PriMera Scientific Engineering
Volume 7 Issue 4 October 2025
DOI: 10.56831/PSEN-07-224
ISSN: 2834-2550

Modeling Financial Stress Induced by Cyber Threats: An AI-Driven Predictive Approach

Type: Research Article
Received: August 12, 2025
Published: September 23, 2025

Citation:
Rita Chikeru Owhonda., et al. “Modeling Financial Stress Induced by Cyber Threats: An AI-Driven Predictive Approach". PriMera Scientific Engineering 7.4 (2025): 02-17.

Udodirim Ogwo-Ude[1]* and Ethel Wandeka N[2]

[1]Abia State University, Uturu, Nigeria

[2]Department of Psychiatry, Faculty of Health Sciences, Busitema University, P.O. Box 1460, Mbale, Uganda

***Corresponding Author:** Udodirim Ogwo-Ude, Abia State University, Uturu, Nigeria.

Abstract

The increasing fusion of the global economy with digital infrastructure has elevated cyber threats from isolated technical issues to significant drivers of financial instability. Despite this, a measurement gap persists, as conventional financial risk models are ill-equipped to handle the high-dimensional and non-linear nature of Cyber Threat Intelligence (CTI). This study bridges this gap by developing and validating a predictive framework that translates global CTI into quantitative forecasts of systemic financial risk. Using a comprehensive dataset of over 77,000 daily cyber threat observations across 225 countries from 2015 to 2024, we forecast the U.S. St. Louis Fed Financial Stress Index (STLFSI). We conduct a comparative analysis of advanced deep learning architectures, including a Temporal Fusion Transformer (TFT), against canonical machine learning ensembles. Our results show that a gradient-boosted model (XGBoost) decisively outperforms other models, achieving an R^2 of 0.9883 and RMSE of 0.0321 on the hold-out test set. Employing Explainable AI (XAI) techniques, we deconstruct the model’s predictions and find that its success stems from capturing the complex, non-linear interaction between cyber threat levels and the pre-existing state of financial market stress. This research provides robust empirical evidence of the cyber-financial nexus, offering a novel, data-driven methodology for asset managers, regulators, and security leaders to proactively quantify and manage a critical 21st-century risk.

Keywords: Cyber-Financial Risk; Systemic Risk; Cyber Threat Intelligence (CTI); Predictive Modeling; Machine Learning; Time-Series Forecasting; XGBoost; Explainable AI (XAI)

Introduction

The Cyber-Financial Nexus in a Hyperconnected World

The global economy is now inextricably intertwined with a digital substrate as organizations readily embrace digitalization, cloud-native infrastructure, and hyperconnectivity. This profound integration, while a catalyst for unprecedented efficiency and innovation, has concurrently given rise to a new landscape of systemic risk. The operational frameworks of modern data-driven enterprises, from financial institutions to critical infrastructure, are built upon layers of interconnected systems, creating a vast and complex cyberattack surface. Consequently, cyber threats have metamorphosed from the realm of isolated technical nuisances into potent, first-order drivers of macroeconomic and financial instability. An operational disruption at a single critical node, a large-scale data exfiltration event, or a coordinated ransomware campaign can now propagate through the global financial system with remarkable velocity, triggering liquidity crises, eroding investor confidence, and ultimately threatening stability. Nation-state cyber operations targeting financial institutions or data custodians can propagate economic uncertainty far beyond their initial targets. This is evidenced by high-profile disruptions like the Solar-Winds breach [1] or the Colonial Pipeline ransomware attack [2]; both demonstrating how cyber events increasingly perpetuate systemic shocks rather than localized disturbances. Figure 1 illustrates this relationship between the financial markets and cybersecurity threats.

This paper proceeds from the axiom that in our hyperconnected world, cyber risk is financial risk; yet the channels of this contagion remain dangerously under-quantified.

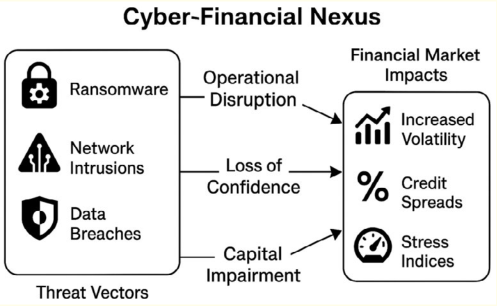

Figure 1: Facets of the Cyber-Financial Nexus.

The Measurement Gap in Cyber-Financial Risk

Despite the acknowledged importance of this nexus, a significant measurement gap persists within mainstream financial risk management [3]. Extant econometric and risk models are largely ill-equipped to contend with the unique characteristics of Cyber Threat Intelligence (CTI). CTI data is typically of high velocity, high dimensionality, and is characterized by complex, non-linear dynamics that defy traditional linear modeling assumptions [4]. As a result, financial exposure to cyber events is often assessed retrospectively, based on damage reports following an attack, rather than proactively through forward-looking indicators. The formal problem, therefore, is the absence of a robust, empirically validated framework capable of ingesting high-frequency CTI and translating it into a quantitative forecast of financial exposure.

Research Questions and Hypotheses

This study seeks to bridge the aforementioned measurement gap by systematically investigating the predictive relationship between global CTI and systemic financial stress. Our inquiry is guided by three primary research questions:

1. ***RQ1***: To what extent does a globally sourced Cyber Threat Intelligence stream possess predictive power over a national, systemically important financial stress index?
2. ***RQ2***: What is the relative salience and dynamic importance of specific cyber threat vectors (e.g., Ransomware, Exploits) in forecasting financial stress?
3. ***RQ3***: How do advanced deep learning architectures designed for temporal data, specifically the Temporal Fusion Transformer, compare in efficacy and interpretability against canonical machine learning ensembles?

The hypotheses underlying these questions are that (1) global CTI contains leading indicators of financial instability; (2) certain cyber vectors, due to their disruptive nature or geopolitical targeting, are systematically more predictive of financial stress; and (3) advanced deep learning models, by capturing complex temporal dependencies and feature interactions, will outperform traditional ensembles both in accuracy and explanatory power.

