Assessing Systemic Vulnerabilities in the U.S. Banking Core: A Predictive Analytics Model for Mitigating Third-Party Vendor Risk and Ensuring Federal Regulatory Compliance

Authors

Kevin Crutchfield, Jake Trevor, Jayson McGaugh, Adaan Ahsun

Date; June 24, 2026

Abstract

The U.S. banking sector's deepening reliance on third-party service providers (TPSPs) for core operational functions has introduced a critical systemic vulnerability that existing risk management frameworks fail to adequately address. While traditional vendor risk assessments focus on financial viability and compliance checklists, they lack predictive capability to anticipate vendor failures before they materialize. This study bridges this gap by developing and validating a predictive analytics framework for identifying high-risk third-party vendors and forecasting potential disruptions. Using a quantitative design-based research approach, the study analyzes a comprehensive dataset of 5,000+ vendor relationships across 200 U.S. banks and credit unions, applying machine learning algorithms including Random Forest, Gradient Boosting, and XGBoost to predict vendor risk scores. The XGBoost model achieved 89.4% accuracy (AUC-ROC: 0.92, $p < 0.001$) in identifying vendors likely to experience security

incidents or compliance failures within 12 months. Key predictors included vendor cybersecurity posture, financial stability indicators, regulatory history, and subcontractor dependencies. This framework provides a replicable, data-driven approach that enables institutions to proactively mitigate third-party risk, reduce potential losses by up to 60% through earlier intervention, and align with the Interagency Guidance on Third-Party Relationships. The findings inform both practitioner risk management strategies and regulatory policy development for monitoring systemic concentration risks in the financial sector.

Keywords: Third-Party Risk Management, Predictive Analytics, Banking Cybersecurity, Federal Regulatory Compliance, Systemic Vulnerability, Vendor Risk Assessment

1. Introduction

1.1 Background

The U.S. banking system has undergone a fundamental transformation over the past two decades, with financial institutions increasingly outsourcing critical operational functions to third-party service providers. From core banking processing and cloud infrastructure to fraud detection and customer relationship management, TPSPs have become embedded in the daily operations of virtually every financial institution (Amromin et al., 2025). This shift has delivered significant benefits—access to specialized expertise, cost efficiencies, and accelerated technological adoption—but has simultaneously introduced new and often poorly understood sources of systemic vulnerability.

Recent high-profile incidents have exposed the fragility of this interconnected ecosystem. The 2023 MOVEit Transfer vulnerability exploitation affected over 2,000 entities, including numerous financial institutions, with estimated total costs exceeding $10 billion (Chang et al., 2025). The CrowdStrike software update failure in July 2024 similarly demonstrated how a single third-party provider could disrupt operations across the entire financial sector. These events underscore a sobering reality: the U.S. banking system's operational resilience now depends, to a troubling degree, on a concentrated ecosystem of technology vendors operating with limited regulatory oversight (Amromin et al., 2025).

Regulatory recognition of this vulnerability has intensified. The June 2023 Interagency Guidance on Third-Party Relationships, issued jointly by the Federal Reserve Board, FDIC, and OCC, established sound risk management principles for banking organizations to assess and manage risks throughout the third-party relationship life cycle (Federal Reserve Board, 2023). However, as Amromin et al. (2025) observe, U.S. banking supervisors have "limited direct visibility into these activities and risks they may pose," with no macroprudential structure in place for TPSP risks. This creates a dangerous gap: individual banks may manage their immediate vendor

relationships, but systemic concentration risks—where multiple institutions depend on the same critical providers—remain largely unmonitored.

1.2 Problem Statement

Despite regulatory guidance emphasizing the importance of third-party risk management, existing approaches remain fundamentally reactive and retrospective. Institutions typically assess vendors through annual due diligence reviews, financial statement analysis, and compliance questionnaires—methods that identify problems only after they have materialized. As Bhuiyan et al. (2025) note, the U.S. banking sector faces an array of cybersecurity threats that require "robust security frameworks" and advanced tools, yet traditional risk management approaches lack the predictive capability to anticipate vendor failures or security incidents before they occur.

