Duopoly Business Competition in Cross-Silo Federated Learning

Chao Huang, Shuqi Ke, and Xin Liu, *Fellow, IEEE*

***Abstract*—In cross-silo federated learning, clients (e.g., organizations) collaboratively train a global model using private local data. In practice, clients may be not only collaborators but also business competitors. This article studies the overlooked but practically important problem of business competition in cross-silo FL. We formulate the clients' market competition as a three-stage game, where the clients decide FL training strategies in Stage I and the pricing strategies in Stage II, and then heterogeneous customers decide purchasing strategies in Stage III. The game analysis is highly challenging, as clients' collaborations and competitions are complexly coupled. We manage to characterize the equilibrium properties and find that market competition always reduces the clients' profits and can further lead to a worse global model when clients' costs are high. To mitigate this issue, we propose a general framework that enables proper revenue (profit plus cost) sharing among clients. Both theoretical and numerical results with MNIST and CIFAR-10 show that revenue sharing can greatly improve the global model accuracy and clients' profits. Counter-intuitively, even if market competition limits clients' profits, it can lead to a better global model when clients' costs are low, as clients strive to survive in the market by contributing more training data.**

***Index Terms*—Federated learning, business competition, data sharing, game theory, machine learning.**

I. Introduction

FEDERATED Learning (FL) is a decentralized machine learning scheme proposed by Google, where multiple clients collaboratively train a global model under the coordination of a central server [1]. Clients train models locally and only need to upload model updates (e.g., gradients or parameters). More specifically, clients train models using their private local data and then upload the model updates to the central server. The central server aggregates all clients' uploaded model updates (e.g., using a weighted average algorithm) to generate a global model. It then sends back the global model to clients for further local training. The iteration between local training and global aggregation terminates until the global model converges.

Based on the scale and participating clients, FL can be classified into two types: cross-device FL and cross-silo FL [2]. In cross-device FL, clients are distributed entities, e.g., smartphones, wearables, and edge devices. Cross-device clients are unlikely to participate in the whole training process due to limited computation and communication resources. Hence, cross-device FL may require many clients (e.g., from thousands to millions) to succeed. In cross-silo FL, however, clients are companies or organizations (e.g., hospitals and banks). The number of participants is small (e.g., 2 to 100), and each client is expected to participate in the entire training process.

While prior related work emphasizes cross-device FL, this article focuses on cross-silo FL. Practical industrial applications of cross-silo FL abound. In the financial domain, WeBank and Swiss Re cooperatively conduct FL analysis and offer financial and insurance services [3]. In the medical and health care domain, networked pharmaceutical institutions collaborate to train models for drug discovery based on private and highly sensitive screening datasets [4].

The success of cross-silo FL requires the clients to contribute sufficient resources (e.g., local data) for model training. To this end, proper cooperation mechanisms are needed to encourage contributions. The authors in [5] proposed a mechanism to enable the organizations' long-term cooperation. The authors in [6] studied the mechanism design to encourage institutions to contribute training resources. However, in practical cross-silo FL applications, the clients may be not only model training collaborators but also business competitors. On one hand, the cross-silo clients collaborate to train FL models and each can benefit from the well-trained global model. On the other hand, the clients may compete in selling model-related services (e.g., loan interest determination in digital banking [7], [8]) to potential customers in the market (see Fig. 1). The clients may worry that contributing to FL will benefit their business competitors, who do not contribute as much but receive the same final model nonetheless. Business competition is a critical component in cross-silo FL, and it makes the collaboration among clients more challenging (and intriguing), due to the complex coupling and conflicting objectives between collaboration and competition. However, this important problem is under-explored and its impact is poorly understood, which motivates us to ask the first key question:

Key Question 1: How will business competition affect cross-silo FL in terms of the global model accuracy and clients' profits (from selling model-related services to customers)?

Manuscript received 8 February 2023; revised 7 June 2023; accepted 13 July 2023. Date of publication 31 July 2023; date of current version 8 January 2024. This work was supported in part by USDA/NIFA under Grant GS012020-67021-32855, in part by the NSF under Grant IIS1838207, in part by CNS under Grant 1901218, and in part by OIA - under Grant 2134901. Recommended for acceptance by Dr. Hongbo Jiang. *(Corresponding author: Chao Huang.)*

Chao Huang and Xin Liu are with the Department of Computer Science, University of California, Davis, CA 95616 USA (e-mail: fchhuang@ucdavis.edu; xinliu@ucdavis.edu).

