Federated Learning for Privacy-Preserving Financial Risk Forecasting Across Distributed Capital Markets

Gruce Yeber
Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL, USA.
weber2002@uab.edu

Francis M. Hayes
School of Computing, Clemson University, Clemson, SC, USA.
francismhayes582@clemson.edu

Damien A. Dunn
School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR, USA.
damien.dunn853@oregonstate.edu

Jack Kones
Department of Electrical Engineering and Computer Science, University of Kansas, Lawrence, KS, USA.
jackj@ku.edu

Abstract

The integration of machine learning into financial risk forecasting has historically required the aggregation of sensitive trading and portfolio data into centralized repositories, raising substantial privacy concerns and regulatory barriers. This paper proposes a federated learning framework that enables distributed capital markets to collaboratively train risk forecasting models without exposing raw client-level or institutional data. We examine the architectural trade-offs between model accuracy, communication efficiency, and statistical heterogeneity across participating entities. The study further analyzes how differential privacy mechanisms, secure aggregation protocols, and cryptographic primitives can be layered onto the federated pipeline to satisfy compliance with regulations such as the General Data Protection Regulation and the California Consumer Privacy Act. Governance challenges, including fair contribution metrics, incentive alignment, and auditability, are discussed in the context of multi-institutional consortia. Deployment sustainability concerns, such as energy consumption of iterative model updates and the robustness of the system to adversarial poisoning attacks, are evaluated through a synthesis of recent empirical results and theoretical guarantees. We also draw cross-domain comparisons from healthcare and mobile sensing applications to highlight structural parallels and transferable solutions. The paper concludes by outlining future research directions, including adaptive communication compression for latency-sensitive trading environments, verifiable computation for model integrity, and the incorporation of alternative data sources under privacy constraints. This work provides a comprehensive system-level roadmap for deploying federated learning in privacy-sensitive financial forecasting infrastructures.

Keywords

federated learning, privacy preservation, financial risk forecasting, capital markets, differential privacy, secure aggregation, distributed machine learning, financial regulation, adversarial robustness, governance.

1. Introduction

The forecasting of financial risk has long relied on statistical models trained on consolidated datasets that aggregate transaction records, portfolio holdings, and market-derived signals from multiple institutions. In traditional centralized machine learning paradigms, data from diverse sources must be transmitted to a single server where model training occurs, exposing proprietary information to potential breaches and regulatory penalties. The increasing stringency of data protection laws, coupled with the growing economic value of financial data, has created an urgent need for training methodologies that preserve privacy while maintaining predictive performance [1]. Federated learning, initially developed for mobile keyboard prediction and healthcare analytics, offers a compelling alternative by enabling multiple participants to collaboratively train a shared model without transferring raw data [2]. Instead, only encrypted or aggregated model updates are communicated to a central coordinator, thereby localizing sensitive information within each institution's infrastructure.

Capital markets present a uniquely challenging environment for federated learning due to the high dimensionality of financial time series, the non-stationary nature of market regimes, and the stringent latency requirements of trading systems. Unlike applications in image classification or natural language processing, financial risk forecasting requires models that can adapt rapidly to evolving volatility patterns and that can generalize across heterogeneous client bases with differing risk exposures [3]. Moreover, the participants in a federated network for risk forecasting are often competitors who are reluctant to share even aggregated gradient information for fear of revealing trading strategies or portfolio compositions [4]. Therefore, the design of a federated system for this domain must simultaneously address statistical heterogeneity, communication overhead, privacy guarantees, and incentive compatibility.

This paper provides a system-level analysis of federated learning applied to privacy-preserving financial risk forecasting across distributed capital markets. We examine the architectural decisions involved in selecting aggregation algorithms, encryption schemes, and differential privacy budgets, and we assess their implications for model utility and convergence speed. The governance of such systems is explored, including mechanisms for measuring each participant's contribution, auditing model behavior for fairness, and establishing contractual frameworks that align institutional incentives. We also consider the sustainability of federated deployments in terms of computational resource consumption, network bandwidth, and resistance to adversarial manipulation. By drawing on case studies from healthcare and mobile computing, we identify transferable design principles and highlight the unique constraints imposed by financial regulations such as the Markets in Financial Instruments Directive and the Basel Accords.