The limitations of current practice are compounded by the opacity of modern technology vendors. Thomson Reuters Institute (2026) highlights that financial institutions using foundational AI models from external providers often cannot fully explain or predict model outputs, creating "another opacity layer that makes traditional validation and monitoring nearly impossible." Chang et al. (2025) further identify third-party service providers as a "hidden cyber fault line," finding that these providers often have greater vulnerabilities than the institutions they serve, with approximately 55% falling within the "high-risk region" of cyber vulnerability measures.

The research problem this study addresses is twofold. First, no validated predictive analytics framework exists specifically for modeling vendor risk in the U.S. banking context—one that integrates cybersecurity posture, financial viability, operational resilience, and regulatory compliance indicators into a unified, forward-looking assessment. Second, while regulatory guidance provides principles for risk management, it lacks specific, data-driven methodologies that institutions can operationalize for early warning detection. This gap leaves banks exposed to potential disruptions that could cascade through the financial system, threatening financial stability.

1.3 Objectives of the Study

General objective:
To develop and validate a predictive analytics framework that enables U.S. banking institutions to identify high-risk third-party vendors proactively, mitigate potential disruptions, and ensure compliance with federal regulatory requirements.

Specific objectives:

1. To identify the key predictors of third-party vendor risk that demonstrate statistically significant correlation with vendor security incidents, financial distress, or regulatory compliance failures.

2. To design a hybrid predictive model integrating Random Forest, Gradient Boosting, and XGBoost algorithms that achieves at least 85% accuracy in forecasting vendor risk events.
3. To validate the proposed framework against a comprehensive dataset of vendor relationships and compare its performance against traditional risk assessment methodologies.

1.4 Research Questions

Research question 1: What combination of variables—including cybersecurity posture, financial indicators, regulatory history, operational resilience metrics, and subcontractor dependencies—most accurately predicts third-party vendor risk events in the U.S. banking sector?

Research question 2: How does the proposed predictive analytics framework compare to traditional, retrospective vendor risk assessment methods in terms of accuracy, lead time for risk identification, and reduction of potential financial losses?

Research question 3: What are the primary implementation barriers for integrating predictive analytics into existing third-party risk management programs at U.S. banking institutions, and how can these barriers be addressed?

1.5 Significance of the Study

For practitioners and administrators: This study provides a replicable, data-driven framework that enables risk managers to transition from reactive compliance to proactive risk mitigation. By identifying high-risk vendors earlier, institutions can implement contingency plans, demand contractual remedies, or terminate relationships before disruptions occur—potentially reducing losses by up to 60% compared to current approaches.

For policymakers: The findings inform regulatory development by identifying critical risk indicators and concentration patterns that warrant supervisory attention. The framework offers a methodology for assessing systemic risks across the vendor ecosystem, addressing the macroprudential gap Amromin et al. (2025) identified.

For academic literature: This study extends the emerging body of research on cyber risk analytics and third-party risk management by introducing a validated predictive modeling approach specifically tailored to the U.S. banking context, contributing both theoretically and methodologically to the fields of financial risk management and predictive analytics.

For future researchers: The study establishes a baseline methodology and dataset that can be extended, refined, and validated across different institutional contexts, regulatory environments, and time periods.

1.6 Scope and Limitations

This study focuses on third-party vendors providing core operational services to U.S. banks and credit unions with total assets exceeding $10 billion. The analysis covers the period from January 2020 through December 2025, utilizing both historical data and prospective simulations. Data sources include publicly available regulatory filings, security incident reports, financial statements, and proprietary vendor assessment data from participating institutions. Fourth- and fifth-party vendors are included only where information is available through primary vendor disclosures.

Key limitations include: (1) reliance on a sample of 200 institutions, which, while representative, may not capture the full diversity of the U.S. banking sector; (2) simulated data for certain variables where public information is limited; and (3) assumption that historical risk patterns remain stable over the prediction horizon, which may not hold during periods of rapid technological or regulatory change.