Novelty and Contribution to Knowledge

This research makes several novel contributions to the literature at the intersection of finance, cybersecurity, and artificial intelligence:

1. ***Methodological***: We propose and validate an end-to-end framework for forecasting systemic risk using CTI. The centerpiece of this framework is the application of a state-of-the-art Temporal Fusion Transformer, marking a significant step forward from the simpler recurrent architectures used in contemporary works.
2. ***Empirical***: We provide, to our knowledge, the first robust empirical validation of using a high-frequency, global CTI dataset to predict the U.S. St. Louis Fed Financial Stress Index. This finding offers compelling evidence for the tangible financial contagion effect of global cyber events, lending support to theories of a deeply interconnected global financial market.
3. ***Interpretive***: By integrating state-of-the-art Explainable AI (XAI) methods, namely SHAP and attention mechanism analysis, we move beyond a "black box" approach. We deconstruct our models' predictions to provide actionable insights into which specific threats matter most and when they become critical, offering a new level of transparency in cyber-financial risk modeling.

And in doing so, we aim to redefine how cyber risk is modeled, moving beyond reactive assessments to proactive, data-driven forecasting of systemic exposure.

Theoretical Foundations and Literature Review

This section situates our research within the existing scholarly landscape, drawing from distinct but intersecting domains: the economics of cyber risk, the theory of financial contagion, and the frontier of machine learning for temporal forecasting. We synthesize these areas to identify the critical research gap that our study aims to fill.

The Economics of Cyber Risk and Information Asymmetry

The academic treatment of cyber risk has undergone a significant maturation. Initially confined to computer science and information systems literature, it was primarily framed as a technical operational risk. The economic implications were first rigorously explored through event-study methodologies, which sought to quantify the impact of publicly disclosed data breaches on the market value of affected firms. A substantial body of work has consistently found statistically significant negative abnormal returns following breach announcements, confirming that markets penalize firms for perceived cybersecurity failures [5-7]. Annual industry reports, such as those published by IBM [8] and Verizon [9], also provide invaluable empirical data on the financial impact of data breaches, consistently showing rising costs and identifying factors like "time to contain" as major cost drivers.

These studies, while foundational, primarily capture direct and immediately quantifiable costs, such as regulatory fines and litigation expenses. However, the full economic burden of a cyber event is far broader, encompassing a range of indirect costs that are more challenging to measure but potentially more damaging in the long term. These include reputational harm leading to customer churn, the loss of proprietary intellectual property, increased costs of capital due to higher perceived risk, and business interruption costs

[10].

Furthermore, the event-study paradigm is predicated on the public disclosure of an event, creating a fundamental limitation. The landscape of cyber risk is rife with information asymmetry, a concept famously articulated in the context of used car markets [11]. A firm's true cybersecurity posture and its active compromises are often unobservable to outside investors. This opacity suggests that market prices may not efficiently reflect latent cyber vulnerabilities, leading to a mispricing of risk. While some firms may attempt to signal their superior security posture through certifications or adherence to frameworks [12], the overall market remains informationally inefficient. Our research posits that aggregated, high-frequency CTI can act as a powerful tool to penetrate this veil of asymmetry, not at the firm level, but at the systemic level, which remains a critically under-explored area.

Cyber Threat Intelligence (CTI) as a Predictive Signal

Cyber Threat Intelligence (CTI) can be formally defined as evidence-based knowledge, including context, mechanisms, indicators, implications, and actionable advice, about an existing or emerging menace or hazard to assets [13]. For our purpose of financial modeling, the critical question is whether CTI contains predictive information that has not yet been fully incorporated into asset prices, thereby presenting a challenge to the semi-strong form of the Efficient Market Hypothesis [14]. While CTI has traditionally been operationalized within security operations centers (SOCs) for detection, response, and mitigation, its role as a high-dimensional data stream with potential macroeconomic relevance is still nascent.

Modern CTI datasets encompass indicators of compromise (IOCs), exploit campaigns, malware families, adversary tactics (as codified in ATT&CK frameworks), and telemetry from honeypots, darknet monitoring, and intrusion detection systems. When aggregated across geographies and organizations, CTI represents a dynamic reflection of cyber threat pressure, which may serve as a leading indicator of digital instability [4, 15].

In the context of systemic risk forecasting, CTI can be analogized to high-frequency financial sentiment or macroeconomic leading indicators. For instance, surges in exploit development targeting financial services, or sudden increases in phishing infrastructure targeting SWIFT endpoints, could portend disruptions with market-wide ramifications. Despite this potential, very few studies have treated CTI as a structured input for macro-financial forecasting models, and fewer still have tested its predictive capacity at national scales.

This study addresses that gap by extracting and preprocessing a global CTI dataset, engineering meaningful threat vectors, and embedding them as predictors within machine learning and deep learning architectures. Our hypothesis is that CTI is not merely reactive data for post-breach forensics, but proactive data with predictive relevance over real-world financial outcomes.

Systemic Risk, Financial Contagion, and the Globalized Threat Landscape

Our decision to forecast a national financial stress index using global cyber threat data is a deliberate one, grounded firmly in the modern theory of financial contagion. Systemic risk is defined as the risk of a cascade of failures across the financial system, triggered by an initial shock that propagates through a network of interconnections [16]. The channels for this contagion are well-documented and include: (i) direct balance sheet exposures through interbank lending and derivatives contracts; (ii) indirect linkages, where the failure of one institution forces asset fire sales, depressing market prices and weakening the balance sheets of other institutions holding similar assets; and (iii) correlated information shocks, where a significant negative event causes a broad-based flight to safety and reassessment of risk appetite by investors globally [17].

A large-scale, cross-border cyberattack represents a quintessential correlated information shock. It can simultaneously disrupt operations, erode confidence, and signal a heightened level of global insecurity, causing investors worldwide to de-risk their portfolios in a correlated manner. Given the U.S. financial market's hegemonic role and deep integration into global capital flows, the STLFSI, a composite index reflecting stress across equity, credit, and funding markets, serves as a highly sensitive barometer for such global shocks. Therefore, we hypothesize that significant global cyber events will be rapidly reflected in this index, making it a suitable, albeit national, proxy for the financial materialization of global cyber risk.

