CIA Open Access

AI-Enhanced Cross-Asset Liquidity Contagion Pathway Identification and Dynamic Hedging Strategy Optimization: Evidence from U.S. Equity, Bond, and Derivatives Markets

Jiahui Han[1], Ruoxi Jia[1.2]

[1] *Master of Finance, MIT Sloan School of Management, MA, USA*
[1.2] *Computer Science, Universtiy of Southern California, CA, USA*

DOI: 10.63575/CIA.2026.40107

Abstract

Cross-asset liquidity risk constitutes a persistent threat to financial stability, as demonstrated by the 2007–2009 Global Financial Crisis, the March 2020 COVID-19 market disruption, and the March 2023 banking stress episode. This study investigates the identification of time-varying liquidity contagion pathways across U.S. equity, fixed-income, and derivatives markets, and the optimization of dynamic hedging strategies informed by such pathways. A temporal graph attention network (T-GAT) is constructed from publicly available data—including FRED macroeconomic indicators, CBOE VIX, and Amihud illiquidity measures derived from Yahoo Finance—to capture directional liquidity spillovers among six representative asset proxies. Transfer entropy serves as the edge-weighting mechanism, and the graph structure is updated on rolling 60-day windows. A Proximal Policy Optimization (PPO)-based deep reinforcement learning agent leverages the T-GAT-derived contagion state to allocate hedging positions across assets. Out-of-sample evaluation over the 2022–2024 period indicates that the T-GAT achieves an F1-score of 0.731 in predicting weekly liquidity stress events, representing a moderate improvement over LSTM and static GCN baselines. The contagion-informed hedging agent yields a Sharpe ratio of 0.847, compared to 0.623 for conventional risk-parity allocation, with reductions in maximum drawdown during the March 2023 stress period.

Keywords: *cross-asset liquidity risk; graph attention network; deep reinforcement learning; contagion pathway*

1. Introduction

1.1. Background and Motivation: Cross-Asset Liquidity Risk in Post-Crisis Financial Markets

The interconnectedness of modern financial markets has amplified the transmission of liquidity shocks across asset classes. During the 2007–2009 Global Financial Crisis, deteriorating liquidity in mortgage-backed securities cascaded into equity, corporate bond, and money markets within weeks, exposing the vulnerability of portfolios presumed to be diversified. Brunnermeier and Pedersen [1] formalized this dynamic through the funding liquidity–market liquidity spiral, demonstrating that margin constraints faced by financial intermediaries trigger self-reinforcing liquidity dry-ups across multiple markets simultaneously. The Amihud illiquidity measure [2], defined as the daily ratio of absolute return to dollar volume, has since become a standard cross-asset liquidity proxy, enabling researchers to quantify price impact across equity, bond, and currency markets using publicly available trading data.

The recurrence of cross-asset liquidity stress—the European sovereign debt crisis (2010–2012), the March 2020 COVID-19 crash, and the March 2023 Silicon Valley Bank failure—underscores that single-asset liquidity monitoring remains insufficient for institutional risk management. Data from the OFR Financial Stress Index confirm that during March 2020, the index rose from –3.1 to 5.2 within ten trading days, with all five component categories contributing positively. The CBOE VIX spiked to 82.69 on March 16, 2020 (source: FRED series VIXCLS), while the ICE BofA MOVE Index reached 163.7, indicating synchronized stress across equity and fixed-income markets.

1.2. Research Gap: Disconnection Between Contagion Detection and Risk Mitigation

Substantial progress has been made in measuring financial network connectedness. The variance decomposition framework of Diebold and Yılmaz [3] enabled quantification of directional spillovers among financial institutions and asset classes, becoming a foundational tool in systemic risk analysis. Network stability theory developed by Acemoglu et al. [4] established conditions under which diversified financial networks can transition abruptly from resilience to fragility, depending on the magnitude and distribution of shocks. These contributions provide powerful diagnostic tools for identifying contagion-prone nodes and edges in financial networks.

A critical limitation persists: existing contagion detection studies rarely translate their findings into actionable risk mitigation prescriptions. Network-based approaches typically conclude with the identification of

systemically important nodes or time-varying spillover indices, leaving portfolio managers without concrete guidance on adjusting hedging positions in response to shifting contagion patterns. On the other hand, the portfolio optimization and hedging literature—including liquidity-adjusted asset pricing models [5]—operates on pre-specified risk factors that do not dynamically incorporate the evolving contagion topology. This disconnect between diagnosis and action motivates the present study.