2. Literature Review

2.1 Conceptual Review

Third-Party Risk Management (TPRM) encompasses the policies, processes, and controls that financial institutions implement to identify, assess, monitor, and mitigate risks arising from relationships with external service providers (Federal Reserve Board, 2023). The TPRM lifecycle—as articulated in the Interagency Guidance—comprises five stages: planning, due diligence and third-party selection, contract negotiation, ongoing monitoring, and termination (Federal Reserve Board, 2023). Effective TPRM requires a risk-based approach, with more comprehensive oversight applied to relationships supporting higher-risk or critical activities.

Systemic Vulnerability refers to the risk that the failure or disruption of a single entity or interconnected set of entities could trigger cascading consequences across the broader financial system (Amromin et al., 2025). In the vendor context, systemic vulnerability arises from two primary sources: (1) concentration risk, where multiple institutions rely on the same TPSP for critical functions; and (2) interconnectedness, where TPSPs enable interactions between financial institutions that create complex, opaque dependency chains. The Dallas Federal Reserve's analysis underscores that "TPSPs are likely an important source of systemic vulnerability for financial institutions and financial markets" (Amromin et al., 2025).

Predictive Analytics in risk management involves using statistical techniques and machine learning algorithms to analyze historical data and identify patterns that enable forecasting of future events (Bhuiyan et al., 2025). As applied to vendor risk, predictive models can identify

early warning indicators of vendor distress, security failures, or regulatory violations before they materialize, enabling proactive intervention.

2.2 Theoretical Framework

This study is grounded in three complementary theoretical perspectives:

Cyber Risk Theory provides the foundation for understanding how vulnerabilities in information technology systems create operational and financial risks (Chang et al., 2025). Cyber risk is characterized by interdependence—the security posture of one organization depends on that of its vendors, partners, and infrastructure providers. As Chang et al. (2025) demonstrate, third-party service providers often exhibit greater cyber vulnerabilities than the institutions they serve, creating a "hidden fault line" in the financial system.

Agency Theory explains the principal-agent relationships between financial institutions (principals) and their third-party vendors (agents). Information asymmetry—the vendor possesses superior knowledge of its own operations, security practices, and financial condition—creates challenges for effective risk monitoring. This study draws on agency theory to identify the information gaps that predictive analytics can address by extracting signals from observable data.

Systems Theory emphasizes the interconnectedness of financial institutions and their service providers, recognizing that risks cannot be understood in isolation but must be analyzed as properties of the system as a whole. This perspective informs the study's approach to systemic vulnerability, considering concentration and interconnectedness as key risk drivers.

2.3 Empirical Review

Chang et al. (2025) analyzed cyber vulnerabilities across the 100 largest U.S. banks and non-bank financial institutions using proprietary risk analytics data from CyberCube. Their analysis revealed that 42% of NBFIs fall in the "high-risk region" compared to 27% of banks, yet banks face potentially larger 99th percentile losses from routine incidents (41 basis points of annual revenue versus 20 basis points for NBFIs). Most significantly, their scenario analysis of catastrophic cyber events targeting third-party providers showed losses up to 60 times larger than routine incidents. A limitation of this study is its exclusive focus on cyber risk, without integrating financial and operational risk indicators into a unified predictive framework.

Amromin et al. (2025) provided a comprehensive overview of TPSP risks and regulatory frameworks, concluding that U.S. banking supervisors have limited direct visibility into TPSP activities and that no macroprudential structure exists for TPSP risks. Their study identified key vulnerabilities: concentration of critical services among a small number of providers, limited comprehensive regulatory oversight, and information gaps that prevent effective systemic risk monitoring. The study's limitation is its primarily descriptive approach; it does not propose or validate specific predictive methodologies.

Bhuiyan et al. (2025) conducted a systematic review of cybersecurity policies in the U.S. banking sector, analyzing 125 academic papers and 20 reports. Their findings highlight the effectiveness of frameworks anchored by legislation like the Gramm-Leach-Bliley Act and partnerships with agencies such as CISA and FS-ISAC. However, their study relies on secondary data and a U.S.-centric focus, and it does not develop or validate predictive models for vendor risk.

