When Federated Learning Meets Oligopoly Competition: Stability and Model Differentiation

Chao Huang, Justin Dachille, and Xin Liu, *Fellow, IEEE*

***Abstract*—Federated learning (FL) is decentralized machine learning framework that finds various applications in health, finance, and the Internet of things. This article studies the under-explored business competition in FL, where organizations are both collaborators in training a shared model and competitors in providing model-based services to a continuum of customers. We focus on an oligopoly case with three organizations. To understand how competition affects FL collaboration, we start with a benchmark case where organizations are not competitors, and show that they have an incentive to collaborate. However, in the presence of competition, organizations may prefer to train local models instead of collaborating via FL (even if FL incurs zero training costs). The reason is that FL intensifies price competition by improving organizations' model performance to a similar level. To address this issue, we devise a model differentiation mechanism in which organizations adaptively adjust their model performance, enabling differentiated model-based services to customers. We prove that the adaptive mechanism converges in polynomial time and is incentive compatible. Perhaps surprisingly, numerical experiments on CIFAR-10 show that the mechanism can simultaneously improve the model performance, organizations' revenues, and social welfare. The improvement is up to 22.31%, 14.42%, and 19.50%, respectively.**

***Index Terms*—Artificial intelligence, business competition, federated learning (FL), game theory, machine learning, mechanism design.**

I. INTRODUCTION

A. Motivations

FEDERATED learning (FL) is a popular machine learning paradigm that enables decentralized, collaborative model training across multiple participating organizations or devices while keeping the data localized [1]. Unlike traditional centralized approaches where raw data is pooled together (e.g., into a single server) for model training, FL allows each participant to train the model on their local data set and share only the model updates, thereby preserving data privacy. More specifically, FL typically consists of several steps.

1) *Local Training:* Each participant trains a shared model on their own data set.

Manuscript received 30 March 2024; accepted 4 May 2024. Date of publication 13 May 2024; date of current version 8 August 2024. This work was supported in part by United States Department of Agriculture under Grant 2020-67021-32855, and in part by NSF under Grant OIA-2134901. *(Corresponding author: Chao Huang.)*

The authors are with the Department of Computer Science, The University of California at Davis, Davis, CA 95616 USA (e-mail: fchhuang@ucdavis.edu; jtdachille@ucdavis.edu; xinliu@ucdavis.edu).

This article has supplementary downloadable material available at https://doi.org/10.1109/JIOT.2024.3400216, provided by the authors.

Digital Object Identifier 10.1109/JIOT.2024.3400216

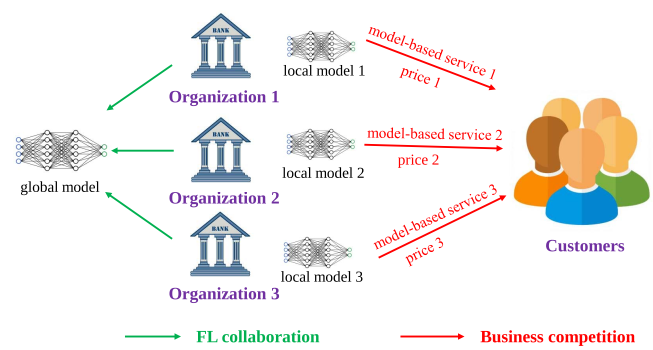

Fig. 1. FL collaboration and business competition. The organizations collaborate in FL training, and each can obtain an improved local model. Meanwhile, the organizations may also compete in selling model-based services at different prices to customers.

2) *Model Update Sharing:* Instead of sending raw data, participants share only the model updates, such as weights or gradients, with a coordinating server or among themselves.
3) *Aggregation:* A server or some designated device aggregates the model updates to generate a new shared model.
4) *Distribution:* The updated shared model is then distributed back to the participants for further local training.

The above iterative steps stop once the shared model converges, or the overall time (e.g., for training and communication) exceeds a predefined threshold.

Prior FL studies focused on algorithm development to improve training performance in the presence of data and system heterogeneity, which makes FL more promising in various applications, such as healthcare, finance, and the Internet of Things [2], [3]. However, the interaction between FL and business competition is not well understood, which is the focus of this article. To be more concrete, participating organizations, although aligned in developing a shared model, are usually direct competitors in the market for model-based services [2], [4]. We discuss some potential examples as follows (see also Fig. 1).

1) *Healthcare [5]:* Multiple healthcare organizations might collaborate to develop a more accurate diagnostic model via FL, but they also compete for patients who seek high-quality diagnostic services.
2) *Finance [6]:* Several reinsurance firms could collaboratively train risk assessment models that better predict large-scale financial losses due to catastrophic events. However, they may also compete in underwriting contracts for customers who seek to mitigate risk exposure.

3) *Retail [7]:* A group of retail companies could jointly train a customer behavior model to understand purchasing patterns better, yet fiercely compete in offering personalized marketing or discounts based on the insights drawn from the shared FL model.

The above examples motivate our first question below.

Question 1: How does business competition affect organizations' FL collaboration?

We focus on an oligopoly setting with three organizations and a continuum of distributed customers, which provides a manageable yet rich setting to capture the complexity of collaborative and competitive dynamics. Note that our analysis can be generalized to where there are three groups of service providers and each group may contain several organizations.[1] Each organization aims to maximize its own revenue via strategically forming FL training coalitions and deciding market prices for providing model-based services to the customers. We use the concept of *core stability* from coalitional game theory to analyze how organizations form stable FL coalitions [8]. Core stability provides a theoretical framework for evaluating the sustainability/stability of collaborative arrangements, particularly in settings where competing incentives among organizations can undermine coalition formation. To answer Question 1, we study two cases.

1) *Noncompetitive:* Organizations are not competitors and each has an exclusive pool of customers.
2) *Competitive:* Organizations compete for the same pool of customers.

For the noncompetitive case, not surprisingly, organizations have an incentive to collaborate via FL. This is because FL collaboration leads to improved model performance, which in turn, drives higher revenues as customers in their exclusive markets are more likely to procure enhanced model-based services. For the competitive case, however, organizations vying for the same pool of customers tend to avoid FL collaboration despite its potential to improve model performance. The reason is that FL harmonizes model performance among the organizations, which intensifies price competition facing the same pool of customers. This can lead to a lower revenue for certain organizations and hence they tend to train local models instead of collaborating via FL.

The above results unfortunately imply that business competition serves as a barrier in organizations' willingness to engage in FL collaborations. This motivates our second question.

Question 2: How to encourage FL collaboration among competing organizations?

To answer Question 2, we propose a model differentiation mechanism, in which organizations adaptively modify the performance of their local models when providing services to customers. In this mechanism, each organization can apply alterations to its local model, e.g., via customized noise or reduced model size. The goal is to enable each organization to offer model-based services that possess differentiated qualities, thereby reducing competition and enhancing revenues. Note that prior FL mechanisms typically require monetary transfer among the central server and the organizations. Our proposed model differentiation mechanism differs in that it does not require any monetary transfer among organizations, which is easier to implement in practice.

While it might be intuitive to think that such model differentiation could harm customers by reducing competition and potentially raising prices, our findings suggest a more nuanced picture. Interestingly, we show that the differentiation mechanism can enhance both organizational revenues and social welfare (defined as the summation of organizations' revenues and consumer surplus). Although model differentiation may result in higher prices, it incentivizes organizations to collaborate through FL, leading to the provision of higher quality services. This creates a win-win situation, where organizations generate higher revenues and customers receive more value, ultimately enhancing overall social welfare.

B. Key Contributions

We summarize the key contributions of this article as follows.

1) *Oligopoly Competition in FL:* To our best knowledge, this work is among the first studies on oligopoly competition in FL. Our work offers implications for the overlooked organization strategy and social welfare in competitive FL systems.
2) *Stability Analysis:* We analyze the stability of FL coalitions. In the absence of competition, we show that grand coalition (all organizations forming an FL coalition) is core stable. In the presence of competition, however, separation (no organizations performing FL) tends to be core stable.
3) *Model Differentiation Mechanism:* We propose a model differentiation mechanism to encourage FL collaboration in the presence of competition. The mechanism incentivizes FL collaboration without requiring monetary transfer among organizations, and it is proven to be incentive compatible.
4) *Numerical Experiments:* We conduct numerical experiments on CIFAR-10. We show that compared to the case without model differentiation, our mechanism can significantly improve the model performance, organizations' revenues, as well as the social welfare. The improvement is up to 22.31%, 14.42%, and 19.50%, respectively.

The remainder of this article is organized as follows. Section II reviews related work and Section III introduces the system model. Section IV analyzes the noncompetitive case, while Section V studies the competitive case and further answers Question 1. Section VI presents the model differentiation mechanism and answers Question 2. Section VII provides numerical results. Section VIII discusses the model extension and Section IX concludes this article.

