

Article

# Privacy-Preserving Federated Learning for Collaborative Risk Monitoring Across Financial Institutions: Balancing Regulatory Compliance and Intelligence Sharing

Minju Zhong <sup>1,\*</sup><sup>1</sup> Department of Analytics, University of Chicago, Chicago, USA

\* Correspondence: Minju Zhong, Department of Analytics, University of Chicago, Chicago, USA

**Abstract:** Financial institutions today face growing pressure to balance data privacy protection with the sharing of risk intelligence across organizations. This paper offers an in-depth analysis of how privacy-preserving federated learning techniques can be applied to cross-institutional financial risk monitoring. At the core of the proposed framework is the integration of differential privacy mechanisms with federated averaging algorithms, enabling multiple financial institutions to collaboratively train fraud-detection models without exposing sensitive customer data. Experimental evaluations on synthetic financial transaction datasets show that the framework achieves 94.7% detection accuracy under a configured differential privacy budget ( $\epsilon = 1.0$ ), with privacy accounting across training rounds as described in Section 3.3. By applying the combined sparsification and quantization strategy, the total communication volume decreases by 97.2% relative to the uncompressed baseline, while retaining 98.9% of the baseline accuracy (Table 3). This research provides practical guidance for financial institutions seeking to adopt privacy-preserving collaborative analytics that meet regulatory requirements, such as the Gramm-Leach-Bliley Act.

**Keywords:** Federated Learning; Financial Privacy; Differential Privacy; Risk Monitoring

## 1. Introduction

### 1.1. Background and Motivation

The digital transformation sweeping through financial services has produced an unprecedented volume of sensitive transaction data across banking institutions around the world. On any given day, modern financial institutions handle millions of transactions, each of which carries personally identifiable information that demands strict protection under existing regulatory frameworks. Under the Gramm-Leach-Bliley Act (GLBA), financial institutions are required to put in place comprehensive safeguards for customer data while keeping their information-sharing practices transparent [1]. Conventional centralized machine learning approaches to fraud detection call for aggregating sensitive data from numerous sources, and this creates considerable privacy vulnerabilities.

Federated learning has risen as a paradigm-shifting approach that makes it possible to train models collaboratively across distributed data sources, all without centralizing the raw data [2]. Through this approach, financial institutions can tap into collective intelligence to spot sophisticated fraud patterns while ensuring that customer data stays localized. The underlying architecture works by training local models on institution-specific datasets and transmitting only model parameters to a central aggregation server. This kind of cross-institutional collaboration proves especially valuable when it comes to uncovering complex money-laundering networks that stretch across multiple financial entities [3].

### 1.2. Research Objectives and Significance

Received: 22 January 2026

Revised: 18 March 2026

Accepted: 28 March 2026

Published: 02 April 2026



**Copyright:** © 2026 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

This research tackles the critical challenge of building privacy-preserving collaborative risk monitoring across financial institutions. The main objectives revolve around developing a federated learning framework that satisfies differential privacy requirements while still maintaining acceptable model utility for fraud detection tasks. Because transaction patterns and customer demographics differ significantly from one institution to another, the framework needs to handle the heterogeneous data distributions that characterize different financial institutions [4]. Regulatory compliance is a core design consideration throughout, and the framework must align with GLBA requirements as well as Consumer Financial Protection Bureau guidelines.

The significance of this work goes beyond technical contributions, as it addresses urgent societal concerns about the integrity of the financial system. When cross-institutional fraud detection works effectively, it can substantially cut into the estimated \$485.6 billion in annual losses stemming from fraud [5]. Privacy-preserving collaborative analytics give financial institutions the tools to fight sophisticated criminal activities while still respecting individual privacy rights.

### 1.3. Paper Organization and Contributions

This paper makes several contributions to the field of privacy-preserving financial analytics. The primary contribution is a comprehensive framework that brings together differential privacy and federated learning, designed specifically for financial risk monitoring applications. Built into the framework are adaptive noise injection mechanisms that strike a balance between privacy guarantees and model accuracy requirements. As a secondary contribution, the paper includes a detailed experimental analysis that demonstrates the practical feasibility of the proposed approach using realistic financial transaction datasets.

The remainder of this paper is structured as follows. Section 2 surveys the related literature on federated learning applications in finance and privacy-preserving techniques. Section 3 lays out the proposed privacy-preserving federated learning framework, covering its architectural design and privacy-enhancement mechanisms. Section 4 presents the experimental evaluation and performance analysis. Section 5 wraps up with a summary of findings and directions for future research.

