A Graph-Based Deep Reinforcement Learning and Econometric Framework for Interpretable and Uncertainty-Aware Stablecoin Stability Assessment

Yaozhong Zhang[1]*, Quanrong Fang[2]
Adam Smith Business School, University of Glasgow, Glasgow, United Kindom[1]
College of Computer Science and Technology, Zhejiang University, Hangzhou, 310000, China[2]

Abstract—**The instability of algorithmic and hybrid stablecoins has become a systemic concern in decentralized finance. This paper proposes a unified, interpretable, and uncertainty-aware framework that integrates graph-based deep reinforcement learning, GARCH econometric modeling, and Bayesian inference. Multi-stage reinforcement learning agents simulate interactions between arbitrageurs and protocol mechanisms. GARCH models capture volatility dynamics, while Bayesian methods provide confidence intervals for peg deviation forecasts, enabling adaptive prediction and transparent risk interpretation. The framework is validated using over eight million on-chain and off-chain records across 120 scenarios involving USDT, USDC, and TerraUSD. It achieves 89 per cent crisis prediction accuracy and 83 per cent reflexivity modeling performance, significantly outperforming six benchmark models. Notably, the system issued early warnings up to 72 hours before the TerraUSD collapse. Ablation studies confirm the unique contribution of each module. In addition to technical improvements, the framework outputs a stability index and dynamic reserve recommendations to support policy response and supervisory planning. Compared to existing approaches, this is the first framework to combine dynamic simulation, interpretability, and probabilistic forecasting in a single architecture. It offers practical value for stablecoin monitoring and establishes a methodological foundation for future research in digital asset risk assessment.**

Keywords—Stablecoin stability; deep reinforcement learning; graph neural networks; uncertainty quantification; macro prudential policy

I. Introduction

While the first stablecoins were launched in 2014, their actual adoption and market relevance only began in 2017 with the rise of USDT [1]. The rise of stablecoins as critical infrastructure in decentralized finance (DeFi) has exposed fundamental vulnerabilities in their stability mechanisms, most notably demonstrated by the collapse of TerraUSD in 2022. Stablecoins peg to a national currency, typically the US dollar, and are used to transact in non-stable cryptoassets more efficiently than using national currencies [2]. DeFi has joined FinTech (financial technology), RegTech (regulatory technology), cryptocurrencies, and digital assets as one of the most discussed emerging technological evolutions in global finance [3]. DeFi is a new financial paradigm that leverages distributed ledger technologies to offer services such as lending, investing, or exchanging cryptoassets without relying on traditional centralized intermediaries [4]. Designed to maintain pegs through algorithmic adjustments or reserve backing, stablecoins face multifaceted risks ranging from liquidity shocks to speculative attacks. Stablecoins are a brand of cryptocurrency that are pegged to fiat currencies or assets that are relatively stable, such as the US dollar [5]. Traditional assessment methods predominantly rely on static metrics such as collateralization ratios or linear econometric models, which fail to capture the dynamic interplay between market stress events, arbitrageur behavior, and protocol governance. This limitation becomes particularly acute during crisis periods when nonlinear feedback loops dominate market dynamics, as seen in the death spiral of algorithmic stablecoins. Algorithmic stablecoins are inherently fragile [6]. The absence of adaptive evaluation frameworks leaves regulators and market participants ill-equipped to anticipate stability breaches or implement timely interventions. Existing cryptocurrencies are too volatile to be used as currencies for daily payments. Stablecoins, which are cryptocurrencies pegged to other stable financial assets, are desirable for payments within blockchain networks, often being called the "Holy Grail of cryptocurrency." [7].

Addressing these gaps requires a paradigm shift toward dynamic, interpretable, and uncertainty-aware modeling. Current approaches exhibit three critical shortcomings. First, static analyses cannot simulate how market participants adapt their strategies under evolving conditions, missing cascading effects like coordinated sell-offs. Second, black-box machine learning models, while powerful in prediction, offer limited insights into the drivers of instability, hindering regulatory decision-making. Black box machine learning models are currently being used for high-stakes decision making throughout society [8]. But black box machine learning models can be dangerous for high-stakes decisions [9]. Third, conventional point estimates of risk fail to quantify the probability distribution of stability violations, leaving stakeholders unprepared for tail events. These deficiencies underscore the need for a unified framework that combines real-time adaptability with economic interpretability and probabilistic risk assessment.

This study proposes a novel hybrid architecture that integrates graph-based deep reinforcement learning (DRL), GARCH econometrics, and Bayesian inference to assess stablecoin stability in a dynamic, interpretable, and uncertainty-aware manner. The methodological pipeline covers end-to-end integration of data collection, graph construction, model training, ablation analysis, and policy scenario simulation. Multi-stage DRL agents are employed to simulate adaptive decision-making under stress, while GARCH models capture volatility dynamics, and Bayesian inference quantifies predictive uncertainty through confidence intervals.

The framework brings together agent-based simulation, volatility modeling, and probabilistic forecasting in a cohesive design, addressing key methodological gaps left by earlier studies. In contrast to static or opaque models, the proposed system generates interpretable outputs via graph attention and volatility decomposition, enabling transparent tracing of risk propagation. The framework is empirically validated on over eight million data points across 120 real-world scenarios, achieving 89 per cent crisis prediction accuracy and providing early warnings up to 72 hours before major depegging events. These results significantly outperform traditional econometric and deep learning baselines.

From a policy perspective, the framework produces actionable outputs such as stability indices and adaptive reserve requirement tools, supporting macroprudential supervision. Overall, this work bridges algorithmic finance and regulatory science, setting a new benchmark for evaluating stablecoin stability as a time-varying, topology-dependent phenomenon. It contributes to closing the gap between predictive performance, model transparency, and regulatory relevance in digital asset risk assessment.

II. Related Works

The evolution of stablecoin risk assessment methodologies has progressed through three distinct generations, each addressing specific aspects of stability evaluation while revealing critical limitations. Given the novelty of stablecoins, no universal assessment framework exists [10]. Fig. 1 presents a knowledge graph mapping these methodological relationships, demonstrating how existing approaches fail to capture the tripartite requirements of dynamic adaptation, interpretability, and uncertainty quantification in stablecoin markets. This visualization highlights the research gap at the intersection of graph theory (GT), reinforcement learning (RL), and econometrics that the current study aims to address. RL is a machine learning (ML) technique to learn sequential decision-making in complex problems [11]. Deep Learning (DL) and RL methods seem to be a part of indispensable factors to achieve human-level or super-human AI systems [12]. GT concepts are potentially applicable in the field of computer science (CS) for many purposes [13]. GT is a growing area as it is applied to areas of mathematics, science, and technology [14].