II. Related Work

A. Coalition in FL

There are some excellent recent studies on the stability of coalitions in FL systems, e.g., [9], [10], [11], [12], [13],

[1] We discuss the extension to more than three organizations in Section VIII.

[14], [15]. For example, Donahue and Kleinberg [9], and [10] gave a comprehensive stability analysis of FL coalitions where participants aim to minimize model loss. Jiang and Wu [11] considered a tradeoff between model performance and communication/computation costs and presented a merge-and-split algorithm to find the stable coalition. Bao et al. [15] studied coalition structures that alleviate negative transfers among clients. Ray Chaudhury et al. [12] incorporated the notion of fairness into core stability and analyzed the corresponding FL coalition.

Our work differs from this stream of studies in that none of them studied business competition in FL. The consideration of business competition in FL is practically important but also technically challenging. The main reason is that organizations need to simultaneously optimize its FL coalition and prices facing the same base of strategic customers.

B. Incentive Design in FL

There has been an increasing volume of studies on the incentive mechanism design for FL. Refer to [16], [17], and [4] for a few excellent surveys. To name a few, [18], [19], [20] designed contract mechanisms to encourage data sharing and training participation in FL. Lu et al. [21] used auctions to appropriately select clients for model training. Luo et al. [22] devised a Stackelberg game approach to incentivize client participation.

Our work differs from prior incentive studies in that prior studies focused on monetary incentives, which may be unviable for competing organizations in practice. Instead, we propose a model differentiation mechanism that does not require any monetary payment. The mechanism design, however, is highly nontrivial due to clients' strategic decisions on FL learning collaboration.

C. Business Competition in FL

Business competition is a critical component in practical FL systems and until recently there are several related papers, e.g., [2], [6], [23], [24], [25]. Wu and Yu [6] studied a fully competitive market where organizations are price takers, and hence did not model the important business strategy related to price design. Wang et al. [23] introduced a mechanism that aims to maximize the overall system profit with the help of a coordinator. Huang et al. [2] analyzed a duopoly case where two organizations aim to maximize their own profit. Sun et al. [24] discussed the implications of competition in an electric vehicle market. However, the studies in [23] and [2] used monetary transfer among organizations to incentivize collaboration, which may not be viable in practice. Instead, our work used a model differentiation mechanism that does not require any monetary transfer.

The model differentiation mechanism in our work is motivated by the vertical differentiation literature in economics, e.g., [26], [27], [28]. Vertical differentiation refers to variations in services and qualities that distinguish higher end offerings from lower end alternatives in a competitive market. We note that the analysis of model differentiation in FL is more challenging than prior vertical differentiation studies. That is, prior work typically assumed that organizations independently adjust the qualities without collaboratively improving service qualities. In FL, however, organizations need to jointly consider the collaborative training coalition, which in turn affects the feasible space of all organizations' model differentiation.

III. System Model

In Section III-A, we first introduce the FL process among organizations. In Sections III-B and III-C, we define the strategies and objectives functions of the customers and organizations, respectively. In Section III-D, we formulate the game-theoretical interactions between customers and organizations.

A. Federated Learning Process

We consider a set $\mathcal{N} = \{1, 2, 3\}$ of three organizations who aim to collaboratively train a global model without exchanging raw data. Each organization possesses a private local data set $\mathcal{D}_n$ with size $D_n = |\mathcal{D}_n|$. The organizations can decide whether to participate in FL training, and we use $\mathcal{S} \subseteq \mathcal{N}$ to denote the set of participating organizations. For example, $\mathcal{S} = \{1, 2\}$ means that organizations 1 and 2 collaborate in FL training, while organization 3 does not participate in FL and instead trains a local model itself.

If an organization n does not participate in FL, it trains a local model using $\mathcal{D}_n$ and has a model accuracy denoted as $A_n^l \in [0, 1]$. If organizations in $\mathcal{S}$ participate in FL, the FL process consists of two phases: 1) global training and 2) fine-tuning, which we discuss as follows.

1) Global Training: Consider a global model parameterized by $\boldsymbol{w}$. We use $f_n(\boldsymbol{w}; \zeta_n)$ to denote the loss of model $\boldsymbol{w}$ over organization n's mini-batch instance ζ_n, which is randomly sampled from $\mathcal{D}_n$. Let $f_n(\boldsymbol{w}) \triangleq \mathbb{E}_{\zeta_n \sim \mathcal{D}_n}[f_n(\boldsymbol{w}; \zeta_n)]$ denote the expected loss of model $\boldsymbol{w}$ over organization n's data set. The goal of FL is to minimize the expected loss of the model $\boldsymbol{w}$ over the data sets of all participating clients

$$\min_{\boldsymbol{w}} F(\boldsymbol{w}) \triangleq \sum_{n \in S} a_n f_n(\boldsymbol{w}) \tag{1}$$

where coefficient $a_n \geq 0$ denotes the weight assigned to organization n's data set and $\sum_{n \in \mathcal{S}} a_n = 1$. We typically have $a_n = D_n / (\sum_{s \in \mathcal{S}} D_s)$.

To derive the optimal weights $\boldsymbol{w}^* \triangleq \arg\min_{\boldsymbol{w}} F(\boldsymbol{w})$, FL proceeds in multiple training rounds. In each round r, the organizations execute the following steps.

1) Each organization n downloads the global model $\boldsymbol{w}^{\boldsymbol{r-1}}$ obtained from the last training round.
2) Each organization n trains the model $\boldsymbol{w}^{\boldsymbol{r-1}}$ multiple times, and each time it uses a mini-batch instance ζ_n randomly sampled from $\mathcal{D}_n$.
3) Each organization n sends the model updates $\boldsymbol{\omega}_{\boldsymbol{n}}^{\boldsymbol{r}}$ to a (trusted) server for synchronization, who produces an updated global model $\boldsymbol{w}^{\boldsymbol{r}}$ to be downloaded in the following round. One widely used synchronization algorithm is FedAvg [31], [32], where $\boldsymbol{\omega}^{\boldsymbol{r}} = \sum_{n \in \mathcal{S}} a_n \boldsymbol{w}_{\boldsymbol{n}}^{\boldsymbol{r}}$.

The above iterative training stops when the global model converges. We use $A^{\mathcal{S}} \in [0, 1]$ to denote the global model accuracy with participating organizations in $\mathcal{S}$.

TABLE I
COMPARISON WITH RELATED WORK

References	[9]–[14]	[18]–[20]	[6], [23]–[25]	[29], [30]	**This work**
Data heterogeneity	✓	×	✓	×	✓
Voluntary participation	✓	✓	✓	×	✓
Business competition	×	×	✓	×	✓
Monetary incentive	✓	✓	✓	✓	×
Non-monetary incentive	×	×	×	×	✓
Social welfare analysis	×	×	×	×	✓

2) Fine-Tuning: Organization $n \in \mathcal{S}$ may further fine-tune the global model post-convergence by retraining some or all parameters using local data [33]. Fine-tuning can enhance the global model accuracy, particularly when organizations hold nonidentically and independently (non-IID) data. The main reason is that when data across organizations are non-IID, the global model's performance can suffer on each organization's specific data distribution, and fine-tuning is a promising way to rectify this and enhance performance. After fine-tuning is completed, each organization $n \in \mathcal{S}$ obtains a personalized local model, and we use $A_n^{\mathcal{S}} \in [0, 1]$ to indicate the accuracy.

3) Model-Based Service: Organizations use their models to generate model-based services, such as loan predictions by banks or disease diagnoses by hospitals [34]. For ease of presentation, we use $A_n(\mathcal{S})$ to denote organization n's final model accuracy, and

$$A_n(\mathcal{S}) = \begin{cases} A_n^l, & \text{if} \quad n \notin \mathcal{S} \\ A_n^{\mathcal{S}}, & \text{if} \quad n \in \mathcal{S}. \end{cases} \tag{2}$$

Specifically, if organization n does not participate in FL (i.e., $n \notin \mathcal{S}$), it uses locally trained model A_n^l to generate service. If organization n participates in FL (i.e., $n \in \mathcal{S}$), it uses fine-tuned global model $A_n^{\mathcal{S}}$ to generate service. Then, organizations enter the market competition to sell these services to prospective customers. We next model the strategies and objective functions for both customers and organizations.

B. Customer Strategy and Objective

1) Customer Valuation Type: We consider a continuum population of customers with size normalized to 1. Each individual's valuation of the model-based service is represented by θ [35], e.g., willingness-to-pay. A larger value of θ means a greater utility derived from the service. Valuations vary across the customer population and are modeled as a random variable θ with a probability density function (PDF) $h(\theta)$ and a cumulative distribution function (CDF) $H(\theta)$ on the support $[0, \theta_{\max}]$. While individual valuations are unknown, their distribution (i.e., $h(\theta)$, $H(\theta)$, and $\theta_{\max}$) is assumed to be known to the organizations, often through market research [36].