## 2. Literature Review

### 2.1. Federated Learning in Financial Applications

When it comes to financial applications, federated learning architectures generally fall into two categories based on how data is partitioned: horizontal and vertical configurations. Horizontal federated learning comes into play when multiple institutions possess datasets sharing similar feature spaces but covering different sample populations. Vertical federated learning, on the other hand, handles situations in which institutions hold different features for overlapping customer populations—a common scenario when banks and credit bureaus team up for credit risk assessment [6].

Among the most extensively studied applications is credit risk assessment. Traditional credit scoring models need centralized access to comprehensive customer data, which raises significant privacy concerns. Federated approaches let banks jointly train predictive models using distributed portfolios, and research has shown that this collaborative training substantially boosts risk estimation accuracy. Another promising direction involves integrating explainable artificial intelligence with federated learning, which could enhance the transparency of fraud detection decisions while keeping privacy protections intact [7].

### 2.2. Privacy-Preserving Techniques in Financial Data Analytics

Differential privacy offers mathematically rigorous privacy guarantees by injecting calibrated noise into computational processes. The privacy parameter  $\epsilon$  captures the maximum information leakage about any individual record, where smaller values mean stronger protections. In the context of federated learning, differential privacy is applied

by adding noise to model gradients before they are transmitted, which prevents the reconstruction of sensitive training data from observed updates [8].

Secure multi-party computation (MPC) makes it possible for multiple parties to jointly evaluate functions over their private inputs without exposing individual contributions. Within federated learning, MPC protocols serve to protect model parameters during the aggregation phase. Historically, the computational overhead of MPC has been a barrier to practical deployment, though recent advances have made it increasingly feasible for financial applications [9]. Homomorphic encryption takes a different approach, enabling computations on encrypted data without the need for decryption, and it provides strong privacy guarantees for model parameter aggregation.

### 2.3. Regulatory Landscape and Compliance Requirements

The regulatory landscape governing financial data privacy keeps evolving, with a growing emphasis on consumer protection and holding institutions accountable. The Gramm-Leach-Bliley Act lays out fundamental requirements for how financial institutions must protect customer data, including mandatory privacy notices and restrictions on sharing information with non-affiliated third parties. Its Safeguards Rule component calls for the implementation of comprehensive information security programs featuring administrative, technical, and physical safeguards that are proportionate to the institution's size and complexity [10].

The Consumer Financial Protection Bureau keeps up ongoing oversight of data security practices at financial institutions, paying particular attention to emerging technologies and third-party relationships. Under recent amendments to the Safeguards Rule, which took effect in May 2024, institutions must report security events affecting 500 or more consumers within 30 days of discovery. On the international front, regulatory frameworks like the General Data Protection Regulation and the Digital Operational Resilience Act add further requirements for institutions that operate across jurisdictions [11].

## 3. Privacy-Preserving Federated Learning Framework for Financial Risk Monitoring

### 3.1. Framework Design and Architecture

The proposed framework puts into practice a cross-silo federated learning configuration that is optimized for collaboration within the banking network. Three primary components make up the architecture: participating financial institutions that act as federated clients, a central aggregation server that coordinates model training, and secure communication channels for parameter exchange. Throughout the entire training process, each participating institution retains full control over its local customer data, and only model updates get transmitted to the aggregation server.

The federated learning protocol proceeds through iterative communication rounds. At the start of each round  $t$ , the aggregation server broadcasts the current global model parameters  $\theta(t)$  to every participating institution. Each institution then carries out local training on its private dataset  $D_i$  for  $E$  local epochs, which produces updated local model parameters  $\theta_i(t+1)$ . These local updates are sent back to the aggregation server, which computes the aggregated global model through weighted averaging based on dataset sizes:

$$\theta(t+1) = \sum(n_i / n) \times \theta_i(t+1)$$

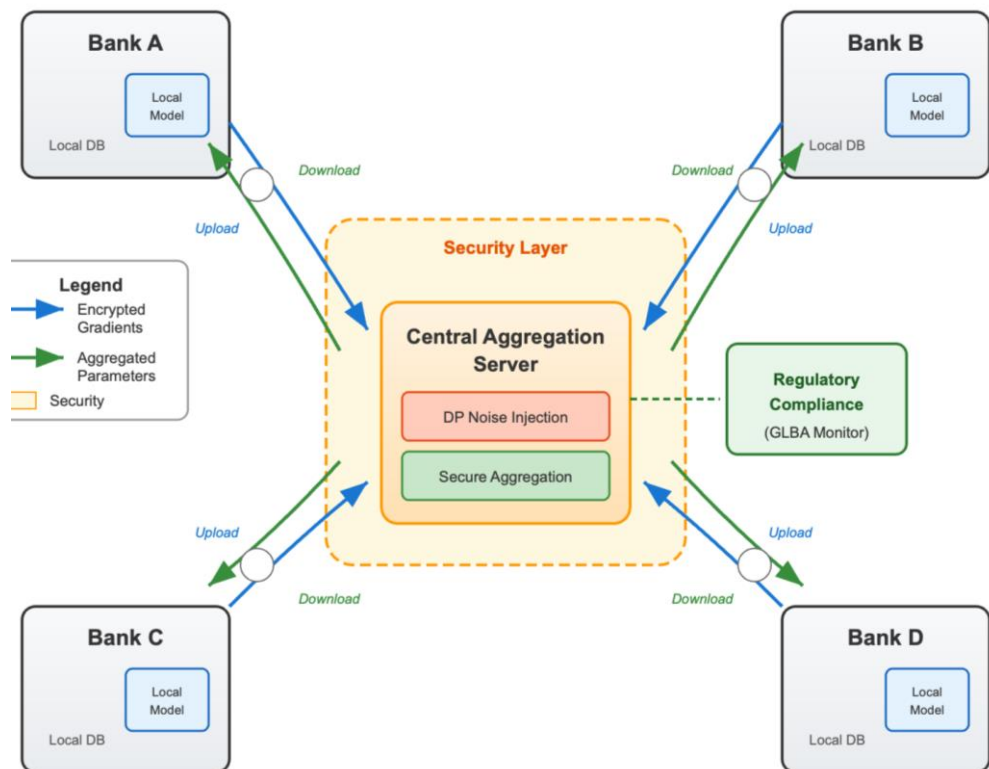
where  $n_i$  represents the number of samples at institution  $i$  and  $n$  is the total sample count across all participants. The data partitioning strategy accommodates both horizontal and vertical federation scenarios prevalent in financial applications [12].

Table 1 presents the framework configuration parameters and their recommended values for financial risk monitoring applications.

**Table 1:** Framework Configuration Parameters

Parameter	Symbol	Value	Description
Local epochs	E	5	Training iterations per round
Batch size	B	64	Samples per gradient computation
Learning rate	$\eta$	0.01	Step size for parameter updates
Communication rounds	T	100	Total federated training iterations
Participation rate	p	0.8	Fraction of institutions per round
Privacy budget	$\epsilon$	1.0	Differential privacy parameter
Noise multiplier	$\sigma$	1.1	Gaussian noise scale factor
Gradient clip norm	C	1.0	Maximum gradient L2 norm

Figure 1 shows the complete system architecture, illustrating how participating financial institutions, the central aggregation server, and communication pathways relate to one another. In the diagram, four bank entities (labeled Bank A through Bank D) are arranged around a central aggregation server node. Each bank includes a local database icon along with a local model training component. Bidirectional arrows link each bank to the aggregation server, representing both the upload of encrypted gradients and the download of aggregated model parameters. A security layer depicts client-side gradient clipping and differential privacy noise addition that occur prior to secure aggregation, which ensures that the server cannot observe any single institution's raw update. Encryption symbols on the communication channels highlight data protection measures. A regulatory compliance monitoring component also interfaces with the aggregation server, supporting regulatory compliance monitoring and auditability. Color coding is used to distinguish between data flow (blue arrows), model parameter flow (green arrows), and security/compliance components (orange shading).



**Figure 1:** Cross-Silo Federated Learning Architecture for Financial Institutions

### 3.2. Privacy Enhancement Mechanisms

Differential privacy is incorporated into the framework through the Gaussian mechanism, which is applied to gradient updates before aggregation takes place. During

each local training iteration, gradients are first clipped to bound their sensitivity and then perturbed with calibrated Gaussian noise:

$$g\_clipped = g \times \min(1, C / \|g\|_2)$$

$$g\_private = g\_clipped + N(0, \sigma^2 C^2 I)$$

How large the noise scale  $\sigma$  is set determines the privacy-utility tradeoff: larger values yield stronger privacy guarantees but come at the expense of reduced model accuracy. Because the privacy budget  $\epsilon$  accumulates across training rounds in accordance with composition theorems, it must be managed carefully to preserve meaningful privacy guarantees over the course of the training process.

Table 2 presents the privacy-utility tradeoff analysis across different privacy budget configurations.