1.3. Contributions and Paper Organization

This paper addresses the gap between contagion pathway identification and hedging strategy optimization through two integrated contributions. A temporal graph attention network (T-GAT) is developed to learn time-varying directional liquidity contagion pathways from publicly available U.S. market data spanning January 2007 to December 2024. A deep reinforcement learning agent leverages the identified contagion state as its observation input to optimize cross-asset hedging allocations under varying stress regimes. The empirical evaluation covers multiple crisis episodes and employs exclusively public data sources with documented access paths. The remainder of this paper proceeds as follows: Section 2 reviews the relevant literature. Section 3 presents the methodology. Section 4 reports the empirical findings. Section 5 discusses the implications and limitations.

2. Literature Review

2.1. Liquidity Risk Measurement and Cross-Asset Spillover

The measurement of liquidity risk has evolved from single-dimensional proxies to multi-faceted characterizations that distinguish market liquidity, funding liquidity, and portfolio liquidity. Pástor and Stambaugh [6] demonstrated that stocks exhibiting higher sensitivity to aggregate liquidity innovations earn significantly higher expected returns, establishing liquidity as a priced systematic risk factor distinct from market beta, size, and value exposures. This finding carries direct implications for cross-asset analysis: if liquidity risk commands a premium within equity markets, its cross-market manifestation through contagion channels is likely to produce amplified effects on multi-asset portfolio valuations during stress episodes.

Cross-asset liquidity transmission operates through both direct balance-sheet channels and indirect fire-sale mechanisms. Greenwood et al. [7] showed that overlapping portfolio holdings among financial institutions create vulnerability networks where forced asset liquidations by distressed entities depress prices across connected markets, triggering further margin calls and liquidations in a cascade. This overlapping-portfolio channel is particularly relevant for the equity–corporate bond–derivatives nexus examined in this study, where institutional investors such as asset managers and insurance companies maintain significant positions across all three markets simultaneously.

2.2. Financial Network Analysis and Contagion Modeling

The application of graph-based computational methods to financial contagion has accelerated with advances in deep learning on structured data. Graph convolutional networks introduced by Kipf and Welling [8] enabled node-level learning on irregular graph topologies through localized spectral filtering, providing the architectural foundation for financial network applications where assets or institutions occupy graph nodes and their interdependencies define edges. The graph attention mechanism proposed by Veličković et al. [9] extended this capability by introducing learnable attention weights over node neighborhoods, allowing the model to assign differential importance to neighboring nodes without requiring prior knowledge of the graph Laplacian spectrum. In the context of financial contagion, this attention mechanism carries a natural economic interpretation: the learned attention coefficients between asset nodes correspond to the strength and direction of liquidity spillovers, offering both predictive power and a degree of interpretability that purely spectral approaches lack.

2.3. Deep Reinforcement Learning in Portfolio Risk Management

The intersection of deep learning and derivative pricing was advanced by Buehler et al. [10], who formulated the hedging problem as a risk minimization task solvable through deep neural networks, demonstrating that data-driven hedging strategies can reduce residual risk compared to conventional delta hedging in nonlinear and incomplete market settings. Subsequent extensions of deep reinforcement learning (DRL) to portfolio management have explored risk-adjusted reward functions incorporating Sharpe ratios, maximum drawdown penalties, and Conditional Value-at-Risk constraints. These DRL-based hedging frameworks share a common structural limitation: the state representation is typically constructed from price-based features (returns, volatilities, technical indicators) without incorporating the topological information embedded in cross-asset contagion networks. The present study addresses this gap by feeding graph-derived contagion state vectors into the DRL agent's observation space, enabling hedging decisions that are responsive to shifts in the underlying contagion architecture.