2) Consumer Strategy: Customers face choices among services of different qualities and prices offered by competing organizations. For a customer with valuation θ, the decision to purchase is denoted as $d_\theta = n \in \{1, 2, 3\}$, where n corresponds to organization n. Here, we assume organizations offer substitutable services (e.g., insurance or medical diagnosis), and a customer will buy from only one organization.

3) Consumer Payoff: The payoff of a customer is defined as the utility derived from the service minus the unit price paid. If a customer chooses a service from organization n, the payoff (i.e., objective) function is

$$u_\theta(d_\theta; \mathcal{S}, \boldsymbol{p}) = \theta \cdot A_n(\mathcal{S}) - p_n \ \forall \ d_\theta = n \in \{1, 2, 3\} \tag{3}$$

where $\boldsymbol{p} \triangleq \{p_1, p_2, p_3\}$. A higher $A_n(\mathcal{S})$ is associated with a better service quality, resulting in a higher utility $\theta \cdot A_n(\mathcal{S})$ for the customer. Here, we assume that a customer's utility is a linear function in $A_n(\mathcal{S})$ [35]. We will discuss the use of a more general utility function in Section VIII. Information about the service quality $A_n(\mathcal{S})$ from different organizations can be obtained through feedback systems [37], [38].

C. Organizations' Decisions and Revenues

1) Organizations FL Coalition and Pricing Strategies: The organizations first need to decide the FL coalition strategy represented by $\mathcal{S} \subseteq \mathcal{N}$. In addition, each organization n decides the unit price $p_n \geq 0$ for providing model-based service to a customer. A customer needs to pay p_n if it purchases a service from organization n.

2) Organization Revenue: An organization n's revenue from an individual customer is the charged price p_n. Therefore, the total revenue $R_n(\mathcal{S}, \boldsymbol{p})$ organization n receives from the entire customer pool is

$$R_n(\mathcal{S}, \boldsymbol{p}) = \int_0^{\theta_{\max}} p_n \cdot \mathbb{1}_{d_\theta = n}(\mathcal{S}, \boldsymbol{p}) \cdot h(\theta) d\theta \tag{4}$$

where $\mathbb{1}$ is an indicator, i.e., $\mathbb{1}_{d_\theta = n} = 1$ if and only if $d_\theta = n$.[2]

Note that while organizations participating in FL incur communication costs (for model transmission) and computation costs (for model training), we normalize them to be zero [2], [39]. This is because organizations usually have strong computational capabilities, such as high-performance servers, and dependable communication infrastructure, such as high-speed wired networks. Nevertheless, one can easily extend our analysis by subtracting a cost term in (4).

D. Three-Stage Game Formulation

We formulate the interactions between organization and customers as a three-stage game. In Stage I, organizations choose the FL coalition strategy $\mathcal{S}$, followed by the unit prices $\boldsymbol{p}$ in Stage II. Each organization seeks to optimize its revenue in (4). In Stage III, customers decide their purchasing strategies to maximize their individual payoffs in (3). The game is analyzed using backward induction.

[2]In this work, we normalize the unit service cost when serving each customer to be zero. Our analysis and conclusions can be easily extended to where such cost is a nonzero constant.

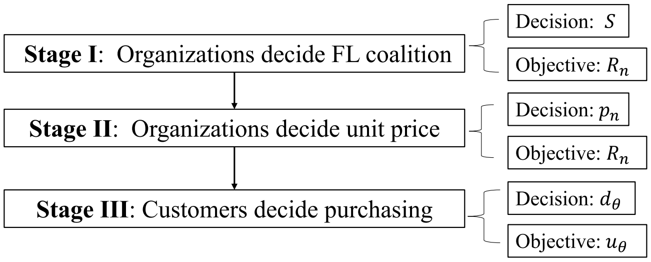

Fig. 2. Three-stage Stackelberg game.

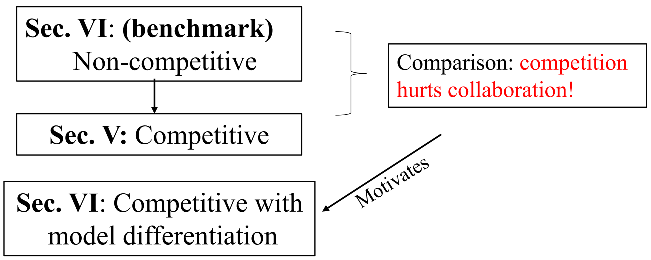

Fig. 3. Functional diagram connecting Sections IV–VI.

TABLE II
KEY NOTATIONS

Decision Variables	
d_θ	purchasing strategy of type-θ customer
$\boldsymbol{p} = \{p_1, p_2, p_3\}$	pricing strategies of organizations
$\mathcal{S}$	FL coalition strategy of organizations
Parameters	
D_n	size of organization n's local data set
θ	customer type (i.e., valuation)
$h(\cdot)$, $H(\cdot)$	PDF and CDF of customer type
Functions	
$u_\theta(d_\theta; \mathcal{S}, \boldsymbol{p})$	payoff of type-θ customer
$A_n(\mathcal{S})$	organization n's final model accuracy
$R_n(\mathcal{S}, \boldsymbol{p})$	revenue of organization n

We summarize the key notations in Table II, and further include a diagram in Fig. 3 showing the connections among the following sections.

IV. Noncompetitive Case

We start with solving the noncompetitive case, which constitutes a crucial component to answer Question 1, i.e., how does business competition affect organizations' FL collaboration? We first discuss the modeling for the noncompetitive case, and then present the organizations' optimal prices and FL coalitions.

A. Model for Noncompetitive Case

In the noncompetitive case, we consider that each organization n has an exclusive customer pool (i.e., market n) represented as a continuum of size $1/3$ with PDF $h(\theta)$ and CDF $H(\theta)$ on support $[0, \theta_{\max}]$. A customer within market n can only purchase service from organization n. We define $d_{\theta,n}(\mathcal{S}, p_n) \in \{n, 0\}$ as customer θ's purchase decision in market n, where $d_{\theta,n}(\mathcal{S}, p_n) = n$ means purchasing, and $d_{\theta,n}(\mathcal{S}, p_n) = 0$ means not purchasing. It is easy to show that the optimal decision is

$$d^*_{\theta,n}(\mathcal{S}, p_n) = n \cdot \mathbb{1}_{\theta A_n(\mathcal{S}) - p_n \geq 0}. \tag{5}$$

B. Organizations' Optimal Pricing and FL Coalition

1) Organization Optimal Pricing: We first analyze the optimal pricing in Stage II. Based on (5), we can write each organization n's revenue $R_n^{\mathrm{NC}}(\mathcal{S}, p_n)$ (in the non-competitive case) as follows:

$$R_n^{\mathrm{NC}}(\mathcal{S}, p_n) = \frac{p_n}{3} \cdot \left[1 - H\left(\frac{p_n}{A_n(\mathcal{S})}\right)\right]. \tag{6}$$

An arbitrary choice of $H(\cdot)$ (or equivalently $h(\theta)$) can render the optimization of $R_n^{\mathrm{NC}}(\mathcal{S}, p_n)$ over p_n intractable. To tackle this, we make an assumption below.

Assumption 1: The customers' valuation follows a uniform distribution, i.e., $h(\theta) = 1/\theta_{\max}\ \ \forall \theta \in [0, \theta_{\max}]$.

We use a uniform distribution mainly for analytical tractability [6], [35], [40]. We will discuss how to relax this assumption in Section VIII. We will also use other distributions (e.g., truncated normal distribution) in the numerical experiments in Section VII.

Next, we present the organizations' optimal price $p_n^{\mathrm{NC}*}$ in the noncompetitive case in Proposition 1.

Proposition 1: Under Assumption 1, the optimal price of each organization $n \in \{1, 2, 3\}$ in the noncompetitive case is

$$p_n^{\mathrm{NC}*}(\mathcal{S}) = \frac{\theta_{\max}}{2} A_n(\mathcal{S}). \tag{7}$$

As a result, organization n's revenue is

$$R_n^{\mathrm{NC}}(\mathcal{S}) = \frac{\theta_{\max}}{12} A_n(\mathcal{S}). \tag{8}$$

Due to space limitation, we defer all the technical proofs to the Appendix in the supporting document.

Proposition 1 implies that if organization n has a higher model accuracy, it will set a larger price to optimize the revenue. Importantly, in the noncompetitive case, each organization's optimal price only depends on its own model accuracy $A_n(\mathcal{S})$. We will show next that this incentivizes organizations to form FL coalitions.

2) Organization Optimal FL Coalition: Even if the organizations do not compete in selling model-based services, they still interact in a game-theoretical fashion in terms of FL coalition. We model the interactions among organizations as an FL coalition game below.

Game 1: The FL coalition game among organizations is as follows.

