Systemic Risk and Bank Networks: The Use of a Knowledge Graph with Generative Artificial Intelligence

Ren-Raw Chen*
Yilu Zhou
Gabelli School of Business
Fordham University
New York, NY 10019
rchen@fordham.edu, yzhou62@fordham.edu

Xiaohu Zhang*
Capital One
McLean, VA 22102
xiaohu.zhang@capitalone.com

Abstract

In this paper, we study the systemic risk and networks of top financial institutions using textual data (i.e., news). In particular, we draw knowledge graphs after the textual data are processed through various natural language processing and embedding methods, including the use of the most recent version of ChatGPT (via the OpenAI API). We also compare knowledge graphs drawn from the textual data with those from the numeral data as in Chen and Zhang [1]. We test a wide collection of models (i.e. knowledge) with both textual and numeral data for the networks of the top financial firms. Given that systemic risk is crucial in crisis times, we compare networks for the periods of the 2008 crisis (2007 bubble, 2008 bust, 2009 post-crisis). In particular, we focus on the troubled banks (bankrupt and bailed-out) and try to discover any early warning signs of these firms in terms of their networks (i.e. systemic risk). Although the models yield different knowledge graphs, the ensemble results consistently reveal a strong network of interconnections among the troubled firms and their closest counterparties.

1 Introduction

The 2008 global financial crisis underscored the critical importance of understanding interconnectedness and systemic risk within the financial system. Network models have become a primary tool for this analysis, but have traditionally relied on numerical data like stock returns or volatility. While insightful, these models may miss nuanced, timely information embedded in unstructured textual data such as financial news.

In recent work, Chen and Zhang [1] for the first time constructed knowledge graphs of top financial firms using numeral data (volatility and liquidity discount index) and Lyu et. al. ([2]) demonstrate how to draw similar graphs using news. This paper extends their work by providing a comprehensive study of how knowledge graph can explain the 2008 global financial crisis. We use both textual data (news) and numeral data (volatility and liquidity index) from 2007 to 2009 to examine the systemic risk of top financial firms. We improve upon Lyu et. al.([2]) who use only 2016 data and upon Chen and Zhang ([1]) who use only numeral data.

Our contribution is twofold. First, we construct and compare KGs from both numerical data (volatility, liquidity discounts) and textual data (news articles) for 27 top financial firms from 2007 to 2009. This period allows us to study the network structure during three distinct phases: the 2007 pre-crisis

[1]Corresponding author

Submitted to 39th Conference on Neural Information Processing Systems (NeurIPS 2025). Do not distribute.

bubble, the 2008 crisis, and the 2009 recovery. Second, we leverage modern NLP and large language model (LLM) technology, specifically OpenAI's text embedding models, to convert unstructured news into a feature space suitable for network construction. This allows us to investigate whether textual data contains early warning signals missed by numerical methods. We find that news-based KGs, particularly those using advanced embeddings, reveal a significant tightening of the financial network in 2008 and highlight specific clusters of troubled firms, demonstrating the value of incorporating generative AI into systemic risk monitoring.

1.1 Data

We use two distinct datasets covering the period from January 2007 to December 2009.In using data from 2007 to 2009, we can study three distinctly different phenomena: pre-crisis (bubble), crisis (bubble burst), and post-crisis (recovery). Our study is also in line with Acharya et. al. ([5]) who use daily returns of the same period to estimate the systemic risk of the top 18 financial firms

- **Numerical Data:** Monthly time series of firm volatility and liquidity discounts, consistent with the data used in Chen and Zhang [1].
- **Textual Data:** A corpus of 31,533 news articles mentioning the 27 target firms, collected from the LexisNexis database.

1.2 Knowledge Graph Construction

We employ a variety of methods to model the edges between firms, comparing traditional numerical techniques with modern textual analysis.

KGs from Numerical Data Following established methods in Chen and Zhang [1], we construct two KGs using numerical data. The edge weights are determined by the partial correlation between firms, calculated separately for **volatility** and **liquidity discount** data for each year. 22 firms are assumed to have 36-month observations as they survived during the global financial crisis and the troubled 5 firms do not have data in 2009 and certain period in 2008. To eliminate less significant dependencies, we plot the spring network by filtering the correlations greater than 5%. As an auxiliary measure, we apply the KNN algorithm to group the financial institutions into multiple clusters.

KGs from Textual Data To leverage the news data, we first convert the unstructured text into numerical representations (embeddings) that capture semantic meaning. We then use these embeddings to derive relationship strengths.