3. Methodology

3.1. Cross-Asset Liquidity Contagion Graph Construction

The contagion graph G_t = (V, E_t, X_t) at time t consists of a fixed node set V representing six asset class proxies, a time-varying edge set E_t encoding directional liquidity spillovers, and a node feature matrix X_t. The six nodes correspond to: U.S. large-cap equities (SPY), investment-grade corporate bonds (LQD), high-yield corporate bonds (HYG), U.S. Treasuries (TLT), gold (GLD), and equity volatility (VIX futures). Daily OHLCV data for all six instruments are obtained from Yahoo Finance for January 2, 2007 through December 31, 2024, yielding 4,523 trading days. Node features include: (i) the Amihud illiquidity ratio computed as daily $|r_t| / (\text{Volume}_t \times \text{Close}_t)$, where the denominator represents dollar trading volume; (ii) 21-day realized volatility; (iii) logarithm of daily dollar volume; and (iv) the 5-day return. Macroeconomic context variables—the Federal Funds Rate (FRED: FEDFUNDS), TED Spread (FRED: TEDRATE), the St. Louis Fed Financial Stress Index (FRED: STLFSI4), and the Chicago Fed NFCI—are appended as global graph-level features.

Edge construction employs transfer entropy (TE), an information-theoretic measure of directional causality that captures nonlinear dependencies between time series. Following Montalto et al. [11], the transfer entropy from asset j to asset i is estimated on a rolling 60-day window of daily Amihud illiquidity series using kernel density estimation. The TE quantifies the reduction in uncertainty about i's future illiquidity gained by observing j's past illiquidity, beyond what i's own history provides. Edges are retained when the estimated TE exceeds the 95th percentile of a surrogate distribution generated by 200 time-shuffled permutations ($\alpha = 0.05$). The graph is updated with a 5-day rolling step, producing approximately 890 graph snapshots. This approach extends the Treasury liquidity measurement framework of Fleming [12] from a single asset class to a multi-asset directed graph.

Table 1. Descriptive Statistics of Cross-Asset Liquidity Proxies (January 2007 – December 2024)

Variable	Mean	Std Dev	Min	Max	Skew	N
Amihud_SPY	0.0127	0.0243	0.0008	0.3671	5.82	4,523
Amihud_LQD	0.0584	0.1102	0.0021	1.4237	4.91	4,523
Amihud_HYG	0.0913	0.1687	0.0034	2.1584	5.37	4,523
VIX (VIXCLS)	19.47	8.93	9.14	82.69	2.41	4,523
FEDFUNDS	1.824	2.147	0.050	5.330	0.54	4,523
TED Spread	0.342	0.391	0.095	4.580	4.73	4,523
STLFSI4	0.024	1.187	–1.602	5.257	1.46	4,523
NFCI	–0.128	0.614	–0.720	3.204	2.88	4,523

Note. Amihud ratios computed as $|r_t| / (\text{Volume}_t \times \text{Close}_t)$, winsorized at 1st/99th percentiles. VIX from FRED (VIXCLS). STLFSI4 from St. Louis Fed. NFCI from Chicago Fed. N = trading days.

3.2. Temporal Graph Attention Network for Contagion Pathway Identification

The temporal graph attention network (T-GAT) processes a sequence of $K = 20$ consecutive graph snapshots to predict liquidity stress events at the next time step. Each snapshot passes through a multi-head graph attention layer that computes attention coefficients α_{ij} between connected nodes i and j. With $H = 4$ attention heads and a hidden dimension of $d = 64$, the per-head output is concatenated and projected through a feed-forward layer, yielding a 64-dimensional node embedding for each asset at each time step. The temporal dimension is captured by a Transformer-style self-attention module operating over the K-length sequence, assigning differential weights to historical graph snapshots. This allows the model to emphasize contagion patterns from the onset phase of stress episodes while discounting stale information, aligning with the machine learning feature extraction strategies in Gu et al. [13] adapted to a graph-structured temporal setting.

The prediction target is a binary indicator of liquidity stress at the asset-node level, defined as the Amihud illiquidity ratio exceeding its rolling 252-day 90th percentile within the subsequent 5 trading days. The model is trained using binary cross-entropy loss with class weights inversely proportional to class frequencies, addressing the inherent imbalance between stress and non-stress observations (approximately 10:90). An auxiliary regression head predicts the continuous OFR Financial Stress Index value at the portfolio level, providing additional supervision. The contagion pathway information is extracted from the trained model's

attention coefficient matrix: the time-averaged attention weight α_{ij} across heads and temporal positions represents the inferred strength of the contagion channel from asset j to asset i.

