

Federated and Privacy-Preserving Machine Learning for Uncertainty-Aware Credit Default Risk Prediction in Distributed Financial Systems

Zilong Liu and Dedy Dwi Prastyo

University of Illinois Urbana-Champaign, USA

Abstract

Credit-default prediction is an important function of modern financial institutions because inaccurate assessment of borrower risk can increase loan losses, capital requirements, and systemic exposure. The growing availability of financial data has created opportunities for machine-learning-based credit-risk prediction, but centralized data architectures introduce significant privacy, security, and governance challenges. Financial institutions frequently maintain borrower information in separate databases because of regulatory restrictions, commercial confidentiality, and institutional policies. Federated learning provides an alternative architecture by enabling machine-learning models to be trained across distributed datasets without requiring the underlying data to be transferred to a central repository. However, distributed prediction also requires mechanisms for evaluating the reliability and uncertainty of model outputs. This study examines federated and privacy-preserving machine learning for uncertainty-aware credit-default risk prediction in distributed financial systems. A conceptual literature-based methodology was employed to synthesize research on multivariate financial-risk prediction, ensemble learning, federated credit-risk assessment, homomorphic encryption, privacy-preserving learning, and probabilistic AI diagnostics. Mahajan et al. (2024) demonstrate the relevance of multivariate data analytics and ensemble learning to financial-risk prediction. Konatham et al. (2025) extend this perspective to distributed credit-risk assessment through federated learning and homomorphic encryption. Barvaliya (2025) contributes a probabilistic-AI perspective through convergence assurance and rank-aware, tail-sensitive diagnostics for multi-chain inference. The findings indicate that an integrated architecture can combine local financial-risk modelling, federated learning, cryptographic protection, and uncertainty diagnostics. The proposed framework retains sensitive borrower information within institutional environments while supporting collaborative model development and reliability-aware credit-risk prediction. The study identifies data heterogeneity, class imbalance, communication overhead, privacy-utility trade-offs, computational cost, model security, and uncertainty assessment as important challenges for distributed financial machine learning.

Keywords Federated learning, credit-default prediction, privacy-preserving machine learning, financial risk, ensemble learning, uncertainty-aware prediction, probabilistic AI, homomorphic encryption, distributed financial systems

Received : 05.03.2026

Acceptance : 10.03.2026

Publication : 12.03.2026

1. INTRODUCTION

Credit-risk assessment is fundamental to banking and financial decision-making. Financial institutions must estimate whether borrowers are likely to repay loans while managing the financial and regulatory consequences associated with default. Machine-learning techniques have increasingly been applied to credit-risk prediction because they can identify nonlinear relationships among financial, behavioral, and borrower characteristics. Mahajan et al. (2024), for example, examined financial-risk prediction through multivariate data analytics and ensemble learning, demonstrating the value of combining multiple financial variables and predictive models.

However, increasingly sophisticated credit-risk models require access to sufficiently diverse information. A single financial institution may possess only a limited representation of borrowers, industries, geographic regions, and default events. Combining information across institutions could potentially improve predictive modelling, but direct data sharing can expose sensitive customer information and create regulatory and governance concerns.

Federated learning provides an alternative approach. Instead of transferring borrower records to a central database, participating institutions retain their datasets locally and collaboratively train a shared model. Model parameters or protected updates can be exchanged between institutions while the underlying datasets remain within their respective environments. Konatham et al. (2025) demonstrated this approach through federated learning for credit-risk assessment in distributed financial systems using BayesShield with homomorphic encryption.

Privacy preservation, however, is only one requirement of a reliable distributed credit-risk system. A model may protect sensitive data while still producing predictions that vary in reliability across institutions, borrower groups, or risk levels. This creates a need for uncertainty-aware assessment. Barvaliya (2025) examined convergence assurance for probabilistic AI using rank-aware and tail-sensitive diagnostics for multi-chain inference. This perspective is relevant to distributed credit-risk prediction because probabilistic outputs require appropriate diagnostics before they are used for risk-sensitive decisions.

The combination of multivariate financial modelling, federated learning, privacy-preserving computation, and uncertainty diagnostics therefore provides a broader framework for credit-default prediction. Mahajan et al. (2024) provide the financial-risk and ensemble-learning foundation, Konatham et al. (2025) provide the distributed and cryptographic architecture, and Barvaliya (2025) provides a reliability-oriented probabilistic perspective.

This study therefore examines federated and privacy-preserving machine learning for uncertainty-aware credit-default risk prediction in distributed financial systems. The objective is to develop a conceptual framework that combines collaborative financial-risk modelling, privacy protection, probabilistic assessment, and diagnostics for model reliability.