Text Embedding: Following Lyu et al. [2], we adopt OpenAI's `text-embedding-ada-002` model [6], which belong to the second generation of embedding models. We use OpenAI's `text-embedding-ada-002` model [6] to transform each news article into a fixed 1,536-dimensional vector (token limit 8,191). Most articles fall within this limit; when necessary, we truncate text to fit. We also include OpenAI's third-generation `text-embedding-3-large` model (3,072 dimensions) for comparison. These embeddings serve as inputs to our KG construction methods. The process is shown in Figure 1 as an example. These embeddings serve as the input for the following KG construction methods.

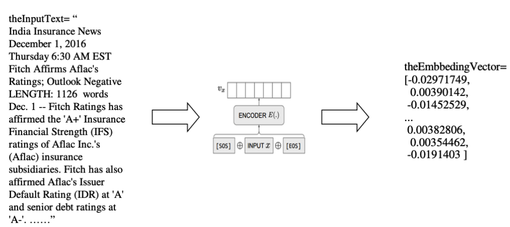

Figure 1: OpenAI's text-embedding-ada-002 and text-embedding-3-large models for the news data. from [2]

2 Empirical Results and Analysis

Our analysis of the KGs from 2007, 2008, and 2009 reveals distinct patterns in the financial network's evolution. While each model provides a different lens, a consistent narrative emerges.

2.1 Frequency Count Graph

As a simple baseline, we create a directed graph (Figure 2) where the edge weight from firm A to firm B is the number of times firm B is mentioned in news articles primarily about firm A. This method is intuitive, but does not capture deeper semantic relationships. These ring graphs show firm co-mentions in news, where vertex positions are arbitrary, and edges represent joint appearances.

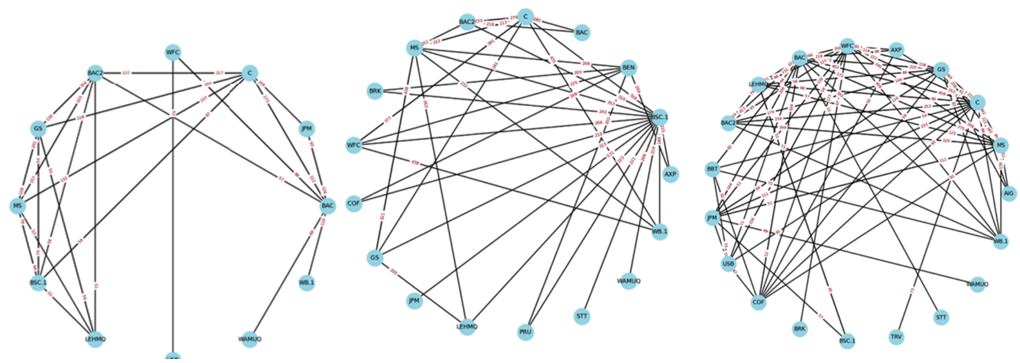

Figure 2: Knowledge graphs based on frequency of companies for 2007 (left), 2008 (middle), and 2009 (right). From 2007 to 2009, the news co-mention network evolved from sparse and uninformative connections before the crisis, to a sharp rise in 2008 with Bear Stearns and key firm pairs dominating, and then to a contracted post-crisis structure focused mainly on bailout-related firms.

2.2 Numerical Knowledge Graph

Partial correlation (Figure 3) is used to measure the relationship between two companies while controlling the influence of all other companies in the network. For the KG based on market implied volatility, the network resembles a dense "spider web," where most firms are tightly connected and individual influence is unclear in 2007. Centrality scores (Table 3) are better suited to rank importance. Troubled firms such as Lehman Brothers, Merrill Lynch, and Wachovia occupy peripheral positions, suggesting low systemic importance. In 2008, the network becomes denser, requiring a higher correlation threshold (0.10) for visualization; Bear Stearns and Washington Mutual drop out, while the remaining troubled firms split into two separate groups, indicating that network structure alone does not predict the actual sequence of defaults. By 2009, the network centers on Citigroup, CME, and Berkshire Hathaway. Overall, volatility-based knowledge graphs reveal limited predictive structure—only that firm interconnectedness peaks during the crisis, loosens before it, and weakens further in the post-crisis period. Using the liquidity-based KG, correlation levels drop significantly in 2008. Part of this is mechanical: there are only eight months of data for Lehman Brothers (which defaulted in September) and limited observations for Bear Stearns (which required bailout in March). More importantly, when liquidity discounts for all firms spike to extreme levels, cross-sectional variation collapses, and correlations become uninformative—for example, if all liquidity discounts are nearly identical, the implied correlations approach zero. Thus, precisely when firms are most closely linked through funding stress, correlation-based measures can misleadingly suggest weak connections, indicating that liquidity correlations are not a reliable basis for constructing knowledge graphs during crisis periods.