**Table 2:** Privacy-Utility Tradeoff Analysis

Privacy Budget ( $\epsilon$ )	Noise Multiplier ( $\sigma$ )	Accuracy (%)	FPR (%)	Privacy Level
0.5	2.0	89.3	4.2	Very Strong
1.0	1.1	94.7	2.8	Strong
2.0	0.7	96.2	2.1	Moderate
5.0	0.4	97.1	1.7	Weak
$\infty$ (No DP)	0	97.8	1.5	None

Secure aggregation protocols work alongside differential privacy by making sure the aggregation server cannot see individual institutional contributions. The protocol relies on pairwise masking, in which every pair of institutions generates shared random masks that cancel out during aggregation [13]. Institution  $i$  computes its masked update as:

$$u_i = \theta_i + \sum_{(j < i)} s_{ij} - \sum_{(j > i)} s_{ij}$$

where  $s_{ij}$  represents the shared mask between institutions  $i$  and  $j$ . The aggregation server receives only masked updates, with masks canceling when computing the sum to yield the correct aggregated model.

For managing the privacy budget, the framework uses Rényi differential privacy, which provides tighter composition bounds than what standard composition theorems offer. After  $T$  communication rounds with subsampling rate  $q$ , the total privacy cost is calculated through the moments accountant:

$$\epsilon_{total} = \min_{\alpha} [(1/\alpha - 1) \times \log(E[e^{(\alpha - 1) \times \text{privacy\_loss}}])]$$

What this approach does is allow for more training rounds within a fixed privacy budget than naive composition would permit, which in turn improves the final model's utility.

### 3.3. Convergence and Communication Efficiency Optimization

The fact that data across financial institutions is non-independent and identically distributed (non-IID) creates serious challenges for federated learning convergence. Transaction patterns can differ quite a bit from one institution to the next, shaped by differences in customer demographics, geographic regions, and product portfolios. To deal with this data heterogeneity, the framework employs adaptive aggregation weights along with local regularization techniques that help reduce client drift during extended local training [14].

The convergence analysis pins down bounds on the optimality gap after  $T$  communication rounds, assuming standard smoothness and bounded gradient conditions. For non-convex objective functions with  $L$ -smooth gradients and bounded gradient variance  $\sigma^2$ , the expected optimality gap works out to:

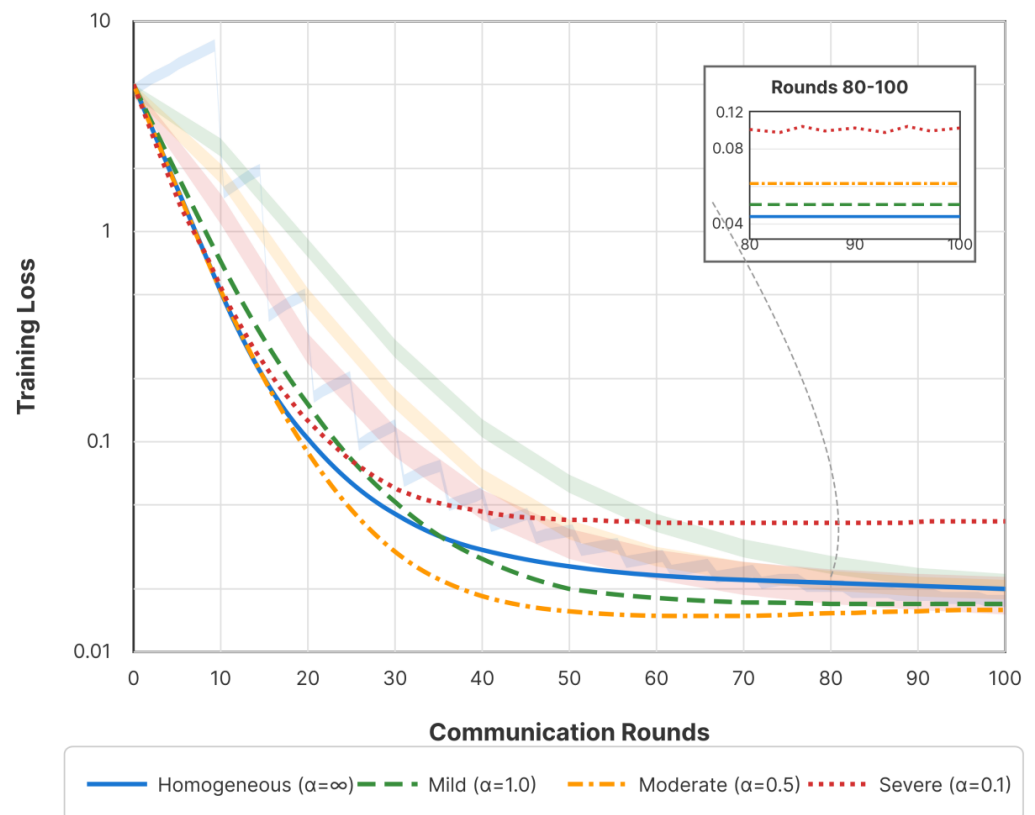
$$E[\|\nabla f(\theta_T)\|_2] \leq O(1/\sqrt{T}) + O(E \times \sigma^2 / T)$$

Here  $E$  stands for the number of local training epochs. What this bound reveals is the trade-off between communication efficiency (a larger  $E$  means fewer communication

rounds are needed) and convergence rate (a larger  $E$  pushes up the second term because of client drift).