2. BACKGROUND OF THE STUDY

Traditional credit-risk modelling generally relies on centralized datasets. Financial institutions collect borrower information, train predictive models, and use the resulting models to support lending decisions. Centralized architectures can simplify model development but create a single environment in which sensitive financial information is aggregated.

Borrower information may include personally identifiable information, income records, transaction histories, repayment histories, credit characteristics, and other sensitive information. Sharing such information between institutions can therefore create substantial privacy and governance challenges.

Federated learning addresses this problem by moving computation toward the data rather than moving all data toward a central computational repository. Each participating institution trains a local model using its own borrower data and contributes protected model updates to a coordinating system.

The predictive foundation for such systems can be derived from multivariate financial-risk modelling. Mahajan et al. (2024) demonstrate the relevance of multivariate data analytics and ensemble learning for financial-risk prediction. Their work provides a basis for combining multiple financial indicators and predictive models within a broader distributed architecture.

Konatham et al. (2025) extend this concept by integrating federated learning with homomorphic encryption for credit-risk assessment in distributed financial systems. Their approach demonstrates how collaborative credit-risk modelling can be combined with cryptographic protection.

The reliability of distributed predictions represents another important issue. Federated learning can generate a shared model, but the output of that model may still contain uncertainty. Institutions may have different borrower populations, different sample sizes, and different distributions of default events. Consequently, a global prediction may not have the same reliability across all participating environments.

Barvaliya (2025) provides a relevant methodological perspective through convergence assurance for probabilistic AI and rank-aware, tail-sensitive diagnostics for multi-chain inference. Applied conceptually to financial risk, this perspective emphasizes that probabilistic outputs should be evaluated for inference reliability rather than interpreted solely from their numerical predictions.

Therefore, a distributed credit-risk architecture should address four interconnected requirements: predictive modelling, data privacy, distributed collaboration, and uncertainty/reliability assessment.

3. LITERATURE REVIEW

3.1 Multivariate Financial-Risk Prediction and Ensemble Learning

Machine learning has become an important approach to financial-risk prediction. Financial institutions frequently need to process multiple indicators simultaneously because credit and financial risk are influenced by interacting variables rather than a single factor.

Mahajan et al. (2024) examined financial-risk prediction through multivariate data analytics and ensemble learning. Their work demonstrates the relevance of combining multiple financial variables with multiple predictive models. This provides an important foundation for federated credit-risk modelling because distributed institutions may each contribute different financial information to a collaborative learning process.

Ensemble learning can also reduce dependence on one algorithm. Different predictive models may capture different relationships within financial data, and combining their outputs can provide a broader predictive representation.

In a federated environment, this multivariate and ensemble perspective can be extended across institutions. Each institution can process its local financial information while contributing model information to a shared predictive architecture.

3.2 Federated Learning for Credit-Risk Assessment

Federated learning extends conventional machine learning by allowing multiple institutions to participate in model development without directly pooling their underlying datasets.

Konatham et al. (2025) provide a direct application of this approach to credit-risk assessment in distributed financial systems. Their work incorporates BayesShield with homomorphic encryption, demonstrating how federated learning can be combined with cryptographic protection.

Federated credit-risk prediction is particularly relevant where institutions possess complementary borrower information but cannot freely exchange raw financial records. The architecture allows institutions to participate in collaborative modelling while retaining local control over their datasets.

However, distributed financial datasets may differ considerably in size and statistical characteristics. Institutions may serve different populations, operate in different geographic regions, or specialize in different forms of lending. These differences can affect the performance of the resulting global model.

3.3 Privacy-Preserving Financial Machine Learning

Federated learning reduces the need to centralize raw data, but it does not automatically eliminate privacy risks. Model updates may still contain information derived from local training data.

Privacy-preserving mechanisms such as differential privacy and homomorphic encryption can provide additional protection. Homomorphic encryption is particularly relevant to the present study because Konatham et al. (2025) explicitly integrate it into a federated credit-risk architecture.

The combination of federated learning and cryptographic protection therefore creates a multi-layer privacy architecture. Local data remain within institutional environments while selected computations or model exchanges can be protected through encryption mechanisms.

3.4 Uncertainty and Probabilistic Diagnostics

Privacy and predictive accuracy do not fully describe the reliability of a financial-risk system. Credit-risk models can produce predictions with different degrees of uncertainty depending on the quality and distribution of available information.

Probabilistic AI provides a mechanism for representing uncertainty around model outputs. Instead of producing only a categorical classification, a probabilistic model can characterize the likelihood of alternative outcomes.