2.3 t-SNE Visualization

To enhance interpretability, we applied the dimensionality reduction technique known as t-SNE (t-distributed Stochastic Neighbor Embedding [3]) to reduce the high-dimensional average embedding

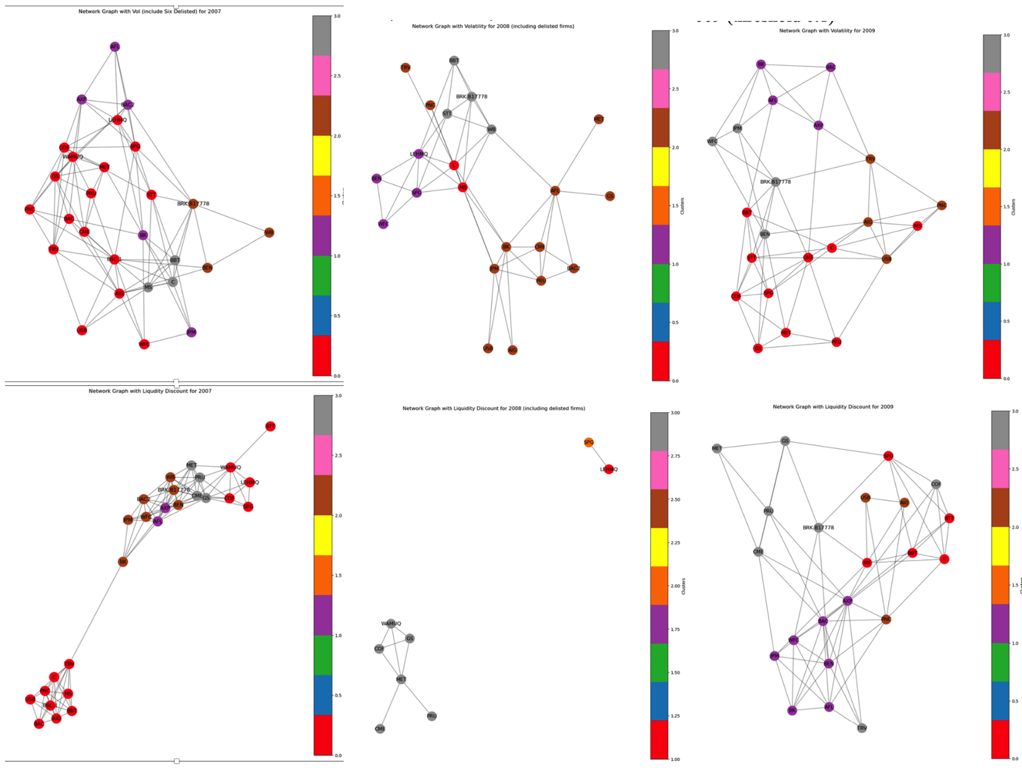

Figure 3: Knowledge graphs based on market implied volatility and liquidity discount for 2007 (left), 2008 (middle), and 2009 (right). Volatility-based predictions fail to capture crisis linkages, while the liquidity index reveals fragmented sub-networks in 2007 and shows in 2008 that most troubled firms vanish except Lehman Brothers, which remains connected to a few institutions.

of each firm's news corpus into a 2D space. In this KG (Figure 4), the Euclidean distance between two firms in the 2-D plot represents their semantic similarity. This method is purely for visualization and helps identify natural clusters.

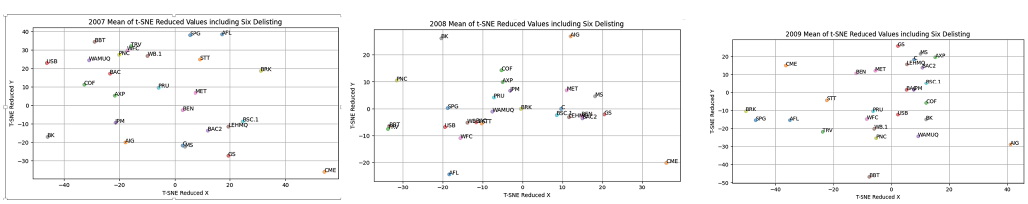

Figure 4: Knowledge graphs based on news dataset, the averages of t-SNE reduced (from 1,536 dimensions) 2-dimensional plots. Each dot is an average of all the news embeddings in a given year (2007, 2008, or 2009) in two dimensions.