1) *Player:* Organization $n \in \{1, 2, 3\}$.
2) *Strategy:* Organizations jointly decide a coalition $\mathcal{S} \subseteq \mathcal{N}$.
3) *Objective:* Organization n's revenue in (8).

Note that the definition in Game 1 differs from prior coalitional games where the organizations are allowed to share the benefits (e.g., revenues) associated with each possible coalition [41]. The main reason is that in the context of competitive FL (to be shown in Section V), revenue sharing may cause a collusion issue which is prohibited in many countries. Instead, we consider where organizations can collaboratively decide FL coalitions but cannot share revenues obtained from customers.

We apply core stability as the solution concept to Game 1. Before defining core stability, we differentiate two terms, i.e., coalition and coalition structure, as follows.

Definition 1:

1) A coalition S is a subset of $\mathcal{N}$, i.e., $S \subseteq \mathcal{N}$.
2) A coalition structure Π is a partition of $\mathcal{N}$, e.g., Π can be $\{1, 2\}, \{3\}$ containing two coalitions $\{1, 2\}$ and $\{3\}$. Moreover, a) the grand coalition Π^G is the coalition structure that contains a single coalition $\mathcal{N}$ and b) the separation Π^R is the coalition structure that contains three coalitions $\{1\}$, $\{2\}$, and $\{3\}$.

Now we are ready to define core stability.

Definition 2: A coalition structure Π is core stable if there does not exist a coalition $\mathcal{S} \notin \Pi$ so that each organization in $\mathcal{S}$ obtains a larger revenue in $\mathcal{S}$ than that obtained in its current coalition in Π.

Note that without any assumptions on $A_n(\mathcal{S})$, it is difficult to analyze the core stable coalition structure. To this end, we make a minor assumption.

Assumption 2: For each n, $A_n(\cdot)$ is monotonic. That is, for any $\mathcal{S}_1 \subseteq \mathcal{S}_2$, we have $A_n(\mathcal{S}_1) \le A_n(\mathcal{S}_2)$.

If $n \notin \mathcal{S}_2$ (and hence $n \notin \mathcal{S}_1$), Assumption 2 holds with equality. If $n \in \mathcal{S}_1$ (and hence $n \in \mathcal{S}_2$), Assumption 2 means that participating organizations achieve a larger accuracy with more FL collaborators. If $n \notin \mathcal{S}_1$ but $n \in \mathcal{S}_2$, Assumption 2 means that an organization achieves a larger accuracy from participating in FL than training a local model. We will show in Section VII that the numerical results are consistent with Assumption 2.

Next, we present the core stable coalition structure below.

Theorem 1: Under Assumptions 1 and 2, the grand coalition Π^G is a core stable coalition structure of Game 1.

Theorem 1 implies that when organizations are not competitors, they have an incentive to form FL coalitions. Forming a larger FL coalition leads to a better model, and organizations can set a higher price to maximize their revenues (see Proposition 1). This demonstrates the benefits of FL. However, we will show in the next section that even if FL improves all organizations' model accuracies, they tend to not collaborate in the presence of business competition.

V. Competitive Case

In this section, we solve the competitive case where organizations face the same pool of customers.

A. Stage III—Customer Purchasing

Lemma 1 computes the customer's optimal purchasing strategies, given the FL coalition $\mathcal{S}$ and the prices $\boldsymbol{p}$.

Lemma 1: Given $\mathcal{S}$ and $\boldsymbol{p}$, a type-θ customer's optimal purchasing $d^*_\theta(\mathcal{S}, \boldsymbol{p})$ is

$$d^*_\theta(\mathcal{S}, \boldsymbol{p}) = \arg\max_n(\theta A_n(\mathcal{S}) - p_n). \tag{9}$$

Lemma 1 follows directly from (3). It shows that in a competitive oligopoly market, a type-θ customer will purchase service from organization n if organization n brings the customer the largest payoff.

B. Stage II—Organization Pricing

In Stage II, given FL coalition $\mathcal{S}$, each organization chooses the pricing p_n to optimize its own revenue in (4), anticipating the customers' optimal purchasing behaviors.

For ease of presentation, we assume without loss of generality that $A_1(\mathcal{S}) > A_2(\mathcal{S}) > A_3(\mathcal{S})$. Based on this, we reindex the organizations in which organization 1 has the best model and organization 3 has the worse model. Next, we define the neutral customer type, which will be useful in the Stage II analysis.

Definition 3: Denote $\sigma_{m,n}$ as the neutral customer type. A type-$\sigma_{m,n}$ customer obtains the same payoff by purchasing service from either organization m or organization n, i.e., $u_{\sigma_{m,n}}(m; \mathcal{S}, \boldsymbol{p}) = u_{\sigma_{m,n}}(n; \mathcal{S}, \boldsymbol{p})$. We can derive $\sigma_{m,n}$ below

$$\sigma_{m,n}(\mathcal{S}, \boldsymbol{p}) = \frac{p_m - p_n}{A_m(\mathcal{S}) - A_n(\mathcal{S})}. \tag{10}$$

Based on Definition 3, we can write organizations' revenues as a function of $\boldsymbol{p}$ as follows:

$$\begin{cases} R_1(\mathcal{S}, \boldsymbol{p}) = p_1\left[1 - H\left(\max\{\sigma_{1,2}(\mathcal{S}, \boldsymbol{p}), \sigma_{1,3}(\mathcal{S}, \boldsymbol{p})\}\right)\right] \\ R_2(\mathcal{S}, \boldsymbol{p}) = p_2 H\left(\sigma_{1,2}(\mathcal{S}, \boldsymbol{p}) - \sigma_{2,3}(\mathcal{S}, \boldsymbol{p})\right) \\ R_3(\mathcal{S}, \boldsymbol{p}) = p_3 H(\min\{\sigma_{2,3}(\mathcal{S}, \boldsymbol{p}), \sigma_{1,3}(\mathcal{S}, \boldsymbol{p})\}. \end{cases} \tag{11}$$

We formulate the price competition game in Stage II as follows.

Game 2: Given $\mathcal{S}$, the three organizations' pricing competition in Stage II can be modeled as the following game.

1) *Player:* Organization n for $n \in \{1, 2, 3\}$.
2) *Strategy:* Each organization n decides $p_n \ge 0$.
3) *Objective:* Each organization n receives a revenue in (11).

We aim to solve the Nash equilibrium (NE) of Game 2.

Definition 4: Given $\mathcal{S}$, a profile $\boldsymbol{p}^*(\mathcal{S}) \triangleq \{p^*_n(\mathcal{S})\}_{n\in\mathcal{N}}$ constitutes an NE of Game 2 if for all $n \in \{1, 2, 3\}$

$$R_n\left(p^*_n(\mathcal{S}), \boldsymbol{p}^*_{-n}(\mathcal{S})\right) \ge R_n\left(p'_n(\mathcal{S}), \boldsymbol{p}^*_{-n}(\mathcal{S})\right) \quad \forall p'_n(\mathcal{S}) \ne p^*_n(\mathcal{S}) \tag{12}$$

where $\boldsymbol{p}^*_{-n}(\mathcal{S}) = \{p^*_m(\mathcal{S})\}_{m\in\mathcal{N}\setminus\{n\}}$.

At an NE, each organization's price is a best response to the price set by other organizations, i.e., the equilibrium is the fixed point of all organizations' best response prices [8].

Solving the NE of Game 2 can be challenging, as the max and min functions in (11) make $R_n(\mathcal{S}, \boldsymbol{p})$ nondifferentiable in p_n. To address this issue, we first show that $R_n(\mathcal{S}, \boldsymbol{p})$ is piece-wise concave. Then, we decompose $R_n(\mathcal{S}, \boldsymbol{p})$ into several concave segments, compute the optimal solution in each segment, and compare across multiple segments to obtain the global optimal solution.

We present the price equilibrium in Proposition 2.

Proposition 2: Under Assumption 1, the NE of Game 2 is

$$\begin{cases} p^*_1(\mathcal{S}) = \frac{(A_1(\mathcal{S}) - A_2(\mathcal{S}))(3A_1(\mathcal{S}) + A_2(\mathcal{S}) - 4A_3(\mathcal{S}))\theta_{\max}}{6(A_1(\mathcal{S}) - A_3(\mathcal{S}))} \\ p^*_2(\mathcal{S}) = \frac{(A_1(\mathcal{S}) - A_2(\mathcal{S}))(A_2(\mathcal{S}) - A_3(\mathcal{S}))\theta_{\max}}{3(A_1(\mathcal{S}) - A_3(\mathcal{S}))} \\ p^*_3(\mathcal{S}) = \frac{(A_1(\mathcal{S}) - A_2(\mathcal{S}))(A_2(\mathcal{S}) - A_3(\mathcal{S}))\theta_{\max}}{6(A_1(\mathcal{S}) - A_3(\mathcal{S}))}. \end{cases} \tag{13}$$

TABLE III
ACCURACY RESULTS

Coalition structure	Acc. of A	Acc. of B	Acc. of C
$\{A,B,C\}$	**88.03%**	**88.21%**	**88.17%**
$\{A,B\},\{C\}$	79.76%	80.07%	70.91%
$\{A,C\},\{B\}$	85.50%	61.16%	86.08%
$\{B,C\},\{A\}$	50.04%	87.32%	87.62%
$\{A\},\{B\},\{C\}$	50.04%	61.16%	70.91%

TABLE IV
PRICE RESULTS

Coalition structure	Price of A	Price of B	Price of C
$\{A,B,C\}$	0.519	2.519	1.037
$\{A,B\},\{C\}$	9.984	20.492	4.992
$\{A,C\},\{B\}$	18.883	9.442	38.442
$\{B,C\},\{A\}$	4.960	9.920	19.960
$\{A\},\{B\},\{C\}$	**86.584**	**173.168**	**574.084**

We further summarize the price comparison as follows.