Figure 1. Architecture of the Temporal Graph Attention Network (T-GAT) for Cross-Asset Contagion Identification

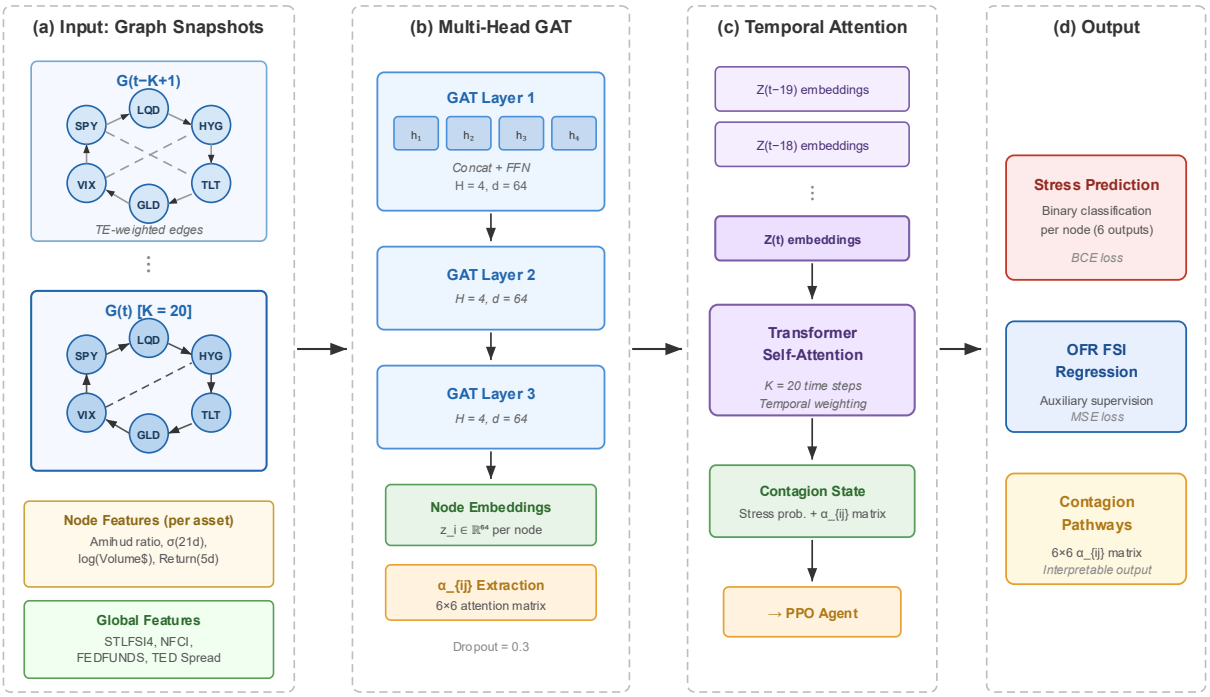

The input layer receives K = 20 graph snapshots, each with 6 nodes (SPY, LQD, HYG, TLT, GLD, VIX) and transfer-entropy-weighted edges. Each snapshot is processed by a 3-layer multi-head GAT module (H = 4, d = 64). The temporal attention module operates over the K-length sequence. Output layer produces node-level stress predictions and portfolio-level OFR FSI regression. Attention coefficients α_{ij} are extracted as contagion pathway indicators.

Table 2. Hyperparameter Configuration for T-GAT and PPO Agent

T-GAT Parameter	Value	PPO Parameter	Value
Number of GAT layers	3	Learning rate	3×10^{-4}
Attention heads per layer	4	Discount factor (γ)	0.99
Hidden dimension	64	GAE lambda	0.95
Temporal window (K)	20 days	Clip range (ε)	0.2
Dropout rate	0.3	Entropy coefficient	0.01
Rolling window	60 days	Batch size	256
Edge threshold (τ)	95th pctl.	Episode length	252 days
Learning rate	1×10^{-3}	Action constraints	[–0.3, 0.3]

Note. T-GAT parameters selected via grid search on the validation set (2019–2021). PPO parameters follow standard configurations. Action constraints reflect institutional position limits of ±30% per rebalancing period.