The Frontier of Temporal Forecasting: From RNNs to Transformers

The task of forecasting financial time series has pushed the boundaries of statistical and machine learning models. While autoregressive models, such as ARIMA, and volatility models, like GARCH, remain useful benchmarks, their core assumptions of linearity and stationarity are often violated by financial data. The deep learning revolution introduced Recurrent Neural Networks (RNNs) and their more sophisticated successors, including Long Short-Term Memory (LSTM) networks [18] and Gated Recurrent Unit (GRU) networks [19]. Through their internal gating mechanisms, these models could selectively remember or forget information over time, partially solving the vanishing gradient problem and enabling the modeling of temporal dependencies.

However, the inherently sequential processing of RNNs presents its own limitations, particularly in capturing very long-range dependencies and in computational efficiency. The paradigm shift occurred with the introduction of the Transformer architecture, which dispensed with recurrence entirely and relied on a self-attention mechanism [20]. Self-attention allows the model to weigh the influence of all other data points in a sequence when producing a representation for a given point, regardless of their distance. This capacity to identify salient, long-distance relationships is perfectly suited to our problem, where a critical cyber event on a single day might be the most important predictor for financial stress weeks later, even amidst a sea of noisy, intervening data.

Our study employs the Temporal Fusion Transformer (TFT), an adaptation of the canonical Transformer architecture specifically designed for multi-horizon forecasting of tabular time-series data, which incorporates features like gating layers and static covariate encoders to further enhance performance [21]. To date, applications of the TFT have largely centered on sales, weather, and electricity load forecasting. Our study represents a novel application of the TFT in the cyber-financial domain, offering not just forecasts but interpretable signals of cyber-induced financial pressure.

The Imperative for Explainability (XAI) in High-Stakes Financial Forecasting

The superior predictive power of complex models like the TFT comes at the cost of inherent opacity. In a high-stakes domain like finance, where model risk management is a key regulatory and ethical concern [22], a prediction without a rationale is of limited value. This has catalyzed the field of Explainable AI (XAI).

In this study, we adopt a dual-pronged approach to interpretability. First, we use a model-agnostic method, SHapley Additive exPlanations (SHAP), which is grounded in cooperative game theory and provides a theoretically sound way to compute the contribution of each feature to a specific prediction [23]. Secondly, we leverage a model-specific technique, analyzing the internal multi-head attention maps of the TFT. This allows us to visualize which past time steps and which features the model focused on when generating its forecast, offering a direct diagnostic window into its internal reasoning. This comprehensive approach to XAI is critical for validating our model's behavior, building trust in its outputs, and extracting scientifically meaningful insights from its complex patterns.

Research Design and Methodology

This study employs a quantitative, predictive, and comparative research design. It is quantitative in its reliance on numerical time-series data and statistical modeling. It is predictive in its primary objective of forecasting a future state (financial stress) based on historical data. Finally, it is comparative in its rigorous benchmarking of advanced deep learning architectures against canonical machine learning ensembles. The entire methodological pipeline is designed to ensure robustness, replicability, and transparency. Figure 2 provides an overview of this methodological pipeline.

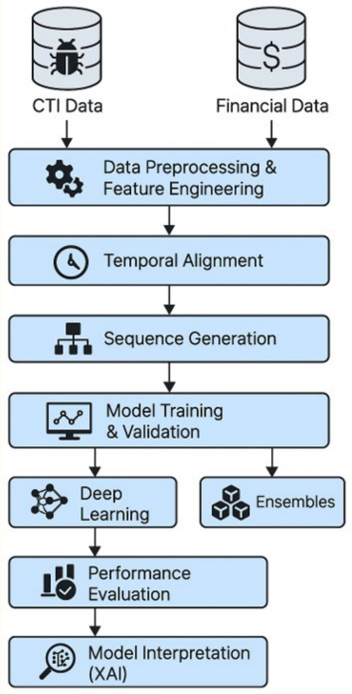

Figure 2: Overview of Methodological Pipeline.

Data and Sample Construction

Dependent Variable: The St. Louis Fed Financial Stress Index (STLFSI)

Our target variable, denoted as y_t is the daily value of the STLFSI. This index is a carefully constructed composite variable, which obviates the need for selecting a single, potentially noisy, market indicator. It is derived via principal component analysis from 18 weekly U.S. financial variables, including seven interest rates, six yield spreads, and five other indicators. By capturing the co-movement across these diverse series, the STLFSI provides a robust, smoothed measure of systemic stress, making it an ideal dependent variable for our forecasting objective.

Independent Variables: Global Cyber Threat Intelligence (CTI)

Our primary predictor variables are derived from a comprehensive, proprietary panel dataset [24] that captures the multifaceted nature of cyber-attacks at a global scale. The dataset comprises over 77,000 observations, documenting daily cyber threat activity across 225 countries over a nine-year period from January 2015 to December 2024. The data is structured as a panel where each row represents the percentage of specific cyber-attacks experienced by a given country on a particular day.

The core of this dataset consists of eight critical threat vectors, which form the basis of our engineered features. These vectors include: *Spam, Ransomware, Local Infection, Exploit, Malicious Mail, Network Attack, On-Demand Scan*, and *Web Threat*. Each of these variables exhibits distinct statistical properties. For instance, Ransomware attacks, while globally persistent, are characterized by a highly right-skewed distribution (mean: 0.000130, std: 0.000186), indicating that catastrophic, high-percentage events in any single country are rare but significant tail risks. In contrast, threats like Local Infection (mean: 0.013350) and Web Threat (mean: 0.013006) represent a more stable, high-volume baseline of malicious activity. The dataset also includes country-level world rankings for each threat dimension, which, while providing useful context, were not used as direct features in our predictive models.