2. Background and Motivation

Financial risk forecasting encompasses a broad set of tasks including value-at-risk estimation, expected shortfall computation, volatility prediction, and drawdown probability modeling. Traditional approaches rely on historical returns and covariance matrices estimated from pooled data, which implicitly assume that all relevant information is captured in a single centralized repository. However, in practice, risk exposures are distributed across multiple

brokerages, asset managers, and clearinghouses, each holding partial views of the market. The aggregation of these fragmented datasets could improve forecast accuracy by capturing cross-institutional correlations and systemic risk factors [5]. Yet such aggregation is often prohibited by regulatory frameworks that mandate data minimization and purpose limitation, such as the General Data Protection Regulation in Europe and the California Consumer Privacy Act in the United States.

Federated learning directly addresses this tension by allowing models to be trained on decentralized data while keeping raw observations within the boundaries of each participant [2]. In the canonical federated averaging algorithm, each client computes gradients on its local data using a shared global model, encrypts or masks these gradients, and transmits them to a central server that updates the model by weighted averaging [1]. This process repeats over multiple communication rounds. Crucially, the server never has access to the raw data or even the full gradient vectors if secure aggregation is employed. For financial risk forecasting, this paradigm enables consortiums of banks, hedge funds, and regulated exchanges to jointly train models for volatility forecasting, tail risk prediction, and portfolio optimization without exposing individual positions or trading patterns.

Nevertheless, the application of federated learning to financial time series presents several technical challenges. Financial data is inherently non-independent and identically distributed across clients, with each institution observing different asset classes, market regimes, and return distributions. This heterogeneity can cause the global model to perform poorly on underrepresented clients or to converge slowly due to conflicting gradient directions [6]. Furthermore, the temporal correlation of financial observations violates the assumption of sample independence that underlies standard federated optimization, potentially leading to biased updates when clients train on consecutive time steps [3]. Additional complexity arises from the need to update forecasts at high frequency; in algorithmic trading environments, risk models must be refreshed within milliseconds, while federated learning rounds typically require seconds or minutes due to communication delays. These constraints motivate the investigation of adaptive aggregation strategies, local fine-tuning, and asynchronous communication protocols.

3. Federated Learning Architecture for Financial Risk Forecasting

A federated learning system for financial risk forecasting consists of three primary components: client nodes representing financial institutions, a central aggregation server (potentially operated by a neutral third party such as a clearinghouse or regulatory authority), and a communication network that supports encrypted payloads. Each client holds a private dataset of historical prices, returns, volumes, and perhaps alternative data such as news sentiment or macroeconomic indicators. The global model to be trained is typically a deep neural network or gradient-boosted tree ensemble designed to predict forward-looking risk measures such as conditional value-at-risk or volatility over multiple horizons [7]. In each communication round, the server broadcasts the current model parameters to all participating clients. Each client then performs one or more local training epochs on its own data, computes the gradient update, and applies a local differential privacy mechanism by adding calibrated noise to the update vector before transmission [8]. The server collects these perturbed updates and uses secure multi-party computation techniques to aggregate them without learning individual contributions [9]. The aggregated update is then applied to the global model, and the process repeats until convergence.

The choice of aggregation rule significantly influences model quality and convergence speed. The standard federated averaging algorithm weights each client's update by the size of its local dataset, but this may not be optimal when clients have vastly different data distributions or when some clients are more reliable than others [10]. In financial settings, a small number of institutions may hold disproportionately informative data, such as those with broad market exposure or access to high-frequency order book data. Weighted aggregation schemes that incorporate a measure of data quality, such as the predictive power of local gradients, can improve global model performance while mitigating the risk of overfitting to noisy or outlier clients [11]. Alternatively, more sophisticated algorithms such as FedProx introduce a proximal term to penalize large deviations from the global model, which helps stabilize training when data distributions are heterogeneous [6].

Communication efficiency is a critical concern because financial risk forecasting often requires near-real-time updates. The transmission of full gradient vectors for deep networks can consume substantial bandwidth, especially when hundreds of clients participate. Techniques such as gradient compression, sparsification, and quantization reduce the size of each update without severely degrading accuracy [12]. For example, clients can send only the top-k largest gradient components, or they can use stochastic quantization to represent each element with fewer bits. The server then reconstructs an approximate aggregate. In latency-sensitive environments, asynchronous federated learning allows clients to send updates independently without waiting for all participants to finish a round, though this introduces staleness and can slow convergence if not carefully managed [13]. A hybrid approach that combines synchronous rounds for model quality with asynchronous micro-updates for rapid adaptation to market shocks may offer the best trade-off.