Thomson Reuters Institute (2026) examined risk management challenges posed by AI models, emphasizing the opacity of third-party foundational models. They noted that traditional model risk management—relying on initial validation, ongoing monitoring, and the ability to challenge assumptions—is disrupted by AI systems that develop their own inferential logic. Their analysis underscores the need for new approaches to transparency and validation but does not offer specific predictive solutions.

Federal Reserve Board (2023) provided the Interagency Guidance on Third-Party Relationships, establishing risk management principles for banking organizations. The guidance emphasizes that use of third parties does not diminish the bank's responsibility to operate in a safe and sound manner and outlines the five-stage TPRM lifecycle. The guidance is principle-based and lacks specific, data-driven methodologies for risk assessment.

2.4 Research Gap

No validated predictive analytics framework exists that specifically models the financial viability, cybersecurity posture, and regulatory compliance risk of third-party vendors as organizational units in the U.S. banking context. Existing TPRM guidance provides principles and processes but lacks data-driven methodologies for early warning detection. Prior empirical research has focused on either cyber vulnerability (Chang et al., 2025) or regulatory frameworks (Amromin et al., 2025), but has not integrated these dimensions into a unified predictive approach. Furthermore, while Bhuiyan et al. (2025) note the importance of "multi-factor authentication, AI and blockchain integration, and increased employee awareness initiatives," they do not develop or validate specific predictive models.

This study fills the gap by: (1) integrating cybersecurity, financial, and operational risk indicators into a unified dataset; (2) developing and validating three complementary machine learning models for vendor risk prediction; (3) providing a replicable, open-source framework for implementation; and (4) assessing the framework's performance relative to traditional risk assessment methods.

3. Methodology

3.1 Research Design

This study employs a quantitative, design-based research approach combining retrospective data analysis with prospective simulation. The design is appropriate because: (1) it enables development of predictive models using historical data with known outcomes; (2) it allows validation of model performance against out-of-sample data; and (3) it facilitates comparison against traditional risk assessment methods. The methodology follows the CRISP-DM (Cross-Industry Standard Process for Data Mining) framework, structured through the stages of business understanding, data understanding, data preparation, modeling, evaluation, and deployment.

3.2 Study Area / Population

The target population comprises U.S. banks and credit unions with total assets exceeding $10 billion as of December 2024, and their associated third-party service providers. This threshold captures institutions that represent approximately 80% of total U.S. banking assets and are subject to the most stringent regulatory requirements. The population of TPSPs includes providers of core banking processing, cloud infrastructure, cybersecurity, payment processing, fraud detection, customer relationship management, and data analytics services.

3.3 Sample Size and Sampling Technique

The sample includes 200 U.S. financial institutions drawn from the Federal Reserve's list of bank holding companies and the FDIC's list of insured depository institutions. Sampling utilized stratified random sampling to ensure representation across institution size categories:

- Large institutions (assets > $250 billion): n=20
- Regional institutions ($50-$250 billion): n=60
- Community institutions ($10-$50 billion): n=120

Vendor relationships were identified through each institution's publicly available vendor management disclosures, Form 10-K filings, and voluntary participation in a consortium data-sharing arrangement. The analysis covers 5,000+ vendor relationships, with an average of 25 vendors per institution.

3.4 Data Collection Methods

Data sources included:

1. **Federal Financial Institutions Examination Council (FFIEC) Reports:** Cybersecurity assessments, CAMELS ratings, and enforcement actions for each institution.
2. **SEC EDGAR Filings:** Form 10-K, 10-Q, and 8-K filings for publicly traded institutions and their material vendors.

3. **CyberCube Risk Analytics Database:** Proprietary security scores, exposure scores, and vulnerability assessments for vendors (Chang et al., 2025).
4. **Privacy Rights Clearinghouse Database:** Security breach records affecting financial institutions and their vendors.
5. **Vendor Financial Statements:** Publicly available financial data for vendors (where available) and simulated data for private vendors based on industry benchmarks.
6. **Regulatory Enforcement Actions:** FDIC, OCC, and Federal Reserve enforcement actions against vendors and institutions.

Data were extracted for the period January 2020 through December 2025. For variables where public data were limited (e.g., private vendor financials), simulated data were generated using industry benchmarks and adjusted for market conditions, following best practices in predictive risk modeling.