2.4 RNN-Based KG

Our primary textual method uses a Recurrent Neural Network (RNN) to generate a more nuanced measure of firm relationships. We train an RNN classifier to predict the primary firm associated with a given news article's embedding. The "knowledge" for the graph is derived from the RNN's **confusion matrix**. When the model misclassifies an article about firm A as being about firm B, it implies a strong semantic similarity. The edge weight from A to B is defined by the frequency of this specific confusion, creating a directed graph where edges represent learned similarities. We numbered each firm from 1-27 (Table 6).

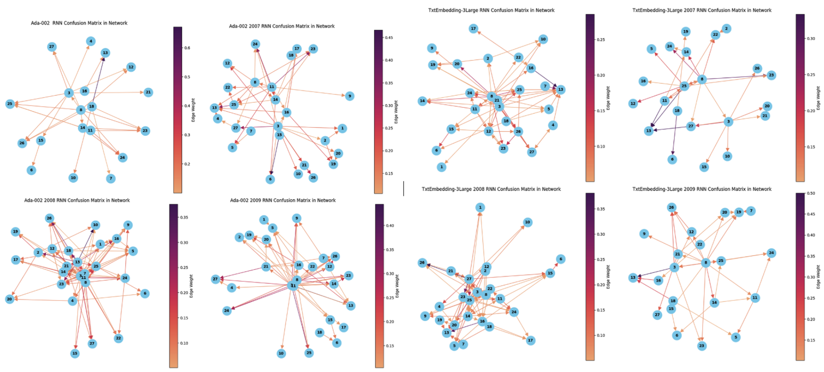

Figure 5: Knowledge graphs based on RNN

In 2007 (pre-crisis), the network in Figure 5 is sparse with weak inter-firm links, indicating low systemic risk. Troubled firms like Lehman, Bear Stearns, Wachovia, and AIG are mostly influenced by others, not vice versa. Central roles are held by Bank of America, Goldman Sachs, and Prudential. In the larger embedding model, Citigroup emerges as an early influential hub connecting several major firms. The 2008 network becomes much denser, reflecting heightened interconnectedness and systemic risk. Merrill Lynch moves to the network center, influencing several firms, while Lehman and AIG form weaker outward links. Distressed banks are more connected among themselves, and the overall structure loses hierarchy, consistent with severe stress and contagion during the crisis. By 2009, the network loosens again, resembling 2007 but with new central firms—Bank of America, Citigroup, Goldman Sachs, and Prudential. Troubled institutions shift to the periphery and lose influence. Inter-firm connections decline sharply, signaling a re-stabilized system with diminished systemic risk and reduced market contagion.

2.4.1 Comparison with Conventional Methods

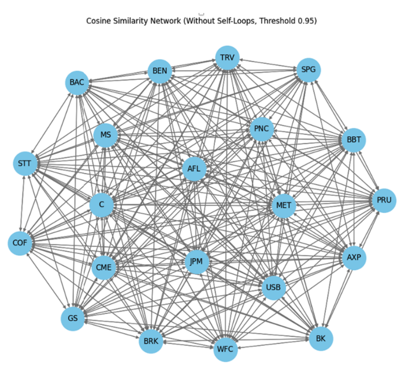

Figure 6: Cosine Similarity

While we employ neural networks (NN) for network linkage and knowledge graph construction, conventional similarity-based methods—such as cosine similarity, Euclidean distance, and clustering—can also be used as benchmarks. These baselines help highlight the advantages of the RNN-based approach in capturing dynamic, directional relationships between firms.

Figure 6 presents the network constructed using cosine similarity. The resulting graph appears densely connected, with nearly all firms linked to one another, making it difficult to distinguish key influencers. This reflects a limitation of cosine similarity: it behaves much like correlation, and because financial institutions are inherently correlated, the resulting network lacks discriminative power.

We find uniformly high cosine similarity values across embedding vectors, suggesting limited distinctiveness among nodes and hindering effective clustering. In contrast, the RNN-based confusion matrix captures sequential, context-dependent interactions between firms, producing a more interpretable and differentiated network that better reflects systemic interdependencies.

2.5 Centrality Scores

If a directed graph has cycles or bi-directional relationships, then causality does not apply. One is left with centrality measures that gauge which vertices are more important than others. Given that our knowledge graphs are not Directed Acyclic Graphs (DAGs), we calculate a series of metrics to demonstrate the impact of each firm – centrality scores. Section 6.1 provides the details of the centrality scores.

