0.042	0.903
Sentiment dispersion	0.218	0.107	0.019	0.681
Daily sentiment change	0.003	0.214	-0.781	0.834
Bitcoin daily return (%)	0.13	2.94	-12.47	14.26
Ethereum daily return (%)	0.16	3.57	-15.81	17.64
Crypto-market volatility index	0.034	0.021	0.008	0.139

As shown in Table 1, the overall mean stablecoin price remained close to the intended USD 1.00 peg, with a mean closing price of USD 1.0007 and a relatively small overall standard deviation of 0.0094. However, the minimum observed price of USD 0.9276 and the maximum price of USD 1.0583 indicated that economically meaningful deviations occurred within the observation period. The mean absolute deviation from the peg was USD 0.0049, although the maximum absolute deviation reached USD 0.0724, demonstrating that the dataset included episodes of substantial temporary instability rather than representing only near-perfectly pegged conditions. Considerable heterogeneity was also observed in market activity. Mean daily trading volume was USD 18.64 billion, with values ranging from USD 0.41 billion to USD 87.93 billion. Similar dispersion was observed in on-chain indicators, including transaction volume, transaction count, active addresses, and exchange-related flows. In particular, net exchange flow ranged from -USD 5.82 billion to USD 6.47 billion, suggesting periods of pronounced changes in the balance between stablecoin inflows to and outflows from exchanges. Market sentiment also showed substantial temporal variation. Although the average overall sentiment score was slightly positive at 0.086, the observed range extended from -0.842 to 0.791. Negative sentiment intensity reached a maximum of 0.914 during periods of heightened market concern, whereas positive sentiment intensity reached 0.903 during more favorable market conditions. The relatively large variability in sentiment measures, together with the observed variation in on-chain activity and stablecoin price deviations, provided sufficient heterogeneity for estimating nonlinear predictive relationships using deep learning techniques.

To evaluate the principal research objective, the predictive performance of the proposed integrated deep learning model was compared with several benchmark models. The comparison included Support Vector Regression, Random Forest, Extreme Gradient Boosting, a conventional Long Short-Term Memory model using price and market information, a Gated Recurrent Unit model, and the final multimodal attention-based LSTM model integrating historical market indicators, on-chain data, and deep learning-derived sentiment information. All performance indices reported in Table 2 were calculated using the chronologically separated test dataset, which was not used during model estimation or hyperparameter optimization.

Table 2. Out-of-Sample Predictive Performance of the Proposed Model and Benchmark Algorithms

Model	MAE	RMSE	MAPE (%)	R ²	Directional Accuracy (%)
Support Vector Regression	0.00482	0.00741	0.481	0.711	64.33
Random Forest	0.00421	0.00672	0.420	0.758	67.68
Extreme Gradient Boosting	0.00376	0.00605	0.375	0.801	70.43
Conventional LSTM	0.00329	0.00537	0.328	0.836	73.48
GRU	0.00311	0.00508	0.310	0.851	74.70
Attention-LSTM without sentiment	0.00274	0.00453	0.273	0.881	77.13
Proposed integrated Attention-LSTM	0.00196	0.00324	0.195	0.939	83.23

The results presented in Table 2 demonstrate that the proposed integrated Attention-LSTM model produced the highest predictive accuracy among all evaluated approaches. The final model achieved a Mean Absolute Error of 0.00196 and a Root Mean Squared Error of 0.00324, indicating that the average prediction error was extremely small relative to the approximately USD 1.00 reference value of stablecoins. The Mean Absolute Percentage Error was 0.195%, which was considerably lower than those of all benchmark algorithms. The model also achieved an R^2 value of 0.939, indicating that approximately 93.9% of the variation in test-period stablecoin prices was captured by the integrated predictive framework. Directional accuracy reached 83.23%, demonstrating that the model was also effective in identifying whether the subsequent movement represented an increase or decrease in price and peg deviation. Among the conventional machine learning models, Extreme Gradient Boosting performed better than Support Vector Regression and Random Forest, with an RMSE of 0.00605 and an R^2 of 0.801. Nevertheless, all deep learning models provided superior predictive performance compared with these conventional algorithms. The conventional LSTM produced an RMSE of 0.00537, while the GRU model reduced this error to 0.00508. The inclusion of an attention mechanism further improved performance, as the Attention-LSTM model without sentiment information achieved an RMSE of 0.00453 and an R^2 of 0.881. Most importantly, incorporating deep learning-derived sentiment information into the final multimodal architecture reduced the RMSE from 0.00453 to 0.00324, representing an approximately 28.48% reduction in prediction error relative to the otherwise comparable attention-based model without sentiment. These findings support the central premise of the study that stablecoin price behavior cannot be optimally predicted solely from historical price patterns and that the simultaneous incorporation of blockchain activity and market sentiment provides substantial incremental predictive information.