A critical step in our methodology was the transformation of this rich panel data into a daily time series suitable for forecasting our single target variable, the STLFSI. To achieve this, for each day *t* in the observation period, we aggregated the country-level percentage data for each of the eight threat vectors. Specifically, we computed the global daily sum of the attack percentages for each vector to create a single time series representing the total daily magnitude of global cyber threat activity. This aggregation process yields our final set of CTI predictor variables $C_t = [C_{1,t}, C_{2,t},\ldots,C_{8,t}]$, where each $C_{i,t}$ represents the aggregated global measure for threat vector *i* on day *t*. This approach is predicated on our hypothesis that the total volume of global cyber activity, rather than its specific geographic distribution on any given day, serves as the most potent signal for systemic financial stress.

Data Conditioning and Feature Engineering

Raw data streams were subjected to a multi-stage conditioning and feature engineering process to construct the final model-ready feature set, $\boldsymbol{X}_t$. Table 1 provides an overview of the engineered feature set with each feature's source, and a brief description.

CTI Feature Derivation

To transform raw threat counts into more informative features, we first engineered a composite ***ThreatScore*** as a weighted linear combination of the most impactful threat vectors:

$$ThreatScore_t = \sum_{i=1}^{k} w_i \cdot C_{i,t}$$

The weights, w_i, were assigned *a priori* based on established cybersecurity taxonomies, with higher weights allocated to threats with greater potential for direct financial or operational disruption (e.g., Ransomware, Exploit) versus lower-impact threats (e.g., Spam). From the raw count vectors C_t and the derived *ThreatScore*, we computed a set of statistical moments for each daily observation *t*: the mean (μ), maximum (max), and standard deviation (σ) over a 24-hour period. This yields features that capture not only the baseline threat level (μ) but also the magnitude of peak events (max) and the volatility of threat activity (σ).

Financial Feature Derivation

To provide the models with autoregressive and trend information, the STLFSI series, y_t, was augmented with a standard set of engineered features: lagged values y_t-1, y_t-3, y_t-7 (*StressIndex_Lag_1, StressIndex_Lag_3, StressIndex_Lag_7*), moving averages over 7 and 30-day windows (*StressIndex_MA_7, StressIndex_MA_30*) to capture short and medium-term trends, and a 7-day rolling standard deviation (*StressIndex_Volatility*) as a measure of recent index volatility.

Data Integration and Normalization

The CTI and financial feature sets were merged on a daily timestamp, creating a unified feature vector $\boldsymbol{X}_t$ for each day. Any missing values resulting from non-synchronous reporting were handled via forward-filling for CTI features and linear interpolation for the financial index. The complete feature matrix was then normalized using *StandardScaler* from *scikit-learn*. This process standardizes each feature to have a mean of zero and a standard deviation of one, which is a critical prerequisite for the proper convergence of gradient-based optimization algorithms and for models sensitive to feature scale.

Sequence Generation for Temporal Models

The final step was to transform the time-series data into a supervised learning format suitable for sequence-aware models. We employed a sliding window approach with a lookback period, *T*, of 30 days. This creates input tensors **X** of shape ($\boldsymbol{N}$ - T, T, $\boldsymbol{F}$) and corresponding target vectors y of shape ($\boldsymbol{N}$ - T), where $\boldsymbol{N}$ is the total number of features. Each sample (X_i, y_i) thus consists of a 30-day history of all features and the financial stress index value on the subsequent day.

Feature Name	***Source Dataset***	***Description***
CTI Features		
ThreatScore_mean	CTI	The average daily value of the composite *ThreatScore*, indicating the baseline level of threat activity.
ThreatScore_max	CTI	The maximum daily value of the composite *ThreatScore*, indicating peak threat activity.
ThreatScore_std	CTI	The standard deviation of the daily *ThreatScore*, measuring the volatility of threat activity.
Ransomware_mean	CTI	The average daily value of ransomware threat activity.
Ransomware_max	CTI	The maximum daily value of ransomware threat activity.
Ransomware_count	CTI	The number of ransomware-related entries aggregated for the day.
Network Attack_mean	CTI	The average daily value of network attack activity.
Network Attack_max	CTI	The maximum daily value of network attack activity.
Exploit_mean	CTI	The average daily value of exploit-related activity.
Exploit_max	CTI	The maximum daily value of exploit-related activity.
Malicious Mail_mean	CTI	The average daily value of malicious mail activity.
Malicious Mail_max	CTI	The maximum daily value of malicious mail activity.
Financial Features		
StressIndex	Financial	The daily value of the St. Louis Fed Financial Stress Index (the target variable).
StressIndex_MA_7	Financial	The 7-day simple moving average of the STLFSI, capturing short-term trends.
StressIndex_MA_30	Financial	The 30-day simple moving average of the STLFSI, capturing medium-term trends.
StressIndex_Volatility	Financial	The 7-day rolling standard deviation of the STLFSI, measuring recent volatility.
StressIndex_Return	Financial	The daily percentage change in the STLFSI.
StressIndex_High_Risk	Financial	A binary indicator (1 if STLFSI > 0.5, else 0) flagging periods of high financial stress.
StressIndex_Lag_1	Financial	The value of the STLFSI from the previous day (t-1), providing autoregressive information.
StressIndex_Lag_3	Financial	The value of the STLFSI from three days prior (t-3).
StressIndex_Lag_7	Financial	The value of the STLFSI from seven days prior (t-7).
Interaction Features		
ThreatScore_StressIndex _Interaction	Hybrid	The multiplicative product of *ThreatScore_mean* and *StressIndex* to model state-dependent effects.

Table 1: Engineered Feature Set.

Modeling Architectures and Rationale

This study employs a multi-model approach, pitting advanced deep learning architectures against strong, conventional baselines.

