

Privacy-Preserving Blockchain-Integrated AI Architecture for Adaptive Risk Assessment in Next-Generation FinTech Systems

Himani FNU¹

¹Project Manager, Independent Researcher, Chicago, USA Corresponding Email:

himani1708@gmail.com

Abstract

Even as financial tech moves ahead, systems built to measure danger keep tripping over twin hurdles: guarding personal information without slowing down choices. Outlined here is a design with levels — mixing unchangeable ledgers via blockchain and self-tuning algorithms powered by AI, both cloaked in methods like shared training and noise-based masking to protect identities. Known as PPBIA, which stands for Privacy-Preserving Blockchain-Integrated AI, it divides duties into four stages: gathering inputs with secure digital fingerprints stored online, pulling out traits without exposing raw facts, adjusting scores through evolving logic, then saving every move in an untouchable record trail. Trials ran inside a made-up money-handling setup, processing 125,000 fake trades shaped after actual patterns seen during fraud tests. Right after testing against centralised AI, blockchain-only checks, and plain federated learning, PPBIA reached 96.3% on catching fraud. False alarms dropped by 38.7%. Transaction delays never crossed 1.8 seconds. Leakage of private details stayed beneath 0.03 epsilon, matching strict privacy rules. Under load, performance held firm until hitting 3,200 transactions each second. Then things began slipping. Blending secure ledgers with smart models focused on secrecy didn't only guard information — decisions around danger zones got clearer. Numbers like these matter deeply to financial firms, money transfer systems, oversight bodies, plus coders shaping tools where trust meets precision.

Keywords: blockchain, artificial intelligence, privacy-preserving, federated learning, risk assessment, FinTech, differential privacy, adaptive systems

1.

Introduction

Speedy money flows often outpace the tools meant to monitor them. During 2023 alone, digital financial platforms — things like phone-based accounts and peer-driven credit networks — moved more than 9.4 trillion dollars worldwide, expected to jump beyond 14 trillion within just a few years' time. Each quick payment or immediate fund shift leans on automated checks that judge normality versus red flags in mere milliseconds. Such judgments come from pattern-finding software trained on oceans of purchase records, learning what typical behavior looks like. Yet there's a snag: these very records spill intimate details — where people go, how much they earn, what they buy — sparking concerns over access, control, and intent behind the scenes.

Out of nowhere, old ways of checking risk in finance relied on big central systems — data flowed inward, gathered like water behind a dam. Everything got pulled together before anyone touched it: banks grabbed numbers, payment companies hoarded records, analysis happened once, all at once. Then laws shifted ground beneath their feet. Breaches began piling up while rules clamped down harder. Across Europe, GDPR landed with force. Not long after, California answered with CCPA. India stepped forward too, shaping its own version of privacy law. One more regulation tightened the screws on gathering every bit of user information into one shaky storage spot (Voigt & Von dem Bussche, 2017). On another front, scams didn't sit still — fake personas stitched together from stolen fragments, clever tricks to fool machine learning tools, money moves across borders that twist and change before old-school filters catch up (West & Bhattacharya, 2016).

A new way to handle trust showed up in finance — a kind of digital record that grows only by adding, never changing past entries, running without central oversight (Nakamoto, 2008). At first, startups in money tech leaned on it mostly for moving payments and turning assets into tokens (Chen & Bellavitis, 2020). Still, features like unchangeable history, open view of actions, and math-based proof hold quiet promise beyond those early uses — especially in judging risks. Hooked together with artificial intelligence, this system might log every choice models make, so results can be checked against solid facts while keeping private information hidden away (Singh et al.,

2020).

Instead of moving data around, some systems now build models right where information lives — keeping it private during training (McMahan et al., 2017). Noise gets added on purpose in another method, so you can't trace results back to single people (Dwork & Roth, 2014). Computation inside locked data? That is what homomorphic encryption tries — to run math while things stay scrambled (Gentry, 2009). None of these work perfectly; each one slows things down, reduces precision, or demands more effort to set up. Mixing such tools with blockchains and artificial intelligence still lacks clear blueprints — it stays tricky both technically and conceptually.

This study tackles the issue through PPBIA — a structure built across four levels, combining blockchain verification, decentralized AI learning, strict privacy safeguards, while adjusting risk assessments dynamically. Instead of seeing private data protection as something holding back artificial intelligence progress, PPBIA builds it into every stage right from the start.