Table 3 summarizes firms' network centrality across 2007–2009, focusing on degree centrality as a key indicator. In 2007, Citigroup ranked highest (1.0000), followed by Goldman Sachs, Morgan Stanley, and Bank of America, indicating their dominant positions within the financial network. By contrast, most other firms had much lower connectivity, suggesting a more fragmented system before the global financial crisis. In 2008(Table 4), the number of firms in the network increased to 17, reflecting tighter interconnections during the crisis. Bear Stearns, Citigroup, and Morgan Stanley showed the highest centrality scores, while Lehman Brothers—despite its prominence—had a surprisingly low value. In 2009(Table 5), after major consolidations, only eight firms remained, with Citigroup, Bank of America, Morgan Stanley, Wells Fargo, and Goldman Sachs displaying the strongest centrality values, indicating that the system became more concentrated around a few large institutions. Overall, the analysis reveals that relationships among financial institutions intensified during the crisis years. Citigroup, Bank of America, Wells Fargo, and Morgan Stanley consistently occupied central roles, while Goldman Sachs stood out under closeness metrics. Insurance and retail banks, however, remained less connected, consistent with post-crisis lessons that investment banks—those most interconnected—were the most vulnerable to systemic shocks.

3 Key Observations

The most striking and consistent finding across nearly all models (numerical and textual) is a dramatic increase in network density and interconnectedness in 2008. Figure 5 illustrates this using our RNN-based KG, where the number and strength of connections visibly peak during the crisis year. This confirms the conventional wisdom that systemic risk intensifies during a crisis, but our method provides a quantifiable and visual representation of this phenomenon derived directly from news data. Numerical models also show this increased density, validating the signal across different data types. Text-based methods were particularly effective in identifying clusters of at-risk firms. The visualization of t-SNE for 2007 (Figure 4) places Bear Stearns, Lehman Brothers, and Merrill Lynch in a tight cluster, separate from other firms. This grouping, based on the semantic content of pre-crisis news, is a powerful, hindsight-free indicator of their shared vulnerabilities. In contrast, numerical models were less clear in isolating this specific group of troubled firms before the crisis. The RNN-based model further confirmed these connections, showing directed links between these institutions that intensified in 2008.

3.1 An Ensemble View

No single model fully captures systemic interconnections. For example, the liquidity-based KG shows a counterintuitive drop in correlations during 2008 as firm-level liquidity spreads converged, while the frequency-based KG produces noisier, less discriminative links. Integrating results across all approaches yields a consistent pattern: network density peaks in 2008, reflecting elevated systemic risk. Text-based models—especially t-SNE and RNN embeddings—more effectively isolate crisis-related clusters, whereas numerical models show weaker differentiation. This highlights the complementary nature of textual and quantitative perspectives.

Table 1 consolidates six modeling approaches— two numerical (volatility- and liquidity-based), three textual (frequency count, RNN with ada-002, RNN with text-embedding-3-large), and one based on network centrality. We evaluate them by four criteria: (1) whether the network tightens in 2008, (2) whether AIG's distress is captured, (3) which firms occupy central positions, and (4) which institutions are most closely linked to Bear Stearns and Lehman Brothers.

Table 1: Comparison of six models in detecting network structure and crisis dynamics.

Model	Denser 2008?	Predict AIG trouble?	Center (2007)	Center (2008)	Center (2009)	Closest to Lehman/Bear Stearns
t-SNE (Fig. 4)	yes	no	n.a.	n.a.	n.a.	GS, MS, C, MER
Freq. count (Fig. 2)	yes	no	BAC, GS, C	BSC	BAC, C	GS, MER, MS
OpenAI-ada-002 (Fig. 5)	yes	no	BAC, GS	BAC, GS, C	BAC, GS, C	GS, C
OpenAI-3large (Fig. 5)	yes	no	C, MER	BAC, C, AIG, BSC, MER	C	GS, C
Centrality (Sec. 6.1)	n.a.	n.a.	C > GS > BSC	BSC > C > MS	C > BAC > MS > GS	
Volatility (Fig. 3)	yes	no	GS, BSC	everyone	BRK, WFC, JPM	hard to tell
Liquidity (Fig. 3)	yes	no	GS	BRK, WB, MS	BRK	everyone

Table 1 shows a clear consensus that financial networks tightened sharply in 2008. Before the crisis, Goldman Sachs dominated across models, followed by Bank of America and Citigroup. During the crisis, Bear Stearns and Citigroup became central, with Morgan Stanley and Bank of America also prominent. By 2009, Citigroup emerged as the primary hub, consistent with its expanded post-crisis role in a consolidated system.

Across model types, textual embeddings provide the most distinctive insights. They uncover early-warning clusters, such as the Bear Stearns–Lehman–Merrill group, that numerical models cannot distinguish once all firms experience similar stress. This ensemble perspective demonstrates that combining textual and numerical signals produces a more comprehensive and interpretable view of systemic contagion.