Barvaliya (2025) contributes to this area by examining convergence assurance for probabilistic AI using rank-aware and tail-sensitive diagnostics for multi-chain inference. The methodological importance of this work lies in emphasizing diagnostics for probabilistic inference.

For distributed credit-risk prediction, this perspective can be used to conceptualize an additional validation layer. A federated model may produce probability-based risk estimates, while diagnostic procedures can assess whether those probabilistic outputs demonstrate sufficient stability and reliability.

Tail-sensitive assessment is also relevant to credit-default prediction because default represents an adverse outcome that may be relatively infrequent but financially significant. Therefore, evaluating model behavior in risk-sensitive regions can complement conventional average predictive-performance measures.

3.5 Integrating the Three Perspectives

The three principal references provide complementary components of a unified architecture.

Mahajan et al. (2024) support the predictive foundation through multivariate financial-risk analytics and ensemble learning.

Konatham et al. (2025) support the distributed architecture through federated learning and homomorphic encryption.

Barvaliya (2025) supports the uncertainty and inference-reliability dimension through convergence assurance and rank-aware, tail-sensitive diagnostics.

The combination of these perspectives allows distributed financial institutions to move from simple collaborative prediction toward privacy-preserving and reliability-aware financial intelligence.

4. METHODOLOGY

This study adopts a conceptual literature-synthesis methodology. The selected literature was analyzed according to five dimensions: financial-risk prediction, ensemble learning, federated learning, privacy-preserving computation, and uncertainty-aware probabilistic assessment.

Particular attention was given to the complementary contributions of Mahajan et al. (2024), Konatham et al. (2025), and Barvaliya (2025).

Methodological Dimension	Primary Function	Representative Reference
Multivariate financial analytics	Integrate multiple financial-risk indicators	Mahajan et al. (2024)
Ensemble learning	Combine predictive models	Mahajan et al. (2024)
Federated learning	Enable collaborative modelling across institutions	Konatham et al. (2025)
Homomorphic encryption	Protect selected computations involving sensitive information	Konatham et al. (2025)
Probabilistic prediction	Represent uncertainty in risk outputs	Barvaliya (2025)
Convergence diagnostics	Assess reliability of probabilistic inference	Barvaliya (2025)
Tail-sensitive diagnostics	Examine important regions of predictive distributions	Barvaliya (2025)

No original multi-bank financial dataset was supplied for the study. Consequently, the Results section represents a conceptual synthesis of the literature rather than newly generated empirical performance measurements.

5. RESULTS

5.1 Multivariate and Ensemble Financial Prediction

The literature synthesis indicates that credit-risk prediction benefits from the consideration of multiple financial variables and predictive models. Mahajan et al. (2024) demonstrate the relevance of multivariate data analytics and ensemble learning to financial-risk prediction.

Within a distributed architecture, these principles can be applied locally by participating institutions and incorporated into collaborative model development. This approach allows financial-risk systems to consider multiple indicators rather than relying on a single predictor.

5.2 Federated Credit-Risk Assessment

The second major finding concerns the role of federated learning. Konatham et al. (2025) demonstrate that federated learning can be applied to credit-risk assessment in distributed financial systems while incorporating homomorphic encryption.

This provides a foundation for collaborative credit-risk modelling in environments where financial institutions cannot freely centralize borrower information. The approach shifts the architecture from centralized data aggregation toward distributed model collaboration.

5.3 Privacy-Preserving Computation

The synthesis further indicates that federated learning should be complemented by privacy-preserving mechanisms. Keeping data locally can reduce direct data-sharing requirements, but model exchanges may still expose information.

Homomorphic encryption provides an additional protection mechanism within the architecture described by Konatham et al. (2025). This suggests that secure credit-risk modelling should treat privacy as an architectural component rather than as a separate post-processing step.

5.4 Uncertainty-Aware Credit Prediction

A fourth finding concerns the reliability of risk predictions. A federated model can produce a credit-risk probability, but the probability should not automatically be treated as equally reliable across all institutions or borrower populations.

Barvaliya (2025) emphasizes convergence assurance and rank-aware, tail-sensitive diagnostics for probabilistic AI. Conceptually, this supports adding a probabilistic diagnostic layer to federated credit-risk systems.

Such a layer could evaluate the stability of probabilistic predictions and provide additional information about uncertainty. This is particularly relevant to adverse-risk outcomes where relatively rare events may have significant financial consequences.