Ablation analysis was subsequently conducted to identify the incremental contribution of different categories of predictors. In this analysis, the architecture of the principal prediction model was maintained as consistently as possible while individual information domains were either included or excluded. The purpose of this procedure was to determine the extent to which price history, conventional market variables, on-chain indicators, and sentiment features contributed to prediction accuracy. Table 3 presents the performance of the alternative information configurations.

Table 3. Ablation Analysis of the Contribution of Market, On-Chain, and Sentiment Information

Predictor Configuration	MAE	RMSE	MAPE (%)	R^2	Directional Accuracy (%)
Historical price only	0.00451	0.00694	0.450	0.741	66.16
Price + conventional market variables	0.00368	0.00586	0.367	0.814	71.34
Price + on-chain variables	0.00313	0.00502	0.312	0.853	74.39
Price + sentiment variables	0.00339	0.00541	0.338	0.834	73.17
Price + market + on-chain variables	0.00274	0.00453	0.273	0.881	77.13
Price + market + sentiment variables	0.00258	0.00427	0.257	0.894	78.66
Price + on-chain + sentiment variables	0.00229	0.00378	0.228	0.916	80.79
Full integrated model	0.00196	0.00324	0.195	0.939	83.23

According to Table 3, each additional information domain contributed to an improvement in stablecoin price prediction, although the magnitude of the improvement varied across predictor categories. A model relying exclusively on historical prices achieved an RMSE of 0.00694 and an R^2 of 0.741, indicating that price history alone explained a substantial portion of subsequent behavior but left considerable prediction error unexplained. Adding conventional market variables reduced RMSE to 0.00586, whereas combining historical prices with on-chain indicators produced a larger improvement, reducing RMSE to 0.00502 and increasing R^2 to 0.853. Sentiment information also contributed predictive value when added to price history, resulting in an RMSE of 0.00541 and an

R^2 of 0.834. The stronger isolated contribution of on-chain variables suggests that blockchain-level activity contained particularly important information concerning stablecoin supply, liquidity movements, and transactional pressure. However, the most important result was the complementary effect observed when on-chain and sentiment indicators were combined. The configuration including price, on-chain, and sentiment variables achieved an RMSE of 0.00378 and an R^2 of 0.916, outperforming models that excluded either information domain. The complete integrated model, which simultaneously incorporated historical prices, conventional market characteristics, on-chain indicators, and market sentiment measures, produced the lowest prediction error and highest explanatory power, with an RMSE of 0.00324, MAPE of 0.195%, R^2 of 0.939, and directional accuracy of 83.23%. Relative to the historical-price-only specification, the full integrated model reduced RMSE by approximately 53.31%. These results indicate that the information contained in blockchain activity and market sentiment was neither redundant nor fully captured by historical price and trading variables. Instead, each data domain provided complementary information, supporting the use of a multimodal predictive architecture.

In addition to overall predictive accuracy, the model was evaluated specifically under periods of stable market conditions and episodes of relatively pronounced de-pegging. Observations were categorized as normal-peg observations when absolute deviation from USD 1.00 was less than 0.50%, moderate-deviation observations when the absolute deviation ranged from 0.50% to less than 1.50%, and high-deviation observations when absolute deviation was 1.50% or greater. Model performance was additionally estimated separately for USDT, USDC, and DAI to determine whether predictive accuracy generalized across alternative stablecoin structures. The results are shown in Table 4.

Table 4. Predictive Performance Across Stablecoins and Different Degrees of Peg Deviation

Evaluation Subsample	n	MAE	RMSE	R^2	Directional Accuracy (%)
USDT	110	0.00151	0.00246	0.947	85.45
USDC	109	0.00179	0.00291	0.942	84.40
DAI	109	0.00259	0.00417	0.921	79.82
Normal-peg observations	241	0.00131	0.00188	0.914	82.16
Moderate-deviation observations	56	0.00273	0.00406	0.927	83.93
High-deviation observations	31	0.00526	0.00791	0.903	90.32
Total test sample	328	0.00196	0.00324	0.939	83.23