Corollary 1: Under Assumption 1, $p_1^*(\mathcal{S}) > p_2^*(\mathcal{S}) > p_3^*(\mathcal{S})$.

There are important implications behind Proposition 2 and Corollary 1. First, as organization 1 has the best model, it will set the highest price at NE. Second and more importantly, each organization's price depends on (and tends to decrease in) others' model accuracies. This is different from the noncompetitive case where an organization's price only relies on its own model (see Proposition 1). We will show next that such interdependency discentivizes organizations to form FL collaborations.

C. Stage I—Organization FL Coalition

The organizations play an FL coalition game similar to Game 1, except that the revenues are replaced by $R_n(\mathcal{S},\boldsymbol{p}^*)$ in (11) with $\boldsymbol{p}^*$ given in (13). We summarize the stable coalition structure in Theorem 2.

Theorem 2: Under Assumptions 1 and 2, there exist scenarios where the separation Π^R is the core stable coalition structure.

Theorems 1 and 2 answer Question 1: When organizations are not competitors, they have a natural incentive to collaborate via FL due to improved model accuracy. When organizations are competitors, however, even if FL enhances model performance, it also intensifies the price competition. This decentivizes the organizations to form FL coalitions. Next, we use a concrete example to illustrate the rationale behind Theorem 2.

Example 1: Consider three organizations A, B, and C who have non-IID data. We train models using different FL coalitions and report the accuracy (acc.) values in Table III.[3] We use $\theta_{\max} = 10^4$, calculate the equilibrium price using Proposition 2 and report the values in Table IV. We further calculate the revenues using (11) and report them in Table V.

From Table III, we observe that the grand coalition Π^G leads to the best models among all possible coalition structures, which is consistent with Assumption 2. In Tables IV and V, however, it is the separation Π^R that achieves the highest price and revenues. Even if grand coalition obtains the best model due to learning from all organizations' data, it results in very similar accuracy levels among the organizations. According to Proposition 2, the equilibrium prices are significantly reduced due to intensified competition. Unfortunately, the separation (without FL collaboration) ends up being most beneficial for revenue-maximizing organizations in a competitive market.

[3]Detailed data preparation and hyperparameters are given in Section VII.

TABLE V
REVENUE RESULTS

Coalition structure	Revenue of A	Revenue of B	Revenue of C
$\{A,B,C\}$	0.019	1.586	0.346
$\{A,B\},\{C\}$	3.328	13.546	0.028
$\{A,C\},\{B\}$	6.294	0.037	25.479
$\{B,C\},\{A\}$	0.007	3.307	13.280
$\{A\},\{B\},\{C\}$	**6.742**	**57.722**	**338.023**

So far we have shown that business competition can be a barrier that prevents organizations from collaborating via FL (even if FL incurs zero cost). Next, we present a model differentiation mechanism to encourage FL collaboration among competing organizations.

VI. Competitive Case With Model Differentiation

In this section, we present an adaptive model differentiation mechanism to enhance FL collaboration. In Section VI-A, we present the mechanism. In Section VI-B, we solve the three-stage game considering model differentiation.

A. Adaptive Model Differentiation Mechanism

To enhance FL collaboration, we present an adaptive model differentiation mechanism where organizations can iteratively adjust their model performance. Our design is inspired by *vertical differentiation* in economics [26], [27], [28]. However, our design is more challenging than prior related literature due to two reasons. First, prior work focused on the duopoly case with two organizations whereas our work studies an oligopoly case with three organizations. Second and more importantly, the unique characteristic of FL coalition introduces an additional technical challenge on model differentiation. That is, prior work typically assumed that organizations can independently adjust the qualities/accuracies. In the context of FL, however, organizations need to jointly consider the FL coalition behavior which in turn determines the feasible space of all organizations' model differentiation.

We present the adaptive model differentiation mechanism $\mathcal{M}$ in Algorithm 1, which has a time complexity of $O(N^2)$. Note that model differentiation occurs after Stage I but before Stage II, and hence $\mathcal{S}$ decided in Stage I can be regarded as given. The model differentiation mechanism proceeds in multiple iterations. In each iteration t, each organization n updates its model accuracy (lines 3-11) in a noncooperative (noncollusive) manner. The mechanism terminates until the relative difference of organizations' model accuracies between two consecutive iterations is small. Note that Algorithm 1 requires information exchange of $\tilde{A}_n(t)$, which is easy to implement in practice through feedback and reputation systems.

Algorithm 1 Adaptive Model Differentiation Mechanism

Initialization Let iteration index $t = 0$. Each organization $n \in \mathcal{N}$ starts with $\tilde{A}_n(t = 0) = A_n(\mathcal{S})$.
repeat
  for each organization $n \in \mathcal{N}$ **do**
    if $\tilde{A}_n(t) \geq \max\{\tilde{A}_j(t)\}_{j \neq n}$ **then**
      $\tilde{A}_n(t) = A_n(\mathcal{S})$
    else if $\min\{\tilde{A}_j(t)\}_{j \neq n} < \tilde{A}_n(t) < \max\{\tilde{A}_j(t)\}_{j \neq n}$ **then**
      $\tilde{A}_n(t) = \min\{A_n(\mathcal{S}), \frac{\min\{\tilde{A}_j(t)\}_{j \neq n} + \max\{\tilde{A}_j(t)\}_{j \neq n}}{2}\}$
    else if $\tilde{A}_n(t) \leq \min\{\tilde{A}_j(t)\}_{j \neq n}$ **then**
      $\tilde{A}_n(t) = \max\{\min\{A_n(\mathcal{S}), 2\min\{\tilde{A}_j(t)\}_{j \neq n} - \max\{\tilde{A}_j(t)\}_{j \neq n}\}, 0\}$
    end if
  end for
  Update strategy profile: $\tilde{\boldsymbol{A}}(t + 1) \leftarrow \tilde{\boldsymbol{A}}(t)$.
  Update iteration index: $t \leftarrow t + 1$.
until $\tilde{\boldsymbol{A}}(t)$ converges.

An important question of interest is whether the mechanism converges, which we discuss in Proposition 3.

Proposition 3: Under Assumption 1, the model differentiation mechanism in Algorithm 1 converges in polynomial time.

Having established the convergence, we next solve the three-stage game incorporating mechanism $\mathcal{M}$, where we use $\tilde{A}_n(\mathcal{S})$ to denote the converged model accuracy of organization n.

B. Three-Stage Solution With Model Differentiation

1) Stage III—Customer Purchasing With Model Differentiation: Note that model differentiation does not affect the feasible decisions or the payoff functions of customers. Hence, given $\mathcal{S}$, mechanism $\mathcal{M}$, and pricing $\boldsymbol{p}$, the customers' optimal decisions can be characterized similar to Lemma 1, except that $A_n(\mathcal{S})$ is replaced by $\tilde{A}_n(\mathcal{S})$.

2) Stage II—Organization Pricing With Model Differentiation: Similar to Section V-B, without loss of generality, we reindex the organizations using $\tilde{A}_1(\mathcal{S}) > \tilde{A}_2(\mathcal{S}) > \tilde{A}_3(\mathcal{S})$. In this case, after model differentiation, organization 1 has the best model and organization 3 has the worse model. Since model differentiation is completed before Stage II, one can regard $\tilde{A}_n(\mathcal{S})$ as given constants. Hence, we are able to derive the equilibrium price $\tilde{p}_n(\mathcal{S})$ similar to Proposition 1, except that $A_n(\mathcal{S})$ is replaced by $\tilde{A}_n(\mathcal{S})$.

Before we analyze Stage I, we present another important property of our proposed model differentiation mechanism.

Proposition 4: Given $\mathcal{S}$, and let Assumption 1 hold. Then, the model differentiation mechanism in Algorithm 1 leads to a no smaller revenue for each organization than that where there is no model differentiation.