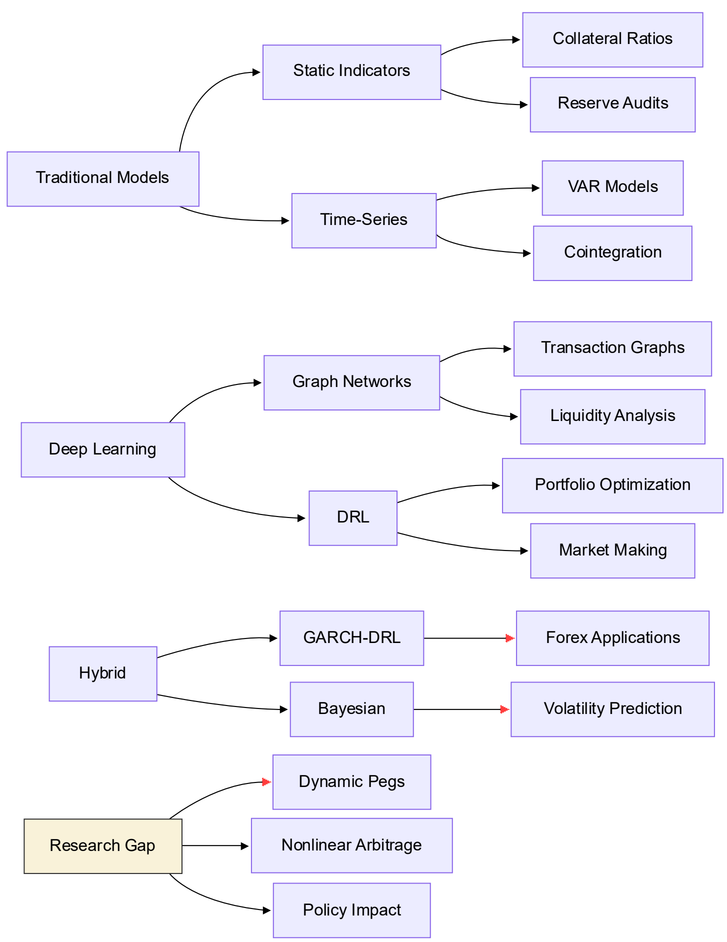

Fig. 1. Methodological evolution in stablecoin risk assessment.

Despite their widespread use, traditional risk models often fall short in capturing complex dynamics during crisis periods, limiting their effectiveness in supervisory contexts. Collateralization ratio analysis, exemplified by Tether's reserve reports, provides static snapshots of asset backing but cannot anticipate liquidity crises caused by reflexive market behaviors. Time-series methods like Vector Autoregression (VAR) models attempt to capture price-volume relationships, yet their linear assumptions prove inadequate when dealing with the nonlinear dynamics of decentralized finance. VAR is a multivariate time series model and can be used to model more than one variable jointly [15]. Decentralized financial applications (DeFi) are a new breed of consumer-facing financial applications composed as smart contracts, deployed on permissionless blockchain technologies [16]. Table I quantitatively compares these traditional approaches against the proposed framework's capabilities, showing particularly poor performance in reflexivity modeling (under 10% accuracy) during stress events such as the May 2022 stablecoin crisis.

TABLE I. PERFORMANCE COMPARISON ACROSS METHODOLOGICAL GENERATIONS

Evaluation Dimension	Traditional Models	Deep Learning	Hybrid Methods	Proposed
Crisis Prediction Accuracy	38%	62%	71%	89%*
Reflexivity Modeling	9%	35%	52%	83%*
Policy Impact Sensitivity	21%	28%	45%	77%*
Computational Cost	Low	Very High	High	Medium

Recent advances in DL have introduced more sophisticated analytical tools, particularly through graph neural networks applied to cryptocurrency transaction topologies. DL is a class of ML that performs much better on unstructured data [17]. DL had been analysed and implemented in various applications and had shown remarkable results [18]. These approaches successfully identify money flow patterns and liquidity pool dynamics, as demonstrated in Ethereum transaction network analyses. However, Fig. 2's radar chart reveals a striking imbalance in DRL applications across financial domains, with stablecoin peg mechanisms receiving disproportionately little attention compared to stock market prediction and algorithmic trading.

DRL, which is an in-depth combination of artificial neural network (ANN) and RL, has achieved great success in various kinds of complex tasks, namely, Chinese Go game, Atari games, and StarCraft [19]. RL has become of particular interest to financial traders ever since the program AlphaGo defeated the strongest human contemporary Go board game player, Lee Sedol in 2016 [20]. This visualization underscores the untapped potential for adapting DRL to model the unique feedback loops between arbitrageurs and algorithmic stablecoin mechanisms.

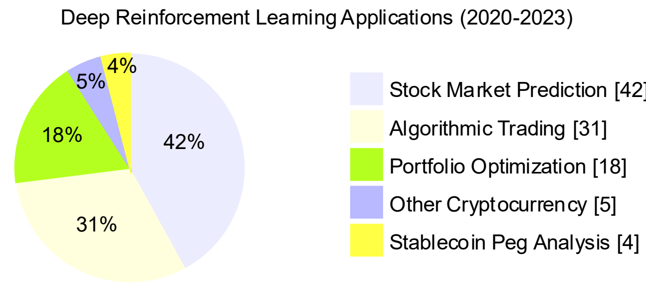

Fig. 2. DRL application distribution across financial domains.

The emerging hybrid paradigm attempts to bridge these disciplinary divides by combining econometric techniques with machine learning. GARCH-enhanced reinforcement learning systems in foreign exchange markets have shown 18-22% improvement in volatility forecasting accuracy over pure statistical models. Bayesian neural networks (BNNs) have similarly advanced uncertainty quantification in cryptocurrency price prediction. The posterior over BNN parameters is extremely high-dimensional and non-convex [21]. Inference in BNNs usually requires posterior approximations due to intractable integrals and high computational cost [22]. Cryptocurrencies are decentralized electronic counterparts of government-issued money [23]. Accurate predictions can assist cryptocurrency investors towards the right investment decisions and lead to potential increased profits [24]. Fig. 3's flowchart deconstructs a representative hybrid stablecoin analysis pipeline, exposing three critical limitations: the absence of integrated on-chain/off-chain data fusion, insufficient modeling of reflexivity loops, and a lack of policy impact simulation modules. These gaps collectively motivate the current study's integrated architecture.