As presented in Table 4, the proposed integrated model maintained strong predictive performance across all three stablecoins, although some differences were observed. Prediction accuracy was highest for USDT, for which the model achieved an MAE of 0.00151, an RMSE of 0.00246, an R^2 of 0.947, and directional accuracy of 85.45%. USDC showed similarly strong performance, with an RMSE of 0.00291 and an R^2 of 0.942. Prediction errors were somewhat higher for DAI, with an MAE of 0.00259 and an RMSE of 0.00417, while its R^2 remained high at 0.921. This result suggests that the model generalized effectively across stablecoins while encountering greater prediction difficulty in a decentralized and collateral-sensitive stablecoin such as DAI, whose price dynamics may reflect more complex interactions among collateral conditions, decentralized finance activity, and market-level expectations. The model also remained informative across different levels of peg instability. As expected, absolute prediction errors were lowest during normal-peg observations, for which RMSE was only 0.00188. Prediction error increased to 0.00406 during moderate deviations and 0.00791 during high-deviation periods because these observations represented more volatile and difficult market conditions. However, directional accuracy increased to 90.32% during high-deviation periods. Thus, although exact point prediction became more difficult during severe market stress,

the model was particularly effective in identifying the direction of abnormal stablecoin movements. The R^2 value of 0.903 within the high-deviation subsample further indicates that the model retained substantial explanatory capacity even during the most economically important periods of peg instability.

Explainability analysis was performed using SHapley Additive exPlanations to determine which variables contributed most strongly to the predictions generated by the integrated deep learning model. Mean absolute SHAP values were calculated to provide an overall ranking of feature importance. The directional interpretation was derived from the relationship between feature values and their SHAP contributions to predicted absolute peg deviation. Positive directional effects indicate that increases in the variable were generally associated with larger predicted deviations from the USD 1.00 peg, whereas negative effects indicate a stabilizing relationship. The principal explainability results are presented in Table 5.

Table 5. SHAP-Based Importance and Direction of the Principal Predictors of Stablecoin Price Deviation

Rank	Predictor	Mean Absolute SHAP Value	Relative Importance (%)	Predominant Relationship With Predicted Peg Deviation
1	Negative sentiment intensity	0.00241	13.74	Positive
2	Net exchange inflow	0.00216	12.31	Positive
3	Previous-day peg deviation	0.00194	11.06	Positive
4	On-chain transaction volume	0.00171	9.75	Nonlinear
5	Sentiment dispersion	0.00152	8.67	Positive
6	Large-holder transaction volume	0.00139	7.93	Positive
7	Crypto-market volatility	0.00127	7.24	Positive
8	Token supply change	0.00111	6.33	Nonlinear
9	Exchange outflow	0.00096	5.47	Predominantly negative
10	Bitcoin return	0.00084	4.79	Asymmetric
11	Active-address change	0.00076	4.33	Nonlinear
12	Trading-volume change	0.00068	3.88	Nonlinear
13	Positive sentiment intensity	0.00062	3.54	Predominantly negative
14	Ethereum return	0.00055	3.14	Asymmetric
15	Market-capitalization change	0.00032	1.82	Predominantly negative

The SHAP analysis presented in Table 5 showed that negative market sentiment was the most influential variable in the final prediction framework, accounting for 13.74% of the importance among the 15 principal predictors. Higher negative sentiment intensity was generally associated with larger predicted deviations from the one-dollar peg. Net exchange inflow was the second most important factor, with a relative contribution of 12.31%, and increased net inflows were predominantly associated with greater predicted peg deviation. The previous day's deviation from the peg ranked third, confirming a degree of short-term persistence in stablecoin instability. On-chain transaction volume showed a substantial but nonlinear relationship, indicating that both unusually low and unusually high blockchain activity could contain relevant information depending on concurrent market conditions. Sentiment dispersion ranked fifth, suggesting that disagreement or heterogeneity in market expectations was independently informative beyond the average polarity of sentiment. Large-holder transaction activity and overall cryptocurrency-market volatility were also important predictors, contributing 7.93% and 7.24%, respectively. The direction of large-holder activity was predominantly positive, indicating that unusually large transfers tended to increase predicted

instability, particularly when they coincided with negative sentiment or elevated exchange inflows. Conversely, exchange outflows and positive sentiment intensity were predominantly associated with lower predicted peg deviation. The asymmetric relationships observed for Bitcoin and Ethereum returns indicated that large negative movements in major cryptocurrencies had a stronger destabilizing contribution than equivalently sized positive movements had a stabilizing contribution. Taken together, the explainability findings demonstrated that stablecoin price instability was not determined by a single class of variables but reflected an interaction between behavioral sentiment, blockchain flows, prior price instability, and broader cryptocurrency-market conditions.