Proposition 4 means that the proposed mechanism is incentive compatible, i.e., each organization has an incentive to execute the mechanism since it leads to a larger revenue.

3) Stage I—Organization FL Coalition: With model differentiation, the organizations decide the FL coalition, anticipating the equilibrium prices in Stage II and customers' optimal purchasing in Stage III. For convenience, define:

1) $\mathcal{S}^*, \tilde{\mathcal{S}}^*$: The FL coalition within the core stable structures without model differentiation and with model differentiation, respectively.
2) $A^{\max}(\mathcal{S}) \triangleq \max\{A_1(\mathcal{S}), A_2(\mathcal{S}), A_3(\mathcal{S})\}$: The best model accuracy among three organizations under coalition $\mathcal{S}$.
3) $W(\mathcal{S})$: Social welfare under coalition $\mathcal{S}$. Social welfare is defined as the summation of all customers' payoffs in (3) and the three organizations' revenues in (4).

Next, we discuss how model differentiation affects the model performance and social welfare.

Theorem 3: Under Assumptions 1 and 2, there exist scenarios where $A^{\max}(\tilde{\mathcal{S}}^*) > A^{\max}(\mathcal{S}^*)$ and $W(\tilde{\mathcal{S}}^*) > W(\mathcal{S}^*)$.

Theorem 3 together with the model differentiation mechanism answers Question 2. Model differentiation can mitigate the impact of price competition, enabling organizations to jointly improve model performance via FL without compromising their competitive advantages. Contrary to the intuition that this might disadvantage customers by lowering competition, our results show that the mechanism can boost both the model performance and social welfare. The key reason is that while differentiation might elevate prices, it promotes collaboration in FL and thereby enhances the service quality. This increases the value received by customers [see (3)] and hence can improve the social welfare.

VII. Numerical Results

We conduct numerical experiments to validate our assumptions and analysis. Specifically, in Section VII-A, we introduce the simulation setup. In Section VII-B, we validate Assumption 2 and Theorem 3. In Section VII-C, we relax some technical assumptions (e.g., Assumption 1) and discuss results that go beyond our theoretical framework.

A. Simulation Setup

We use train FL models on CIFAR-10 [42] using FedAvg. The CIFAR-10 data set features 10 classes and 60 000 data points. For non-IID scenarios, we employ the established Dirichlet distribution with a controlling parameter $\beta > 0$ [43]. A smaller β indicates a greater dissimilarity in clients' label distributions, which corresponds to a higher degree of non-IID. In our experiments, we consider two values of $\beta \in \{0.1, 0.01\}$, where $\beta = 0.1$ corresponds to a mildly non-IID scenario and $\beta = 0.01$ corresponds to an extremely non-IID scenario. For the results shown in Tables III–V, we use $\beta = 0.1$ and organizations $A, B,$ and, C have 3000, 5000, and 8000 data points, respectively.

The key hyperparameters are as follows [32]. The model architecture is based on ResNet-18 [44]. The local epoch count is set to 5 and the batch size to 64. Local and global learning rates are 0.1 and 1, respectively. We conduct 50 communication rounds, after which each organization n fine-tunes the global model using $\mathcal{D}_n$ for 5 additional epochs to obtain their final local models.

B. Validating Assumption 2 and Theorem 3

We first consider $\beta = 0.1$ and a uniform distribution with $\theta_{\max} = 10^4$. Organizations A, B, C have

TABLE VI
RESULTS WITHOUT MODEL DIFFERENTIATION UNDER $\beta = 0.1$

Coalition structure	A's accuracy | price | revenue	B's accuracy | price | revenue	C's accuracy | price | revenue	$A^{\max}(\mathcal{S}^*)$	$W(\mathcal{S}^*)$
$\{A,B,C\}$	86.81% | 13.15 | 7.67	86.68% | 3.76 | 1.25	86.55% | 1.88 | 0.16	-	-
$\{A,B\},\{C\}$	84.40% | 7.28 | 2.43	84.53% | 14.62 | 9.73	59.32% | 3.64 | 0.00	-	-
$\{A,C\},\{B\}$	83.01% | 31.14 | 10.38	63.94% | 15.57 | 0.09	83.59% | 63.88 | 42.23	-	-
$\{B,C\},\{A\}$	60.85% | 7.40 | 0.02	83.47% | 29.90 | 19.86	83.20% | 14.79 | 4.93	-	-
$\{A\}$, $\{B\}$, $\{C\}$	60.85% | 41.50 | 13.83	63.94% | 213.55 | 118.27	59.32% | 20.75 | 2.34	63.94%	2.005e7

TABLE VII
RESULTS WITH MODEL DIFFERENTIATION UNDER $\beta = 0.1$

Coalition structure	A's accuracy | price | revenue	B's accuracy | price | revenue	C's accuracy | price | revenue	$A^{\max}(\tilde{\mathcal{S}}^*)$	$W(\tilde{\mathcal{S}}^*)$
$\{A,B,C\}$	86.81% | 13.15 | 7.67	86.68% | 3.76 | 1.25	86.55% | 1.88 | 0.16	-	-
$\{A,B\}$, $\{C\}$	71.93% | 299.88 | 99.96	84.53% | 1136.00 | 654.37	59.32% | 149.94 | 13.59	84.53%	2.224e7
$\{A,C\},\{B\}$	73.77% | 240.51 | 80.17	63.94% | 120.25 | 10.69	83.59% | 893.26 | 516.11	-	-
$\{B,C\},\{A\}$	60.85% | 135.19 | 12.15	83.47% | 1015.27 | 585.61	72.16% | 270.37 | 90.12	-	-
$\{A\},\{B\},\{C\}$	60.85% | 62.64 | 20.88	63.94% | 224.12 | 130.26	57.76% | 31.32 | 2.68	-	-

3000, 3000, and 2500 data points, respectively. We report the results without model differentiation in Table VI, and with differentiation in Table VII, where we use bold texts to highlight the corresponding core stable coalition structure.

1) Validating Assumption 2: From Table VI, we observe a weak increase in accuracy for all organizations as the size of the coalition increases. Specifically, when all three organizations A, B, and C form a grand coalition, the accuracy levels are the highest for all clients, e.g., 86.81% for organization A. Conversely, the accuracy tends to drop when the organizations form smaller coalitions, suggesting a synergistic effect on accuracy in larger coalitions. Similar observations can be made in Table VI, which are consistent with Assumption 2.

However, the price and revenue metrics show a more complex relationship with coalition structure. For instance, in Table VI, organization B achieves its highest revenue of 118.27 when in the separation structure, even though its accuracy is not maximized. This implies that in the presence of competition, accuracy and revenues are not always positively correlated, and hence one needs to design a proper mechanism to encourage FL collaboration. We summarize the above observations as follows.

Observation 1: 1) Larger coalitions tend to improve model accuracy for all organizations and 2) an organization's accuracy and revenue are not always positively correlated.

2) Validating Theorem 3: The comparison between Tables VI and VII indicates a substantial positive impact of model differentiation. There are three important observations. First, without model differentiation, the separation structure is core stable. FL coalition intensifies price competition and reduces the organization revenue. As a result, organizations prefer to train local models (see Table VI). With model differentiation, however, organizations tend to collaborate and both organizations A and B choose to form FL coalitions (see Table VII).

Second, we observe that given the same coalition structure (e.g., $\{A, B\}, \{C\}$), each organization yields a higher revenue with model differentiation than without it. This implies that our mechanism is incentive compatible, which also validates Proposition 4.

Third and most importantly, from Tables VI and VII, we observe that model differentiation can greatly improve the model accuracy $A^{\max}$, the organizations' total revenues, and the social welfare. In particular, the improvement is 20.59%, 14.42%, and 10.92%, respectively. These observations are consistent with Theorem 3. We summarize the above key observations as follows.

Observation 2: Model differentiation can simultaneously improve the model accuracy, organizations' revenues, as well as the social welfare.

However, we must note that the customer surplus (i.e., the summation of all customers' payoffs) experiences a drop by 9.8%. This is due to the increased prices from model differentiation albeit the quality improvement. Nevertheless, model differentiation improves the social welfare. A more sophisticated mechanism design to ensure a low or zero drop of customer surplus is left to future work.

C. Results With Relaxed Assumptions

So far our results are based on the assumptions that 1) the customers' valuation θ follows a uniform distribution and 2) each customer's utility linearly depends on the model accuracy [see (3)]. Now, we relax these assumptions and study whether previous results/observations continue to hold. To this end, we consider $\beta = 0.01$, customer's valuation follows a truncated normal distribution with mean 5000 and standard deviation 2000 on support $[0, 10^4]$. We further consider a quadratic utility function i.e., $\forall\, d_\theta = n \in \{1, 2, 3\}$

$$u_\theta(d_\theta; \mathcal{S}, \boldsymbol{p}) = \theta \cdot A_n^2(\mathcal{S}) - p_n. \tag{14}$$

The quadratic utility can model the scenario where customers gain marginally increasing benefits as the quality of service (e.g., in terms of medical diagnosis and financial insurance) increases. We report the results without model differentiation in Table VIII, and with differentiation in Table IX, where we use bold texts to highlight the corresponding core stable coalition structure.