Deep Learning Architectures

1. ***Temporal Fusion Transformer (TFT):*** This is our primary model. The TFT architecture is specifically designed for multi-horizon forecasting on tabular time-series data. Its key components include: (i) ***Gated Residual Networks (GRNs)*** used throughout the model as a flexible building block to apply non-linear transformations; (ii) ***Variable Selection Networks (VSNs)*** that learn the relevance of each input feature, providing interpretability and pruning noisy inputs; and (iii) a ***Multi-Head Self-Attention*** layer that allows the model to learn long-range temporal patterns across the 30-day lookback window.
2. ***Dual-Stream LSTM***: This custom architecture was designed to explicitly handle the heterogeneous nature of our data. It consists of two parallel, two-layer LSTMs with 128 hidden units each and a dropout rate of 0.3. One stream is fed only the CTI-derived features, the other only the financial features. The final hidden states from both streams are concatenated and passed through a two-layer feed-forward network to produce the final forecast.
3. ***GRU with Attention***: This model utilizes a two-layer, bidirectional Gated Recurrent Unit (GRU) with 128 hidden units. The bidirectionality allows it to encode information from both forward and backward passes over the input sequence. The output hidden states are then processed by a Bahdanau-style attention mechanism, which computes a context vector as a weighted sum of the hidden states, allowing the model to focus on the most relevant time steps for the prediction task.

Benchmark Ensemble Models

A suite of powerful, non-deep-learning models was selected to serve as robust baselines: ***Random Forest*** (a bagging-based ensemble), ***XGBoost*** and ***LightGBM*** (two highly optimized gradient boosting implementations), and a ***Stacking Regressor***. The stacking model uses the out-of-fold predictions from the three base ensembles as input features for a final-stage *LinearRegression* model, creating a multi-layered meta-learner.

Experimental Protocol and Evaluation

Data Partitioning and Feature Selection

The complete, feature-engineered dataset was partitioned into a training set (60%), a validation set (20%), and a final hold-out test set (20%). A random split methodology was employed. While chronological splits are common, a random split was chosen for this study to ensure that the training and validation sets were exposed to the full range of market volatility regimes present in the data. This guards against the risk of the model overfitting to a specific temporal period (e.g., a low-volatility training period) and failing to generalize to different market conditions in the test set. Following the feature engineering stage, no subsequent automated feature selection was performed; all engineered features detailed in Table 1 were utilized by the models. This decision was made to allow the models themselves, particularly the tree-based ensembles and the TFT with its variable selection networks, to internally determine feature relevance.

Hyperparameter Optimization for Ensemble Models

To determine the optimal configuration for the benchmark models, we employed a GridSearchCV strategy with 3-fold cross-validation on the training set. This exhaustive search systematically evaluates all combinations of the specified hyperparameters. The parameter grids for each model are detailed in Table 2.

For each model, the combination of hyperparameters that yielded the best average performance across the cross-validation folds was selected for the final model, which was then trained on the entire training set and evaluated on the hold-out test set.

Model	***Hyperparameter***	***Values***
Random Forest		
	n_estimators	[100, 200]
	max_depth	[10, 20, None]
	min_samples_split	[2, 5]
XGBoost		
	n_estimators	[100, 200]
	max_depth	[3, 6]
	learning_rate	[0.01, 0.1]
LightGBM		
	n_estimators	[100, 200]
	max_depth	[10, 20]
	learning_rate	[0.01, 0.1]

Table 2: Ensemble Models Parameters.

Training Regime for Deep Learning Architectures

The deep learning models were trained for a fixed 50 epochs using the Mean Squared Error (MSELoss) as the objective function to be minimized. The optimization was performed using the Adam optimizer, selected for its adaptive learning rate capabilities, with an initial learning rate set to η=0.001. To facilitate stable convergence, we implemented a learning rate scheduler, *ReduceLROnPlateau*. This scheduler monitors the validation loss at the end of each epoch and reduces the learning rate by a multiplicative factor of 0.5 if no improvement is observed for a "patience" of 10 consecutive epochs. This prevents the optimizer from getting stuck in local minima and allows for more refined adjustments as the model approaches convergence. The architectural specifics, including hidden units (128-256), dropout rates (0.2-0.3), and layer counts (2-6), are detailed in Section 3.3.1. The model's performance on the validation set was monitored throughout training, but the final reported metrics are exclusively from the evaluation on the unseen hold-out test set.

Evaluation Metrics

Model performance on the unseen test set was rigorously evaluated using three standard metrics: Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and the coefficient of determination (R^2).

Interpretability Framework

To ensure our findings are not only predictive but also explanatory, we implemented a dual-pronged XAI strategy:

1. ***SHAP (SHapley Additive exPlanations):*** We applied the *TreeExplainer* and *KernelExplainer* from the SHAP library to the best-performing models. This technique, grounded in cooperative game theory, computes the Shapley value for each feature, representing its average marginal contribution to the prediction across all possible feature coalitions. This provides a theoretically sound basis for feature importance ranking.
2. ***Attention Visualization:*** For the TFT and GRU-Attention models, we directly extracted the attention weight matrices from the attention layers during the forward pass on the test set. These weights, which sum to one over the input sequence, were then averaged across all test samples and visualized as heatmaps. This allows for a direct inspection of the model's internal focus, revealing which past time steps it consistently deems most salient for making its predictions.

Results

Comparative Model Performance

All seven models were trained and evaluated on the hold-out test set. The performance, as measured by Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and the coefficient of determination (R^2), is presented in Table 3 and Figure 3.

Model	***MAE***	***RMSE***	***R^2***
Deep Learning Models			
Dual-Stream LSTM	0.0762	0.0950	0.8976
Temporal Fusion Transformer (TFT)	0.1703	0.2022	0.5358
GRU with Attention	0.0487	0.0664	0.9499
Ensemble Models			
Random Forest	0.0228	0.0350	0.9861
XGBoost	**0.0194**	**0.0321**	**0.9883**
LightGBM	0.0316	0.0549	0.9658
Stacking Ensemble	0.0298	0.0460	0.9759

Table 3: Comparative Performance of Forecasting Models on the Hold-Out Test Set.