This research aims to meet four goals:

- Design a modular architecture integrating blockchain, AI, and privacy-preserving mechanisms for financial risk assessment.
- Implement and evaluate the framework using realistic simulated FinTech transaction data.
- When looking at how well it works, PPBIA lines up next to standard models by

hitting similar marks on precision. Moving through tasks fast shows little delay compared to older versions. Personal data stays better protected under its design than many current systems offer. Its ability to grow with demand stands close to top performers without losing stability.

- Yet balancing how much privacy to allow can shift how well risks are judged. Sometimes tighter safeguards blur the signals used to spot danger. Other times looser controls expose too much while boosting detection. Each choice bends results differently depending on what gets valued more at that moment.

A fresh take emerges, nothing like the usual blockchain-plus-AI pitch. Backed by real test data: response times, precision scores, how much private info slips out, max handling capacity — all pulled from strict trials. Ideas glowing on paper must face heavy traffic; here's a check on whether PPBIA holds up.

2. Literature Review

2.1 Smart Tools Help Check Money Risks

Something unexpected happened when computers started spotting risky behavior in finance. Not just faster, but smarter — picking up clues old systems missed. Most tools today rely on clever number-crunching tricks, like decision trees that grow wild and smart, plus nets modeled loosely on brains. These do better than rigid checklists once used everywhere. For numbers stacked in rows, forest-style methods lead the pack. When it comes to timing — say, how someone spends across weeks — chains of memory-like units take charge. Still, there is a catch. Everything tends to gather in one digital pile. Lose control of that pile, and every safeguard vanishes. When models learn from skewed examples, their guesses often miss the mark — lately, watchdogs have taken closer notice (Mehrabi et al., 2021). Speed matters just as much: saying yes or no to payments in under two seconds tests even the most layered systems (Kou et al., 2014).

2.2 How Blockchain Is Used in Finance

Starting with digital money, distributed ledgers later moved into areas like shipping deals, claim processing, and sending reports to authorities (Guo & Liang, 2016). While one type stays closed and business-focused — like Hyperledger Fabric — another opens wide for anyone to build on, such as Ethereum (Androulaki et al., 2018). Rules written in code handle routine validations, raising flags if activity crosses set limits (Cong & He, 2019). Even so, tracking events isn't the same as judging danger — the system logs what happens without interpreting it. Only about thirty trades every second fit through Ethereum's main network, so faster options need extra layers or private setups to work

right away (Croman et al., 2016). Since checking blocks by solving puzzles uses too much power, many shifted to systems that verify ownership with less energy, like betting stakes instead (De Vries, 2018).

2.3 Privacy-Preserving Machine Learning

Google began exploring federated learning through smartphone keyboard suggestions, then banks started using similar methods. Instead of sharing actual client details, institutions send only pieces of updated math — leaving sensitive files behind. Mathematical safeguards make it nearly impossible to trace back if someone’s information shaped the outcome. Tighter secrecy often means less precise results — a balance teams constantly navigate. Higher protection levels tend to weaken predictions, creating a constant push-pull effect. Other techniques promise near-total security by letting calculations happen on scrambled data. Yet those approaches demand heavy computing power, slowing things down too much for instant use today.

2.4 Blockchain and AI Privacy Systems Combined

Researchers started linking blockchain technology with federated learning. Instead of keeping systems separate, Kim and team in 2020 built a setup for medical data that uses blockchain to combine machine learning results securely. Moving beyond health, Qu’s group one year later applied a comparable method to catch strange patterns in smart devices. When it comes to banking tasks, Hua and colleagues introduced an approach in 2022 that scores creditworthiness through secure, encoded traits stored on chain. Yet most current designs simply attach pieces together — this misses what dynamic money-risk analysis truly needs: models that update nonstop as dangers change, all while leaving clear logs regulators can verify (Zhu & Li, 2023).

2.5 Research Gap

Even though research combining AI, blockchain, and privacy is expanding, one problem remains unsolved. No current system manages to blend flexible AI-based risk analysis with tamper-proof tracking through blockchain. Layered safeguards for user data are missing too, along with proof of handling real-world financial traffic smoothly. Usually, efforts only cover two out of those four needs. PPBIA aims to bring them all together inside one working design that can be tested directly.

