Fed-RNBS: A Federated Risk-adjusted Nash Bargaining Solution for Financial Applications

Subeen Choi
School of Computer Science & Engineering
Soongsil University
Seoul, Korea
schoi@soongsil.ac.kr

Bong Jun Choi
School of Computer Science & Engineering
Soongsil University
Seoul, Korea
davidchoi@soongsil.ac.kr

***Abstract*—The practical application of Multi-Task Federated Learning (MTFL) in the financial domain is challenged by a combination of class imbalance, Non-IID data, and tasks with asymmetric failure costs. To address these compounded challenges, this paper proposes a game-theoretic algorithm, the Federated Risk-adjusted Nash Bargaining Solution (Fed-RNBS). Fed-RNBS negotiates an optimal model update by dynamically incorporating task-specific risks and learning states. Experiments on a real-world mortgage dataset show that the proposed model improves the F1-Score by 46.6% on the highly imbalanced delinquency prediction task compared to baseline models. This demonstrates the effectiveness of Fed-RNBS in building robust and reliable models for complex real-world financial environments.**

***Index Terms*—Multi-Task Federated Learning, Financial Risk Management, Gradient Conflict, Game Theory, Nash Bargaining Solution**

I. Introduction

Federated Learning (FL) [1] is a distributed learning paradigm designed for collaboratively training models in environments where data is decentralized across multiple institutions or devices, ensuring privacy is preserved. Instead of transferring local data externally, each client shares only the model parameters learned locally with a central server. This approach not only fundamentally guarantees data privacy but also offers the advantage of mitigating persistent issues inherent in traditional centralized learning, such as communication overhead, computational bottlenecks at the central server, and storage constraints.

The utility of Federated Learning can be further extended when combined with Multi-Task Learning (MTL) [2]. MTL is a methodology wherein a single model learns multiple related tasks concurrently. By sharing knowledge across tasks, it can achieve superior overall performance and efficiency compared to training separate models for each task. The combination of these two paradigms, Multi-Task Federated Learning (MTFL), aims to build a single, powerful model capable of performing multiple tasks simultaneously, without requiring clients to share their privacy-sensitive data.

The financial sector represents one of the most ideal yet challenging environments for applying the MTFL methodology. Financial institutions possess vast amounts of data essential for building accurate and robust AI models. However, they face a dilemma: this data is fragmented within each institution's firewall, and strict privacy regulations render inter-institutional sharing virtually impossible [3]. Furthermore, the financial domain is characterized by two significant data challenges that impede the performance of federated learning. The first is the Non-Identically and Independently Distributed (Non-IID) nature of the data [4]. Each bank has a unique customer base and product portfolio, leading to disparate data distributions among clients (banks), which hinders the convergence of the global model. The second is the problem of severe class imbalance. In financial risk prediction, risk events manifest as a very small number of minority samples. For instance, in the delinquency prediction data (Task 1) addressed in this study, actual delinquency samples constitute only 1.96% of the entire dataset, making it extremely difficult for the model to learn the patterns of the minority class.

Our prior work [5] validated the feasibility of MTFL in finance using a shared encoder architecture with static loss weighting. While this approach yielded promising results, we also identified two critical limitations that impede its practical application: (1) performance degradation caused by conflicting gradients among different tasks, and (2) training instability stemming from task-specific data imbalances.

To address these challenges, we propose Game Theory as a solution. Game theory is a mathematical framework for analyzing strategic interactions among multiple decision-makers [6]. Although prior research has applied the Nash Bargaining Solution (NBS) to multi-task learning to mitigate gradient conflicts [7], this approach has a fundamental limitation: it treats all tasks symmetrically. This symmetric treatment makes it ill-suited for the financial domain, where tasks often have asymmetric failure costs. For example, the cost of failing to predict a loan delinquency is substantially greater than that of incorrectly predicting a prepayment. Consequently, a symmetric approach can lead to an unstable model from a risk management perspective. More recently, an asymmetric approach (AuxiNash) that automatically learns task preferences through bi-level optimization has also been proposed [8], but it assumes a general-purpose auxiliary learning environment.