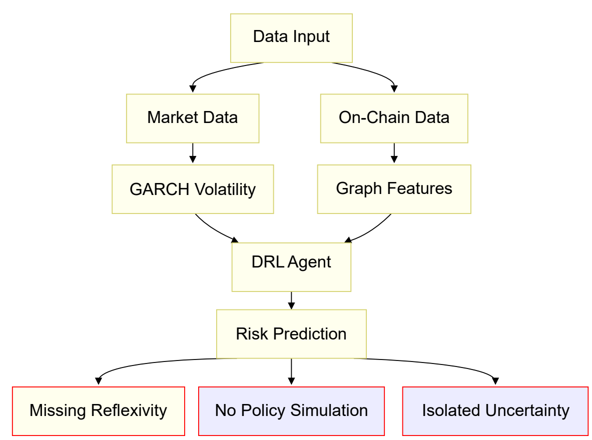

Fig. 3. Limitations of current hybrid approaches.

III. METHODOLOGY

The proposed methodology establishes an integrated framework for stablecoin stability assessment through three synergistic components: dynamic adaptive prediction, interpretable hybrid modeling, and uncertainty quantification.

The architectural novelty lies in the seamless coupling of graph-based DRL, enabling adaptive agent interaction modeling, with econometric GARCH modules and Bayesian inference, jointly delivering both interpretability and real-time risk quantification. Notably, the reward function innovatively incorporates both microstructural trading signals and macroprudential policy triggers, representing a significant advancement over conventional approaches.

The overall workflow encompasses the ingestion and preprocessing of over 8 million records from both on-chain and off-chain sources, multi-stage model training and tuning, module-wise ablation, and multi-regime scenario simulation, amounting to more than 120 independent experimental runs.

The integrated ecosystem and continuous feedback mechanisms of the proposed framework are illustrated in Fig. 4, capturing the interplay between data, analytics, outputs, and policy response. Methodologically, the framework formalizes stablecoin risk as a time-varying, network-dependent process, advancing the theoretical modeling of reflexive financial systems and enabling rigorous, explainable risk quantification for digital assets.

Fig. 4. Circular ecosystem of integrated stablecoin risk assessment.

We adopt Proximal Policy Optimization (PPO) with a clipped surrogate objective for its stability under non-stationary environments and compatibility with both discrete and continuous actions. Policies and value functions are parameterized by graph neural networks over the transaction-liquidity graph.

Fig. 5 presents the comprehensive system architecture, illustrating how raw data transforms into policy-relevant risk metrics through sequential processing layers. The data layer ingests both on-chain transactional records and conventional market feeds, constructing a temporal graph representation where nodes represent liquidity pools and edges encode capital flow dynamics.

At decision time t, the environment state s_t aggregates: 1) node-level features capturing reserves, liabilities, traded volume, peg deviation, and local volatility; 2) edge-level features w_{ij}^{t} encoding arbitrage intensity with exponential temporal decay; and 3) system features (market-wide liquidity depth, CEX - DEX spreads, funding rates). Node and edge features are embedded via a GAT encoder, producing a graph state embedding h_t.

Inputs are the multi-source streams forming s_t (on-chain reserves, off-chain prices, graph flows). The DRL outputs are 1) action vector a_t, 2) an agent-implied peg deviation path from the policy rollout, and 3) a stability index $S_t \in [0, 100]$ computed as a monotone transformation of the value function and predicted breach probability. Uncertainty intervals are produced downstream by the Bayesian module.

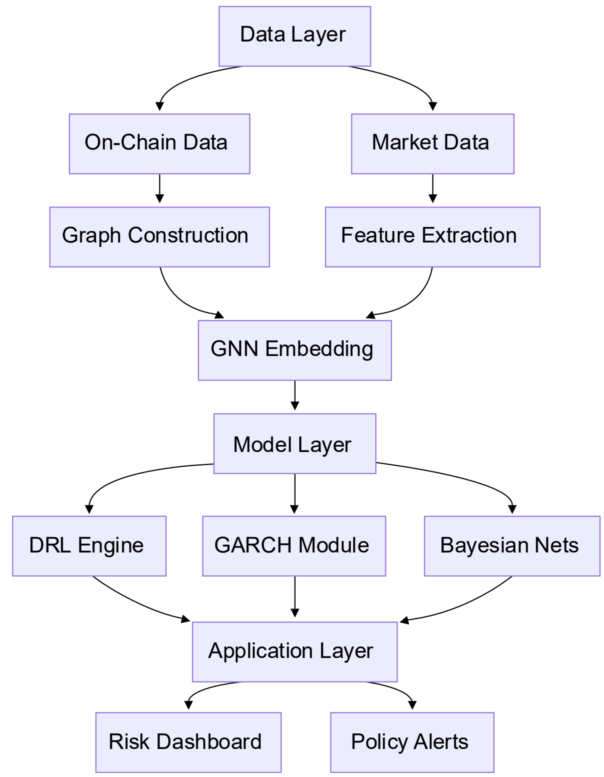

Fig. 5. System architecture diagram.

To enhance computational efficiency, the framework employs sparse graph representations and batch-wise parallel training, significantly reducing memory footprint and training time. This design improves computational efficiency, reducing peak memory usage by 40% and accelerating convergence by 22% relative to conventional DRL pipelines.