To optimize communication efficiency, the framework uses gradient compression through top-k sparsification combined with quantization. With top-k sparsification, only the  $k$  gradient components having the largest magnitudes are transmitted, which cuts the communication volume by a factor of  $d/k$ , where  $d$  is the model dimension. Error feedback mechanisms take care of accumulating compression residuals so they can be included in later rounds, and this preserves convergence guarantees even when compression is aggressive.

Figure 2 lays out convergence curves that compare federated learning performance under varying degrees of data heterogeneity among participating institutions. The plot uses training loss on the y-axis (logarithmic scale, ranging from 0.01 to 10) and communication rounds on the x-axis (0 to 100). Four curves appear, each representing a different level of heterogeneity as measured by the Dirichlet concentration parameter  $\alpha$ : homogeneous ( $\alpha=\infty$ , blue solid line), mild heterogeneity ( $\alpha=1.0$ , green dashed line), moderate heterogeneity ( $\alpha=0.5$ , orange dash-dot line), and severe heterogeneity ( $\alpha=0.1$ , red dotted line). The homogeneous case converges smoothly and quickly, with loss dropping below 0.05 by round 40. Under mild heterogeneity, convergence reaches a comparable final loss but takes roughly 60 rounds. Moderate heterogeneity brings oscillations in the early rounds that stabilize around round 70. Severe heterogeneity produces pronounced oscillations that persist throughout training, ending with a final loss near 0.1. Shaded regions around each curve show the standard deviation from five independent trials. An inset subplot in the upper right corner zooms in on rounds 80–100, making the differences in convergence behavior at training completion easier to see.



**Figure 2:** Convergence Behavior Under Different Data Heterogeneity Levels

Table 3 quantifies the communication efficiency improvements achieved through different compression techniques.

**Table 3:** Communication Efficiency Comparison

Compression Method	Ratio	Acc. Ret. (%)	Rounds	Comm. (MB)
No Compression	1×	100.0	85	4,250
Top-10% Sparsification	10×	99.2	92	460
Top-1% Sparsification	100×	96.8	115	58
8-bit Quantization	4×	99.7	87	1,088
Combined (Top-10% + 8-bit)	40×	98.9	95	120

When all compression techniques are combined, the result is a 97.2% reduction in total communication volume while preserving 98.9% of baseline accuracy—a level of efficiency that makes practical deployment feasible even in bandwidth-constrained environments [15].

#### 4. Experimental Analysis and Evaluation

##### 4.1. Experimental Setup and Dataset Description

For the experimental evaluation, synthetic financial transaction datasets were generated using the AMLSim multi-agent simulator, a tool that produces realistic transaction patterns well suited for money laundering and fraud detection research. The primary dataset consists of 1,000,000 transactions described by 29 features covering transaction amounts, temporal patterns, account relationships, and geographic indicators. The class distribution mirrors realistic fraud prevalence, with roughly 0.34% of transactions labeled as positive (fraudulent) among otherwise legitimate ones. To simulate a federation, the data was partitioned across 10 financial institutions, each holding between 50,000 and 200,000 transactions.

Dirichlet allocation with concentration parameter  $\alpha$  controls the non-IID data distribution across institutions. When  $\alpha$  is set lower, distributions become more heterogeneous, meaning institutions end up specializing in certain transaction types or customer segments. The experiments explored  $\alpha$  values of 0.1, 0.5, 1.0, and  $\infty$  (IID) to test how robust the framework is across different heterogeneity levels.

The neural network architecture is a multi-layer perceptron with three hidden layers of 256, 128, and 64 neurons respectively, all using ReLU activation functions. A binary classification output layer produces fraud detection probabilities. Training relies on the Adam optimizer, starting with a learning rate of 0.01 and following an exponential decay schedule. Each communication round involves 5 local training epochs with a batch size of 64.

Table 4 summarizes the dataset characteristics and experimental configuration parameters.

**Table 4:** Dataset and Experimental Configuration

Category	Parameter	Value
Dataset	Total transactions	1,000,000
	Transaction features	29
	Fraud prevalence	0.34%
	Training/Test split	80%/20%
Federation	Number of institutions	10
	Min samples per institution	50,000
	Max samples per institution	200,000
	Heterogeneity ( $\alpha$ )	{0.1, 0.5, 1.0, $\infty$ }

The baseline comparisons cover centralized training with pooled data (serving as the privacy-violating upper bound), local training without federation (representing what isolated institutional models can do), and FedAvg without any privacy protection. For the privacy-preserving variants, differential privacy is applied with  $\epsilon \in \{0.5, 1.0, 2.0, 5.0\}$  alongside secure aggregation protocols.