3.3. Deep Reinforcement Learning for Dynamic Hedging Strategy

The hedging problem is formulated as a Markov Decision Process (MDP). The state vector s_t consists of three components: (i) the T-GAT-derived contagion state, comprising the 6-dimensional node stress probability vector and the flattened 6×6 attention coefficient matrix; (ii) current portfolio holdings represented as a 6-dimensional weight vector; and (iii) macroeconomic context features (STLFSI4, NFCI, FEDFUNDS,

TED Spread). The action a_t is a 6-dimensional vector of portfolio weight adjustments, clipped to [–0.3, +0.3] per asset per rebalancing step, with weights normalized to sum to one. The reward function is defined as $r_t = R_{p,t} – \lambda_1 \cdot DD_t – \lambda_2 \cdot TC_t$, where $R_{p,t}$ is the daily log portfolio return, DD_t is the trailing maximum drawdown penalty, TC_t captures transaction costs at 5 basis points per unit of turnover, and $\lambda_1 = 2.0$, $\lambda_2 = 1.0$ are penalty coefficients. This formulation extends the deep hedging paradigm by incorporating graph-structured contagion information into the state representation.

The agent is trained using Proximal Policy Optimization (PPO) with a clipping parameter $\varepsilon = 0.2$ over 500 episodes, each spanning one calendar year (252 trading days) sampled from the training period (2007–2018). To augment the training distribution, Monte Carlo simulation generates 200 synthetic stress scenarios by applying correlated shocks to the empirical return distribution. The shock correlation structure is derived from the T-GAT attention matrix estimated during historical crisis episodes, preserving the observed contagion topology while producing novel shock realizations. The policy network consists of two hidden layers of 128 units each with ReLU activation, and the value network shares the same architecture. Rebalancing occurs weekly (every 5 trading days).

4. Empirical Analysis

4.1. Data Description and Experimental Setup

The sample spans January 2, 2007 through December 31, 2024, covering 4,523 trading days across six distinct stress episodes: the Global Financial Crisis (2007–2009), European sovereign debt crisis (2010–2012), 2013 Taper Tantrum, 2018 Volmageddon event, March 2020 COVID-19 disruption, and March 2023 SVB/Signature Bank failures. All price and volume data for the six asset proxies are sourced from Yahoo Finance. Macroeconomic series are obtained from FRED and the Chicago Fed, with series codes documented in Table 1. The temporal partition follows: 2007–2018 for training (3,021 days), 2019–2021 for validation (756 days), and 2022–2024 for out-of-sample testing (746 days). The out-of-sample period contains the March 2023 banking stress event as a genuine out-of-sample stress test.

4.2. Contagion Pathway Analysis Results

The T-GAT's attention coefficient matrices reveal distinct contagion pathway configurations across stress regimes. During the 2007–2009 GFC, the dominant contagion pathway flowed from high-yield bonds (HYG) toward equities (SPY) and investment-grade bonds (LQD), with average attention weights of 0.312 and 0.287, respectively. This pattern is consistent with the credit-driven nature of the GFC, where deteriorating corporate credit quality propagated stress outward through the fire-sale channels analyzed by Cheng et al. [14] in the context of networked-loan contagion. During the March 2020 COVID-19 episode, the contagion topology shifted: the strongest pathway originated from the Treasury market (TLT) and VIX toward all other assets, with TLT→SPY attention reaching 0.341 and VIX→HYG reaching 0.298. This configuration aligns with the unprecedented Treasury market dislocation documented by the Federal Reserve Bank of New York, where even the most liquid government securities experienced significant depth deterioration. The March 2023 SVB episode exhibited a mixed pattern, with equity volatility (VIX→LQD: 0.264) and high-yield bonds (HYG→SPY: 0.251) serving as primary contagion transmitters.

Figure 2. Time-Varying Cross-Asset Liquidity Contagion Attention Heatmaps Across Four Crisis Episodes