4. Privacy and Security Considerations in Distributed Capital Markets

Privacy guarantees in federated learning are typically achieved through a combination of differential privacy, secure aggregation, and occasionally homomorphic encryption. Differential privacy ensures that the contribution of any single client's data does not have a significant effect on the final model, thereby protecting against inference attacks that attempt to reconstruct training samples from the model weights [8]. In the financial context, a client's trading data could be used to deduce proprietary strategies or client positions, so a strong privacy budget is often required. However, adding noise to updates reduces model accuracy, especially for rare events that are critical for tail risk forecasting [14]. Therefore, the selection of the privacy budget, typically denoted by epsilon, involves a trade-off between privacy protection and predictive utility. Recent work has proposed adaptive noise calibration that allocates more noise to less informative features and less noise to features that are essential for risk prediction [9].

Secure aggregation protocols prevent the server from observing any individual client's update. Using techniques such as secret sharing or pairwise masking, the server can compute the sum of all client updates without ever seeing the individual contributions [15]. This is particularly important in capital markets, where even the gradient vector may reveal information about portfolio composition or trading intensity. For instance, a gradient value associated with a specific asset weight could indicate whether a client holds a long or short position in that asset. Secure aggregation ensures that only the aggregate update is disclosed, thereby preventing the server from acting as a single point of privacy compromise. However, secure aggregation adds computational overhead and requires that all clients participate in the same round; dropouts must be handled through threshold schemes or robust protocols.

Beyond privacy, security concerns include model poisoning, where a malicious client submits crafted updates to corrupt the global model, and free-riding, where a client benefits from the global model without contributing meaningful updates. In financial risk forecasting, a poisoned model could lead to incorrect risk estimates that might be exploited by an adversary, potentially causing large losses or systemic instability. Defenses against poisoning include anomaly detection on gradient norms, cryptographic proof of computation, and reputation systems that discount contributions from untrusted clients [16]. The federated learning framework should also incorporate accountability mechanisms so that any malicious behavior can be traced back to the responsible institution, though this must be balanced against privacy requirements.

5. Structural Trade-offs and System Design

The design of a federated learning system for financial risk forecasting involves several interdependent trade-offs that must be resolved according to the specific operational context. One fundamental trade-off is between model accuracy and communication cost. Deep neural networks that achieve high predictive accuracy often require millions of parameters, leading to large communication overhead per round. In contrast, simpler models such as linear factor models or shallow trees are cheaper to communicate but may capture less complex risk dynamics [7]. The system architect must decide on the model capacity based on the available bandwidth and the tolerance for latency. For instance, a consortium of high-frequency trading firms with dedicated fiber connections can afford larger models than a group of retail brokers with limited network infrastructure.

Another trade-off concerns the degree of data heterogeneity that the system can tolerate. In federated learning with highly heterogeneous data, the global model may not perform well for any single client, leading to poor local risk forecasts. One solution is to train a personalized model for each client by allowing local fine-tuning or by using multi-task learning formulations that learn a shared representation as well as client-specific parameters [17]. Personalization, however, increases the total number of parameters and may require more communication or local computation. An intermediate approach is to cluster clients with similar data distributions and train a separate global model per cluster, which balances shared learning with customization.

The choice between synchronous and asynchronous training also affects system performance. Synchronous rounds guarantee that the model update is based on contributions from all clients, which tends to produce more stable convergence, but it introduces a straggler problem: if one client takes longer to compute its local update, all others must wait, increasing total training time [13]. Asynchronous updates eliminate waiting but can cause the model to oscillate when clients with stale updates participate. In financial markets where data arrives in real time, asynchronous methods may be appealing because they allow continuous incorporation of new market information. However, careful staleness control mechanisms, such as weighting updates inversely to their delay, are necessary to maintain stability.

6. Governance, Fairness, and Policy Implications

The deployment of a federated learning network across competing financial institutions raises complex governance questions. Who owns the global model? How are contributions measured and rewarded? What mechanisms prevent a large institution from dominating the training process? These questions must be addressed through formal consortium agreements that specify data contribution rights, model usage licenses, and dispute resolution procedures.