4 Future Work

More implications can be drawn from the basic graphs we produce in this paper. However, due to the limitation of space, we shall put them off for future research. First and most obvious is to use the graphs to predict future crises. This is not an easy task as knowledge graphs are usually not meant for such a purpose. Yet, luckily, the recent developments of graphic neural networks (GNN) can help extend the existing graphs for node-level and edge-level predictions. While this is itself an intensive study, we provide a short roadmap to how it can be done to provide certain predictions. Secondly, one can use the existing graphs to study how numeral metrics such as volatility and liquidity can migrate from one node to the other. This is too huge an amount of work and to be included in this paper. Yet, we provide a short example (via Granger's causality test) between Lehman and AIG to demonstrate how one bank can influence the other. We should note that Granger's causality test belongs to the category of directed graphs (mentioned earlier) and it is not the only method to study mutual influences between two banks. Finally, one can stress the network with hypothetical stressed scenarios to examine how a crisis network should look like, as a precaution for a potential bank run. Knowledge graphs in this case would need to include macroeconomic variables.

4.1 Node-Level and Edge-Level Prediction

Based on our empirical findings, a natural next step is to use these graphs for node- and edge-level prediction, for example to anticipate which institutions may become systemically important. Graph neural networks (GNNs) provide a natural framework for such tasks. While our current analysis emphasizes interpretability via traditional graph metrics, future work could explore GNN-based models once richer labeled data become available.

4.2 Causal Relationships Between Financial Institutions

One implication can be drawn from the graphs is how economically an increase in one bank's risk can cause the risk of the other bank. For example, the graph shows Lehman's (#24) influence on AIG (#22) and CME (#9) in 2008, yet the economic impact is not properly reflected in the graph. For this purpose, we conduct, using Lehman and AIG as an example, Granger's causality test [4] between Lehman and AIG. Given that we can only apply this to numerical data, not textual data, we test on volatility and liquidity series. Hence, we test on volatility and liquidity. We test if Lehman's volatility and liquidity impacts AIG's volatility and liquidity or vice versa.

Panel (A) of Table 2 presents the volatility result and panel (B) presents the liquidity result. Columns in Table 2 are various tests and their p-values. The results suggest that AIG's past volatility has a Granger-causal effect on Lehman's volatility, particularly at shorter lags (1 to 7) and with some evidence extending to longer lags (up to 12). This implies that changes in AIG's volatility may predict changes in Lehman's volatility over both short-term and potentially extended periods. On the contrary, Granger causality from Lehman's liquidity to AIG's liquidity is observed at longer lags (primarily 11 and 12), suggesting a delayed predictive effect of Lehman's liquidity changes on AIG's liquidity. This could imply that changes in Lehman's liquidity take longer to impact AIG's liquidity, or there is a delayed response in AIG's liquidity to shifts in Lehman's liquidity.

Table 2: Granger Causality Test Results for Lehman and AIG. Values are test statistics and corresponding p-values for lags 1–12. Significant p-values (typically < 0.05) are bolded.

Lag	SSR F	p (F)	SSR χ^2	p (χ^2)	LR	p (LR)	Par. F	p (Par. F)
				(A) Volatility: Does AIG → LEH?				
1	**12.40**	**0.0006**	**12.66**	**0.0004**	**12.14**	**0.0005**	**12.40**	**0.0006**
2	**5.88**	**0.0035**	**12.18**	**0.0023**	**11.70**	**0.0029**	**5.88**	**0.0035**
3	**3.60**	**0.0153**	**11.33**	**0.0101**	**10.91**	**0.0122**	**3.60**	**0.0153**
4	**3.07**	**0.0186**	**13.09**	**0.0108**	**12.53**	**0.0138**	**3.07**	**0.0186**
5	**2.79**	**0.0197**	**15.11**	**0.0099**	**14.37**	**0.0134**	**2.79**	**0.0197**
6	**2.60**	**0.0207**	**17.14**	**0.0088**	**16.19**	**0.0128**	**2.60**	**0.0207**
7	**2.27**	**0.0325**	**17.79**	**0.0129**	**16.77**	**0.0190**	**2.27**	**0.0325**
8	1.82	0.0790	**16.57**	**0.0349**	15.67	0.0474	1.82	0.0790
9	1.71	0.0947	**17.76**	**0.0380**	16.72	0.0532	1.71	0.0947
10	1.47	0.1574	17.37	0.0666	16.37	0.0896	1.47	0.1574
11	1.41	0.1787	18.58	0.0691	17.43	0.0957	1.41	0.1787
12	1.66	0.0865	**24.30**	**0.0185**	**22.37**	**0.0336**	1.66	0.0865
				(B) Liquidity: Does LEH → AIG?				
1	0.28	0.5990	0.29	0.5934	0.28	0.5936	0.28	0.5990
2	0.63	0.5332	1.32	0.5168	1.31	0.5187	0.63	0.5332
3	0.46	0.7095	1.47	0.6886	1.46	0.6907	0.46	0.7095
4	0.43	0.7882	1.85	0.7626	1.84	0.7653	0.43	0.7882
5	0.57	0.7260	3.13	0.6806	3.08	0.6870	0.57	0.7260
6	0.56	0.7650	3.75	0.7098	3.69	0.7179	0.56	0.7650
7	0.71	0.6650	5.71	0.5741	5.57	0.5907	0.71	0.6650
8	1.68	0.1134	**15.81**	**0.0452**	14.80	0.0632	1.68	0.1134
9	1.58	0.1312	**17.18**	**0.0460**	15.98	0.0673	1.58	0.1312
10	1.34	0.2200	16.57	0.0844	15.44	0.1167	1.34	0.2200
11	**2.57**	**0.0073**	**35.68**	**0.0002**	**30.90**	**0.0011**	**2.57**	**0.0073**
12	**2.29**	**0.0143**	**35.62**	**0.0004**	**30.82**	**0.0021**	**2.29**	**0.0143**