There are three key observations from Tables VIII and IX. First, larger coalitions generally improve model accuracy for all organizations. For instance, in both tables, when the

TABLE VIII
RESULTS WITHOUT MODEL DIFFERENTIATION UNDER $\beta = 0.01$

Coalition structure	A' accuracy | price | revenue	B' accuracy | price | revenue	C' accuracy | price | revenue	$A^{\max}(\mathcal{S}^*)$	$W(\mathcal{S}^*)$
$\{A,B,C\}$	86.89% | 2992.96 | 2095.71	86.84% | 0.00 | 0.00	86.84% | 0.00 | 0.00	-	-
$\{A,B\},\{C\}$	84.11% | 2.47 | 0.56	84.21% | 8.42 | 6.50	57.90% | 1.24 | 0.00	-	-
$\{A,C\},\{B\}$	84.27% | 9.27 | 7.15	60.83% | 1.36 | 0.00	84.16% | 2.72 | 0.62	-	-
$\{B,C\},\{A\}$	61.96% | 1521.89 | 1065.65	82.89% | 0.00 | 0.00	82.89% | 0.00 | 0.00	-	-
$\boldsymbol{\{A\}, \{B\}, \{C\}}$	61.96% | 64.91 | 48.87	60.83% | 14.23 | 3.04	57.90% | 6.66 | 0.01	61.96%	2.179e7

TABLE IX
RESULTS WITH MODEL DIFFERENTIATION UNDER $\beta = 0.01$

Coalition structure	A' accuracy | price | revenue	B' accuracy | price | revenue	C' accuracy | price | revenue	$A^{\max}(\tilde{\mathcal{S}}^*)$	$W(\tilde{\mathcal{S}}^*)$
$\{A,B,C\}$	86.89% | 2992.96 | 2095.71	86.84% | 0.00 | 0.00	86.84% | 0.00 | 0.00	-	-
$\{A,B\},\{C\}$	71.06% | 129.18 | 25.77	84.21% | 898.34 | 659.68	57.90% | 56.86 | 0.28	-	-
$\boldsymbol{\{A,C\}, \{B\}}$	84.27% | 809.48 | 594.77	60.83% | 51.93 | 0.25	72.55% | 117.83 | 23.54	84.27%	2.604e7
$\{B,C\},\{A\}$	61.96% | 1521.89 | 1065.65	82.89% | 0.00 | 0.00	82.89% | 0.00 | 0.00	-	-
$\{A\},\{B\},\{C\}$	61.96% | 109.80 | 80.93	59.93% | 17.03 | 3.43	57.90% | 7.56 | 0.03	-	-

coalition includes all three organizations (i.e., $\{A, B, C\}$), the accuracy is the highest for each organization. This again validates Assumption 2 and is consistent with Observation 1. Second, given the coalition structure, model differentiation leads to a no smaller revenue for each organization, which supports Proposition 4. Third, model differentiation can indeed yield simultaneous benefits in terms of model accuracy and social welfare. The improvement is 22.31% and 19.5%, respectively. This is consistent with Observation 2. We summarize the above key observations as follows.

Observation 3: Observations 1 and 2 continue to hold when we consider a truncated normal distribution for customers' valuation and a quadratic customer utility function.

VIII. DISCUSSIONS

In this section, we discuss some possible extensions of our model and analysis.

A. More General Valuation Distribution

The analysis (e.g., previous propositions and theorems) relies on the assumption that customers' valuation follows a uniform distribution. We note that our analysis can be extended to a broader class of distributions where $h(\theta)/[1 - H(\theta)]$ is nondecreasing in θ. This condition is satisfied by many commonly used distributions, such as uniform distribution, normal distribution, and gamma distribution. In this case, deriving a closed-form solution to Stage II can be challenging, but the price optimization can be shown to be a piece-wise concave problem, which can be solved in polynomial time.

B. More General Utility Function

The theoretical results assumed that customers' utility function is linear in the model accuracy. We note that our results can be extended to where the utility function is concave and strictly increasing in the accuracy. Although a closed-form solution to Stage II may no longer be available, the price optimization is still a piece-wise concave problem, and one can solve it in polynomial time. In addition, in case of nonconcave and nonlinear utility functions, our experiments in Tables VIII and IX using a quadratic function show that the key results (e.g., Theorem 3) continue to hold.

C. More Than Three Organizations

When there are more than three organizations, one can derive the customers' optimal purchasing similar to (9). In this case, assuming that $A_1 > A_2 > \cdots > A_N$, one can write each organization's revenue function as follows $R_n = p_n \cdot M_n$, where

$$\begin{aligned} M_1 &= H(\theta_{\max}) - H\big(\max(\sigma_{1,2}, \sigma_{1,3}, \ldots, \sigma_{1,N-1}, \sigma_{1,N})\big) \\ M_N &= H\big(\min(\sigma_{1,N}, \sigma_{2,N}, \ldots, \sigma_{N-1,N}\theta_{\max})\big) - H(0) \end{aligned} \quad (15)$$

and for all $n = 2, \ldots, N-1$

$$M_n = H\big(\sigma_{n-1,n} - \max(\sigma_{n,n+1}, \sigma_{n,n+2}, \ldots, \sigma_{n,N}, 0)\big). \quad (16)$$

One can use algorithms, such as gradient ascent, to find the solutions to Stage II.

However, we note that the solutions to Stage I are empirically infeasible under a large N due to there being a total number of Bell(N) of possible coalition structures. Here, Bell($\cdot$) is the Bell number that increases exponentially in N [45]. For example, Bell(5) = 52 but Bell(10) = 115975. One possible remedy to this is to establish a closed-form correspondence between the coalition and the model accuracy, which is left to future work.

IX. CONCLUSION

This article studies the under-explored problem of business competition in FL using an oligopoly framework with three organizations. We introduce a multistage game to model these interactions, accounting for both collaboration in FL training and competition in selling model-based services. Our findings reveal a paradox: although organizations have an incentive to collaborate for FL training in a noncompetitive setting, the presence of competition discourages such collaboration. This is mainly because FL tends to harmonize model performance across organizations, intensify price competition, and hence reduce revenues. To encourage FL collaboration, we present an adaptive model differentiation mechanism in which the organizations iteratively adjust their model performance. Perhaps counter-intuitively, both theoretical analysis and numerical experiments show that our mechanism not only enhances model accuracy and organizations' revenues, but also improves the overall social welfare.

For the future work, it is interesting to study the case where organizations may offer supplementary services and a customer can buy services from more than one organization. It is also interesting to study how the asymmetry among the organizations, in terms of data quantity or quality, affects the FL coalition and price competition. Considering that organizations interact repeatedly within a long time window is another promising research avenue.

REFERENCES

[1] P. Kairouz et al., "Advances and open problems in federated learning," *Found. Trends® Mach. Learn.*, vol. 14, nos. 1–2, pp. 1–210, 2021.

[2] C. Huang, S. Ke, and X. Liu, "Duopoly business competition in cross-silo federated learning," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 1, pp. 340–351, Jan./Feb. 2024.

[3] Y. Li et al., "VARF: An incentive mechanism of cross-silo federated learning in MEC," *IEEE Internet Things J.*, vol. 10, no. 17, pp. 15115–15132, Sep. 2023.

[4] C. Huang, M. Tang, Q. Ma, J. Huang, and X. Liu, "Promoting collaborations in cross-silo federated learning: Challenges and opportunities," *IEEE Commun. Mag.*, vol. 62, no. 4, pp. 82–88, Apr. 2024.

[5] D. C. Nguyen et al., "Federated learning for smart healthcare: A survey," *ACM Comput. Surv.*, vol. 55, no. 3, pp. 1–37, 2022.

[6] X. Wu and H. Yu, "MarS-FL: Enabling competitors to collaborate in federated learning," *IEEE Trans. Big Data*, early access, Jun. 28, 2022, doi: 10.1109/TBDATA.2022.3186991.

[7] Y. Wang, M. Jia, N. Gao, L. Von Krannichfeldt, M. Sun, and G. Hug, "Federated clustering for electricity consumption pattern extraction," *IEEE Trans. Smart Grid*, vol. 13, no. 3, pp. 2425–2439, May 2022.

[8] M. J. Osborne and A. Rubinstein, *A Course in Game Theory*. Cambridge, MA, USA: MIT press, 1994.

[9] K. Donahue and J. Kleinberg, "Model-sharing games: Analyzing federated learning under voluntary participation," in *Proc. AAAI*, 2021, pp. 5303–5311.