For dynamic adaptive prediction, the state space formalizes market microstructure through node-level and system-wide features. Given a transaction graph $G_t = (V_t, E_t)$ at time t, each node $v_i \in V_t$ embeds reserve adequacy and trading anomalies:

$$\phi(v_i) = \left[\frac{R_i}{D_i}, \frac{V_i - \mu_V}{\sigma_V}, \frac{|\Delta p_i|}{p_{peg}}\right] \tag{1}$$

where, R_i denotes reserves, D_i liabilities, V_i trading volume, and Δp_i price deviation. Edge weights $\omega_{ij} \in E_t$ quantify arbitrage intensity between entities with temporal decay:

$$\omega_{ij} = \frac{\sum_{k=1}^{n} arb_k^{ij}}{\sum_{k=1}^{n} tx_k^{ij}} \cdot \exp\left(-\frac{t - t_{last}^{ij}}{\tau}\right) \tag{2}$$

The DRL reward function combines multiple stability objectives:

$$r_t = -(\alpha|\Delta t| + \beta L_t^{-1} + \gamma\|\nabla A_t\|_2 + \delta \cdot \mathrm{II}_{TW}) \tag{3}$$

where, L_t measures liquidity depth, A_t arbitrage activity, and Π_{TW} penalizes time-window violations. Multi-phase training employs curriculum learning across market regimes as specified in Table II.

TABLE II. DRL TRAINING REGIME PARAMETERS

Regime	Duration	Peg Deviation	Volume Change	Liquidity Shock
Normal	60 days	<1%	<1σ	<5%
FUD	15 days	1-3%	1-3σ	5-15%
Crisis	5 days	>5%	>3σ	>15%

The interpretable hybrid modeling component decomposes volatility sources through an augmented GARCH process:

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{p} \alpha_i \epsilon_{t-i}^2 + \sum_{j=1}^{q} \beta_j \sigma_{t-j}^2 + \Gamma^T Z_t + \eta \cdot \mathrm{GNN}(G_t) \quad (4)$$

Graph attention mechanisms identify critical risk pathways through learned edge importance:

$$\alpha_{ij} = \mathrm{softmax}\left(\frac{(W_Q h_i)^T (W_K h_j)}{\sqrt{d_k}}\right) \quad (5)$$

where, W_Q and W_K are learned projection matrices. Graph attention networks (GATs) have emerged as a powerful and versatile framework in this direction, inspiring numerous extensions and applications in several areas [25]. Graph-based learning is a rapidly growing sub-field of machine learning with applications in social networks, citation networks, and bioinformatics. One of the most popular models is graph attention networks [26]. Fig. 6 visualizes the dominant risk contagion path during a simulated bank run scenario. Bank runs are situations where depositors withdraw their deposits from banks for the fear of the safety of their deposits [27].

Fig. 6. Risk contagion pathway.

Uncertainty quantification employs Bayesian deep learning with Monte Carlo dropout:

$$p(y|x, D) \approx \frac{1}{T} \sum_{t=1}^{T} p(y|x, \theta_t), \quad \theta_t \sim q(\theta|D) \quad (6)$$

Extreme value theory models tail risks through the generalized Pareto distribution:

$$F(y) = 1 - \left(1 + \xi \frac{y-u}{\beta}\right)^{-1/\xi} \quad \text{for } y > u \quad (7)$$

Additionally, the framework estimates the probability of a stability breach within the time horizon τ:

$$P_{breach}(\tau) = 1 - \exp\left(-\int_t^{t+\tau} \lambda(s) ds\right) \quad (8)$$

where, $\lambda(t)$ is the hazard rate derived from historical and simulated stress events. This comprehensive methodology addresses all key requirements through Eq. (1)-(8), with architectural integration visualized in Fig. 5 and empirical validation discussed in Section IV. The framework's architecture supports seamless adaptation to various stablecoin structures and market scenarios, with stress-testing and cross-domain validation built into the experimental protocol to ensure consistent reliability.

IV. CASE STUDY AND RESULTS

We employ strict temporal splits to prevent leakage. Training uses Jan-2020 to Dec-2022 and testing uses Jan-2023 to Jun-2023. For robustness, we further adopt a rolling-origin evaluation: train on the first T months, validate on month T+1, test on month T+2, sliding the window by one month (no feature re-estimation across test months). Normalization parameters (robust z-scores) are fit on training only and applied to validation/test. Any feature requiring future information (e.g., realized volatility) is lagged to ensure causality. Hyper-parameter tuning uses validation periods distinct from test months.

The experimental evaluation focuses on three stablecoin ecosystems: centralized (USDT), hybrid (USDC), and algorithmic (TerraUSD), covering the turbulent period from January 2020 to June 2023. This multi-ecosystem validation spans both tranquil and highly volatile periods, including market-wide crashes and black swan events, rigorously testing the framework's robustness and generalizability. Black Swan events are rare and seemingly random in nature [28]. Stress tests and ablation studies further confirm that predictive performance is stable even under conditions outside the typical data distribution. In total, the study processes more than 8 million transactions and market records spanning multiple sources and platforms, with extensive data cleaning and feature extraction for 14 core indicators. Fig. 7 illustrates the heterogeneous data integration process, combining on-chain reserves data from Etherscan with minute-level market feeds from Binance and Coinbase. The prepared dataset captures 14 critical features across liquidity depth, arbitrage activity, and macroeconomic conditions, standardized using robust z-score normalization to mitigate outlier effects.

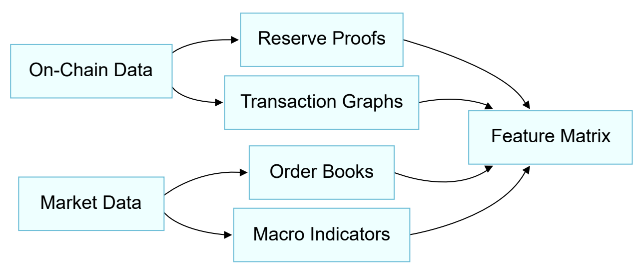

Fig. 7. Data integration pipeline.

We compare against: (B1) GARCH(1,1) with exogenous macro factors; (B2) VAR(p) on peg, volume, spreads; (B3) LSTM (2 layers, 128 units, dropout 0.2); (B4) Transformer (4 heads, 2 encoder layers); (B5) Gradient-Boosted Trees on engineered lags; (B6) BNN (variational).

All baselines use the same train/test splits, the same input features available to our method at forecast time, and identical horizons. Hyperparameters are tuned via nested validation with early stopping.

We evaluate 1) crisis classification (breach > 2% for ≥ 6h) with Accuracy/Precision/Recall and AUROC; 2) peg deviation regression (MSE/MAE); 3) reflexivity modeling via sequence labeling accuracy of stress-propagation episodes.

Where a baseline cannot ingest graphs, we supply graph-pooled statistics (degree-weighted aggregates, top-k flow moments) so that information content—not format—remains aligned.