A common approach is to establish a neutral governing body, such as an industry association or a regulatory sandbox, that manages the central server and enforces the privacy and security protocols [18].

Fairness in federated learning for financial risk forecasting has two dimensions: distributive fairness, which ensures that all participants gain approximately proportional benefit from the model, and procedural fairness, which ensures that the training process does not systematically disadvantage certain types of institutions (e.g., smaller firms with less data). If the global model consistently underperforms for small banks while benefiting large ones, the smaller participants may withdraw, reducing the network's overall data diversity. Contribution metrics based on the reduction in global loss or the improvement in validation accuracy per client can be used to allocate usage rights or financial incentives [19]. Moreover, the model should be regularly audited for bias against particular asset classes or market segments, especially if the training data is imbalanced.

Policy implications are significant. Regulatory bodies such as the European Securities and Markets Authority and the U.S. Securities and Exchange Commission have expressed interest in how artificial intelligence is used in risk management, particularly regarding model interpretability and accountability. Federated learning introduces additional opacity because the model is trained on data that the regulator cannot directly inspect. To satisfy compliance, the system must provide a mechanism for auditing the global model's behavior without violating client privacy. Techniques such as differential privacy with a public audit trail, or cryptographic zero-knowledge proofs of correct training, could offer a path forward. The integration of federated learning into financial regulation could also reduce systemic risk by enabling more accurate and timely risk forecasts across institutions without requiring a central repository of all positions, which itself presents a single point of failure or attack.

7. Deployment Challenges and Sustainability

Deploying a federated learning system in production financial environments involves numerous practical challenges. One is the engineering effort required to integrate the federated learning client software into existing trading and risk management platforms, which are often built on legacy technologies with strict reliability requirements. Each institution must maintain a local data pipeline that extracts relevant features, trains the model on a rolling window of recent data, and securely communicates updates. The system must also handle client dropouts, network interruptions, and data staleness gracefully, as financial markets operate around the clock and any downtime can lead to missed risk updates.

Sustainability in the context of federated learning refers to the long-term viability of the training infrastructure. The computational cost of local training for each client can be substantial, especially when models are updated frequently. For a bank with millions of accounts, running multiple epochs of a deep neural network on daily data may require significant GPU or TPU resources. The energy consumption of federated learning networks has been studied in the context of mobile devices, but financial institutions have higher performance requirements and may rely on energy-intensive cloud computing [20]. Optimizations such as early stopping, quantization-aware training, and selective participation (where only a subset of clients is chosen each round) can reduce the resource footprint. Additionally, the carbon footprint of the entire consortium should be considered, and institutions may wish to adopt renewable energy sources for their data centers.

Robustness to adversarial attacks is another sustainability concern. Beyond poisoning, attackers could launch inference attacks to infer whether a particular trade was included in the training data, violating privacy. Defenses such as differentially private training with tight privacy accounting and robust aggregation (e.g., using trimmed mean or median instead of average) can mitigate these risks, but they often come at the cost of reduced model accuracy [14]. The system must also be resilient to changes in market structure, such as the introduction of new asset classes or regulatory changes that alter the definition of risk. Retraining the global model in response to such shifts requires coordinated action among participants, which can be slow. Therefore, the federated learning framework should be designed to support continuous learning and adaptation without disrupting ongoing operations.

8. Case Illustrations and Cross-Domain Comparisons

To better understand the structural trade-offs and potential pitfalls of federated learning in finance, it is instructive to examine analogous implementations in other privacy-sensitive domains. In healthcare, federated learning has been used to train models for medical image diagnosis across hospitals that cannot share patient records due to regulations such as the Health Insurance Portability and Accountability Act [1]. Studies have shown that federated models can achieve accuracy comparable to centralized models when the data distributions are similar across hospitals, but performance degrades significantly when distributions differ, as when one hospital specializes in rare diseases [6]. This mirrors the financial situation where a bank focusing on emerging market equities has a very different risk profile than one trading primarily in government bonds. Solutions developed in healthcare, such as transfer learning with local adaptation layers, have been successfully transferred to financial federated settings.

In mobile computing, federated learning has been deployed for next-word prediction and smart reply suggestions on smartphones [2]. That domain emphasizes communication efficiency because devices operate on metered cellular connections. Techniques such as gradient compression and on-device model caching have been directly adopted in financial systems, albeit with different hardware constraints. Mobile federated learning also deals with highly unbalanced and non-representative data distributions, similar to the financial case. However, mobile applications typically tolerate longer training intervals (e.g., overnight updates), whereas financial risk forecasting often requires near-continuous updates. This difference necessitates the development of asynchronous and streaming federated learning protocols that have not yet been fully explored in either domain.