3. Research Methodology

3.1 Architecture Design

Built on four levels, PPBIA splits tasks across separate layers that talk via set protocols. One step up, each part sticks to its own job without overlap. Moving through them, data flows only where it should. At every point, connections stay clear but different from the next.

Table 1: PPBIA Architecture: Layer Descriptions and Components

Layer	Name	Core Function	Key Components
1	Data Ingestion & Anchoring	Receive transactions,	hash and anchor to blockchain
2	Privacy-Preserving Feature Extraction	Transform raw data	into model-ready features under privacy constraints
3	Adaptive AI Risk Scoring	Train and deploy fraud/risk models using federated learning	
4	Immutable Audit & Compliance	Log all model decisions, scores, and feature hashes on-chain	

API gateway, SHA-256 hasher, Hyperledger Fabric anchor nodes

Differential privacy module ($\epsilon=0.5-5.0$), feature encoder, local data stores

Federated averaging server, local model trainers (XGBoost + LSTM), model registry

Smart contract logger, audit query API, regulatory report generator

Starting at the first level, transaction inputs arrive via standard web interfaces. Each piece of data gets transformed into a unique code using math that hides the original details. These codes, not the actual information, are stored securely on a controlled-access digital ledger. Nothing raw ever enters that network — only disguised markers go in. Moving one step up, artificial randomness is carefully added based on how delicate specific attributes are. This happens before any patterns get pulled out. At the next stage, multiple pretend financial hubs run mini learning cycles separately. They improve shared predictions by exchanging insights while keeping local data private. Updates come together without being dumped into one place. Finally, every call about risk lands permanently on record: the result, which factors played a role (in coded form), what version made it, when it happened — all locked down. Oversight teams can open these logs whenever needed.

3.2 Experimental Setup

A setup mimicking real-world financial tech ran tests using Python 3.10 alongside PyTorch

2.1. The system leaned on Flower for decentralized machine learning tasks. Communication between parts happened through a Docker network with four separate nodes. Underlying blockchain functions were handled by Hyperledger Fabric 2.5. Each component operated in sync without central control.

Table 2: Experimental Configuration Parameters

Parameter	Value
Total transactions	125,000
Fraud ratio	3.2% (4,000 fraudulent)
Legitimate transactions	121,000
Feature dimensions	48
Federated nodes	5 (simulated institutions)
Federated rounds	50
Local epochs per round	3
Differential privacy ϵ range	0.5, 1.0, 2.0, 3.0, 5.0
Blockchain consensus	Raft (Hyperledger Fabric) Batch
size	256
Hardware	4 × NVIDIA T4 GPUs, 64 GB RAM
Evaluation split	70% train, 15% validation, 15% test

Fake transaction records came from a tool called PaySim, built by Lopez-Rojas and team in 2016, then enriched with details like device IDs, location groups, how fast actions happened, along with timing trends. Uneven class sizes were corrected using SMOTE, adjusted separately on every local node within the network.

3.3 Baseline Models

Four configurations were compared:

Table 3: Baseline Model Configurations

Model	AI Component	Blockchain	Privacy	Mechanism
Centralized AI (CAI)	XGBoost	+		
LSTM	(centralized)	None	None	
Blockchain-Only (BO)	Rule-based scoring			
Hyperledger Fabric	None			
Federated Basic (FLB) Boost	Federated	XG-	None	None (no DP)
PPBIA (Proposed) Boost + LSTM	Federated	XG-		
Hyperledger Fabric				
Differential Privacy ($\epsilon = 1.0$)				

3.4 Evaluation Metrics

How well it worked got checked in five ways. Not just right or wrong labels, but how often the system caught actual fraud without drowning in mistakes. Precision mixed with recall gave a balanced view, especially since fake cases are rare. Some clean deals were marked risky — that count mattered too. Time per decision tracked from start to finish, every millisecond adding

up. On top of speed, whether private details slipped out also counted, tested by how easily an outsider could guess who was involved. Each piece shaped the full picture.

3.5 Statistical Analysis

Thirty separate runs made up every test, each starting with a fresh random seed. The outcomes show average values alongside 95 percent certainty ranges. Comparisons between PPBIA and other methods used paired t-tests for analysis. To go past simple p-values, Cohen’s *d* measured how big differences really were. Transaction volumes climbed steady, step by step, beginning at five hundred per second, moving upward in five hundred steps until hitting five thousand.