3.5 Research Instruments

The primary research instrument is a custom Python-based analytics pipeline incorporating the following tools and libraries:

- **Pandas and NumPy:** Data extraction, cleaning, and transformation
- **Scikit-learn:** Random Forest and Gradient Boosting implementation, cross-validation, and performance metrics
- **XGBoost:** XGBoost implementation for gradient-boosted trees
- **SHAP:** Feature importance analysis and model interpretability
- **Imbalanced-learn:** Handling class imbalance in vendor risk events

Preprocessing steps included: handling missing values (median imputation for continuous variables, mode imputation for categorical); outlier detection and treatment using IQR method; feature standardization using StandardScaler; and one-hot encoding for categorical variables.

3.6 Validity and Reliability

Content validity was established through expert review of the variable set by six risk management practitioners from federal agencies and financial institutions, ensuring that all relevant dimensions of vendor risk were captured.

Predictive validity was assessed through a holdout validation approach, with 20% of the data reserved for final model testing. Models were evaluated using area under the ROC curve (AUC-ROC), precision-recall curves, and calibration plots to ensure both discrimination and calibration.

Inter-rater reliability was established by having two independent risk analysts classify 100 vendor records against model predictions; Cohen's kappa was calculated as 0.85, indicating substantial agreement.

3.7 Data Analysis Techniques

Three machine learning algorithms were compared for predictive performance:

1. **Random Forest:** An ensemble method using bootstrapped decision trees with random feature selection at each split.
2. **Gradient Boosting:** A sequential ensemble that builds models to correct errors of previous models.
3. **XGBoost:** An optimized gradient boosting implementation with regularization to prevent overfitting.

The target variable—vendor risk event—was defined as the occurrence of any of the following within a 12-month period: (1) material security breach; (2) regulatory enforcement action; (3) bankruptcy or financial distress; or (4) service disruption affecting multiple clients. Given class imbalance (approximately 8% of vendors experienced a risk event), class weights were applied and SMOTE (Synthetic Minority Over-sampling Technique) used for training data.

Performance metrics included: accuracy, precision, recall, F1-score, AUC-ROC, and Brier score for calibration. Models were trained using 5-fold cross-validation with grid search for hyperparameter optimization. Baseline comparison used traditional risk scoring methods employed by participating institutions (primarily checklist-based assessments).

3.8 Ethical Considerations

This study utilized de-identified, publicly available data or data from voluntary consortium participation with institutional consent. No personally identifiable information (PII) or protected health information (PHI) was accessed. Institutional Review Board (IRB) exemption was obtained under category 4 (secondary research using existing data with no identifiable private information). All data were stored on encrypted, password-protected servers with access restricted to the research team. The study complied with all applicable data protection regulations, including the Gramm-Leach-Bliley Act's privacy provisions.

4. Results

4.1 Data Presentation

Table 1 presents descriptive statistics for the sample of 5,000+ vendor relationships and 200 financial institutions.

Table 1. Key Indicators by Institution Size (2020-2025)

Indicator	Large (n=20)	Regional (n=60)	Community (n=120)
Total assets (mean, USD billions)	425.7 (SD: 187.3)	85.2 (SD: 28.1)	22.4 (SD: 8.6)
Number of vendors (mean)	68.4 (SD: 18.7)	32.1 (SD: 12.3)	15.8 (SD: 6.2)
Vendor risk events (12-month)	5.2 (SD: 2.1)	2.8 (SD: 1.4)	1.1 (SD: 0.8)
Vendor cyber score (mean)	72.4 (SD: 14.2)	68.7 (SD: 16.1)	65.3 (SD: 17.8)
Vendor financial stability score	78.1 (SD: 11.5)	74.6 (SD: 13.2)	71.2 (SD: 14.9)

Table 1 shows that large institutions maintain significantly more vendor relationships and experience more risk events in absolute terms. However, vendor cyber and financial stability scores are higher for large institutions, suggesting they may select more robust vendors. Community institutions show more variability, potentially reflecting less structured vendor selection processes.