The external validation of computational sentiment measures using responses from the 384 Tehran-based cryptocurrency participants further supported the validity of the sentiment-analysis component. The deep learning-derived overall sentiment score was positively correlated with participants' reported short-term market confidence, $r = .63$, $p < .001$, and expected stablecoin price stability, $r = .57$, $p < .001$. Computational negative sentiment intensity was positively related to participant-reported de-pegging concern, $r = .66$, $p < .001$, perceived stablecoin risk, $r = .61$, $p < .001$, and liquidity concern, $r = .49$, $p < .001$. Conversely, positive sentiment intensity was negatively associated with perceived de-pegging risk, $r = -.52$, $p < .001$. These moderate-to-strong relationships indicated convergence between behavioral perceptions reported by active market participants and the sentiment signals automatically extracted from large-scale textual sources. The results therefore provided additional evidence that the computational sentiment measures represented meaningful market perceptions rather than merely reflecting linguistic variation unrelated to investor expectations.




Figure 1. Observed and Predicted Stablecoin Prices Across the Chronological Test Period Using the Integrated Attention-LSTM Model

Figure 1 illustrates the correspondence between actual and model-predicted stablecoin prices across the chronological test period. The predicted series closely followed the observed market series throughout most normal trading periods, with particularly small discrepancies when stablecoin prices remained close to the USD 1.00 reference value. More importantly, the predicted series responded to major upward and downward departures from the peg rather than converging mechanically toward the long-run mean. The correspondence between observed and predicted values remained particularly strong at the beginning and intermediate portions of de-pegging episodes, indicating that the model was able to incorporate emerging information from changes in exchange flows, blockchain activity, market volatility, and sentiment before or during the development of abnormal price movements. Prediction errors became larger around the most extreme turning points, which is consistent with the higher RMSE observed for the high-deviation subsample; however, the model generally identified the direction of such movements correctly. After large deviations, the predicted values also tracked the subsequent return toward the peg, suggesting that the architecture was sensitive not only to destabilizing signals but also to indicators associated with market normalization. The visual pattern therefore supported the quantitative evaluation metrics and showed that the proposed model generated predictions that were temporally aligned with actual stablecoin price behavior rather than achieving high aggregate accuracy by simply predicting a constant value close to USD 1.00.

Overall, the findings provided consistent evidence for the effectiveness of combining on-chain information, conventional market indicators, and deep learning-based sentiment measures for stablecoin price prediction. The superiority of the full integrated model was evident across all major evaluation criteria, including MAE, RMSE, MAPE, R^2 , and directional accuracy. The ablation analysis further demonstrated that eliminating either on-chain or sentiment information produced a measurable deterioration in out-of-sample performance. The model retained high explanatory power across USDT, USDC, and DAI and remained informative during moderate and severe de-pegging episodes. Finally, SHAP-based explainability showed that negative sentiment, exchange flows, previous peg deviations, transaction activity, sentiment dispersion, large-holder activity, and overall market volatility were among the principal drivers of model predictions. These findings indicate that stablecoin pricing is characterized by the simultaneous influence of technical blockchain conditions and behavioral market expectations and that integrating these dimensions within an explainable deep learning framework can substantially improve both prediction accuracy and interpretability.

Discussion and Conclusion

The present study aimed to design and evaluate an explainable deep learning model for predicting stablecoin prices and peg deviations through the integration of conventional market indicators, on-chain variables, and deep learning-based market sentiment information. The findings showed that the proposed integrated Attention-LSTM model achieved the highest predictive performance among all evaluated approaches, with an MAE of 0.00196, an RMSE of 0.00324, an MAPE of 0.195%, an R^2 of 0.939, and a directional accuracy of 83.23%. These results indicate that the model was able to capture a very large proportion of the variation in stablecoin prices and accurately identify the direction of short-term changes. Importantly, the proposed framework outperformed Support Vector Regression, Random Forest, Extreme Gradient Boosting, conventional LSTM, GRU, and an Attention-LSTM model without sentiment variables. This pattern confirms the value of combining sequential deep learning with heterogeneous

information sources in cryptocurrency forecasting. Previous studies have similarly emphasized the capacity of deep learning models to identify nonlinear and temporally dependent structures in cryptocurrency prices that cannot be adequately captured by conventional statistical or shallow machine learning methods (2, 3, 5). The superior performance of the proposed model is also consistent with comparative evidence showing that deep learning and ensemble-based approaches can provide substantial forecasting gains when applied to complex and volatile cryptocurrency markets (4, 8).