[10] K. Donahue and J. Kleinberg, "Optimality and stability in federated learning: A game-theoretic approach," in *Proc. NeurIPS*, 2021, pp. 1287–1298.

[11] S. Jiang and J. Wu, "Coalition formation game in the cross-silo federated learning system," in *Proc. IEEE MASS*, 2022, pp. 49–57.

[12] B. Ray Chaudhury, L. Li, M. Kang, B. Li, and R. Mehta, "Fairness in federated learning via core-stability," in *Proc. NeurIPS*, 2022, pp. 5738–5750.

[13] S. Arisdakessian, O. A. Wahab, A. Mourad, and H. Otrok, "Coalitional federated learning: Improving communication and training on non-IID data with selfish clients," *IEEE Trans. Services Comput.*, vol. 16, no. 4, pp. 2462–2476, Jul./Aug. 2023.

[14] M. Hu, D. Wu, Y. Zhou, X. Chen, and M. Chen, "Incentive-aware autonomous client participation in federated learning," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 10, pp. 2612–2627, Oct. 2022.

[15] W. Bao, H. Wang, J. Wu, and J. He, "Optimizing the collaboration structure in cross-silo federated learning," in *Proc. ICML*, 2023, pp. 1–19.

[16] X. Tu, K. Zhu, N. C. Luong, D. Niyato, Y. Zhang, and J. Li, "Incentive mechanisms for federated learning: From economic and game theoretic perspective," *IEEE Trans. Cogn. Commun. Netw.*, vol. 8, no. 3, pp. 1566–1593, Sep. 2022.

[17] Y. Zhan, J. Zhang, Z. Hong, L. Wu, P. Li, and S. Guo, "A survey of incentive mechanism design for federated learning," *IEEE Trans. Emerg. Topics Comput.*, vol. 10, no. 2, pp. 1035–1044, Apr.–Jun. 2022.

[18] J. Kang et al., "Blockchain-empowered federated learning for healthcare metaverses: User-centric incentive mechanism with optimal data freshness," *IEEE Trans. Cogn. Commun. Netw.*, vol. 10, no. 1, pp. 348–362, Feb. 2024.

[19] Y. Zhan, P. Li, Z. Qu, D. Zeng, and S. Guo, "A learning-based incentive mechanism for federated learning," *IEEE Internet Things J.*, vol. 7, no. 7, pp. 6360–6368, Jul. 2020.

[20] D. Ye, X. Huang, Y. Wu, and R. Yu, "Incentivizing semisupervised vehicular federated learning: A multidimensional contract approach with bounded rationality," *IEEE Internet Things J.*, vol. 9, no. 19, pp. 18573–18588, Oct. 2022.

[21] R. Lu et al., "Auction-based cluster federated learning in mobile edge computing systems," *IEEE Trans. Parallel Distrib. Syst.*, vol. 34, no. 4, pp. 1145–1158, Apr. 2023.

[22] B. Luo, Y. Feng, S. Wang, J. Huang, and L. Tassiulas, "Incentive mechanism design for unbiased federated learning with randomized client participation," in *Proc. 43rd ICDCS*, 2023, pp. 545–555.

[23] X. Wang, Q. Le, A. F. Khan, J. Ding, and A. Anwar, "A framework for incentivized collaborative learning," 2023, *arXiv:2305.17052*.

[24] C. Sun, C. Huang, B. Shou, and J. Huang, "Federated learning in competitive EV charging market," in *Proc. IEEE ISGT*, 2023, pp. 1–5.

[25] S. Tan et al., "FedCompetitors: Harmonious collaboration in federated learning with competing participants," in *Proc. AAAI*, 2024, pp. 15231–15239.

[26] X. Wauthy, "Quality choice in models of vertical differentiation," *J. Ind. Econ.*, vol. 44, no. 3, pp. 345–353, 1996.

[27] H. K. Bhargava and V. Choudhary, "Information goods and vertical differentiation," *J. Manag. Inf. Syst.*, vol. 18, no. 2, pp. 89–106, 2001.

[28] D. O. Lauga and E. Ofek, "Product positioning in a two-dimensional vertical differentiation model: The role of quality costs," *Market. Sci.*, vol. 30, no. 5, pp. 903–923, 2011.

[29] M. Tang and V. W. Wong, "An incentive mechanism for cross-silo federated learning: A public goods perspective," in *Proc. IEEE INFOCOM*, 2021, pp. 1–10.

[30] N. Zhang, Q. Ma, and X. Chen, "Enabling long-term cooperation in cross-silo federated learning: A repeated game perspective," *IEEE Trans. Mobile Comput.*, vol. 22, no. 7, pp. 3910–3924, Jul. 2023.

[31] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. Artif. Intell. Stat.*, 2017, pp. 1273–1282.

[32] S. Ke, C. Huang, and X. Liu, "Quantifying the impact of label noise on federated learning," 2023, *arXiv:2211.07816*.

[33] L. Zhang, L. Shen, L. Ding, D. Tao, and L.-Y. Duan, "Fine-tuning global model via data-free knowledge distillation for non-iid federated learning," in *Proc. IEEE/CVF CVPR*, 2022, pp. 10174–10183.

[34] B. Pfitzner, N. Steckhan, and B. Arnrich, "Federated learning in a medical context: A systematic literature review," *ACM Trans. Internet Technol.*, vol. 21, no. 2, pp. 1–31, 2021.

[35] Z. Wang, L. Gao, and J. Huang, "Duopoly competition for mobile data plans with time flexibility," *IEEE Trans. Mobile Comput.*, vol. 19, no. 6, pp. 1286–1298, Jun. 2020.

[36] Z. Zheng, Y. Peng, F. Wu, S. Tang, and G. Chen, "Trading data in the crowd: Profit-driven data acquisition for mobile crowdsensing," *IEEE J. Sel. Areas Commun.*, vol. 35, no. 2, pp. 486–501, Feb. 2017.

[37] M. Jayakumar, R. P. Pradhan, S. Dash, R. P. Maradana, and K. Gaurav, "Banking competition, banking stability, and economic growth: Are feedback effects at work?" *J. Econ. Bus.*, vol. 96, pp. 15–41, Mar./Apr. 2018.

[38] C. Huang, H. Yu, J. Huang, and R. Berry, "An online inference-aided incentive framework for information elicitation without verification," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 1167–1185, Apr. 2023.

[39] C. Huang, J. Huang, and X. Liu, "Cross-silo federated learning: Challenges and opportunities," 2022, *arXiv:2206.12949*.

[40] C. Huang, H. Yu, J. Huang, and R. Berry, "Online crowd learning through strategic worker reports," *IEEE Trans. Mobile Comput.*, vol. 22, no. 9, pp. 5406–5417, Sep. 2023.

[41] Q. Huang, Y. Liu, L. Wang, P. Sun, J. Li, and J. Xu, "A blockchain-enabled coalitional game framework for carbon emission trading," *IEEE Trans. Netw. Sci. Eng.*, early access, Jul. 10, 2023, doi: 10.1109/TNSE.2023.3293453.

[42] "The CIFAR-10 dataset." Accessed: May 24, 2024. [Online]. Available: https://www.cs.toronto.edu/∼ kriz/cifar.html

[43] T.-M. H. Hsu, H. Qi, and M. Brown, "Measuring the effects of non-identical data distribution for federated visual classification," 2019, *arXiv:1909.06335*.

[44] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE CVPR*, 2016, pp. 770–778.

[45] J. Katriel, "On a generalized recurrence for bell numbers," *J. Integer Seq.*, vol. 11, no. 2, p. 3, 2008.

Chao Huang received the Ph.D. degree from The Chinese University of Hong Kong, Hong Kong, in 2021.

From July 2021 to November 2021, he was a Postdoctoral Research Fellow with the Department of Management Sciences, The City University of Hong Kong, Hong Kong. He is currently a Postdoctoral Researcher with the Department of Computer Science, University of California at Davis, Davis, CA, USA. His recent research interests include distributed machine learning and network economics.

Justin Dachille received the B.S. degree in computer science from the University of California at Irvine, Irvine, CA, USA, in 2022. He is currently pursuing the M.S. degree in computer science with the University of California at Davis, Davis, CA, USA.

His research interests include federated learning and split learning.

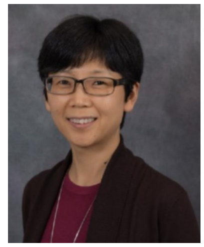

Xin Liu (Fellow, IEEE) received the Ph.D. degree in electrical engineering from Purdue University, West Lafayette, IN, USA, in 2002.

She is currently a Professor of Computer Science with the University of California at Davis, Davis, CA, USA. Her current research interests fall in the general areas of machine learning algorithm development and machine learning applications in human and animal healthcare, food systems, and communication networks. Her research on networking includes cellular networks, cognitive radio networks, wireless sensor networks, network information theory, network security, and IoT systems.