5.5 Integrated Architecture

The combined findings support the following conceptual sequence:

Local Financial Data → Multivariate/Ensemble Modelling → Federated Learning → Privacy-Preserving Computation → Probabilistic Risk Prediction → Convergence and Tail-Sensitive Diagnostics → Credit-Risk Decision Support

Mahajan et al. (2024) primarily support the first two components, Konatham et al. (2025) support the federated and cryptographic components, and Barvaliya (2025) supports the probabilistic diagnostic component.

6. DISCUSSION

The findings demonstrate that federated credit-risk prediction can be strengthened by integrating financial-risk modelling, privacy-preserving distributed learning, and uncertainty assessment.

The contribution of Mahajan et al. (2024) is particularly relevant to the predictive foundation. Multivariate data analytics allows financial-risk models to incorporate multiple interacting indicators, while ensemble learning provides a mechanism for combining different predictive approaches.

Konatham et al. (2025) extend this foundation to distributed financial systems. Their federated-learning architecture with homomorphic encryption demonstrates how institutions can collaborate on credit-risk assessment while incorporating cryptographic protection.

The third component concerns uncertainty. A model that protects privacy and produces accurate predictions may still provide outputs whose reliability varies according to data quality, borrower population, and model configuration. Barvaliya (2025) provides a methodological basis for addressing this issue through convergence assurance and rank-aware, tail-sensitive diagnostics for probabilistic AI.

The integration of these three components suggests that future federated credit-risk systems should not evaluate performance only through accuracy, precision, recall, or similar predictive measures. They should also examine the stability and reliability of probabilistic outputs.

This is especially relevant to credit-default prediction because default events are often relatively uncommon compared with non-default observations. The tail of a predictive distribution may therefore contain information that is highly important for financial-risk management.

Data heterogeneity represents another challenge. Participating institutions may have different customer populations, loan portfolios, underwriting procedures, and historical default rates. A shared model therefore needs to accommodate differences among local datasets.

Privacy-utility trade-offs also remain important. Additional privacy protection can introduce computational overhead or affect predictive information. Homomorphic encryption can strengthen confidentiality but may increase computational requirements. Federated learning can reduce raw-data sharing but may increase communication requirements.

Consequently, the proposed architecture should be viewed as a layered system rather than a single algorithm. Financial-risk modelling addresses predictive information, federated learning addresses collaboration, encryption addresses privacy, and probabilistic diagnostics address uncertainty and inference reliability.

7. PROPOSED INTEGRATED FRAMEWORK

The synthesis supports a six-layer framework for uncertainty-aware federated credit-risk prediction.

Layer	Component	Main Function
1	Local financial data	Maintain borrower and financial information within institutions
2	Multivariate and ensemble modelling	Extract financial-risk patterns from multiple indicators
3	Federated learning	Enable collaborative model development without centralized raw data
4	Privacy-preserving computation	Protect model exchanges and sensitive computations
5	Probabilistic prediction	Produce uncertainty-aware credit-risk estimates
6	Diagnostic and decision layer	Assess inference reliability and support financial decisions

The first two layers reflect the financial-risk modelling perspective of Mahajan et al. (2024). The third and fourth layers reflect the federated and homomorphic-encryption architecture of Konatham et al. (2025). The fifth and sixth layers incorporate the probabilistic inference and diagnostic perspective associated with Barvaliya (2025).

This integrated framework therefore connects financial prediction, distributed learning, privacy protection, and uncertainty assessment within a single conceptual architecture.

8. CHALLENGES AND LIMITATIONS

The proposed framework remains conceptual and does not provide an empirical comparison of federated credit-risk models using real-world multi-institutional data.

Data heterogeneity represents one important challenge. Participating institutions may have different borrower populations and risk distributions, which can influence global model performance.

Class imbalance is another challenge because default events may occur less frequently than non-default events. This issue can become more complex when the degree of imbalance differs among participating institutions.

Privacy-preserving technologies can also introduce computational and operational costs. Homomorphic encryption, for example, can require additional computation compared with conventional unencrypted processing.

Federated learning introduces communication requirements because model updates must be exchanged among participating institutions. Network reliability and institutional interoperability can therefore influence system performance.

Finally, uncertainty diagnostics require careful implementation. Probabilistic outputs should not automatically be interpreted as reliable simply because they are represented as probabilities. Diagnostic procedures are necessary to assess the stability and reliability of the inference process.

9. CONCLUSION

Federated and privacy-preserving machine learning provides a conceptual foundation for collaborative credit-default prediction in distributed financial systems. The integration of multivariate financial-risk prediction, ensemble learning, federated learning, homomorphic encryption, probabilistic modelling, and uncertainty diagnostics can provide a broader architecture for modern credit-risk assessment.