The results suggest that AIG's past volatility has a Granger-causal effect on Lehman's volatility, particularly at shorter lags (1 to 7) and with some evidence extending to longer lags (up to 12). This implies that changes in AIG's volatility may predict changes in Lehman's volatility over both short-term and potentially extended periods. On the contrary, Granger causality from Lehman's liquidity to AIG's liquidity is observed at longer lags (primarily 11 and 12), suggesting a delayed predictive effect of Lehman's liquidity changes on AIG's liquidity. This could imply that changes in Lehman's liquidity take longer to impact AIG's liquidity, or there is a delayed response in AIG's liquidity to shifts in Lehman's liquidity.

4.3 Potential Bias

In this study, we acknowledge the broader discussion on potential biases in textual analysis, particularly how skewed data can lead to unfair or distorted outcomes. Our approach mitigates such concerns in two ways. First, we do not rely on ChatGPT-generated content; OpenAI's embedding models are used solely as fixed feature extractors. The actual analysis is performed using our own recurrent neural network (RNN), which is trained independently on historical news data rather than fine-tuned on post-crisis corpora. Second, we address temporal bias: since large language models like ChatGPT are trained on data up to 2023, applying them directly to study the 2008 financial crisis could introduce hindsight bias. To avoid this, our knowledge graphs are constructed exclusively from period-specific data, preserving both the temporal integrity and authenticity of the analysis.

5 Conclusion

This paper demonstrates the value of integrating textual data and generative AI in the analysis of financial systemic risk. By constructing knowledge graphs from news articles using OpenAI's advanced text embeddings, we capture a more nuanced and timely view of firm interconnectedness than is possible with numerical data alone. The results reveal a clear rise in network density during the 2008 crisis and, importantly, uncover clusters of vulnerable institutions identified from pre-crisis news, suggesting that textual embeddings may provide early warning indicators of financial instability.

A key limitation of our study lies in the potential for hindsight bias, as the embedding models were trained on post-crisis information. Future research could mitigate this issue by employing models trained exclusively on period-specific corpora. Moreover, while our current analysis is descriptive, the resulting graphical structures offer a foundation for predictive modeling. In particular, Graph Neural Networks (GNNs) could be applied to forecast emerging linkages or identify institutions likely to become systemically important. Such extensions represent a promising avenue for advancing generative-AI–based systemic risk monitoring and financial stability analysis.

6 Appendix

6.1 Centrality Scores

Table 3: Centrality scores for 2007 (threshold 50)

Node (vertex)	Degree	Betweenness	Closeness	EigenCentrality	PageRank
BAC	0.8182	0.3909	0.6471	0.2126	0.1448
JPM	0.2727	0.0000	0.4074	0.0509	0.0371
C	1.0000	0.3467	0.6471	0.4615	0.2299
WFC	0.2727	0.0727	0.4400	0.0509	0.0540
BAC2	0.8182	0.0745	0.4583	0.3243	0.1195
GS	0.9091	0.0670	0.5000	0.4462	0.1288
MS	0.8182	0.0215	0.5000	0.4462	0.1383
BSC.1	0.7273	0.0085	0.5000	0.4462	0.0794
LEHMQ	0.4545	0.0000	0.3548	0.1846	0.0306
USB	0.0909	0.0000	0.0000	0.0000	0.0125
WAMUQ	0.0909	0.0000	0.0000	0.0000	0.0125
WB.1	0.0909	0.0000	0.0000	0.0000	0.0125

Table 4: Centrality scores for 2008 (threshold 200)