Shuqi Ke is with the School of Science and Engineering, Chinese University of Hong Kong, Shenzhen, Guangdong 518172, China (e-mail: shuqike@link.cuhk.edu.cn).

This article has supplementary downloadable material available at https://doi.org/10.1109/TNSE.2023.3297880, provided by the authors.

Digital Object Identifier 10.1109/TNSE.2023.3297880

2327-4697 © 2023 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.
See https://www.ieee.org/publications/rights/index.html for more information.

Authorized licensed use limited to: Univ of Calif Davis. Downloaded on January 20,2024 at 01:48:26 UTC from IEEE Xplore. Restrictions apply.

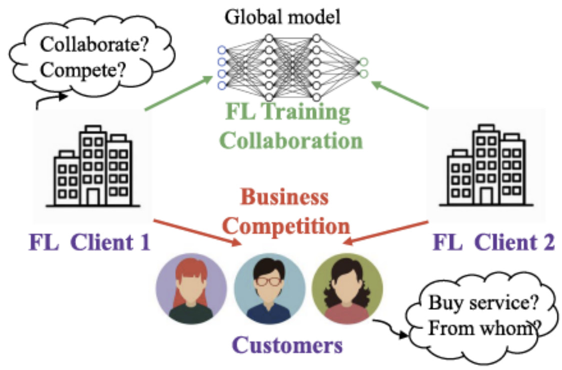

Fig. 1. FL collaboration and business competition in cross-silo FL.

Throughout the article, we use two metrics, i.e., *global model accuracy* and *clients' profits*, to quantify the impact of business competition on cross-silo FL. The global model accuracy is an important indicator of whether and how much clients contribute to FL, while clients' profits reflect from an economic/incentive perspective why clients contribute to FL.

To answer *Key Question 1*, we study the clients' market competition in terms of their model training strategies and pricing decisions. To abstract the interactions among competing clients and heterogeneous customers, we focus on the two-client case, i.e., duopoly competition. Note that duopoly competition exists in practical applications. For example, WeBank and Swiss Re may compete for customers from China who seek reinsurance services [9]. Moreover, the two-client case serves as a first step to understanding the key insights related to business competition. Note that duopoly competition exists in practical applications. Our model can be generalized to where there are two groups of service providers and each group may contain several organizations. When there are more than three clients, an analytical characterization is challenging. However, as will be shown, the (price) competition involves solving a piece-wise concave problem and there exist efficient algorithms (e.g., gradient ascent) that return the optimal solutions We leave the detailed investigation of the case with more than three clients into future work. As a benchmark comparison, we also study the monopoly case where each client is the only service provider in its own market (but it can collaboratively train FL models with the other client in a different market). The comparison between duopoly and monopoly enables us to understand the impact of business competition on cross-silo FL.

An intuitive guess to *Key Question 1* is that, compared to monopoly, duopoly competition leads to a lower client profit and a worse global model. This is because market competition is supposed to limit the clients' profits and hence their incentive to contribute data for FL, resulting in a worse global model. This undesirable outcome turns out to be possible under certain conditions, which motivates our second key question below.

Key Question 2: If business competition damages cross-silo FL, how can we mitigate competition and promote collaboration among clients?

To answer *Key Question 2*, we propose a general revenue-sharing framework to enhance client collaboration. Under the framework, the clients appropriately share revenues (i.e., profits plus clients' costs) from selling model-related services to the customers. We will compare duopoly (without revenue sharing) and duopoly with revenue sharing in terms of the global model accuracy and clients' profits. The comparison provides insights into the practical design and implementations of cooperation mechanisms for competing cross-silo clients.

A. Solutions and Key Contributions

We study the strategic interactions between two competitive clients and a group of heterogeneous customers. We use a three-stage game model to formulate the clients' FL training strategies and pricing decisions, as well as customers' purchasing strategies. More specifically, the two clients decide FL training strategies in Stage I and the corresponding pricing strategies in Stage II. Finally, the customers decide whether and which client to purchase model-related service from in Stage III.

We summarize the key results and contributions below.