The strong performance of the integrated model can be explained by the temporal and nonlinear characteristics of stablecoin price formation. Although stablecoins are designed to remain close to a fixed reference value, their prices do not remain perfectly constant. Instead, short-term deviations emerge from changes in liquidity, redemption pressure, exchange activity, investor confidence, network behavior, and broader crypto-market conditions. The Attention-LSTM architecture was particularly appropriate for this environment because LSTM networks can retain information from prior observations while attention mechanisms can assign greater weight to specific time points that contain unusually informative signals. The results therefore support previous research demonstrating that deep recurrent networks are well suited to cryptocurrency price prediction because they can capture time dependencies, nonlinear dynamics, and interactions among multiple variables (1, 2). The finding is also in line with the work of Akouaouch and Bouayad, who demonstrated that deep learning approaches can be effective for high-frequency and short-term cryptocurrency price prediction, especially when rapid changes in market conditions must be captured (9). Similarly, the use of deep machine learning for forecasting virtual cryptocurrency prices has been supported by prior work showing that such methods can identify complex patterns in digital-asset data that are not easily detected through simpler forecasting techniques (10).

A major finding of the present study was that adding sentiment information to the Attention-LSTM architecture substantially improved prediction accuracy. The Attention-LSTM model without sentiment achieved an RMSE of 0.00453 and an R^2 of 0.881, whereas the full integrated model reduced the RMSE to 0.00324 and increased R^2 to 0.939. This represents a meaningful improvement and indicates that market sentiment contains predictive information that is not fully reflected in historical price, volume, or on-chain variables. This finding is consistent with earlier evidence that textual sentiment can improve cryptocurrency forecasting because digital-asset markets are highly responsive to expectations, online narratives, and rapidly disseminated information (11). Rusydiana also showed that cryptocurrency-related Twitter data contain meaningful positive, negative, and neutral sentiment patterns that can be used to interpret market conditions (12). The present results extend these findings specifically to stablecoins and suggest that sentiment is important even when the underlying asset is designed to maintain a fixed price. This may occur because the stability of a stablecoin depends not only on mechanical stabilization arrangements but also on confidence in reserves, collateral, redemption processes, liquidity, and issuer or protocol reliability.

The ablation analysis provided further evidence that each information domain made an independent contribution to forecasting performance. A model based solely on historical prices achieved an RMSE of 0.00694 and an R^2 of 0.741. Adding conventional market variables improved performance, while adding on-chain variables resulted in a larger improvement. The price-plus-on-chain configuration achieved an RMSE of 0.00502 and an R^2 of 0.853, whereas the price-plus-sentiment model achieved an RMSE of 0.00541 and an R^2 of 0.834. When on-chain and sentiment information were combined, performance improved considerably, and the full integrated model produced the best results. These findings indicate that historical prices, blockchain activity, and sentiment are complementary

rather than redundant information sources. Such a result is theoretically plausible because market prices summarize previous market outcomes, on-chain indicators reflect transactional behavior and liquidity movements at the blockchain level, and sentiment reflects expectations and perceptions before they are fully translated into transactions. The finding is consistent with broader cryptocurrency forecasting research suggesting that the quality and diversity of the feature set can be as important as the forecasting algorithm itself (4, 8). It also supports the argument that cryptocurrency markets require multidimensional modeling because their price dynamics are simultaneously behavioral, technological, and financial.

The importance of on-chain variables was particularly notable. The ablation results showed that on-chain indicators provided greater isolated predictive value than sentiment when each was added separately to historical prices. This suggests that blockchain-native data capture fundamental changes in stablecoin activity that may precede or accompany price instability. Stablecoin systems are especially sensitive to exchange inflows and outflows, token issuance and redemption, large transfers, and sudden changes in transaction volume. These variables can reflect liquidity pressure and changes in market positioning before they become visible in price series. The relevance of such signals is compatible with prior analyses emphasizing that stablecoins have distinct high-frequency behavior and that their stability should not be assumed merely because they target a fixed value (7). Bullmann, Klemm, and Pinna similarly questioned whether stablecoins can consistently provide the stability they are intended to deliver and emphasized the structural mechanisms and risks underlying their operation (6). The current findings reinforce this perspective by showing empirically that blockchain-level variables contribute substantially to forecasting temporary deviations from the peg.