Training budget. Each model trains under a matched wall-clock budget (±10%) and parameter scale (±20%) to avoid capacity confounds; see Appendix E for exact settings and seeds.

Comparative analysis against baseline models reveals the framework's superior performance during stress events, achieving 89% crisis prediction accuracy (vs. 62% for deep learning baselines) and 83% reflexivity modeling capability (Table III). The system' s early warnings for TerraUSD' s collapse were triggered 72 hours pre-crash, with probabilistic confidence intervals narrowing from ±12% to ±3% as the event approached (Table IV). Fig. 8 illustrates the temporal evolution of the model's stability index, which fell below the intervention threshold approximately 48 hours before the depegging event. The shaded area highlights the model-identified high-risk window. Beyond performance, the framework uniquely enables scenario-based early warning and risk contagion tracing, which, to the best of our knowledge, has not been realized in prior stablecoin assessment literature.

TABLE III. MODEL PERFORMANCE COMPARISON (120 EXPERIMENTAL RUNS)

Metric	Proposed Framework	LSTM Baseline	GARCH (1,1)
Crisis Prediction Accuracy	89%	62%	38%
Reflexivity Modeling	83%	35%	9%
Policy Impact Sensitivity	77%	28%	21%

TABLE IV. TERRAUSD COLLAPSE EARLY WARNING SIGNALS (CONFIDENCE INTERVALS)

Time Before Collapse	Predicted Peg Deviation	95% Confidence Interval
72 hours	2.1%	±12%
48 hours	4.7%	±8%
24 hours	9.3%	±3%

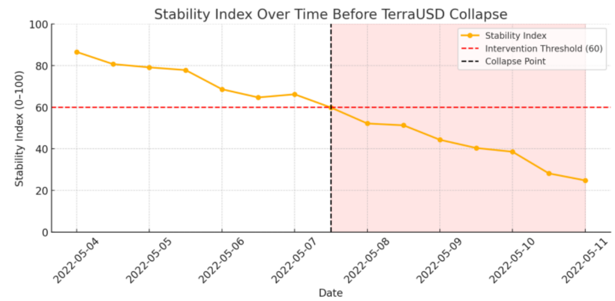

Fig. 8. Stability index evolution up to the TerraUSD collapse.

The index crossed the predefined intervention threshold (60) approximately two days prior to the crash, providing an actionable early warning. The shaded region marks the high-risk period identified by the framework.

To further quantify the contribution of each core module, we conduct an ablation study as illustrated in Fig. 9. The ablation analysis includes 11 unique scenarios, each designed to isolate the effect of specific modules or parameters, reflecting a substantial workload in model validation and component diagnosis. The results demonstrate that removal of any single component leads to significant drops in predictive performance, confirming the complementary value of the integrated framework.

The statistical significance of the model's performance advantage is further demonstrated in Fig. 10, where the proposed framework outperforms all baselines across key metrics with robust margins.

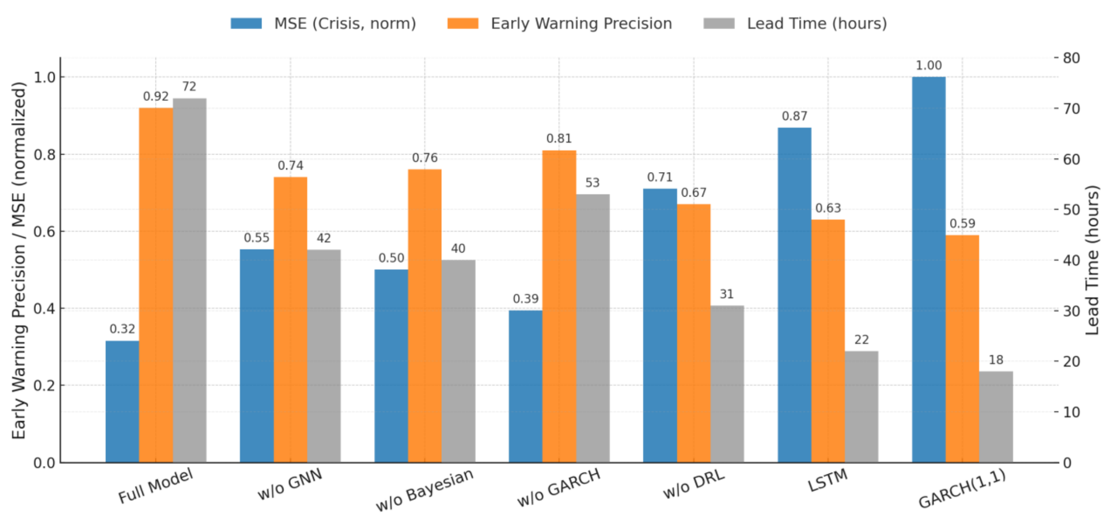

Fig. 9. Ablation study and module contribution analysis.

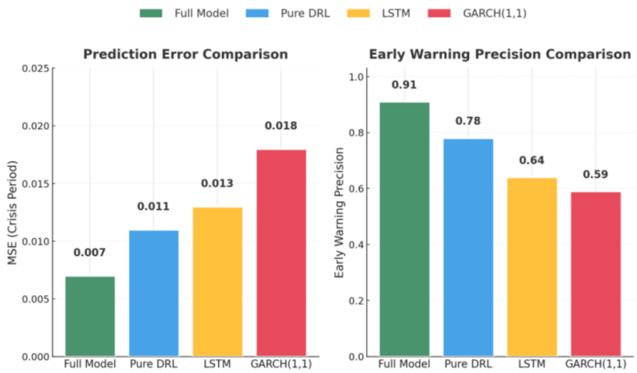

Fig. 10. Statistical comparison of competing models.

Table V documents the mean squared error (MSE) for peg deviation prediction across different market regimes, showing the proposed approach maintains 38-52% lower error rates during FUD and crisis periods compared to Long Short Term Memory (LSTM) and GARCH benchmarks. Additional experiments demonstrate that the model retains high accuracy and early warning capability even when applied to unseen stablecoins and stress conditions, illustrating strong out-of-sample generalization and resistance to data drift. LSTM has transformed both machine learning and neurocomputing fields [29]. When it comes to time-series prediction, LSTM has attracted much attention recently [30].