Another relevant domain is collaborative fraud detection across multiple banks. Several industry consortia have tested federated learning for credit card fraud prediction, where data privacy is paramount and the cost of false negatives is high [15]. In such applications, the aggregation algorithm must be robust to extreme class imbalance, as fraudulent transactions are rare. Financial risk forecasting shares this characteristic because large market drawdowns are infrequent but catastrophic. The lessons learned from fraud detection, such as using focal loss functions and weighted aggregation based on class prior, can be adapted to risk forecasting to improve sensitivity to tail events.

Cross-domain comparisons also highlight the importance of incentive design. In mobile federated learning, participants (phone users) are typically not directly compensated; they receive a better product. In financial consortia, institutions invest significant resources and expect a measurable return on that investment. Therefore, the governance structure must include a mechanism for valuing contributions, for example by adjusting the weight given to

each client's update or by providing differential access to the global model's insights. The success of such incentive schemes depends on the transparency of the measurement process and the trust that all parties have in the central server's neutrality.

9. Future Directions and Forward-Looking Perspectives

The field of federated learning for financial risk forecasting is still in its infancy, and several promising research avenues are emerging. One direction is the use of personalized federated learning to handle client heterogeneity without sacrificing privacy. Approaches such as using meta-learning to initialize client-specific models, or learning a shared representation across clients while allowing each client to have a local head, could improve risk forecasts for smaller or specialized institutions. Another direction is the integration of alternative data sources, such as satellite imagery, news sentiment, or supply chain data, subject to privacy constraints. These data sources are often proprietary and cannot be shared directly; federated learning could allow institutions to jointly train models that leverage these signals without exposing the raw data.

Verifiable computation is another important frontier. In a decentralized setting, clients could maliciously claim to have performed the required local training without actually doing so, or they could submit forged updates. Cryptographic proofs, such as zero-knowledge succinct non-interactive arguments of knowledge, can enable the server to verify that an update was computed correctly without learning anything about the data [16]. While these cryptographic tools are computationally expensive, recent advances in proof systems have made them more practical for machine learning workloads. Their adoption in financial federated learning would increase trust and enable participation from institutions that are otherwise skeptical of the process.

The use of blockchain or distributed ledger technology to replace the central aggregation server is also being explored. A smart contract could coordinate the training rounds, manage contributions, and distribute rewards, eliminating the need for a central authority and reducing the risk of server-side bias or tampering [18]. However, blockchain consensus mechanisms introduce latency and throughput limitations that must be reconciled with the speed requirements of financial risk forecasting. Hybrid architectures that use a central server for aggregation and a blockchain for governance and audit might offer the best of both worlds.

Finally, the alignment of federated learning with financial regulations will require ongoing dialogue between technologists and policymakers. Regulators may mandate that risk models be interpretable, which is challenging for deep neural networks. Federated learning could be extended to train interpretable models, such as generalized additive models or rule-based systems, in a privacy-preserving manner. Alternatively, post-hoc explanation techniques that are privacy-compatible, such as distributing feature importance computations across clients, could satisfy regulatory needs without sacrificing predictive power. As capital markets become increasingly automated and interconnected, the ability to forecast risk accurately while respecting privacy will be a critical competitive and regulatory advantage. The framework presented in this paper provides a foundation for building such systems in a robust, fair, and sustainable manner.

10. Conclusion

This paper has provided a comprehensive analysis of federated learning as a paradigm for privacy-preserving financial risk forecasting across distributed capital markets. We have examined the architectural choices, privacy and security mechanisms, structural trade-offs,

governance requirements, and deployment sustainability of such systems. The unique challenges of financial data, including its high dimensionality, non-stationarity, and heterogeneity, demand careful adaptation of existing federated learning techniques. By integrating differential privacy, secure aggregation, and robust aggregation, a federated framework can satisfy regulatory constraints while maintaining predictive utility. The success of any deployment depends on the establishment of trust among competing institutions, which in turn requires transparent contribution metrics, fair incentive alignment, and auditable processes. Cross-domain insights from healthcare and mobile computing provide valuable guidance, but financial risk forecasting imposes stricter latency and reliability requirements that call for continued innovation in communication-efficient and asynchronous protocols. As the financial industry moves toward greater data collaboration under privacy mandates, federated learning stands out as a viable and necessary approach to harness distributed data for system-wide risk management. Future work should focus on personalized models, verifiable computation, and regulatory alignment to realize the full potential of this technology.