The explainability analysis offered additional insight into the specific factors driving the predictions. Negative sentiment intensity emerged as the most important predictor, followed by net exchange inflow, previous-day peg deviation, on-chain transaction volume, sentiment dispersion, large-holder transaction volume, and broader cryptocurrency-market volatility. The dominant role of negative sentiment suggests that perceived risk and deteriorating confidence may be particularly important during periods of stablecoin stress. This is consistent with evidence that sentiment can become especially influential during extreme market regimes and periods characterized by adverse selection or elevated uncertainty (13). Sakariyahu and colleagues likewise showed that global uncertainty and sentiment factors are significantly related to cryptocurrency market crashes, indicating that behavioral conditions become especially important when markets shift away from normal regimes (14). The present results suggest that similar processes operate in stablecoin markets, where rapidly increasing negative sentiment may signal concerns about liquidity, backing, exchange conditions, or redemption capacity and therefore increase the probability of de-pegging.

The importance of sentiment dispersion is also noteworthy. The results indicated that disagreement or heterogeneity in market perceptions had explanatory value beyond the average level of positive or negative sentiment. This suggests that unstable expectations themselves may represent a risk signal. When investors hold strongly divergent views about the stability of a coin, trading and redemption behavior may become more erratic, contributing to temporary pricing dislocations. Such a mechanism is compatible with the argument that sentiment regimes can affect adverse selection and market quality in cryptocurrency markets (13). It is also consistent with findings showing that investor sentiment influences decision-making and portfolio allocation across both conventional and cryptocurrency markets (15). In the context of stablecoins, heterogeneous expectations may

therefore translate into simultaneous buying, selling, redemption, and exchange movement, producing instability that cannot be explained by average sentiment alone.

The positive contribution of net exchange inflow and large-holder transaction volume further supports the interpretation that liquidity movements and concentrated trading activity are key determinants of stablecoin price instability. Large inflows to exchanges may signal preparation for selling, conversion, arbitrage, or redemption-related activity. Similarly, large-holder transfers can have disproportionate effects on short-term liquidity because stablecoin markets may react strongly when significant quantities are moved by concentrated participants. These findings help explain why the integrated model outperformed price-only models: the model was able to identify changes in underlying market behavior before these changes were fully incorporated into price. This result also aligns with high-frequency stablecoin research showing that short-term pricing behavior is influenced by market microstructure and trading conditions rather than solely by the intended peg (7). From this perspective, peg stability should be understood as an equilibrium outcome maintained by continuous market adjustment rather than as an automatic property of the token.

Another important result was the persistence of prior peg deviation as a major predictor. The previous day's deviation from the peg ranked among the most important explanatory variables, indicating that stablecoin instability has a short-term autoregressive component. Once a coin moves away from its target value, the deviation may persist temporarily because arbitrage, liquidity provision, redemption, and confidence-restoring mechanisms require time to operate. This is consistent with the broader time-series literature on cryptocurrencies, where previous observations often contain substantial information about future movement (2). However, the present study also showed that historical information alone was insufficient, because the price-only model performed substantially worse than models that incorporated on-chain and sentiment variables. Thus, persistence is important but should be interpreted jointly with contemporaneous behavioral and blockchain signals.

The separate analysis across USDT, USDC, and DAI indicated that the model generalized effectively across different stablecoin structures, although performance was somewhat stronger for USDT and USDC than for DAI. This may reflect differences in liquidity, market depth, collateralization design, and stabilization mechanisms. Centralized reserve-backed stablecoins may respond more directly to exchange flows, supply changes, and redemption expectations, whereas decentralized stablecoins can be influenced by more complex interactions among collateral values, decentralized finance activity, smart-contract mechanisms, and protocol-level incentives. The fact that the model still achieved an R^2 of 0.921 for DAI suggests that the integrated framework remained robust even under a more complex stabilization architecture. This finding is compatible with the broader argument that stablecoins should not be treated as a homogeneous asset class because their stability mechanisms differ substantially (6). It also reinforces the importance of incorporating multiple information sources when constructing generalizable forecasting models.

The model's performance during moderate and severe de-pegging episodes is particularly important from a practical perspective. Although absolute prediction errors increased during high-deviation periods, directional accuracy rose to 90.32%. This means that the model became less precise in estimating the exact magnitude of extreme movements but was highly effective in identifying their direction. Such a result is expected because crisis periods involve abrupt nonlinear adjustments that are difficult to predict precisely. However, identifying the direction of a developing de-pegging episode may be more important than minimizing small point-estimation errors in normal periods. The relevance of this finding is supported by research demonstrating that uncertainty, sentiment, and cross-

market spillovers intensify during crisis conditions (14, 16). Soltani and colleagues showed that investor sentiment and cryptocurrency markets can exhibit substantial spillover effects during periods such as the COVID-19 crisis (16). The present findings similarly suggest that stablecoin forecasting models must be evaluated specifically under stressed conditions rather than only on average performance.