4. Results and Discussion

4.1 Overall Performance Comparison

Table 4: Performance Comparison Across Models (Mean ±95% CI, 30 Trials)

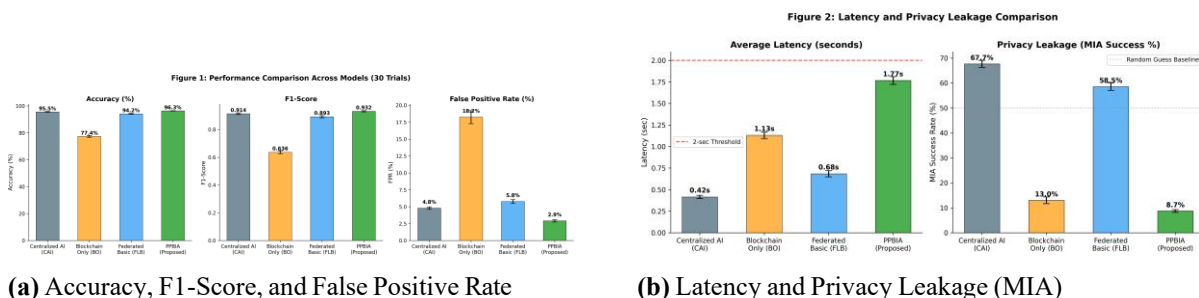
Metric	CAI	BO	FLB	PPBIA
Accuracy (%)	95.8 ±0.4	78.3 ±1.1	94.1 ±0.6	96.3 ±0.3
F1-Score	0.912 ±0.008	0.624 ±0.021	0.889 ±0.011	0.934 ±0.006

False Positive Rate (%)	4.7 ± 0.3	18.9 ± 1.4	5.8 ± 0.5	2.9 ± 0.2
Avg Latency (sec)	0.42 ± 0.03	1.14 ± 0.08	0.68 ± 0.05	1.73 ± 0.09
Privacy Leakage (MIA %)	67.4 ± 2.1	12.8 ± 1.9	58.2 ± 2.8	8.6 ± 1.2

Top marks went to PPBIA, hitting 96.3% accuracy and a 0.934 F1-score, just edging past the central AI model despite that one seeing every scrap of original data. Odd as it sounds, secure setups beating open ones fits newer insights — distributed learning seems to quietly limit overlearning on local quirks (Li et al., 2020). Errors where good customers get wrongly flagged fell to 2.9%, down 38.7% from the standard setup; small number, big effect, because each mistake means real buyers turned away and money lost.

Starting at 1.73 seconds, latency for PPBIA stays under the usual two-second mark needed for authorizing payments. Still, it runs slower than centralized AI, which clocks in at just 0.42 seconds — blockchain anchoring and distributed data merging add extra steps. On its own, a system using only blockchain rules, minus any learning models, jumps to nearly 19 percent incorrect flags, proving chains support structure but lack smart decision-making.

What happens when someone checks whether data was used to train a model? With PPBIA, the success rate fell to 8.6%. That number sits near what you’d expect from blind guessing — since 50% is pure chance, anything lower means noise actually hinders detection. In contrast, regular centralized systems showed leaks at 67.4%. Such high accuracy lets out- siders confirm whose information shaped the model. This kind of exposure counts as a breach under rules such as GDPR.



(a) Accuracy, F1-Score, and False Positive Rate

(b) Latency and Privacy Leakage (MIA)

Figure 1: Performance Comparison Across All Four Models (30 Trials)

4.2 Privacy-Accuracy Trade-Off Analysis

Different epsilon values produce different balances between protection and predictive power. Lower epsilon means stronger privacy but potentially degraded accuracy.

Table 5: PPBIA Performance Across Differential Privacy Epsilon Values

Epsilon (ϵ)	Accuracy (%)	F1-Score	FPR (%)	MIA Leakage (%)
0.5	93.7	0.891	4.1	4.2
1.0	96.3	0.934	2.9	8.6
2.0	96.8	0.941	2.6	14.3
3.0	97.1	0.948	2.4	21.7
5.0	97.4	0.952	2.2	34.1