The key contributions of this paper are as follows:

- We propose a novel algorithm, the Federated Risk-adjusted Nash Bargaining Solution (Fed-RNBS), to si-

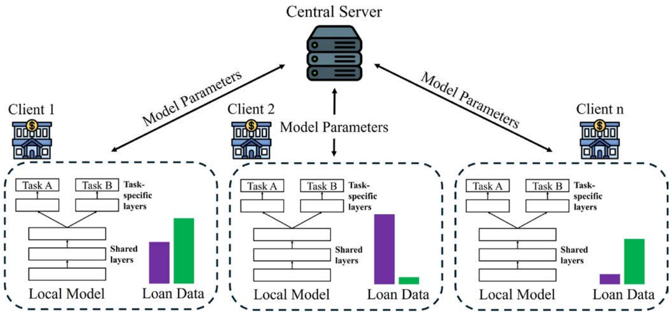

Fig. 1: The overall framework of Fed-RNBS.

multaneously address the challenges of inter-task gradient conflicts and data imbalance. Our method defines the asymmetric failure cost of each task as 'risk' and incorporates it into the bargaining process, thereby ensuring stable gradient updates tailored for the financial environment.

- We demonstrate the effectiveness of Fed-RNBS on a real-world mortgage dataset from Freddie Mac, which exhibits both Non-IID and class imbalance properties. Our results confirm its superiority over existing algorithms in handling extremely imbalanced data found in real-world scenarios.

The remainder of this paper is organized as follows. Section 2 provides a detailed description of our proposed algorithm, Fed-RNBS. Section 3 presents a comprehensive performance evaluation, comparing our proposed algorithm against several baseline models. Finally, Section 4 summarizes our findings and discusses future research directions.

II. PROPOSED METHOD

The overall framework of Fed-RNBS is illustrated in Fig. 1. Each client employs risk-adjusted bargaining to resolve inter-task gradient conflicts during local training, while the central server aggregates the model parameters. Under this coordination, each client (e.g., a financial institution) updates its model using a Non-IID and class-imbalanced dataset.

A. Problem Formulation: MTFL for Financial Tasks

The framework of our study assumes a standard federated learning setting, which consists of multiple clients and a single central server. Each client trains its model independently without sharing raw data externally. We assume that the data held by each client is statistically heterogeneous, exhibiting Non-IID properties, and that severe class imbalance exists within each task.

For multi-task learning, our model adopts a Shared Encoder and Task-specific Decoders architecture. The shared encoder is responsible for learning a common feature representation from the input data that is beneficial for all tasks. The feature vector processed by the encoder is then fed into separate decoders for each task to perform the final predictions. In this work, we define two critical binary classification tasks in financial risk management as follows:

Task 1: Delinquency Prediction. The objective of this task is to predict whether a loan will become delinquent for two or more months at any point during its entire term. It represents one of the most critical challenges in risk management for financial institutions and is considered a high-risk task due to the significant costs associated with prediction failures. The target variable is defined using the 'CURRENT LOAN DELINQUENCY STATUS' feature from the monthly performance file. This task is characterized by a severe class imbalance, with the positive class accounting for only approximately 1.96% of the entire dataset.

Task 2: Prepayment Prediction. The objective of this task is to predict whether a borrower will exhibit a prime repayment pattern, specifically by making an additional payment of $1,000 or more beyond the scheduled monthly installment. This information is crucial for financial institutions in managing liquidity and formulating product strategies. The target variable is defined by the 'EARLY REPAYMENT INDEX', which is calculated by combining data from the origination and monthly files. Compared to delinquency prediction, this is considered a relatively low-risk task.

B. The Challenge of Gradient Conflicts

When training a multi-task model with a shared encoder, each task aims to update the model parameters in a direction that most effectively minimizes its respective loss function. However, a gradient conflict arises when the preferred update directions—that is, the gradients—for each task diverge.

Conventional approaches, including our prior work, typically optimize a simple weighted sum of the task-specific losses. However, this method fails to adequately account for the complex interactions among tasks. It can lead to a phenomenon known as 'negative transfer,' where the gradient of one task interferes with the learning process of another, ultimately degrading the model's overall performance.