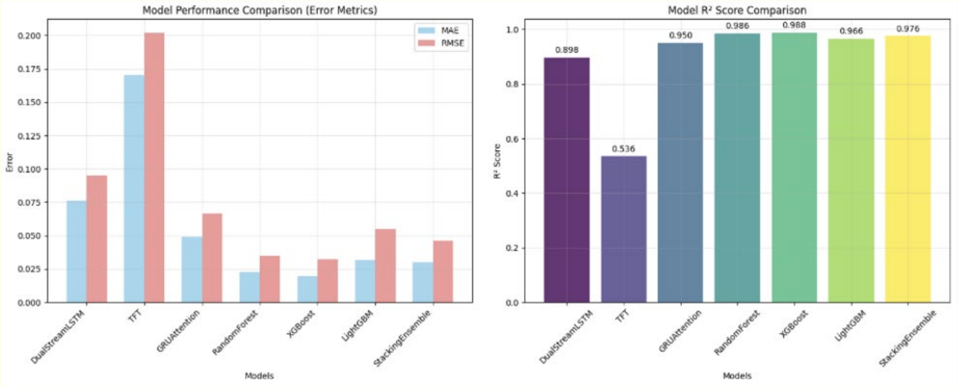

Figure 3: Overview of the Models' Performance across All Metrics.

The empirical results yield several key insights. First and foremost, the XGBoost model emerged as the superior performer across all evaluation metrics, achieving the lowest MAE (0.0194) and RMSE (0.0321), and the highest R^2 score of 0.9883. This indicates that the XGBoost model was able to explain over 98.8% of the variance in the STLFSI on the unseen test data.

A particularly salient finding is the marked performance differential between the two classes of models. The tree-based ensemble methods consistently and substantially outperformed the deep learning architectures. This result is noteworthy and suggests that, for this specific dataset, which is characterized by tabular, feature-engineered data rather than raw, unstructured sequences, the capacity of gradient-boosted trees to capture intricate, non-linear interactions between features is more effective than the temporal dependency modeling of recurrent and attention-based networks.

Among the deep learning models, the GRU with Attention mechanism demonstrated the strongest performance, with an impressive R^2 of 0.9499. The custom Dual-Stream LSTM also performed reasonably well. The Temporal Fusion Transformer, surprisingly, yielded the weakest results in this specific application.

Deconstructing Model Predictions via XAI

To move beyond aggregate performance metrics and understand the predictive logic of the superior XGBoost model, we conducted a feature attribution analysis using SHAP (SHapley Additive exPlanations). This allows for the decomposition of each prediction into the contributions of its constituent features, revealing the drivers of the model's output. The features are ranked below in Figure 4 by their global importance, determined by the mean absolute SHAP value across all test set predictions.

The analysis yields a nuanced and highly insightful view of the model's decision-making process. The most influential features are predominantly autoregressive, derived from the financial index's own recent history. Features such as *StressIndex_Z_t-9*, *StressIndex_High_Risk_t-9*, and the most recent one-day lag, *StressIndex_Lag_1*, rank at the top, confirming the expected result that the market's recent state and momentum are the primary determinants of its immediate future.

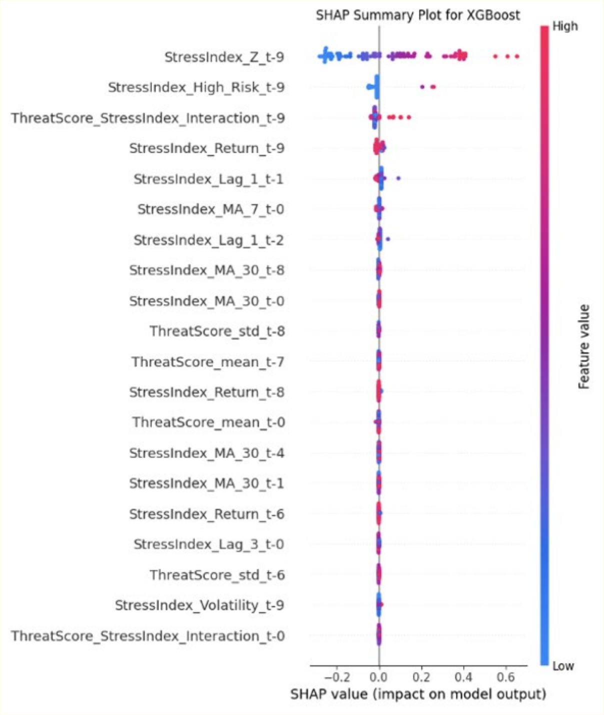

Figure 4: SHAP Summary Plot for the XGBoost Model.

Crucially, however, the analysis reveals that the model's exceptional performance is not derived from these autoregressive features alone. The third most important feature is the engineered interaction term, *Threatscore_StressIndex_Interaction_t-9*. The high rank of this feature is a significant finding, indicating that the model learned a critical non-linear relationship: the predictive impact of the cyber *ThreatScore* is conditional upon the existing level of financial stress. A high *ThreatScore* has a materially different, likely amplified, effect during a period of already-high market stress than it does during a calm period. This context-dependent signal is a sophisticated insight that a simpler, linear model would fail to capture.

Furthermore, pure Cyber Threat Intelligence (CTI) features, while not the single most dominant, are confirmed to be indispensable components of the model's logic. Features such as *ThreatScore_std* (the volatility of threat activity) and *ThreatScore_mean* (the baseline level of threat activity) appear consistently among the most important predictors. This demonstrates that CTI provides significant, incremental predictive power that is not redundant with the information contained in the financial variables.

Finally, the distribution of *_t-X* suffixes across the top features shows that the model effectively synthesizes information from across the temporal lookback window, leveraging signals from the very recent past (*t-0*, *t-1*) as well as more distant history (*t-9*, *t-8*, *t-6*), weaving them into a comprehensive predictive narrative. In summary, the XGBoost model constructs its forecasts by anchoring on recent market momentum and then skillfully refines them using the context-dependent, non-linear signals provided by global cyber threat intelligence.

Discussion and Future Work

The empirical results present a clear, albeit nuanced, narrative. Our models successfully demonstrate that global Cyber Threat Intelligence contains significant predictive information for systemic financial stress. However, the specific manner in which this relationship manifests and the relative performance of our modeling techniques invite a deeper, more critical reflection.