The participant-based validation of computational sentiment also strengthened the credibility of the sentiment component. Deep learning-derived sentiment scores were positively associated with reported market confidence and expected stablecoin stability, while negative sentiment intensity was positively related to perceived de-pegging risk, stablecoin risk, and liquidity concern. This convergence indicates that the textual sentiment model captured meaningful behavioral expectations rather than merely linguistic noise. The finding is compatible with previous research showing that investor sentiment has measurable implications for portfolio decisions and cryptocurrency behavior (15). It also supports the usefulness of social-media and text-based sentiment analysis documented in cryptocurrency research (11, 12). Combining computational sentiment with participant-based validation therefore provided a stronger behavioral foundation for the model than relying on either source independently.

Overall, the results demonstrate that stablecoin price prediction benefits substantially from integrating market, blockchain, and behavioral data within an explainable deep learning architecture. The proposed model achieved superior predictive accuracy, performed robustly across multiple stablecoins, remained informative during de-pegging episodes, and identified economically interpretable predictors of instability. These findings extend previous cryptocurrency forecasting research by showing that high-performing deep learning models can be made more informative through the integration of on-chain and sentiment signals and through the application of explainability techniques. They also support the broader literature indicating that cryptocurrency prices are influenced by nonlinear interactions among market structure, investor behavior, uncertainty, and temporal dynamics (1, 3, 8). The present findings therefore suggest that an explainable multimodal forecasting framework offers a more complete representation of stablecoin behavior than conventional price-based prediction.

Several limitations should be considered when interpreting the findings. First, the predictive model was developed using three major stablecoins and therefore may not fully represent smaller, algorithmic, commodity-backed, or newly issued stablecoins whose liquidity and stabilization mechanisms differ substantially. Second, although the dataset incorporated multiple on-chain, market, and sentiment variables, some potentially relevant information, such as detailed reserve disclosures, exchange-specific order-book depth, regulatory announcements, protocol governance decisions, and intraday collateral liquidation data, was not available consistently across the entire observation period. Third, the use of daily synchronized observations may have obscured very short-lived intraday de-pegging events that emerge and resolve within hours or minutes. Fourth, sentiment data collected from public online sources can contain noise generated by automated accounts, promotional messages, duplicated content, sarcasm, and coordinated communication. Although preprocessing and participant-based validation reduced this concern, such sources cannot be fully eliminated. Finally, the behavioral validation sample was limited to active cryptocurrency participants residing in Tehran, and the resulting perception patterns may not be fully generalizable to investors operating in other countries, regulatory environments, or market structures.

Future studies should extend the proposed framework in several directions. Research could examine a broader range of stablecoins and compare reserve-backed, overcollateralized, algorithmic, commodity-backed, and hybrid stabilization mechanisms using separate model architectures. Higher-frequency datasets based on hourly, fifteen-minute, or transaction-level observations could be used to determine whether on-chain and sentiment signals

provide earlier warnings of de-pegging than daily models. Future models could also incorporate order-book variables, reserve attestations, collateral ratios, decentralized finance liquidation events, cross-chain bridge activity, regulatory news, and macroeconomic indicators. Transformer-based time-series models and multimodal architectures could be compared with Attention-LSTM networks to identify whether newer sequential learning methods improve performance during crisis regimes. Cross-country behavioral validation would also be valuable for examining whether sentiment signals differ according to regulatory environments, financial conditions, or investor sophistication. Finally, future work should place greater emphasis on real-time early-warning systems and probabilistic forecasting so that models provide not only point estimates but also confidence intervals and explicit probabilities of mild, moderate, or severe de-pegging.

The findings have several practical implications for cryptocurrency exchanges, portfolio managers, stablecoin users, decentralized finance platforms, risk analysts, and financial technology firms. Market-monitoring systems should avoid relying exclusively on historical stablecoin prices because price-based indicators may react only after instability has already begun. Instead, real-time surveillance frameworks should integrate exchange flows, on-chain transaction activity, large-holder transfers, supply changes, market volatility, and sentiment indicators. Negative sentiment intensity and rising net exchange inflows should be treated as potential early-warning signals, particularly when they occur simultaneously with prior peg deviations or increasing market volatility. Exchanges and trading platforms could embed explainable forecasting systems into risk dashboards so that users can see not only predicted deviations but also the principal factors contributing to them. Portfolio managers could incorporate model outputs into liquidity management, hedging, and exposure limits, while stablecoin issuers and protocol managers could use such systems to detect emerging confidence or redemption pressures. Practical adoption should nevertheless emphasize transparent model monitoring, frequent retraining, stress testing, and human oversight, especially because stablecoin markets can change rapidly in response to regulatory, technological, and systemic events.