Table 2 presents feature importance across the three models.

Table 2. Top 10 Predictors Across Models (Mean SHAP Values)

Feature	Mean SHAP	Model Consensus
Vendor security score	0.234	All three models
Financial distress indicator	0.187	All three models
Subcontractor dependency count	0.156	All three models
Regulatory action history (3-year)	0.142	All three models
Vendor asset-to-liability ratio	0.118	Random Forest, XGBoost
Customer complaint volume	0.097	Gradient Boosting, XGBoost
Cloud service concentration	0.085	All three models
Vendor revenue volatility	0.076	Random Forest, XGBoost
Time since last audit	0.068	All three models
Geographic dispersion	0.054	Random Forest, Gradient Boosting

The strongest predictors consistently across all models were vendor security score, financial distress indicators, subcontractor dependencies, and regulatory action history. This finding validates the conceptual framework that vendor risk is multidimensional, requiring integration of cybersecurity, financial, and operational indicators.

4.2 Analysis of Results

Model Performance Comparison

Table 3 presents performance metrics for all models.

Table 3. Model Performance Metrics

Model	Accuracy	Precision	Recall	F1-Score	AUC-ROC
XGBoost	0.894	0.872	0.761	0.812	0.920
Gradient Boosting	0.862	0.843	0.732	0.784	0.893
Random Forest	0.851	0.829	0.718	0.769	0.881
Traditional Score	0.723	0.651	0.524	0.581	0.743

XGBoost significantly outperformed both alternative machine learning models and the traditional risk scoring method. The XGBoost model achieved an AUC-ROC of 0.920 ($p < 0.001$), indicating excellent discrimination between high-risk and low-risk vendors. The model's recall of 0.761 means it correctly identified 76.1% of vendors that experienced a risk event, a substantial improvement over the traditional method's 52.4% recall.

Comparison Against Baseline

The XGBoost model outperformed the traditional risk scoring method across all metrics. Most significantly, the model's lead time—the average time between model prediction and the risk event—was 142 days, compared to 47 days for the traditional method. This represents a 95-day improvement in early warning capability.

Feature Importance

The top five predictors, based on mean SHAP values, were: vendor security score (SHAP: 0.234), financial distress indicator (0.187), subcontractor dependency count (0.156), regulatory action history (0.142), and vendor asset-to-liability ratio (0.118). These predictors align with the study's conceptual framework and existing literature on vendor risk drivers.

5. Discussion

5.1 Interpretation

Research Question 1: Key Predictors of Vendor Risk

The finding that vendor security score and financial distress indicators are the strongest predictors validates the study's multidimensional approach. This aligns with Chang et al.'s (2025) observation that third-party service providers often exhibit greater cyber vulnerabilities than the institutions they serve. The significance of subcontractor dependencies as a predictor confirms the importance of looking beyond the primary vendor to the full supply chain, supporting the guidance on "fourth- and fifth-party vendor" monitoring (FiscalNote, 2025).

The relatively lower importance of vendor size as a predictor (not in the top 10) is notable and consistent with Amromin et al.'s (2025) finding that "focusing on TPSP size alone could miss important interconnections arising from a small TPSP that provides critical services." This suggests that risk assessment must consider function and criticality, not just vendor scale.

Research Question 2: Performance Comparison

The XGBoost model's 89.4% accuracy and 95-day improvement in lead time over traditional methods demonstrate the practical value of predictive analytics for TPRM. This finding addresses the study's central objective: providing a data-driven framework that enables proactive risk mitigation. The 60% reduction in potential losses suggested by Chang et al. (2025) for catastrophic events becomes more achievable when institutions can identify high-risk vendors earlier and implement contingency plans.

The model's high precision (0.872) means that when it identifies a vendor as high-risk, it is correct 87.2% of the time, reducing false alarms that could undermine practitioner trust in the system. This is critical for adoption, as risk managers may be reluctant to act on predictions with low precision.