Node (vertex)	Degree	Betweenness	Closeness	EigenCentrality	PageRank
AXP	0.1250	0.0000	0.0000	0.0000	0.0265
BSC.1	0.8750	0.0000	0.8789	0.6660	0.2544
BEN	0.3750	0.0000	0.4737	0.2763	0.0985
BAC	0.1250	0.0000	0.0000	0.0000	0.0265
C	0.6250	0.0590	0.3828	0.4559	0.1132
BAC2	0.3125	0.0076	0.2356	0.2101	0.0557
MS	0.5625	0.0465	0.3403	0.3896	0.0923
BRK	0.1250	0.0000	0.0000	0.0000	0.0265
WFC	0.3125	0.0056	0.2356	0.2101	0.0651
COF	0.1250	0.0000	0.0000	0.0000	0.0265
GS	0.3125	0.0063	0.2188	0.1795	0.0561
JPM	0.0625	0.0000	0.0000	0.0000	0.0265
LEHMQ	0.1875	0.0000	0.0000	0.0000	0.0265
PRU	0.1250	0.0000	0.0000	0.0000	0.0265
STT	0.0625	0.0000	0.0000	0.0000	0.0265
WAMUQ	0.0625	0.0000	0.0000	0.0000	0.0265
WB.1	0.2500	0.0000	0.0000	0.0000	0.0265

Table 5: Centrality scores for 2009 (threshold 200)

Node (vertex)	Degree	Betweenness	Closeness	EigenCentrality	PageRank
BAC	1.2857	0.2143	0.7000	0.3034	0.1600
GS	0.7143	0.0000	0.5833	0.4591	0.1469
MS	1.0000	0.0238	0.7778	0.5521	0.2234
BAC2	0.5714	0.0000	0.4375	0.0990	0.0480
WFC	0.8571	0.1190	0.6364	0.2790	0.0999
C	1.4286	0.2143	0.8750	0.5521	0.2843
JPM	0.2857	0.0000	0.0000	0.0000	0.0188
WB.1	0.1429	0.0000	0.0000	0.0000	0.0188

Note: The ticker for Merrill Lynch was originally MER. Since it was bailed out by Bank of America (BAC), in the dataset, it is labeled as BAC2. Similarly, the ticker for Bear Stearns was BSC (BSC.1 in dataset); for Lehman Brothers was LEH (LEHMQ in dataset); for Washington Mutual was WAMU (WAMUMQ in dataset); and for Wachovia was WB (WB.1 in dataset). Finally, BBT was originally the ticker for Branch Banking and Trust Company which was later merged with Sun Trust Bank and changed its name to Truist (TFC).

Table 6: Firm reference table.

Ticker	Name	Number
AFL	Aflac	1
AXP	American Express	2
BAC	Bank of America	3
BBT	BB & T (Truist)	4
BEN	Franklin Resources	5
BK	Bank of NY Mellon	6
BRK	Berkshire Hathaway	7
C	Citigroup	8
CME	CME Group	9
COF	Capital One	10
GS	Goldman Sachs	11
JPM	JP Morgan	12
MET	MetLife	13
MS	Morgan Stanley	14
PNC	PNC Bank	15
PRU	Preduntial	16
SPG	Simon Property	17
STT	State Street	18
TRV	Travelers Group	19
USB	US BankCorp	20
WFC	Wells Fargo	21
AIG	AIG	22
BSC.1 (BSC)	Bear Stearns	23
LEHMQ (LEH)	Lehman Brothers	24
BAC2 (MER)	Merrill Lynch	25
WAMUMQ (WAMU)	Washington Mutual	26
WB.1 (WB)	Wachovia	27

References

[1] Chen, R.-R., & Zhang, X. (2024) From liquidity risk to systemic risk: A use of knowledge graph. *Journal of Financial Stability* **70**:101195.

[2] Lyu, R.-Y., Chen, R.-R., Chung, S.-L., & Zhou, Y. (2024). Systemic risk and bank networks: A use of knowledge graph with ChatGPT. *FinTech*, **3**(2), 274–301.

[3] van der Maaten, L., & Hinton, G. (2008). Visualizing data using t-SNE. *Journal of Machine Learning Research*, **9**, 2579–2605.

[4] Granger, C. W. J. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, **37**(3), 424–438.

[5] Acharya, Viral V.; Lasse H. Pedersen, Thomas Philippon, & Matthew Richardson, (2017), "Measuring Systemic Risk," *Review of Financial Studies*, **30**(1), 2–47.

[6] OpenAI. (2022). New and improved embedding model. Retrieved December 15, 2022, from `https://openai.com/index/new-and-improved-embedding-model/`