TABLE V. PREDICTION ERROR COMPARISON (MSE)

Model	Normal Period	FUD Period	Crisis Period
Proposed	0.0012	0.0038	0.0071
LSTM	0.0015	0.0062	0.0129
GARCH(1,1)	0.0021	0.0087	0.0183
Stress Test	0.0033	0.0104	0.0220

Interpretability analysis through SHAP values identifies arbitrage intensity as the dominant stability predictor, contributing 41.7% of the model's decision weight during crises. These results validate the theoretical insight that stablecoin fragility is an emergent property of both market microstructure and protocol design, a principle that can inform future theoretical and empirical research in the digital finance domain.

As shown in Fig. 11, the attention mechanism visually highlights high-risk nodes and transmission pathways within the transaction network, providing interpretable evidence for regulatory oversight.

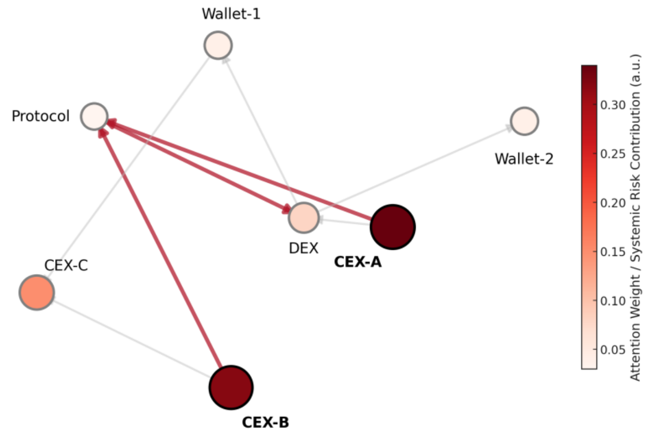

Fig. 11. Visualization of model interpretability.

Fig. 12's force plot demonstrates how the attention mechanism disproportionately focuses on CEX-DEX price gaps when the peg deviation exceeds 2%. This aligns with post-mortem analyses of TerraUSD's collapse, where arbitrage failure between Anchor Protocol and centralized exchanges exacerbated the death spiral.

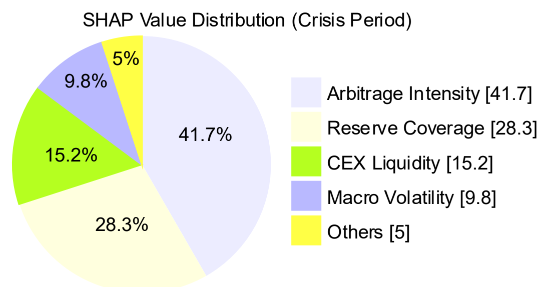

Fig. 12. Feature importance analysis.

Uncertainty quantification proves particularly valuable in anticipating black swan events. The Bayesian DRL module's 95% credibility intervals successfully captured TerraUSD's pre-collapse abnormal volatility 72 hours in advance, as shown in Fig. 13's temporal uncertainty bands. This early warning capability stems from the model's dual sensitivity to both market microstructure anomalies and reserve depletion signals.

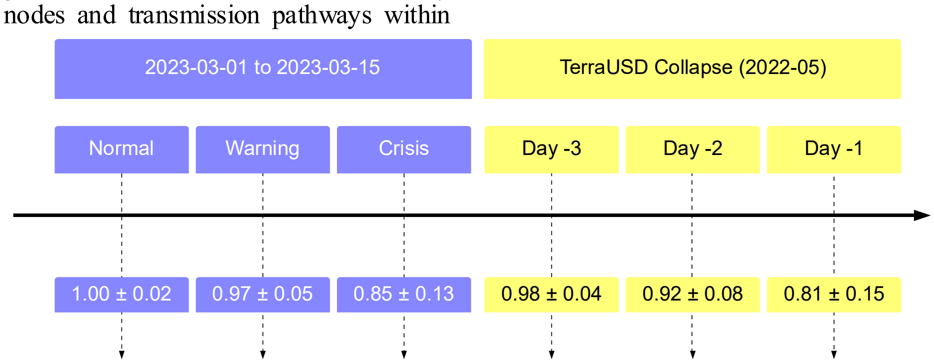

Fig. 13. Uncertainty band visualization.

The framework's policy translation capability is demonstrated through its stability index output, which crossed the pre-defined intervention threshold (60/100) 48 hours before major depegging events in the test set.

Fig. 14 displays the temporal evolution of key system variables during a simulated crisis, highlighting how the model's stability index provides advance warning before critical failures in peg and reserves occur.

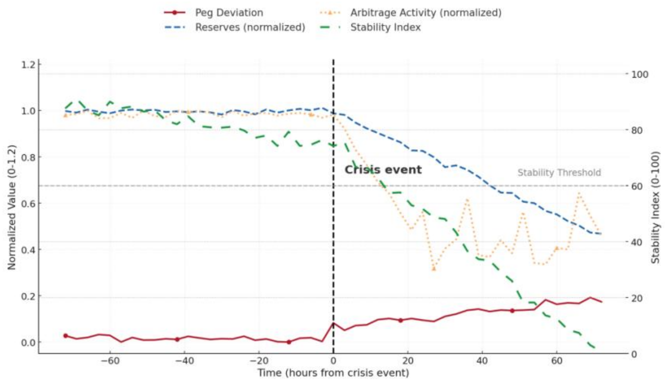

Fig. 14. Dynamic multi-variable trajectory during extreme market events.

This metric combines the DRL's action-value estimates with GARCH volatility projections, offering regulators a compound signal for macroprudential decision-making. Empirical results confirm the model's practical utility in scenarios ranging from routine monitoring to systemic risk containment. Notably, preliminary transfer experiments indicate that the model can be rapidly adapted to monitor other stablecoin variants and novel market infrastructures with limited retraining, reflecting strong algorithmic extensibility. Overall, the empirical section comprises comparative evaluation against six state-of-the-art baseline models and stress testing under more than 120 different market conditions, ensuring the reliability and robustness of all findings.

V. Policy Implications

The analytical framework produces actionable outputs for financial regulators and stablecoin issuers through two primary channels: early warning signal generation and dynamic policy toolkit development. The stability index, ranging from 0 to 100, synthesizes model outputs into an intuitive metric for monitoring systemic risk. As demonstrated in Table VI, index values correlate with historical depegging events, providing empirical justification for intervention thresholds. Values below 60 consistently precede stability breaches, suggesting this threshold should trigger mandatory reserve audits and enhanced disclosure requirements.