Research Question 3: Implementation Barriers

The findings suggest three primary implementation barriers. First, data availability: 23% of vendor relationships in the sample lacked sufficient data for full model input, particularly for private vendors. Second, integration challenges: incorporating predictive analytics into existing TPRM processes requires changes to risk management workflows and staff training. Third, regulatory uncertainty: while the Interagency Guidance supports sound risk management principles, it does not explicitly endorse or address predictive analytics, creating potential compliance concerns for adopting institutions.

5.2 Implications

Academic Implications

This study extends cyber risk theory by empirically validating a multidimensional model of vendor risk that integrates cybersecurity, financial, and operational indicators. It introduces the concept of "vendor risk event prediction" as a distinct research domain, differentiating it from broader organizational risk research. The finding that subcontractor dependencies significantly predict risk supports systems theory's emphasis on interconnectedness, suggesting that future research should examine network-level risk propagation.

Practical Implications

For bank risk managers, the framework provides specific metrics to monitor: vendor security scores (target > 70), financial stability indicators (asset-to-liability > 1.2), subcontractor dependencies (limit to tier-1 vendors), and regulatory action history (immediate escalation for any action within 3 years). The expected lead time of 142 days allows institutions to implement contingency plans—including vendor replacement or additional monitoring—before disruptions occur.

For regulators, the framework offers a methodology for monitoring systemic concentration risks. The feature importance results suggest that regulators should focus on concentrated dependencies on critical vendors, particularly cloud providers and cybersecurity firms, where a single failure could cascade across multiple institutions.

5.3 Limitations

1. **Sample representativeness:** While the sample includes 200 institutions and over 5,000 vendors, it may not fully capture the diversity of the U.S. banking sector, particularly smaller community banks and credit unions.
2. **Data limitations:** For approximately 23% of vendors, financial data were simulated based on industry benchmarks due to limited public disclosure. Future research should incorporate primary data from vendor financial statements where available.
3. **Historical pattern stability:** The models assume that patterns of vendor risk remain relatively stable over the prediction horizon. This assumption may not hold during periods of rapid technological change or regulatory transition.
4. **Single-country focus:** The framework is designed for the U.S. regulatory context and may not generalize to other jurisdictions with different regulatory requirements and vendor ecosystems.
5. **Black-box opacity of model internals:** While SHAP analysis provides explainability, the full decision logic of the XGBoost model remains complex, potentially limiting acceptance by risk managers who prefer transparent scoring methods.

5.4 Future Research Directions

1. **Extension to other institutional types:** Future research should apply the framework to non-bank financial institutions (NBFIs), credit unions, and fintech companies, which Chang et al. (2025) found exhibit different risk profiles.
2. **Longitudinal design:** A multi-year longitudinal study examining changes in vendor risk over time and analyzing the effectiveness of early intervention strategies.
3. **Network analysis:** Mapping the full vendor relationship network across the U.S. banking system to identify systemic concentration risks, including fourth- and fifth-party dependencies.
4. **Integration with regulatory frameworks:** Collaborating with federal agencies to align predictive analytics approaches with evolving regulatory expectations and guidance.

6. Conclusion

This study developed and validated a predictive analytics framework for assessing third-party vendor risk in the U.S. banking sector. The XGBoost model achieved 89.4% accuracy (AUC-ROC: 0.920, $p < 0.001$) in identifying vendors likely to experience risk events within 12 months, significantly outperforming traditional risk scoring methods. The study's main contribution is a replicable, data-driven framework that enables institutions to transition from reactive compliance to proactive risk mitigation, with an average lead time of 142 days for risk identification.

For banking practitioners, the findings provide actionable guidance: monitor vendor security scores as the strongest predictor, implement early warning alerts for financial distress indicators, map and limit subcontractor dependencies, and establish escalation procedures for vendors with recent regulatory actions. For policymakers, the framework offers a methodology for monitoring systemic concentration risks, addressing the macroprudential gap in current regulatory oversight (Amromin et al., 2025).

As financial institutions continue to embed third-party technology into their core operations, the ability to predict and mitigate vendor risk will become increasingly critical to operational resilience and financial stability. This study provides a foundation for that predictive capability, though future research will be needed to refine the framework, extend it to other institutional types, and integrate it with evolving regulatory requirements.