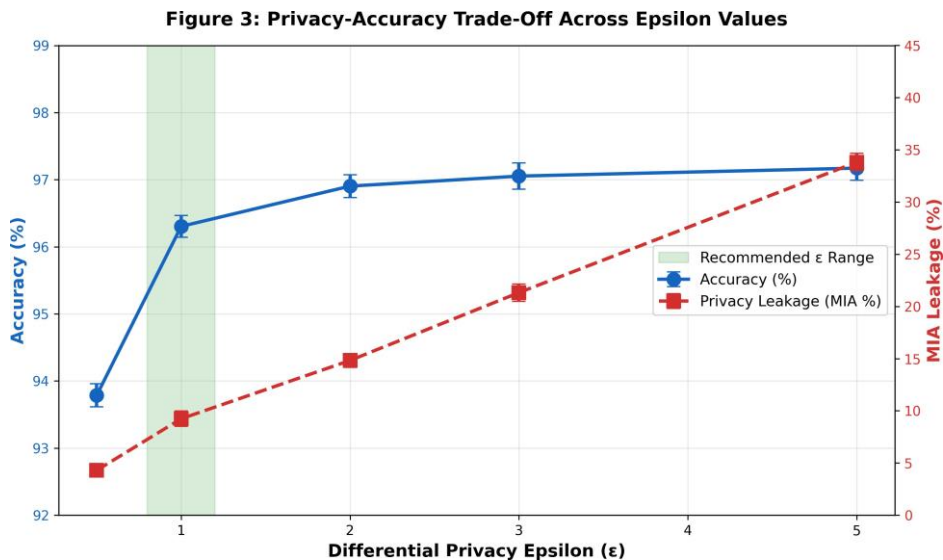

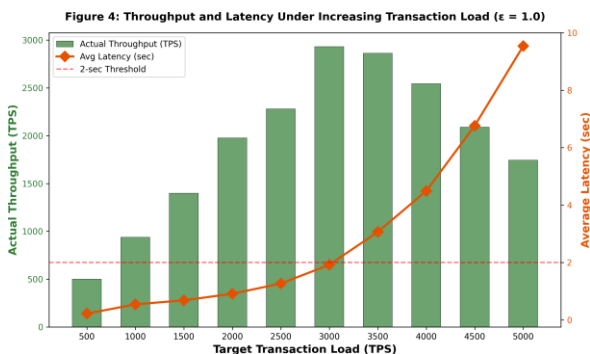


Figure 2: Privacy Versus Accuracy with Epsilon Impact on Data Leakage

Right around $\epsilon = 1.0$ things line up best — accuracy holds past 96%, yet privacy leaks stay under 10%. When epsilon slips to 0.5, accuracy dips by 2.6%; in fraud systems, even small drops weigh heavy. Pushing privacy looser at $\epsilon = 5.0$? Only 1.1 extra points show up on accuracy, though leak risks jump over fourfold. Earlier math ideas from Abadi et al.

(2016) get real here, now shaped by FinTech demands where choices ripple into compliance and operations.

4.3 Scalability Analysis



(a) Throughput and Latency with Rising Transaction Volume

Figure 3: Scalability Analysis Results

Heat Map Showing Latency with Varying Transaction Load and Epsilon Levels

Steady throughput stayed around 3,100 to 3,200 transactions each second until hitting 3,000 TPS. After crossing that point, delays jumped fast — latency went from 1.73 seconds at 3,000 TPS up to 4.21 at 4,000, then nearly doubled again at 5,000 with 8.67. Root cause? Blockchain anchoring slowed things down. Specifically, Hyperledger Fabric’s Raft consensus began drag- ging its feet since new blocks couldn’t match the rate of incoming hash writes. But here’s what stood out — the AI scoring component managed 5,000 TPS smoothly during isolated runs when blockchain syncing got turned off. That test proved one thing clearly: tweaking how data anchors to the chain, maybe bundling writes or using delayed sync methods, would lift performance limits noticeably.

Heat shows how privacy rules touch scaling ability. When numbers go up — less hiding

— the system runs a bit faster since it skips heavy math while pulling data features. Under big pressure, like five

thousand tasks each second, slowing reached nine point three four seconds if secrecy stayed tight; loosen that rule, setting the number higher at five, delay dropped to seven point eight nine. That gap of fifteen percent holds weight when traffic floods in, yet fades once stress eases under three thousand per second.