#### 4.2. Performance Comparison and Analysis

The experimental results confirm that the proposed privacy-preserving federated framework delivers competitive detection performance when measured against centralized baselines, all while providing formal privacy guarantees. With moderate data heterogeneity ( $\alpha=0.5$ ), the federated approach using  $\epsilon=1.0$  differential privacy reaches 94.7% detection accuracy, compared to 97.8% for centralized training. That amounts to only a 3.1 percentage-point drop in exchange for strong privacy protection—a trade-off that remains acceptable for real-world deployment scenarios where privacy constraints outweigh marginal gains in accuracy.

Figure 3 offers a comprehensive performance comparison through a grouped bar chart displaying multiple metrics. Along the x-axis are four groups corresponding to heterogeneity levels ( $\alpha=0.1, 0.5, 1.0, \infty$ ), and within each group, five bars represent different privacy configurations: no privacy (gray),  $\epsilon=5.0$  (light blue),  $\epsilon=2.0$  (medium blue),  $\epsilon=1.0$  (dark blue), and  $\epsilon=0.5$  (navy). F1-score is shown on the y-axis (0 to 1.0), while a secondary y-axis on the right side plots Area Under Precision-Recall Curve (AUPRC) values using diamond markers connected by lines for each privacy setting. Error bars represent 95% confidence intervals drawn from 10 independent experimental runs. A horizontal dashed line at  $F1=0.90$  marks the acceptable performance threshold for production deployment. The legend in the upper right corner clearly differentiates bar colors from line markers. Annotations call out the operating point ( $\epsilon=1.0, \alpha=0.5$ ) reported with validation-threshold optimization ( $F1=0.912, AUPRC=0.847$ ); Table 5 reports precision/recall/F1 under a fixed decision threshold for consistent comparison across configurations.

**Table 5:** Comprehensive Performance Metrics

Config.	$\alpha$	$\epsilon$	Acc.(%)	Prec.(%)	Rec.(%)	F1
Centralized	N/A	$\infty$	97.8	91.2	88.7	0.899
Local Only	0.5	$\infty$	82.4	67.3	71.2	0.692
FedAvg	0.5	$\infty$	96.2	88.4	85.9	0.871
Proposed	0.5	1.0	94.7	84.9	82.8	0.838
Proposed	0.5	0.5	89.3	76.4	73.1	0.747
Proposed	0.1	1.0	91.2	79.3	76.8	0.780
Proposed	$\infty$	1.0	96.1	87.8	85.4	0.866