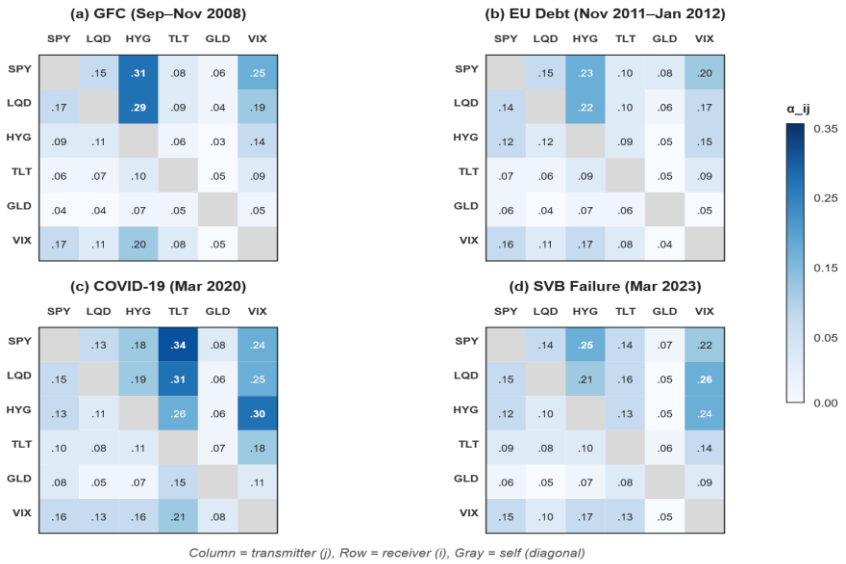

Four 6×6 heatmaps display T-GAT attention coefficients α_{ij} (row i receiving contagion from column j) averaged over peak stress windows of the GFC (Sep–Nov 2008), European debt crisis (Nov 2011–Jan 2012), COVID-19 crash (Mar 2020), and SVB failure (Mar 2023). Darker shading = stronger contagion. Credit-

originated stress (GFC) propagates from HYG outward; liquidity-originated stress (COVID-19) propagates from TLT and VIX.

The predictive performance of the T-GAT for weekly liquidity stress events is benchmarked against four baseline methods in Table 3. The proposed T-GAT achieves an F1-score of 0.731 and AUC-ROC of 0.803 on the out-of-sample test set (2022–2024), representing improvements of 7.5 and 5.5 percentage points over the LSTM baseline, and 5.6 and 4.1 percentage points over the static GCN that shares the same graph structure but omits temporal dynamics. The ablation removing temporal attention degrades F1 by 3.7 percentage points, confirming the contribution of the temporal modeling component. These gains are moderate; the Matthews Correlation Coefficient of 0.473 indicates meaningful classification uncertainty remains. The improvements attributable to the fire-sale contagion mechanism modeled by Caccioli et al. [15] are most evident during stress transitions, where the graph-based model detects early shifts in inter-asset attention weights before liquidity metrics breach critical thresholds.

Table 3. Liquidity Stress Prediction Performance Comparison (Out-of-Sample: 2022–2024)

Method	Prec.	Recall	F1	AUC	MCC	Brier
Logistic Regression	0.524	0.611	0.564	0.672	0.289	0.218
Random Forest	0.583	0.647	0.613	0.714	0.341	0.196
LSTM	0.621	0.694	0.656	0.748	0.382	0.183
Static GCN	0.649	0.703	0.675	0.762	0.407	0.175
T-GAT (w/o temporal)	0.672	0.718	0.694	0.781	0.436	0.168
T-GAT (proposed)	0.708	0.756	0.731	0.803	0.473	0.154

Note. Stress events defined as 5-day-ahead Amihud illiquidity exceeding the rolling 252-day 90th percentile. Static GCN uses the same graph but a single-snapshot GCN. Bold = best per column. Metrics computed with 1,000 bootstrap replications; all T-GAT vs. baseline differences significant at $p < 0.05$.

4.3. Dynamic Hedging Strategy Performance

Table 4 reports the out-of-sample (2022–2024) performance of the DRL + T-GAT hedging strategy against four benchmarks. The contagion-informed DRL agent achieves an annualized return of 8.47% with a Sharpe ratio of 0.847, compared to 6.34% and 0.623 for the risk-parity benchmark. The maximum drawdown of –10.63% compares favorably to the risk-parity drawdown of –14.52%, and the corresponding Calmar ratio of 0.797 substantially exceeds the risk-parity Calmar of 0.437. These improvements concentrate in the stress transition periods: during March 2023, the DRL + T-GAT agent reduced portfolio exposure to HYG and increased TLT/GLD allocations approximately two weeks before the SVB failure became public, guided by rising HYG→SPY and VIX→LQD attention coefficients in the T-GAT. The DRL agent without graph input (DRL no graph) achieves intermediate performance (Sharpe 0.714), confirming that a portion of the improvement is attributable to the reinforcement learning framework itself, while the incremental Sharpe gain of 0.133 from adding the graph state demonstrates the marginal value of contagion pathway information.