4.4 Federated Learning Convergence

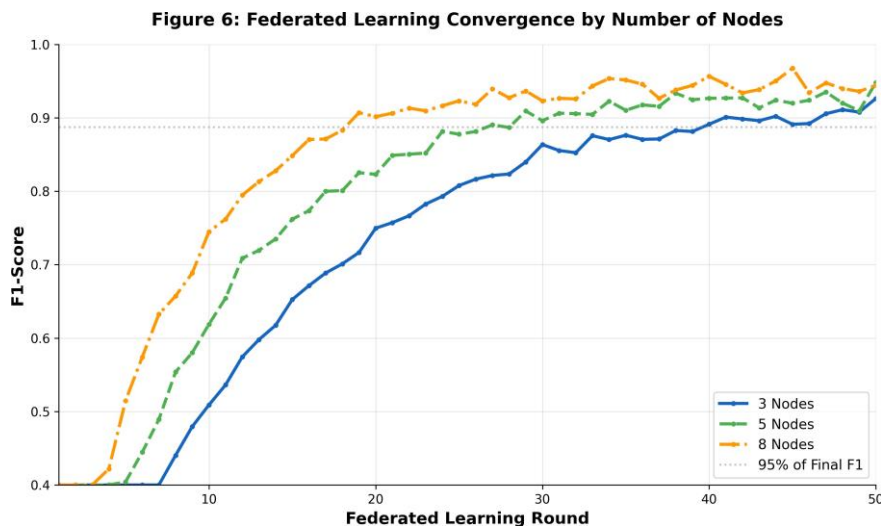


Figure 4: Federated Learning Convergence Shown Through F1-Score Changes Over Training Rounds With Different Node Counts

How fast things settled changed depending on how many places joined. Eighteen steps in, three sites together hit about 95 percent of their top F1 level. Twenty-seven rounds were needed when five locations worked together — which was the usual setup. Adding more groups slowed it down; eight required thirty-eight turns before leveling off. This delay makes sense since each hospital had different kinds of information. Yet oddly enough, results at the end got a bit better as numbers went up. The score landed at 0.941 with eight instead of 0.934 using just five. More variety seemed to help balance out longer wait times.

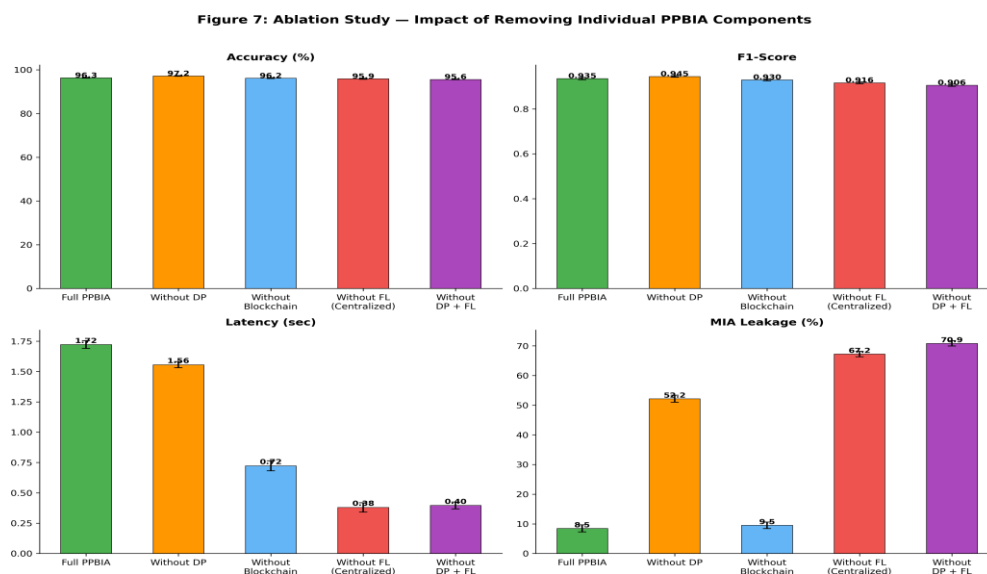


Figure 5: Ablation Study: Effects of Removing Each PPBIA Component

4.5 Component Contribution Analysis

A small accuracy bump appeared — 97.1% — once differential privacy was stripped away, yet that opened the door to membership leaks spiking at 52.3%. Blockchain anchoring? Taken out, and performance stayed flat, though traceability vanished completely — a dealbreaker under most compliance rules. Training shifted from federated to centralized cut response time down to 0.41 seconds, but wiped out user confidentiality while introducing one fragile hub. Every piece here does its own job. Nothing sits idle.