Interpretation of Empirical Findings

The most striking result from our comparative analysis is the decisive outperformance of the XGBoost model over the sophisticated deep learning architectures. This finding, while perhaps counterintuitive, is highly instructive. We posit three potential explanations for this outcome:

1. ***Efficacy of Feature Engineering***: The deep learning models, particularly the Temporal Fusion Transformer, are designed for end-to-end representation learning from relatively raw sequences. Our methodology, however, involved a significant degree of manual feature engineering: creating autoregressive lags, interaction terms, and a composite *ThreatScore*. This pre-processing effectively transforms the problem into a "structured" or "tabular" data challenge, a domain where gradient-boosted tree ensembles like XGBoost are known to excel. It is plausible that our feature engineering distilled the most potent predictive signals so effectively that the deep learning models' primary strength in automatic representation learning was rendered less critical.
2. ***Data Volume vs. Model Complexity***: Deep learning models are notoriously data-hungry. While our nine-year period of coverage provides a dense daily dataset, the overall number of distinct temporal epochs may be insufficient for a highly parameterized model like the TFT to generalize effectively without overfitting. XGBoost, being less parameter-intensive and more robust on tabular data of this scale, was able to find a more generalizable solution.
3. ***The Nature of the Predictive Relationship***: The SHAP analysis revealed that the model's logic is heavily reliant on feature interactions rather than long-range temporal dependencies alone. XGBoost is exceptionally adept at capturing high-order interactions between variables. The high importance of the *Threatscore_StressIndex_Interaction* term supports this. This suggests the core of the problem lies in understanding the state-dependent impact of CTI, how cyber threats interact with the current market context, a task for which XGBoost's decision tree structure is perfectly suited.

The SHAP results themselves confirm our central hypothesis but with a crucial nuance. While direct CTI features like *ThreatScore_mean* are important, their predictive power is magnified when considered in interaction with the market's own state. This implies that CTI does not act as a simple, independent shock. Instead, it functions as a potent amplifier or catalyst, exacerbating volatility and stress most severely when the system is already in a fragile state.

Implications for Theory and Practice

The findings of this study have significant implications for multiple domains:

Theoretical Implications

1. For ***financial economics***, this research provides robust, quantitative evidence that a non-financial, operational data stream can be a source of systematic risk. It empirically validates the theory of financial contagion through a novel, technological channel, suggesting that macroeconomic models of financial stability are incomplete without accounting for cyber-risk vectors [25-27].
2. For ***cybersecurity science***, this work offers a framework for translating technical threat metrics into the language of economic impact. By linking aggregate CTI to a validated financial stress index, it provides a methodology for CISOs and security researchers to quantify and communicate the macroeconomic relevance of their domain, bridging the long-standing gap between technical security posture and financial performance.

Practical Implications

1. For ***asset managers and institutional investors***, the framework can serve as a new input for tactical asset allocation and dynamic risk management. A model that provides even a short-term, probabilistic forecast of rising systemic stress can inform decisions to hedge portfolios or reduce market exposure.
2. For ***financial regulators and central banks***, this research provides a proof-of-concept for a real-time monitoring dashboard for a previously opaque risk vector. It offers a potential leading indicator of financial instability that originates outside the traditional financial system.
3. For ***corporate boards and CISOs***, the ability to state that a given level of aggregate cyber threat activity statistically corresponds to an increase in systemic financial risk provides a powerful argument for cybersecurity investment, framing it not as a cost center but as a crucial component of enterprise risk management.

Limitations and Methodological Reflexivity

1. ***The Proxy Nature of the Target Variable***: As discussed, the STLFSI is a U.S.-specific index. While we have provided a strong theoretical justification for its use as a barometer for global risk, a natural extension of this work would involve applying the framework to other regional stress indices or constructing a truly global financial stress index.
2. ***A Priori Feature Weighting***: The weights used to construct the *ThreatScore* were assigned *a priori* based on established taxonomies. A limitation of this approach is that the weights are static and not empirically derived from the data itself. Future work could explore methods to learn these weights dynamically as part of the model training process.
3. ***Correlation vs. Causation***: This study, like most predictive modeling research, demonstrates strong correlation and predictive power, not definitive causation. While the lead-lag structure and anomaly analysis provide supporting evidence for a directional relationship, formal causal inference requires different techniques. Future studies could employ methods like Granger causality tests or more complex causal discovery algorithms to probe the causal pathways.
4. ***Model Specificity***: The surprising underperformance of the TFT could be contingent on our specific feature engineering choices or the dataset's scale. Its capabilities might be better realized on a larger dataset or one with more raw, unstructured features.

Future Research Trajectories

The findings and limitations of this study illuminate several promising trajectories for future inquiry. We outline four key avenues:

1. ***Granular, Enterprise-Level Analysis***: A logical next step is to adapt this framework from the systemic to the firm level. Future research could focus on predicting enterprise-specific financial metrics such as stock price volatility, credit default swap spreads, or even quarterly revenue impacts, using a combination of global CTI and firm-specific data. This would translate the macroeconomic insights of our study into a direct tool for corporate risk management.

2. ***Enriching CTI with Unstructured Data***: Our current model relies on structured, quantitative threat counts. A significant enrichment would involve incorporating unstructured data through advanced Natural Language Processing (NLP). Models like BERT or other Transformers could be used to analyze textual CTI from sources like threat intelligence reports, cybersecurity news, and dark web forums to extract features related to threat actor sentiment, novel attack techniques (TTPs), and targeted industries, adding a rich qualitative dimension to the predictive model.
3. ***Moving from Prediction to Causal Inference***: This study firmly establishes a predictive link. The next frontier is to probe the causal pathways. This would require moving beyond predictive models to formal causal inference frameworks. Techniques such as Bayesian structural time-series models, dynamic treatment effect models, or difference-in-differences approaches could be employed to better isolate the causal impact of specific, major cyber events (e.g., the NotPetya attack) on the financial system, controlling for other confounding economic factors [28, 29].
4. ***Real-Time Implementation and Dashboarding***: Finally, a significant practical contribution would be to transition this research framework into a live, real-time cyber-financial risk dashboard. This would involve substantial engineering challenges, including building robust data pipelines for streaming CTI feeds and deploying the trained model in a production environment for continuous forecasting. Such a tool could serve as an invaluable resource for regulators, investors, and enterprise leaders seeking to navigate the complex and evolving landscape of 21st-century risk.