Mahajan et al. (2024) provide the foundation for multivariate financial-risk prediction and ensemble learning. Konatham et al. (2025) provide a direct framework for federated credit-risk assessment using homomorphic encryption. Barvaliya (2025) contributes the probabilistic-AI perspective through convergence assurance and rank-aware, tail-sensitive diagnostics.

Together, these references support a framework in which financial institutions retain local borrower data, develop collaborative models through federated learning, protect sensitive computations, generate uncertainty-aware credit-risk predictions, and evaluate the reliability of probabilistic outputs.

The resulting architecture recognizes that effective distributed financial intelligence requires more than predictive performance. Privacy, security, uncertainty, inference reliability, data heterogeneity, and interpretability must also be considered. Future empirical research can evaluate this framework using multi-institutional credit datasets and examine how privacy-preserving federated learning and probabilistic diagnostics affect credit-default prediction under heterogeneous financial conditions.

REFERENCES

1. Barvaliya, K. (2025, November). *Convergence assurance for probabilistic AI: Rank-aware, tail-sensitive diagnostics for multi-chain inference*. In 2025 IEEE International Conference on Emerging Trends in Computing and Communication (ETCOM) (pp. 1–4). IEEE.
2. Konatham, M. R., Uddandaraao, D. P., Vadlamani, R. K., & Konatham, S. K. R. (2025). *Federated learning for credit risk assessment in distributed financial systems using BayesShield with homomorphic encryption*. In 2025 International Conference on Computing Technologies & Data Communication (ICCTDC) (pp. 1–6). IEEE. <https://doi.org/10.1109/ICCTDC64446.2025.11158863>
3. Mahajan, A. S., Yamsani, N., & Konatham, M. R. (2024). *Financial risk prediction through multivariate data analytics and ensemble learning*. Stochastic Modelling and Computational Sciences, 4(2).
4. Anerao, P., Prakash, O., Ikhar, S., Sankaran, N., Avliyakov, A., & Bekbaulievich, A. I. (2025, December). *Federated learning for privacy-preserving credit risk assessment*. In 2025 15th International Conference on Mathematics, Actuarial Science, Computer Science and Statistics (MACS) (pp. 1–7). IEEE.
5. Awadelkarim, A. M. (2026). *Federated and interpretable AI framework for secure and transparent loan default prediction in financial institutions*. Mathematical and Computational Applications, 31(2), 56.
6. Bochner, R. (2026). *Credit risk assessment: A privacy-preserving framework integrating Shapley deep networks with blockchain verification*. International Journal of Automation and Digital Transformation, 5(1), 31–53.
7. Chatterjee, P., & Das, A. (2025, November). *Federated learning for privacy-preserving credit scoring across decentralized financial institutions*. In International Conference on Innovations in Data Science (pp. 323–339). Springer Nature Switzerland.
8. Houshmand, S., & Albadvi, A. (2025). *Credit risk prediction: An application of federated learning*. Journal of Information Systems and Telecommunication, 2(50), 154.
9. Mathpati, S. S., & Chirputkar, A. (2026, March). *Explainable AI and federated learning for privacy-preserving loan default prediction in cooperative banking networks*. In 2026 IEEE International Conference for Convergence in Computing Technology (I3CTCON) (pp. 1–5). IEEE.
10. Naresh, V. S., & Ayyappa, D. (2026). *Privacy-preserving federated credit risk models: Evaluating differential privacy and homomorphic encryption techniques*. Scientific Reports, 16(1), 4379.

11. Pingulkar, S., & Pawade, D. (2024, October). *Federated learning architectures for credit risk assessment: A comparative analysis of vertical, horizontal, and transfer learning approaches*. In 2024 IEEE International Conference on Blockchain and Distributed Systems Security (ICBDS) (pp. 1–7). IEEE.
12. Ramamoorthy, S. (2024). *Federated learning architecture for multi-cloud financial analytics with privacy-preserving data intelligence*. International Journal of Research Publications in Engineering, Technology and Management, 7(6), 11705–11714.
13. Sen, A., Heng, S. H., & Tan, S. C. (2026). *Recent advances in privacy-preserving federated learning for financial applications*. Discover Artificial Intelligence.
14. Zhang, C., & Liu, H. (2025). *A preliminary study on a multi-bank joint credit-risk control model based on federated learning*. Multidisciplinary Research in Computing Information Systems, 5(3), 341–358.
15. Zhang, S., Tay, J., & Baiz, P. (2024). *The effects of data imbalance under a federated learning approach for credit risk forecasting*. arXiv preprint arXiv:2401.07234.