C. Proposed Algorithm: Fed-RNBS

In this paper, we frame the gradient conflict problem as a cooperative game where multiple tasks negotiate to find an optimal update direction. To solve this, we propose the Federated Risk-adjusted Nash Bargaining Solution (Fed-RNBS). Fed-RNBS comprises two key mechanisms: (1) dynamically calculating the bargaining power for each task at every training step to reflect its current state, and (2) an iterative Nash bargaining process that determines the optimal task weights based on the calculated bargaining power.

The Nash Bargaining Solution (NBS) [9] is a concept in cooperative game theory for finding a fair and efficient optimal agreement among players. However, the standard NBS assumes that all players have equal bargaining power, which makes it difficult to apply directly to financial tasks characterized by asymmetric failure costs.

To overcome this limitation, we dynamically adjust the bargaining power, denoted as β, as shown in Equation (1), considering both the current learning state and stability of each task. The bargaining power β_k, defined below, for each task k is determined by the product of the following three components:

$$\beta_k = (I_k \cdot L_{\text{norm},k}) \cdot \left(\frac{1}{\text{Var}(g_k) + \epsilon} \right). \quad (1)$$

1) Pre-defined Importance (I_k): This represents the static importance of a task, which is based on domain knowledge. For instance, in this study, we assign a higher I_k value to the delinquency prediction task (Task 1) due to its greater risk ($I_1 > I_2$). This component reflects the 'Risk-adjusted' aspect of our proposed methodology.
2) Normalized Loss (L_{norm}, k): This metric indicates the difficulty of a task at the current training stage. It functions as an adaptive learning mechanism by temporarily granting higher bargaining power to tasks with higher loss values—that is, those that are lagging and require more training. This encourages the model to focus on underperforming tasks.
3) Inverse of Gradient Variance ($\frac{1}{\text{Var}(g_k)+\epsilon}$): This term serves as a regularizer to control training stability. A large gradient variance ($Var(g_k)$) for a task within a given batch indicates that the training signal is unstable and noisy. By reducing the bargaining power of such tasks, this component prevents the model from updating in inconsistent or erratic directions, thereby enhancing overall training stability.

Based on the dynamically computed bargaining power vector $\beta = [\beta_1, \ldots, \beta_K]$, Fed-RNBS determines the final task weight vector $\alpha = [\alpha_1, \ldots, \alpha_K]$. Instead of approximating this process with a closed-form solution, we employ an iterative negotiation method, as detailed in Algorithm 1, to progressively find the optimal solution. This process repeats for a pre-defined number of iterations ($T_{negotiate}$). It defines an 'error' that measures how much the current weights α deviate from the ideal Nash equilibrium condition ($M \cdot \alpha \propto \beta/\alpha$), and slightly updates α in the direction of the steepest descent (the negative gradient) of this error. Through these incremental updates, the weight vector α converges to an optimal bargaining solution that satisfies the requirements of all tasks.

III. PERFORMANCE EVALUATION

In this section, we present the experimental setup, baseline models, and comprehensive performance analysis conducted to validate the effectiveness of our proposed algorithm, Fed-RNBS.

A. Experimental Setup

To accurately replicate a real-world financial environment, we utilized the Single-Family Loan-Level Dataset provided by the U.S. Government-Sponsored Enterprise (GSE), Freddie Mac [10]. This dataset is constructed based on actual residential mortgage loan information, ensuring a high degree of realism. Furthermore, it includes loan data acquired from various financial institutions, which is advantageous for securing data diversity and generalizability. We used approximately 180,000 loan records collected over about 21 months, starting from the first quarter of 2023. The data consists of an Origination file, containing information at the time of loan inception, and a Monthly Performance file, which includes monthly updated information.