References

1. McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. In Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS) (pp. 1273–1282). PMLR.

2. Konečný, J., McMahan, H. B., Yu, F. X., Richtárik, P., Suresh, A. T., & Bacon, D. (2016). Federated learning: Strategies for improving communication efficiency. arXiv preprint arXiv:1610.05492.

3. Hu, L., & Shen, Y. (2026). A predictive analytics approach for forecasting global stock index returns using deep learning techniques. Decision Analytics Journal, 100685.

4. 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.

5. Gabaix, X. (2009). Power laws in economics and finance. Annual Review of Economics, 1(1), 255–294.

6. Li, T., Sahu, A. K., Zaheer, M., Sanjabi, M., Talwalkar, A., & Smith, V. (2020). Federated optimization in heterogeneous networks. In Proceedings of Machine Learning and Systems (MLSys) (Vol. 2, pp. 429–450).

7. Liu, T. (2026). Interpretable Machine Learning for Volatility Forecasting Under Realistic Walk-Forward Constraints.

8. Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. In Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS) (pp. 308–318). ACM.

9. Liu, T. (2026). Beyond volatility: A leakage-safe residual-stress signal for drawdown risk monitoring. Available at SSRN 6503179.

10. McMahan, H. B., Moore, E., Ramage, D., & y Arcas, B. A. (2016). Federated learning of deep networks using model averaging. arXiv preprint arXiv:1602.05629.

11. Sattler, F., Wiedemann, S., Müller, K. R., & Samek, W. (2020). Robust and communication-efficient federated learning from non-i.i.d. data. IEEE Transactions on Neural Networks and Learning Systems, 31(9), 3400–3413.

12. Alistarh, D., Grubic, D., Li, J., Tomioka, R., & Vojnovic, M. (2017). QSGD: Communication-efficient SGD via gradient quantization and encoding. In Advances in Neural Information Processing Systems (NeurIPS) (Vol. 30, pp. 1707–1718).

13. Xie, C., Koyejo, S., & Gupta, I. (2019). Asynchronous federated optimization. arXiv preprint arXiv:1903.03934.

14. Bagdasaryan, E., Veit, A., Hua, Y., Estrin, D., & Shmatikov, V. (2020). How to backdoor federated learning. In Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS) (pp. 2938–2948). PMLR.

15. Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., … & Seth, K. (2019). Towards federated learning at scale: System design. In Proceedings of Machine Learning and Systems (MLSys) (Vol. 1, pp. 374–388).

16. Blanchard, P., El Mhamdi, E. M., Guerraoui, R., & Stainer, J. (2017). Machine learning with adversaries: Byzantine tolerant gradient descent. In Advances in Neural Information Processing Systems (NeurIPS) (Vol. 30, pp. 119–129).

17. Smith, V., Chiang, C. K., Sanjabi, M., & Talwalkar, A. (2017). Federated multi-task learning. In Advances in Neural Information Processing Systems (NeurIPS) (Vol. 30, pp. 4424–4434).

18. Liu, T. (2022, December). Financial Constraint'Impact on Firms' ESG Rating Based on Chinese Stock Market. In 2022 4th International Conference on Economic Management and Cultural Industry (ICEMCI 2022) (pp. 1085-1095). Atlantis Press.

19. Sim, R., Hsieh, J., & Hong, S. (2023). Fairness and privacy in federated learning: A survey. ACM Computing Surveys, 55(7), 1–35.

20. Qiu, X., Parcollet, T., Beutel, D. J., Topal, T., & Lane, N. D. (2020). Can federated learning save the planet? In NeurIPS 2020 Workshop on Tackling Climate Change with Machine Learning.

21. Xue, P., & Ye, Y. (2026). Attention-enhanced reinforcement learning for dynamic portfolio optimization. Intelligent Systems with Applications, 200622.

22. Liu, T. (2026). Volatility Forecasting and Early-Warning Market Stress Detection: A Leakage-Safe Evaluation with Tree Ensembles and Transformers.