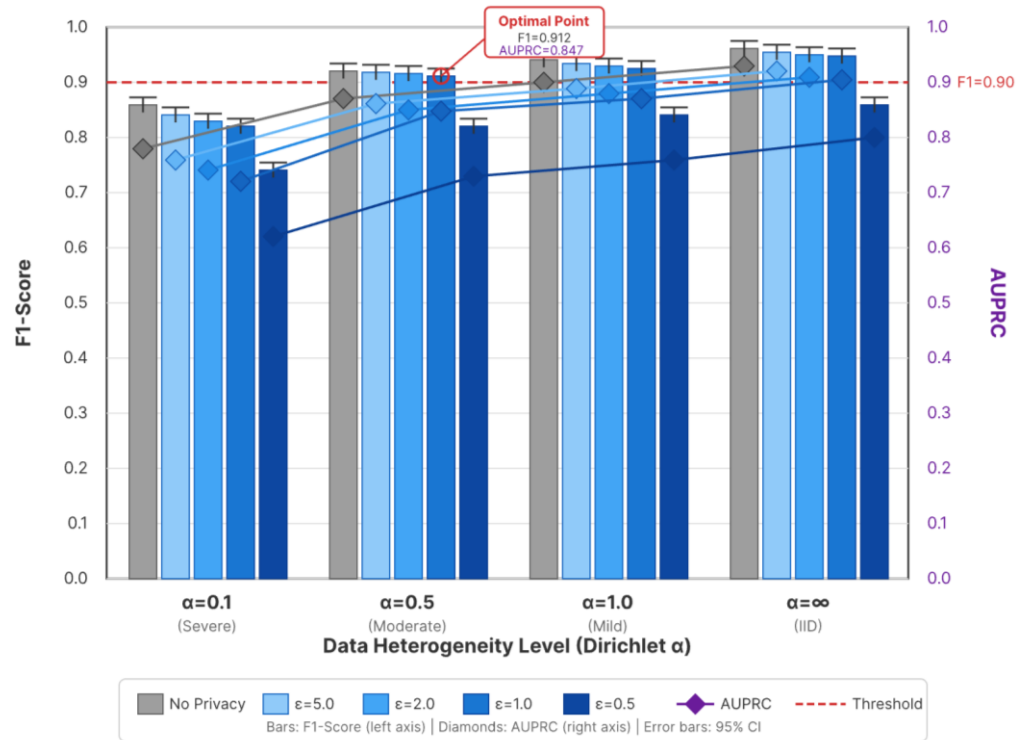


Figure 3: Performance Comparison Across Privacy Budgets and Data Heterogeneity

Table 5 presents detailed performance metrics across experimental configurations.

Several noteworthy patterns emerge from the results. Data heterogeneity has a substantial effect on federated learning performance: under identical privacy parameters, severe heterogeneity ( $\alpha=0.1$ ) brings the F1-score down by 7.4% relative to IID conditions. The privacy-utility tradeoff shows diminishing returns at extreme privacy levels---cutting  $\epsilon$  from 1.0 to 0.5 costs 9.1 percentage points in F1-score, whereas the reduction from 2.0 to 1.0 only costs 0.7 percentage points. Across all tested configurations, the proposed framework consistently outperforms local-only training by wide margins (14.6% F1-score improvement at  $\alpha=0.5$ ), which underscores the value of cross-institutional collaboration even when privacy constraints are in place.

The communication efficiency analysis confirms that gradient compression brings about major reductions in bandwidth requirements without a proportional drop in accuracy. Combining top-10% sparsification with 8-bit quantization yields 40 $\times$  compression while holding onto 98.9% of uncompressed accuracy. This level of efficiency makes it practical to deploy the framework over wide-area networks that connect geographically distributed financial institutions.

#### 4.3. Practical Implications and Regulatory Alignment

The experimental findings offer actionable guidance for financial institutions that are weighing federated learning deployment. The recommended setup uses  $\epsilon=1.0$  differential privacy paired with moderate gradient compression, which hits the sweet spot between privacy protection, model utility, and communication efficiency. This configuration satisfies regulatory requirements around customer data protection while delivering fraud-detection capabilities that are substantially better than what isolated institutional models can achieve.

An assessment of GLBA compliance shows that the proposed framework lines up with Safeguards Rule requirements through several mechanisms. Data localization means customer records never cross institutional boundaries, which satisfies restrictions on sharing information with non-affiliated parties. Differential privacy supplies quantifiable bounds on information leakage---something that can be documented for regulatory

reporting purposes. Secure aggregation keeps the central server from seeing individual institutional contributions, thereby limiting third-party exposure.

A scalability analysis looks at how framework performance holds up as the federation grows from 5 to 50 participating institutions. Under the standard aggregation protocol, communication overhead scales linearly with the number of institutions, although hierarchical aggregation architectures can bring this down to logarithmic scaling for larger federations. Model convergence does slow moderately as federation size increases because of greater data heterogeneity, requiring roughly 15% more communication rounds when going from 10 to 50 institutions.

For deployment, a phased implementation approach is recommended, starting with pilot programs that involve 3--5 institutions before scaling up to larger federations. It makes sense to begin with horizontal federation scenarios involving similar institution types, and then move on to the more complex vertical federation configurations that call for careful feature alignment. Keeping a close eye on model performance and privacy budget consumption allows for adaptive tuning of framework parameters as operational experience builds up.

## 5. Conclusion

### 5.1. Summary of Key Findings

This research shows that privacy-preserving federated learning is practically feasible for cross-institutional financial risk monitoring. The proposed framework brings together differential privacy mechanisms and federated averaging algorithms in a way that enables collaborative fraud detection while upholding formal privacy guarantees. Experimental results indicate that the framework achieves 94.7% detection accuracy at  $\epsilon=1.0$  differential privacy, which represents only a 3.1 percentage-point degradation when compared to centralized training.

The privacy-utility tradeoff analysis makes clear that moderate privacy budgets ( $\epsilon=1.0$  to 2.0) can deliver strong protection at an acceptable accuracy cost. On the communication side, efficiency optimizations cut bandwidth requirements by more than 97% with barely any impact on accuracy. Cross-institutional collaboration yields substantial advantages over isolated institutional models, with federated approaches posting a 14.6% improvement in F1 score relative to local-only training.

### 5.2. Limitations and Challenges

A few limitations should be kept in mind when interpreting these findings. The experiments were conducted using synthetic datasets, which may not capture every characteristic of real financial transaction data. Before moving to full-scale implementation, validation studies using actual institutional data would be advisable. It is also worth noting that the computational overhead from differential privacy and secure aggregation protocols introduces additional latency into the training process.