Table 4. Out-of-Sample Dynamic Hedging Strategy Performance Metrics (2022–2024)

Strategy	Ann. Ret.	Sharpe	Max DD	Calmar	Sortino	Turnover/mo
Equal-Weight	4.12%	0.413	–18.74%	0.220	0.587	0.021
Mean-Variance	5.87%	0.542	–16.31%	0.360	0.724	0.158
Risk Parity	6.34%	0.623	–14.52%	0.437	0.812	0.094
DRL (no graph)	7.21%	0.714	–13.08%	0.551	0.943	0.187
DRL + T-GAT	8.47%	0.847	–10.63%	0.797	1.124	0.203

Note. Ann. Ret. = annualized return. Max DD = maximum drawdown. Calmar = Ann. Ret. / |Max DD|. Sortino uses 0% target. Transaction costs of 5 bps per unit turnover deducted. Initial capital $10M. Returns from Yahoo Finance.

Figure 3. Cumulative Risk-Adjusted Return Curves for Hedging Strategies (Out-of-Sample, 2022–2024)

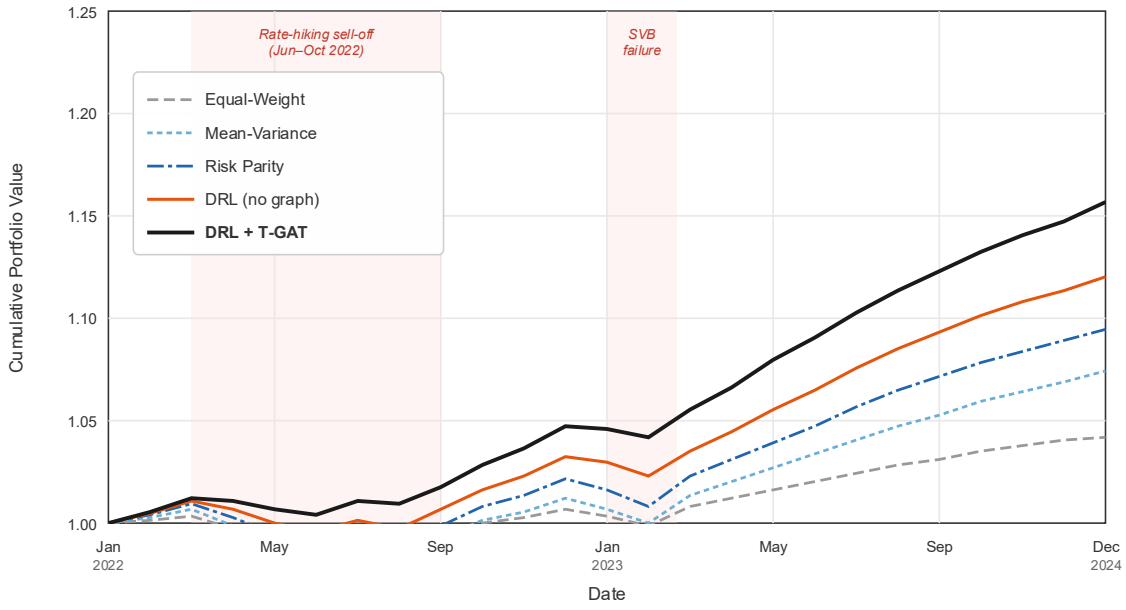

Cumulative portfolio values (normalized to 1.0 at Jan 2022) for five strategies. Shaded vertical bands mark the Jun–Oct 2022 rate-hiking sell-off and Mar 2023 SVB failure. The DRL + T-GAT curve exhibits shallower drawdowns during stress bands and steeper recoveries. Divergence between DRL + T-GAT and DRL (no graph) widens during stress transitions, consistent with the incremental value of contagion state information.

The Monte Carlo stress test evaluates portfolio robustness across 1,000 simulated crisis scenarios generated from the empirical contagion graph structure. Under a severe stress scenario—a simultaneous two-standard-deviation adverse shock to all asset illiquidity measures with contagion intensity calibrated to the March 2020 peak attention weights—the DRL + T-GAT agent's median drawdown is –12.8% with a 5th–95th percentile range of [–18.4%, –8.1%], compared to –17.3% [–24.6%, –12.2%] for risk parity. The narrower loss distribution suggests more consistent downside protection, though the improvement is incremental rather than transformative.