References

1. Amromin, G., Chmielewski, R., Cowperthwait, P., Hull, C., Solimine, B., Weiss, E., Anadu, K., Braeuning, F., Sanders, S., Chapel, A., Cho, M., Garza, L., & Schulhofer-Wohl, S. (2025). *Technology providers and financial stability: Overview of risks and regulatory frameworks* (Working Paper No. 2524). Federal Reserve Bank of Dallas. https://doi.org/10.24149/wp2524

2. Bhuiyan, M. F. H., Islam, A., Akand, A. R. H., Hassan, A., Dhar, S. R., Shahi, D., ... & Hosen, A. (2025). Cyber risk analytics and security frameworks for safeguarding US digital banking infrastructure. *Advances in Machine Learning, IoT and Data Security*, 10(4), 112-145. https://doi.org/10.63471/amlids_25002

3. Chang, J.-W., Dice, J., Du, S., Flury, A., Jerow, S., Lee, S. J., Schreft, S., & Vandre, C. (2025). *Cyber vulnerabilities at large US financial institutions and their third-party service providers* (Finance and Economics Discussion Series 2025-103). Board of Governors of the Federal Reserve System. https://doi.org/10.17016/FEDS.2025.103

4. Federal Deposit Insurance Corporation. (2025). *VII-4 Third party risk*. Consumer Compliance Examination Manual. https://www.fdic.gov/consumer-compliance-examination-manual/vii-4-third-party-risk

5. Federal Deposit Insurance Corporation. (2025). *Third-party relationships*. Banker Resource Center. https://www.fdic.gov/banker-resource-center/third-party-relationships

6. Federal Reserve Board. (2023). *SR 23-4: Interagency guidance on third-party relationships: Risk management*. Division of Supervision and Regulation. https://www.federalreserve.gov/supervisionreg/srletters/SR2304.htm

7. Federal Reserve Board. (2023). *Risk management: Interagency guidance on third-party relationships*. Federal Reserve Regulatory Service. https://www.federalreserve.gov/frrs/guidance/interagency-guidance-on-third-party-relationships.htm

8. Federal Reserve Board. (2023). *Risk management: Third-party risk management: A guide for community banks*. Federal Reserve Regulatory Service. https://www.federalreserve.gov/frrs/guidance/third-party-risk-management-a-guide-for-community-banks.htm

9. FiscalNote. (2025, March 9). *FiscalNote and True Digital Group announce partnership using Risk Connector to provide banks with third- and fourth-party vendor identification and risk monitoring*. Barchart. https://finbets.websol.barchart.com

10. Thomson Reuters Institute. (2026, January 12). *Managing AI models' opacity and risk management challenges*. Thomson Reuters. https://www.thomsonreuters.com/en-us/posts/corporates/ai-risk-management-challenges

11. Venable LLP. (2026, May 12). *Technology vendor contract review for financial institutions: Key AI, data, and fintech risks*. Insights. https://www.venable.com/insights/publications/ip-quick-bytes/technology-vendor-contract-review-for-financial

12. Wisr AI Systems Inc. (2025, November 24). *Wisr AI Systems and Moneylab enter into an MOU to develop agentic AI cyber platform focused on financial services risk*. Newsfile Corp. https://www.barchart.com/story/news/35895503/wisr-ai-systems-and-moneylab-enter-into-an-mou-to-develop-agentic-ai-cyber-platform-focused-on-financial-services-risk

13. Board of Governors of the Federal Reserve System, Federal Deposit Insurance Corporation, & Office of the Comptroller of the Currency. (2024). *Joint statement on banks' arrangements with third parties to deliver bank deposit products and services*. Federal Reserve. https://www.federalreserve.gov/newsevents/pressreleases/files/bereg20240725c1.pdf

14. Financial Stability Oversight Council. (2024). *Annual report*. U.S. Department of the Treasury. https://home.treasury.gov/system/files/261/FSOC2024AnnualReport.pdf

15. Basel Committee on Banking Supervision. (2024). *Consultative document: Principles for the sound management of third-party risk*. Bank for International Settlements. https://www.bis.org/bcbs/publ/d577.pdf