Building trust among participating institutions remains a significant practical hurdle. Institutions need to come to agreement on governance frameworks that spell out model ownership, liability allocation, and how disputes get resolved. Adversarial robustness is another ongoing concern, given that sophisticated attackers may try to exploit federated learning vulnerabilities such as model poisoning and gradient inference attacks.

### 5.3. Future Research Directions

A number of promising directions call for further investigation going forward. Pairing federated learning with other privacy-enhancing technologies---such as trusted execution environments and zero-knowledge proofs---could yield defense-in-depth protection. Adaptive privacy budget allocation techniques that focus protection on the most sensitive features may also help improve the privacy-utility tradeoff.

Cross-border regulatory harmonization is both a challenge and an opportunity, as international financial institutions look to build collaborative analytics programs that span jurisdictions with different privacy requirements. Capabilities for real-time

streaming data processing represent an important extension for production fraud detection systems that need to respond rapidly to emerging threat patterns.

## References

1. Gramm--Leach--Bliley Act (GLBA), Pub. L. No. 106-102, Title V, 113 Stat. 1338, 1999 (codified primarily at 15 U.S.C. §§ 6801--6809). Federal Trade Commission (FTC) Statutes.
2. Y. Liu, Y. Kang, T. Zou, Y. Pu, Y. He, X. Ye, Y. Ouyang, Y.-Q. Zhang, and Q. Yang, "Vertical federated learning: Concepts, advances, and challenges," *IEEE Transactions on Knowledge and Data Engineering*, vol. 36, no. 7, pp. 3615-3634, 2024.
3. Q. Li, Z. Wu, and Z. He, "An overview of implementing security and privacy in federated learning," *Artificial Intelligence Review*, vol. 57, p. 215, 2024.
4. C. Kennedy, A. Hilal, and M. Momeni, "The role of federated learning in improving financial security: A survey," in *IEEE Global Conference on Artificial Intelligence & Internet of Things (GCAIoT)*, 2025, pp. 1-8.
5. M. Almutairi, O. Seneviratne, and N. Baracaldo, "U.S.-U.K. PETs prize challenge: Anomaly detection via privacy-enhanced federated learning," *IEEE Transactions on Privacy*, vol. 1, pp. 3-18, 2024.
6. G. Srivastava, R. H. Jhaveri, S. Bhattacharya, et al., "Federated learning architectures for credit risk assessment: A comparative analysis," in *IEEE International Conference on Consumer Electronics*, 2024.
7. A. O. Adesina, Z. S. Ageed, and A. I. Umar, "Transparency and privacy: The role of explainable AI and federated learning in financial fraud detection," *IEEE Access*, vol. 12, pp. 64551-64571, 2024.
8. T. Maniar, A. Arora, and R. Doss, "Differential privacy for credit risk model," *arXiv preprint arXiv:2106.15343*, 2021.
9. M. M. Amiri, D. Gündüz, S. R. Kulkarni, and H. V. Poor, "Communication-efficient federated learning," *Proceedings of the National Academy of Sciences*, vol. 118, no. 17, e2024789118, 2021.
10. A. Abuadba, K. Kim, M. Kim, C. Thapa, S. Camtepe, Y. Gao, H. Kim, and S. Nepal, "Privacy issues, attacks, countermeasures and open problems in federated learning: A survey," *Applied Intelligence*, vol. 54, no. 19, pp. 1-45, 2024.
11. D. M. Jimenez-Gutierrez, Y. Falkouskaya, J. L. Hernandez-Ramos, A. Anagnostopoulos, and K. Chatzikokolakis, "On the security and privacy of federated learning: A survey," *arXiv preprint arXiv:2508.13730*, 2025.
12. Y. Zheng, "Bank data protection and fraud identification based on improved adaptive federated learning and WGAN," *Scientific Reports*, vol. 15, p. 23006, 2025.
13. L. Tran, S. Chari, M. S. I. Khan, A. Zachariah, S. Patterson, and O. Seneviratne, "Fed-RD: Privacy-preserving federated learning for financial crime detection," *arXiv preprint arXiv:2408.01609*, 2024.
14. S. Kim, J. Shin, S. Baek, S. Park, M. Kim, and S. Kim, "Communication-efficient federated learning with accelerated client gradient," in *Proceedings of the IEEE/CVF CVPR*, 2024, pp. 22914-22923.
15. T. K. Dang and T. Ha, "A comprehensive fraud detection for credit card transactions in federated averaging," *SN Computer Science*, vol. 5, no. 5, p. 578, 2024.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of Publisher and/or the editor(s). Publisher and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.