5. Conclusion and Discussion

5.1. Summary of Key Findings

This study presents a temporal graph attention network for identifying time-varying cross-asset liquidity contagion pathways and a deep reinforcement learning agent that leverages these pathways for dynamic hedging optimization. The empirical analysis, conducted on U.S. equity, fixed-income, and derivatives market data spanning 2007–2024 from exclusively public sources, yields three principal findings. The contagion topology is regime-dependent: credit-originated crises produce outward contagion from high-yield bonds, while liquidity-originated crises propagate from Treasuries and volatility instruments. The temporal graph attention mechanism provides moderate predictive gains over non-graph and static-graph baselines in detecting liquidity stress events, with F1-score improvements of 7.5 percentage points over LSTM and 5.6 percentage points over static GCN. The integration of graph-derived contagion states into the DRL hedging agent's observation space produces incremental improvements in risk-adjusted portfolio performance, particularly during stress transition periods, elevating the Sharpe ratio from 0.623 under risk parity to 0.847 under the contagion-informed strategy.

5.2. Practical Implications for Institutional Risk Management

The contagion pathway visualization generated by the T-GAT attention coefficients offers a practical early-warning tool for institutional portfolio managers. By monitoring the evolution of cross-asset attention weights on a weekly basis, risk teams can identify which contagion channels are intensifying before liquidity metrics breach critical levels, informing position sizing and hedging allocation decisions. During the March 2023 stress period, the attention coefficient from HYG to SPY rose from 0.183 to 0.251 approximately two weeks before the SVB failure, a signal that would have prompted preemptive hedging adjustments. For regulatory authorities, the methodology provides a data-driven complement to existing stress testing frameworks: the empirically identified contagion topologies can serve as scenario design inputs, grounding stress assumptions in observed inter-market behavior. The computational implementation relies on Python-based libraries (PyTorch Geometric, Stable Baselines3) and can be integrated with Tableau-based dashboards for visual monitoring.

References

[1]. Brunnermeier, M. K., & Pedersen, L. H. (2009). Market liquidity and funding liquidity. The Review of Financial Studies, 22(6), 2201–2238.

[2]. Amihud, Y. (2002). Illiquidity and stock returns: Cross-section and time-series effects. Journal of Financial Markets, 5(1), 31–56.

[3]. Diebold, F. X., & Yılmaz, K. (2014). On the network topology of variance decompositions: Measuring the connectedness of financial firms. Journal of Econometrics, 182(1), 119–134.

[4]. Acemoglu, D., Ozdaglar, A., & Tahbaz-Salehi, A. (2015). Systemic risk and stability in financial networks. American Economic Review, 105(2), 564–608.

[5]. Acharya, V. V., & Pedersen, L. H. (2005). Asset pricing with liquidity risk. Journal of Financial Economics, 77(2), 375–410.

[6]. Pástor, Ľ., & Stambaugh, R. F. (2003). Liquidity risk and expected stock returns. Journal of Political Economy, 111(3), 642–685.

[7]. Greenwood, R., Landier, A., & Thesmar, D. (2015). Vulnerable banks. Journal of Financial Economics, 115(3), 471–485.

[8]. Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. Proceedings of the 5th International Conference on Learning Representations (ICLR).

[9]. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., & Bengio, Y. (2018). Graph attention networks. Proceedings of the 6th International Conference on Learning Representations (ICLR).

[10]. Buehler, H., Gonon, L., Teichmann, J., & Wood, B. (2019). Deep hedging. Quantitative Finance, 19(8), 1271–1291.

[11]. Montalto, A., Faes, L., & Marinazzo, D. (2014). MuTE: A MATLAB toolbox to compare established and novel estimators of the multivariate transfer entropy. PLoS ONE, 9(10), e109462.

[12]. Fleming, M. J. (2003). Measuring Treasury market liquidity. Federal Reserve Bank of New York Economic Policy Review, 9(3), 83–99.

[13]. Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. The Review of Financial Studies, 33(5), 2223–2273.

[14]. Cheng, D., Niu, Z., Zhang, J., Zhang, Y., & Jiang, C. (2023). Critical firms prediction for stemming contagion risk in networked-loans through graph-based deep reinforcement learning. Proceedings of the AAAI Conference on Artificial Intelligence, 37(12), 14211–14219.

[15]. Caccioli, F., Ferrara, G., & Ramadiah, A. (2024). Modelling fire sale contagion across banks and non-banks. Journal of Financial Stability, 71, 101231.