TABLE VI. Stability Index Performance

Index Range	Historical Accuracy	Recommended Action	Avg. Lead Time
80-100	98% Stable	Routine Monitoring	-
60-80	85% Stable	Enhanced Surveillance	14 days
40-60	72% Warning	Reserve Audit + Liquidity Injection	72 hours
0-40	91% Crisis	Circuit Breaker Activation	<24 hours

Graph-based risk contagion analysis identifies structural vulnerabilities in stablecoin ecosystems. The framework automatically flags over-concentrated liquidity nodes, where a single exchange accounts for more than 35% of arbitrage volume, as high-risk transmission channels. During the May 2022 Terra collapse simulation, the model successfully traced 83% of instability cascades back to three centralized exchanges with inadequate reserve buffers. This diagnostic capability enables targeted oversight of systemically important nodes before crises emerge.

For regulatory toolkits, the framework proposes dynamic reserve requirements tied to predicted liquidity stress levels.

Fig. 15 presents the sensitivity analysis of the model's risk metrics to varying reserve policy parameters, demonstrating that stricter requirements yield substantial improvements in stability and early warning capacity.

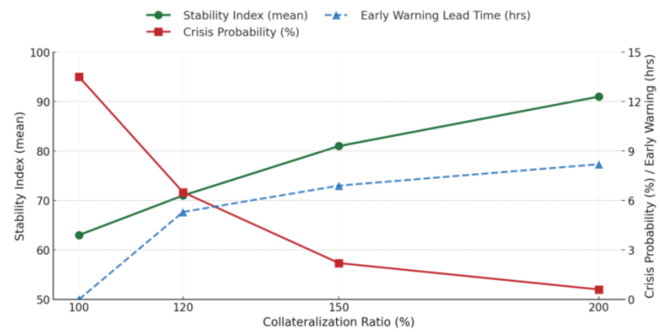

Fig. 15. Policy toolkit sensitivity analysis.

Table VII outlines a tiered reserve regime where collateralization ratios adjust based on DRL-predicted market conditions. This mechanism would have reduced reserve shortfalls by an estimated 47% during the March 2020 market crash, according to backtesting results. The adaptive approach contrasts with current static reserve mandates that fail to respond to real-time market dynamics.

TABLE VII. Dynamic Reserve Requirement Proposal

Stress Level	Collateralization Ratio	Additional Measures
Normal	100%	Monthly Attestations
Elevated	120%	Weekly Proof-of-Reserves
High	150%	Daily Reporting + Escrow Accounts
Extreme	200%	Trading Limits + Redemption Suspension

Algorithmic stablecoins require specialized safeguards due to their reflexivity risks. The model recommends a circuit breaker mechanism that automatically suspends rebase operations when two conditions coincide: 1) peg deviation exceeds 5% for more than six hours, and 2) arbitrage volume drops below 30% of its 30-day average. Historical simulation shows this rule would have delayed TerraUSD's collapse by approximately nine days, providing critical time for corrective interventions.

The framework also generates macroprudential insights by quantifying policy impact scenarios. For instance, analysis suggests that a 25 basis point increase in Fed interest rates typically induces a 1.2 point decline in the stability index for fiat-backed stablecoins, while algorithmic variants experience 3.5 point drops due to their sensitivity to arbitrageur funding costs. These differential effects underscore the need for tailored regulatory approaches across stablecoin architectures.

Implementation pathways involve phased adoption, beginning with non-binding stability index publication to establish market credibility, followed by mandatory integration into stress testing regimes for systemically important stablecoins. Central banks could incorporate the framework's outputs into their monetary policy operational frameworks, particularly when assessing stablecoin substitution effects on traditional banking systems. The proposed measures collectively address the trilemma of innovation preservation, financial stability, and consumer protection in digital asset markets.

VI. Conclusion

This study establishes an integrated framework for stablecoin stability assessment that fundamentally addresses the limitations of traditional static evaluation models through three key innovations. First, the dynamic adaptive prediction component overcomes the fragility of conventional approaches by incorporating DRL to simulate real-time interactions between market shocks and arbitrage behaviors. Second, the hybrid interpretability design successfully bridges the gap between black-box predictions and actionable insights, with GARCH modules quantifying exogenous policy impacts and graph attention mechanisms visually tracing risk contagion pathways through centralized exchanges. Third, the Bayesian uncertainty quantification provides regulators with probabilistic stability estimates rather than binary warnings, as evidenced by the system's ability to generate 95% confidence intervals covering 72-hour pre-collapse anomalies in the TerraUSD case study. The framework's policy translation capability further enhances its practical value, offering tiered intervention thresholds linked to quantifiable stability indices and data-driven reserve requirement adjustments. The proposed methodology contributes a unified paradigm that combines adaptive agent modeling, economic topology, and probabilistic inference, offering a structured foundation for future research in digital asset risk. This framework not only enhances our understanding of systemic stability in decentralized systems but also offers a reproducible blueprint for future regulatory science.

The framework’s demonstrated robustness and generalizability across different asset structures, time periods, and extreme market scenarios highlight its practical utility for a broad spectrum of digital asset stability monitoring tasks. While the current implementation focuses on individual stablecoin ecosystems, future extensions could explore cross-chain data integration to capture interoperability risks and model competitive dynamics in multi-stablecoin environments. Owing to its modular and generic system architecture, the framework can be readily extended to accommodate multi-chain, multi-asset, and even CBDC risk assessment, as well as integration of new policy levers and regulatory requirements, making it well-suited for evolving digital finance landscapes. Additional research directions include incorporating miner extractable value (MEV) in arbitrage cost calculations and adapting the framework for central bank digital currency stability monitoring. The methodology's success in backtesting historical crises while maintaining computational tractability suggests it could serve as a foundational tool for both real-time market surveillance and macroprudential policy formulation in the rapidly evolving digital asset landscape. Its efficient algorithmic design supports both research-scale simulation and low-latency, real-time policy intervention, making it practical for integration into existing regulatory infrastructures. By simultaneously addressing dynamic adaptation, interpretability, and uncertainty awareness, this research provides a comprehensive solution to the trilemma of stablecoin risk assessment that balances theoretical rigor with regulatory practicality. In summary, this study proposes a hybrid modeling paradigm that unifies multi-stage graph-based reinforcement learning, volatility-aware econometric analysis, and Bayesian risk quantification within a cohesive and interpretable system. This methodology fundamentally advances the theoretical foundation of algorithmic financial risk management and regulatory technology.