Algorithm 1 Asymmetric N-Player Nash Bargaining Solver

```
Input: Set of task gradients {g_k}_{k=1}^K.
    Vector of bargaining powers β = [β_1, ..., β_K].
    Negotiation rounds T_neg, learning rate η.
Output: Final weight vector α = [α_1, ..., α_K].
Initialization:
G ← [g_1, g_2, ..., g_K]
M ← G^⊤ G
α ← [1/K, ..., 1/K]^⊤
Iterative Negotiation:
for iteration τ = 1, ..., T_neg do
        ▷ Measure disagreement from the ideal Nash condition
    error ← ||M · α − β/α||^2
        ▷ Calculate gradient of the disagreement
    grad_error ← ∇_α(error)
        ▷ Update weights to reduce disagreement
    Δ_α ← −grad_error
    α ← α + η · Δ_α
        ▷ Project weights back to the simplex (sum to 1)
    α ← α / Σ_{i=1}^K α_i
end for
return α
```

To simulate the real-world scenario where financial institutions possess distinct customer bases, we partitioned the entire training dataset among five clients in a non-uniform (Non-IID) manner using a Dirichlet distribution (α=1.0). Moreover, the persistent issue of class imbalance inherent in financial data was also reflected in our experiments. Table 1 shows the state of severe data imbalance for the two tasks in this study. For Task 1, which predicts the core risk, the proportion of the positive class is a mere 1.96%, making performance evaluation on the minority class a significant challenge.

TABLE I: Dataset Statistics and Class Distribution for Each Task

Statistic	Task 1	Task 2
Total Samples:	185,849	185,849
Negative Samples (Label 0)	182,214 (98.04%)	163,101 (87.76%)
Positive Samples (Label 1)	3,635 (1.96%)	22,748 (12.24%)

The model employed in our experiments follows the shared encoder and task-specific decoder architecture, as described in Section 2. The shared encoder consists of three fully connected layers with 256, 128, and 64 neurons, respectively, each followed by a ReLU activation function. Each task-specific decoder is composed of four hidden layers and an output layer with a Sigmoid function to compute the final binary classification probability. All experiments were conducted for 15 global rounds, with each client performing 3 epochs of local training in each round. To evaluate the performance of our proposed model, we conducted comparisons with the following baseline models.

B. Results and Analysis

Table 2 compares the personalized performance of our proposed Fed-RNBS with that of the baseline models for each task. The experimental results demonstrate that Fed-RNBS outperforms both baseline models on most key metrics, with the exception of Accuracy on Task 2. Furthermore, it proved to generate a much more robust model, particularly in the Non-IID environment and under conditions of severe data imbalance.

TABLE II: Performance Comparison of Fed-RNBS with Baseline Models on Financial Prediction Tasks

Task	Metric	MTFL	Nash-MTL	Fed-RNBS
Task 1 (High Risk)	Accuracy	0.801 ± 0.10	0.807 ± 0.11	$\mathbf{0.834 \pm 0.10}$
	F1-Score	0.266 ± 0.14	0.214 ± 0.22	$\mathbf{0.313 \pm 0.26}$
	Recall	0.411 ± 0.15	0.370 ± 0.26	$\mathbf{0.421 \pm 0.32}$
	Precision	0.204 ± 0.13	0.171 ± 0.19	$\mathbf{0.252 \pm 0.22}$
Task 2 (Low Risk)	Accuracy	$\mathbf{0.762 \pm 0.18}$	0.638 ± 0.10	0.752 ± 0.13
	F1-Score	0.337 ± 0.30	0.425 ± 0.08	$\mathbf{0.583 \pm 0.34}$
	Recall	0.367 ± 0.28	0.406 ± 0.05	$\mathbf{0.553 \pm 0.31}$
	Precision	0.353 ± 0.36	0.512 ± 0.24	$\mathbf{0.626 \pm 0.37}$

Given the severe class imbalance inherent in financial datasets, standard Accuracy can be misleading as it often reflects the correct classification of the majority class rather than the detection of risks. Therefore, we prioritize Recall and F1-Score as the primary indicators of actual risk detection capability. On these critical metrics, Fed-RNBS demonstrated significant gains, achieving a 26.4% improvement in Recall and a 40.2% improvement in F1-Score compared to Nash-MTL. These figures represent the average performance improvement calculated across both tasks. This substantial gain is attributed to the dynamic bargaining power mechanism of Fed-RNBS, which effectively prevents the model from being biased toward the majority class.