Acknowledgments

We would like to express our appreciation and gratitude to all those who helped us carrying out this study.

Authors' Contributions

All authors equally contributed to this study.

Declaration of Interest

The authors of this article declared no conflict of interest.

Ethical Considerations

All ethical principles were adhered in conducting and writing this article.

Transparency of Data

In accordance with the principles of transparency and open research, we declare that all data and materials used in this study are available upon request.

Funding

This research was carried out independently with personal funding and without the financial support of any governmental or private institution or organization.

References

1. Zhang Y, Lee J. Deep Learning for Cryptocurrency Price Prediction: A Survey. *Journal of Financial Data Science*. 2020;2(3):15-35.
2. Zhou J, Li H. Time Series Analysis of Cryptocurrency Prices Using Deep Learning. *Journal of Computational Finance*. 2021;25(1):56-72.
3. Huang Y, Zhao M. Cryptocurrency Price Prediction Using Hybrid Deep Learning Techniques. *Expert Systems with Applications*. 2019;123:195-210.
4. Patel R, Verma A. Cryptocurrency Price Prediction Using Ensemble Deep Learning Models. *Expert Systems with Applications*. 2021;183:115392.
5. Pintelas E, Livieris IE, Stavroyiannis S, Kotsilieris T, Pintelas P. Investigating the Problem of Cryptocurrency Price Prediction: A Deep Learning Approach. *AIAI 2020: Artificial Intelligence Applications and Innovations 2020*. p. 99-110.
6. Bullmann D, Klemm J, Pinna A. In search for stability in crypto-assets: Are stablecoins the solution? *European Central Bank Occasional Paper Series*, 2019.
7. Barucci E, Moncayo GG, Marazzina D. Cryptocurrencies and Stablecoins: A High-Frequency Analysis. *Digital Finance*. 2022;4(2-3):217-39. doi: 10.1007/s42521-022-00055-9.
8. Bouteska A, Abedin MZ, Hajek P, Yuan K. Cryptocurrency price forecasting - A comparative analysis of ensemble learning and deep learning methods. *International Review of Financial Analysis*. 2024;92:103055. doi: 10.1016/j.irfa.2023.103055.
9. Akouaouch I, Bouayad A. A new deep learning approach for predicting high -frequency short -term cryptocurrency price. *Bulletin of Electrical Engineering and Informatics*. 2025;14(1):513-23. doi: <https://doi.org/10.11591/eei.v14i1.7377>.
10. Bakhit S, Rabiei M. Forecasting the Price of the First Virtual Cryptocurrency Based on Deep Machine Learning Techniques: Ghiyasuddin Jamshid Kashani University; 2020.
11. Roberts J. Sentiment Analysis and Cryptocurrency Price Prediction: A Deep Learning Approach. *Journal of Financial Markets*. 2018;37:85-101.
12. Rusydiana AS. Cryptocurrency: A Sentiment Analysis From Twitter Data. *Ift*. 2024;1(1). doi: 10.58968/ift.v1i1.359.
13. Farzulla M. The Extremity Premium: Sentiment Regimes and Adverse Selection in Cryptocurrency Markets. 2026. doi: 10.48550/arxiv.2602.07018.
14. Sakariyahu R, Lawal R, Adigun R, Paterson A, Johan S. One Crash, Too Many: Global Uncertainty, Sentiment Factors and Cryptocurrency Market. *Journal of International Financial Markets, Institutions and Money*. 2024;94. doi: 10.1016/j.intfin.2024.102028.
15. Nourollahi S, Jafari A, Pakmaram A. The Impact of Investor Sentiment on Portfolio Optimization in the Tehran Stock Exchange and Cryptocurrency Markets. *Business, Marketing, and Finance Open*. 2025:1-21.
16. Soltani H, Taleb J, Boujelbene Abbas M. Spillover effects between stock markets, cryptocurrency, and investor sentiment during COVID-19. *European Journal of Management and Business Economics*. 2025;34(1):23-46.