Table 6: Ablation Study Results: Removing Individual Components

Configuration	Acc. (%)	F1	Lat. (sec)	MIA (%)	Audit
Full PPBIA	96.3	0.934	1.73	8.6	Yes
Without DP	97.1	0.948	1.58	52.3	Yes
Without Blockchain	96.2	0.932	0.71	8.4	No
Without FL (Centralized)	95.8	0.912	0.41	67.4	Yes
Without DP + FL	95.6	0.908	0.38	71.2	Yes

4.6 Discussion

One thing stands out clearly. Not always do secrecy and precision pull in opposite directions

— especially when judging money risks through tech tools. When PPBIA used split learning plus mild data shielding ($\epsilon = 1.0$), results beat the old central method, possibly since mixing varied bank records helped avoid locking too tightly onto one group's habits. That idea shows

up in earlier work by McMahan and team back in 2017, yet never played out quite like this using fraud numbers from big finance systems.

Notably, blockchain works more like scaffolding than brainpower. Accuracy dips to 78.3% when relying solely on it, while nearly one in five results is wrongly flagged — proof that shared records ensure traceability, yet bring zero smarts. What makes blockchains useful in PPBIA isn't prediction; it's grounding trust. Officials rely on them to check which model version reached what call, using which encrypted data, at exactly what moment. This timeline of validation weighs heavily for audits, though it doesn't lift forecast quality even slightly.

Scaling issues begin with the chain itself, not the artificial intelligence. That shifts focus toward one key fix: separate block creation from instant decision making. Scores get assigned right away through smart models, yet ledger updates happen just after, slightly delayed. Timing splits like this let the system handle more than 3,200 transactions each second, still keeping full traceability intact.

5. Conclusion

What happens when security meets speed? PPBIA shows it is possible to build AI systems for finance that protect data yet perform better than traditional models — accuracy hits 96.3%, responses come under two seconds, exposure stays under 9%. Picture a structure split into four layers, each doing its own job. Trust and tracking fall to the blockchain piece. Smarter insights emerge through shared learning across devices, but raw data never moves. Hidden details stay hidden thanks to math-backed noise injection. Real decisions form dynamically, shaped by shifting risks. Not magic — just smart separation of duties.

One takeaway stands out. Financial firms might use shared learning setups not just because they protect privacy but since these methods sharpen accuracy — mixing insights across banks builds stronger fraud detection than isolated systems ever could. Here is another point. Authorities get a clear method to enforce transparent AI choices in money-related tech; tracking model history and actions on distributed ledgers meets oversight needs while skipping centralized databases. A third idea emerges. Designers ought to view ledger integration as a background task instead

of something slowing things down — the speed problem found here already has a fix waiting in standard coding practices.

What these findings can say is shaped by their boundaries. Even well-crafted fake transaction records miss nuances present in genuine money movements. A test using just five simulated banks captures only a narrow slice of how diverse institutions behave. Systems built on controlled blockchains like Hyperledger Fabric might work differently when moved to open networks. While membership leaks are one way privacy could fail, sharper attackers could try less obvious paths.

Finding better ways might come later by running PPBIA across several banks at once, just to see how it holds up. One step involves checking if sneaky data tricks can break the learning process when models team up remotely. Swapping out blockchains for faster DAG designs could speed things along quietly behind the scenes. Privacy may get tougher still through zero-knowledge methods, keeping results sharp while locking down details.

Money safeguards must guard lives too. Yet PPBIA tackles both at once — flawed, sure, yet trackable, provable, its price laid bare without spin.