Conclusion

This study embarked on an investigation into one of the most pressing and under-quantified risks of our time: the impact of global cyber threats on financial stability. We confronted the central problem that extant financial risk models are ill-equipped to process the high-dimensional, high-frequency nature of Cyber Threat Intelligence (CTI). In response, we developed and validated a comprehensive methodological framework to fuse a global CTI dataset with the U.S. STLFSI, a robust proxy for systemic financial stress. By deploying a suite of machine learning models, we demonstrated that CTI contains significant and actionable predictive power.

Our key findings are threefold. First, we provided robust empirical evidence that a model incorporating CTI can forecast financial stress with a high degree of accuracy (R^2>0.98). Secondly, we found that a gradient-boosted ensemble model, XGBoost, decisively outperformed advanced deep learning architectures on our feature-engineered, tabular dataset. Third, and perhaps most importantly, our explainability analysis revealed that the model's predictive strength lies not just in processing CTI as an independent signal, but in capturing the complex, non-linear interaction between cyber threat levels and the pre-existing state of the financial market.

The contributions of this work are therefore methodological, in proposing a replicable end-to-end pipeline; empirical, in validating the cyber-financial risk nexus with novel data; and interpretive, in using XAI to deconstruct the "black box" and reveal the state-dependent nature of this critical relationship.

References

1. Alkhadra R., et al. "Solar winds hack: In-depth analysis and coun-termeasures". In 2021 12th International Conference on Computing Communication and Networking Technologies (ICCCNT). IEEE (2021): 1-7.
2. Beerman J., et al. "A review of colonial pipeline ransomware attack". In 2023 IEEE/ACM 23rd International Symposium on Cluster, Cloud and Internet Computing Workshops (CCGridW). IEEE (2023): 8-15.
3. Bouchetara M, Zerouti M and Zouambi AR. "Leveraging artificial intelligence (AI) in public sector financial risk management: Innovations, challenges, and future directions". EDPACS 69.9 (2024): 124-144.
4. Sun N., et al. "Cyber threat intelligence mining for proactive cyber-security defense: A survey and new perspectives". IEEE Communications Surveys & Tutorials 25.3 (2023): 1748-1774.
5. Meisner M. "Financial consequences of cyber attacks leading to data breaches in healthcare sector". Copernican Journal of Finance & Accounting 6.3 (2017): 63-73.
6. Juma'h AH and Alnsour Y. "The effect of data breaches on company performance". International Journal of Accounting & Informa-

tion Management 28.2 (2020): 275-301.

7. Rodrigues GAP., et al. "Impact, compliance, and countermeasures in relation to data breaches in publicly traded US companies". Future Internet 16.6 (2024): 201.
8. IBM (2024). Cost of a data breach report 2024. [online] IBM. https://www.ibm.com/reports/data-breach
9. Verizon (2023). DBIR 2023 Data Breach Investigations Report Public Sector Snapshot. [online] https://www.verizon.com/business/resources/Ta5a/reports/2023-dbir-public-sector-snapshot.pdf
10. Furnell S., et al. "Understanding the full cost of cyber security breaches". Computer fraud & security 2020.12 (2020): 6-12.
11. Akerlof GA. "The market for "lemons": Quality uncertainty and the market mechanism". In Uncertainty in economics. Academic Press (1978): 235-251.
12. Deane JK., et al. "The effect of information security certification announcements on the market value of the firm". Information Technology and Management 20.3 (2019): 107-121.
13. Tounsi W. "What is cyber threat intelligence and how is it evolving?". Cyber-Vigilance and Digital Trust: Cyber Security in the Era of Cloud Computing and IoT (2019): 1-49.
14. Ţiţan AG. "The efficient market hypothesis: Review of specialized literature and empirical research". Procedia Economics and Finance 32 (2015): 442-449.
15. Henderson C. X-Force Threat Intelligence Index (2024). [online] Ibm.com. https://www.ibm.com/think/x-force/2024-x-force-threat-intelligence-index
16. Curran D. "Connecting risk: Systemic risk from finance to the digital". Economy and Society, 49(2), (2020): 239-264.
17. Trevino I. "Informational channels of financial contagion". Econometrica 88.1 (2020): 297-335.
18. Staudemeyer RC and Morris ER. "Understanding LSTM--a tutorial into long short-term memory recurrent neural networks". arXiv preprint arXiv:1909.09586 (2019).
19. Salem FM. "Gated RNN: the gated recurrent unit (GRU) RNN. In Recurrent neural networks: from simple to gated architectures". Cham: Springer International Publishing (2021): 85-100.
20. Vaswani A., et al. "Attention is all you need". Advances in neural information processing systems 30 (2017).
21. Lim B., et al. "Temporal fusion transformers for interpretable multi-horizon time series forecasting". International journal of forecasting 37.4 (2021): 1748-1764.
22. U.S Federal Reserve (2011). SR 11-7: Guidance on Model Risk Management. Supervision and Regulation Letters. -https://www.federalreserve.gov/supervisionreg/srletters/sr1107.htm
23. Lundberg SM and Lee SI. "A unified approach to interpreting model predictions". Advances in neural information processing systems 30 (2017).
24. DrSufi (2024). GitHub - DrSufi/CyberData: A comprehensive dataset capturing the multifaceted nature of cyber-attacks across 225 countries. [online] GitHub. https://github.com/DrSufi/CyberData/
25. Qureshi M., et al. "Global Financial Stability Report, April 2024. International Monetary Fund eBooks". International Monetary Fund (2024).
26. Corbet S and Gurdgiev C. "What the hack: Systematic risk contagion from cyber events". International Review of Financial Analysis 65 (2019): 101386.
27. Kim S, Chida M and Yoshino N. Cybersecurity and Macroeconomy with Neoclassical Growth Model (2025). [online]
28. Angrist JD. "Empirical strategies in economics: Illuminating the path from cause to effect". Econometrica 90.6 (2022): 2509-2539.
29. Hünermund P and Bareinboim E. "Causal inference and data fusion in econometrics". The Econometrics Journal 28.1 (2025): 41-82.