The superiority in F1-Score clearly highlights the key advantages of the proposed model. The significant performance gap in Task 1, in particular, demonstrates that the approach of Fed-RNBS—which goes beyond merely mitigating gradient conflicts to concurrently consider asymmetric task importance and training stability—is highly effective in real-world, imbalanced financial environments. Furthermore, achieving the top performance on the relatively low-risk Task 2 confirms that our model finds a stable trade-off, maximizing performance on the high-risk task without sacrificing the performance of the other.

IV. CONCLUSION

In this paper, we aimed to address the challenges that arise when applying Multi-Task Federated Learning (MTFL) in realistic financial environments, where Non-IID data distributions, extreme class imbalance, and asymmetric risks among tasks coexist. To this end, we proposed a novel game-theoretic algorithm, the Federated Risk-adjusted Nash Bargaining Solution (Fed-RNBS). Fed-RNBS dynamically computes the bargaining power for each task by holistically considering its predefined importance, current learning difficulty, and gradient stability. By applying this power within an iterative Nash bargaining process, it effectively resolves gradient conflicts between tasks. Furthermore, through experiments on a real-world residential mortgage dataset, we demonstrated that the proposed Fed-RNBS outperforms conventional MTFL and symmetric Nash bargaining-based models. It achieved a particularly significant performance improvement on the high-risk and severely imbalanced delinquency prediction task, confirming that our proposed methodology operates effectively and robustly in challenging real-world financial data environments.

ACKNOWLEDGMENTS

This research was supported by MSIT Korea under NRF Korea (RS-2026-00557379, 60%). This research was also supported by the Korea Institute for Advancement of Technology (KIAT) grant funded by the Korean government (MOTIE) (P0017123, The Competency Development Program for Industry Specialist, 10%) and the Convergence Security Core Talent Training Business Support Program (IITP-2026-RS-2024-00426853, 30%) supervised by IITP.

REFERENCES

[1] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, ”Communication-Efficient Learning of Deep Networks from Decentralized Data,” in Proc. Int. Conf. Artif. Intell. Stat. (AISTATS), Fort Lauderdale, FL, USA, Apr. 2017, pp. 1273-1282.

[2] S. Ruder, ”An Overview of Multi-Task Learning in Deep Neural Networks,” arXiv preprint arXiv:1706.05098, 2017.

[3] T. Liu, Z. Wang, H. He, W. Shi, L. Lin, R. An, and C. Li, ”Efficient and Secure Federated Learning for Financial Applications,” Applied Sciences, vol. 13, no. 10, p. 5877, 2023.

[4] J. Bang, S. Woo, J. Lee, “DRL-Empowered Joint Batch Size and Weighted Aggregation Adjustment Mechanism for Federated Learning on Non-IID Data,” ICT Express, Volume 10, no. 4, pp. 863-870, August 2024.

[5] Subeen Choi and Bong Jun Choi, ”Loan Risk Prediction via Multi-Task Federated Learning in a Distributed Environment,” in Proceedings of the Korean Information Science Society Conference, 2025, pp. 1237-1239.

[6] D. Fudenberg and J. Tirole, Game Theory. Cambridge, MA, USA: MIT Press, 1991.

[7] A. Navon, A. Shamsian, I. Achituve, H. Maron, K. Kawaguchi, G. Chechik, and E. Fetaya, ”Multi-task learning as a bargaining game,” arXiv preprint arXiv:2202.01017, 2022.

[8] A. Shamsian, A. Navon, N. Glazer, K. Kawaguchi, G. Chechik, and E. Fetaya, ”Auxiliary learning as an asymmetric bargaining game,” in Proc. Int. Conf. Mach. Learn. (ICML), 2023, pp. 30689-30705.

[9] K. Binmore, A. Rubinstein, and A. Wolinsky, ”The Nash bargaining solution in economic modelling,” The RAND Journal of Economics, vol. 17, no. 2, pp. 176-188, 1986.

[10] Freddie Mac, Single-Family Loan-Level Dataset, 2025. [Online]. Available: https://www.freddiemac.com/research/datasets/sf-loanlevel-dataset. [Accessed: Aug. 24, 2025].