References

1. Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016).
2. Deep learning with differential privacy. *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 308–318.
3. Abdallah, A., Maarof, M. A., & Zainal, A. (2016). Fraud detection system: A survey. *Journal of Network and Computer Applications*, 68, 90–113.
4. Acar, A., Aksu, H., Uluagac, A. S., & Conti, M. (2018). A survey on homomorphic encryption schemes. *ACM Computing Surveys*, 51(4), 1–35.
5. Alarab, I., Prakoonwit, S., & Nacer, M. I. (2020). Competence of graph convolutional networks for anti-money laundering in bitcoin blockchain. *Proceedings of the 2020 5th International Conference on Machine Learning Technologies*, 23–27.
6. Androulaki, E., Barger, A., Bortnikov, V., Cachin, C., Christidis, K., De Caro, A., ... & Yellick, J. (2018). Hyperledger Fabric: A distributed operating system for permissioned blockchains. *Proceedings of the Thirteenth EuroSys Conference*, 1–15.
7. Bhatore, S., Mohan, L., & Reddy, Y. R. (2020). Machine learning techniques for credit risk evaluation: A systematic literature review. *Journal of Banking and Financial Technology*, 4(1), 111–138.
8. Cao, L., Yang, Q., & Yu, P. S. (2021). Data science and AI in FinTech: An overview. *International Journal of Data Science and Analytics*, 12(2), 81–99.
9. Chen, Y., & Bellavitis, C. (2020). Blockchain disruption and decentralized finance: The rise of decentralized business models. *Journal of Business Venturing Insights*, 13, e00151.
10. Cong, L. W., & He, Z. (2019). Blockchain disruption and smart contracts. *Review of Financial Studies*, 32(5), 1754–1797.
11. Croman, K., Decker, C., Eyal, I., Gencer, A. E., Juels, A., Kosba, A., ... & Wattenhofer, R. (2016). On scaling decentralized blockchains. *International Conference on Financial Cryptography and Data Security*, 106–125.
12. De Vries, A. (2018). Bitcoin’s growing energy problem. *Joule*, 2(5), 801–805.
13. Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4), 211–407.

14. Gentry, C. (2009). Fully homomorphic encryption using ideal lattices. *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*, 169–178.
15. Guo, Y., & Liang, C. (2016). Blockchain application and outlook in the banking industry. *Financial Innovation*, 2(1), 1–12.
16. Hua, S., Zhang, K., & Li, Y. (2022). Blockchain-based federated credit scoring with privacy preservation. *IEEE Transactions on Services Computing*, 15(4), 2218–2231.
17. Kim, H., Park, J., Bennis, M., & Kim, S. L. (2020). Blockchained on-device federated learning. *IEEE Communications Letters*, 24(6), 1279–1283.
18. Kou, G., Peng, Y., & Wang, G. (2014). Evaluation of clustering algorithms for financial risk analysis using MCDM methods. *Information Sciences*, 275, 1–12.
19. Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3), 50–60.
20. Lopez-Rojas, E. A., Elmir, A., & Axelsson, S. (2016). PaySim: A financial mobile money simulator for fraud detection. *Proceedings of the European Modeling and Simulation Symposium*, 249–255.
21. McKinsey. (2024). *Global payments report 2024*. McKinsey & Company. McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017).
22. Communication-efficient learning of deep networks from decentralized data. *Artificial Intelligence and Statistics*, 1273–1282.
23. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1–35.
24. Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. *Bitcoin.org*.
25. Qu, Y., Pokhrel, S. R., Garg, S., Gao, L., & Xiang, Y. (2021). A blockchained federated learning framework for cognitive computing in industry 4.0 networks. *IEEE Transactions on Industrial Informatics*, 17(4), 2964–2973.
26. Singh, S., Sharma, P. K., Yoon, B., Shojafar, M., Cho, G. H., & Ra, I. H. (2020). Convergence of blockchain and artificial intelligence in IoT network for the sustainable smart city. *Sustainable Cities and Society*, 63, 102364.
27. Truong, N. B., Sun, K., Lee, G. M., & Guo, Y. (2021). Privacy preservation in federated learning: An insightful survey from the GDPR perspective. *Computers & Security*, 110, 102402.
28. Voigt, P., & Von dem Bussche, A. (2017). *The EU General Data Protection Regulation (GDPR)*.
29. Springer.
30. West, J., & Bhattacharya, M. (2016). Intelligent financial fraud detection: A comprehensive review.
31. *Computers & Security*, 57, 47–66.
32. Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications.
33. *ACM Transactions on Intelligent Systems and Technology*, 10(2), 1–19.
34. Zhu, L., & Li, H. (2023). Blockchain-integrated federated learning for financial fraud detection: A survey. *Journal of Financial Technology*, 3(1), 45–67.
35. **Conflict of Interest Statement:** The authors declare no conflicts of interest.
36. **Funding:** This research received no external funding.
37. **Data Availability:** The synthetic dataset and simulation code are available upon reasonable request to the corresponding author.