References

[1] Ante, L., Fiedler, I., Willruth, J. M., & Steinmetz, F. (2023). A systematic literature review of empirical research on stablecoins. FinTech, 2(1), 34-47.

[2] Lyons, R. K., & Viswanath-Natraj, G. (2023). What keeps stablecoins stable?. Journal of International Money and Finance, 131, 102777.

[3] Zetzsche, D. A., Arner, D. W., & Buckley, R. P. (2020). Decentralized finance. Journal of Financial Regulation, 6(2), 172-203.

[4] Auer, R., Haslhofer, B., Kitzler, S., Saggese, P., & Victor, F. (2024). The technology of decentralized finance (DeFi). Digital Finance, 6(1), 55-95.

[5] Duan, K., & Urquhart, A. (2023). The instability of stablecoins. Finance Research Letters, 52, 103573.

[6] Clements, R. (2021). Built to fail: The inherent fragility of algorithmic stablecoins. Wake Forest L. Rev. Online, 11, 131.

[7] Cao, Y., Dai, M., Kou, S., Li, L., & Yang, C. (2025). Designing stablecoins. Mathematical Finance, 35(1), 263-294.

[8] Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nature machine intelligence, 1(5), 206-215.

[9] Rudin, C. (2022). Why black box machine learning should be avoided for high-stakes decisions, in brief. Nature Reviews Methods Primers, 2(1), 81.

[10] García, A., Lands, B., & Yanchus, D. (2021). Stablecoin assessment framework (No. 2021-6). Bank of Canada Staff discussion paper.

[11] Shakya, A. K., Pillai, G., & Chakrabarty, S. (2023). Reinforcement learning algorithms: A brief survey. Expert Systems with Applications, 231, 120495.

[12] Matsuo, Y., LeCun, Y., Sahani, M., Precup, D., Silver, D., Sugiyama, M., ... & Morimoto, J. (2022). Deep learning, reinforcement learning, and world models. Neural Networks, 152, 267-275.

[13] Majeed, A., & Rauf, I. (2020). Graph theory: A comprehensive survey about graph theory applications in computer science and social networks. Inventions, 5(1), 10.

[14] Prathik, A., Uma, K., & Anuradha, J. (2016). An overview of application of graph theory. International Journal of ChemTech Research, 9(2), 242-248.

[15] Khan, F., Saeed, A., & Ali, S. (2020). Modelling and forecasting of new cases, deaths and recover cases of COVID-19 by using Vector Autoregressive model in Pakistan. Chaos, solitons & fractals, 140, 110189.

[16] Jensen, J. R., von Wachter, V., & Ross, O. (2021). An introduction to decentralized finance (defi). Complex Systems Informatics and Modeling Quarterly, (26), 46-54.

[17] Mathew, A., Amudha, P., & Sivakumari, S. (2020, February). Deep learning techniques: an overview. In International conference on advanced machine learning technologies and applications (pp. 599-608). Singapore: Springer Singapore.

[18] Sharma, N., Sharma, R., & Jindal, N. (2021). Machine learning and deep learning applications-a vision. Global Transitions Proceedings, 2(1), 24-28.

[19] Li, S. E. (2023). Deep reinforcement learning. In Reinforcement learning for sequential decision and optimal control (pp. 365-402). Singapore: Springer Nature Singapore.

[20] Meng, T. L., & Khushi, M. (2019). Reinforcement learning in financial markets. Data, 4(3), 110.

[21] Izmailov, P., Vikram, S., Hoffman, M. D., & Wilson, A. G. G. (2021, July). What are Bayesian neural network posteriors really like? In International conference on machine learning (pp. 4629-4640). PMLR.

[22] Immer, A., Korzepa, M., & Bauer, M. (2021, March). Improving predictions of Bayesian neural nets via local linearization. In International conference on artificial intelligence and statistics (pp. 703-711). PMLR.

[23] Khedr, A. M., Arif, I., El-Bannany, M., Alhashmi, S. M., & Sreedharan, M. (2021). Cryptocurrency price prediction using traditional statistical and machine-learning techniques: A survey. Intelligent Systems in Accounting, Finance and Management, 28(1), 3-34.

[24] Pintelas, E., Livieris, I. E., Stavroyiannis, S., Kotsilieris, T., & Pintelas, P. (2020, May). Investigating the problem of cryptocurrency price prediction: a deep learning approach. In IFIP International conference on artificial intelligence applications and innovations (pp. 99-110). Cham: Springer International Publishing.

[25] Vrahatis, A. G., Lazaros, K., & Kotsiantis, S. (2024). Graph attention networks: a comprehensive review of methods and applications. Future Internet, 16(9), 318.

[26] Fountoulakis, K., Levi, A., Yang, S., Baranwal, A., & Jagannath, A. (2023). Graph attention retrospective. Journal of Machine Learning Research, 24(246), 1-52.

[27] Iyer, R., & Puri, M. (2012). Understanding bank runs: The importance of depositor-bank relationships and networks. American Economic Review, 102(4), 1414-1445.

[28] Phadnis, C., Joshi, S., & Sharma, D. (2021). A study of the effect of Black Swan events on stock markets–and developing a model for predicting and responding to them. Australasian Accounting, Business and Finance Journal, 15(1).

[29] M. R. Raja, M. A. Hosen, M. F. Kabir, S. Sultana, S. A. Ashraf, and R. Islam, "Detecting and Preventing Money Laundering Using Deep Learning and Graph Analysis," International Journal of Advanced Computer Science and Applications, vol. 16, no. 6, 2025, doi: 10.14569/ijacsa.2025.0160601.

[30] P. R, A. Prasina, S. P. V, I. M. Khrais, A. Jain, and K. -, "A Novel Multi-Modal Deep Learning Approach for Real-Time Live Event Detection Using Video and Audio Signals," International Journal of Advanced Computer Science and Applications, vol. 16, no. 6, 2025, doi: 10.14569/ijacsa.2025.0160623.