- *Duopoly Competition in Cross-Silo Federated Learning:* To our best knowledge, this is the first analytical work that studies the cross-silo clients' duopoly business competition considering their FL training strategies and pricing decisions. The problem is of theoretical interest and practical importance for the sustainable development of cross-silo FL.
- *A Three-Stage Competition Model:* We formulate the clients' market competition and customers' purchasing as a three-stage game. The analysis is highly challenging, as clients' FL training strategies (collaboration) and pricing strategies (competition) are complexly coupled. Further, it is difficult to characterize the impact of clients' FL training strategies on the model accuracy.
- *Characterizing Impact of Duopoly Competition:* We show that compared to monopoly, duopoly always leads to a lower client profit. However, despite the profit reduction, it can surprisingly lead to a better global model when clients' costs are low, as clients strive to survive in the competing market by setting lower prices and contributing more training resources.
- *Proposing A General Revenue-Sharing Framework:* When clients' costs are high, on the other hand, duopoly leads to a lower client profit and a worse global model than monopoly. To avoid this undesirable outcome, we propose a general revenue-sharing framework to enhance client collaboration, which is shown to significantly improve the clients' profit and the global model accuracy.
- *Performance Evaluation:* We conduct extensive numerical experiments using the MNIST and CIFAR-10 dataset, and find that the results are consistent with our theoretical analysis. We further compare various benchmark revenue-sharing mechanisms (all of which can be characterized by our proposed framework), and provide insights on adopting the revenue-sharing mechanism in practical cross-silo FL.

The remainder of the article is organized as follows. In Section II, we review the related work. In Section III, we

Authorized licensed use limited to: Univ of Calif Davis. Downloaded on January 20,2024 at 01:48:26 UTC from IEEE Xplore. Restrictions apply.

introduce the model. In Section IV, we study the duopoly competition among cross-silo clients. In Section V, we propose a general revenue-sharing framework and analyze its impact on duopoly competition. We provide numerical results in Section VI and conclude in Section VII.

II. Related Work

Our work studies business competition in cross-silo FL. Hence, we review the related work from two aspects, i.e., cross-silo FL and business competition.

A. Cross-Silo FL

While existing work focuses on cross-device FL, some recent work studies cross-silo FL. For example, the studies in [10], [11], [12] propose personalized approaches to tackle the non-i.i.d. issue. The work in [13] studies the topology design to improve communication efficiency. The studies in [14], [15], [16] propose privacy-preserving mechanisms (e.g., using differential privacy and homomorphic encryption) customized for cross-silo settings. The studies in [5], [6], [17], [18] focus on the incentive mechanism design to promote client cooperation. Interested readers can refer to [2] for a more comprehensive overview of cross-silo FL.

However, none of these studies analyzes business competition in cross-silo FL. To our best knowledge, we propose a first analytical framework to explicitly model duopoly business competition and analyze its impact on cross-silo FL.

B. Business Competition

Business competition exists in many applications. In mobile computing, network operators (e.g., AT&T) strive to attract and keep data plan subscribers [19], [20]. In cognitive radio networks, primary users compete in providing spectrum resources to secondary users [21]. In cross-silo FL, financial companies compete for customers who seek loan and insurance services [22]. Although a recent paper [23] considered FL clients to be price takers in a fully competitive market, it did not model the important business competition part related to pricing design. *Different from [23], our article studies the clients' duopoly business competition in terms of both the clients' FL training and pricing decisions. The complex coupling between FL training (collaboration) and pricing (competition) makes the analysis much more challenging.*

III. Model

In Section III-A, we first introduce a practical cross-silo FL process. In Sections III-B and III-C, we specify the decisions and objectives of the customers and clients, respectively. In Section III-D, we discuss the strategic interactions between customers and clients.

A. A Practical Cross-Silo FL Process

A practical cross-silo FL process consists of two phases [10]:

- *Global iteration:* the clients use (possibly part of) their local data to iteratively train a shared global model.
- *Local fine-tuning:* after the global model converges, each client further fine-tunes the global model using its local data, and generates a final local model.

We explain in more detail the above two phases below.

1) Global Iteration: Consider a set $\mathcal{N} = \{1, 2\}$ of clients (e.g., WeBank and Swiss Re [3]) who aim to collaboratively train a global model. Each client owns a private local dataset and can use a subset of data for global model iteration. Define:

- $\mathcal{D}_n$: each client n's local data set with size $D_n = |\mathcal{D}_n|$.
- $\mathcal{X}_n$: each client n's chosen data set for global model iteration, where $\mathcal{X}_n \subseteq \mathcal{D}_n$ with its size being $x_n = |\mathcal{X}_n|$.

In global iteration, clients use their chosen local data to train a global model represented by a weight vector ω. To derive the optimal weights ω^*, the global iteration proceeds in multiple rounds. In each round r, the clients perform the following steps:

- Each client n downloads the global model ω^{r-1} generated from the previous round.
- Each client n locally trains the model ω^{r-1} over its chosen data set $\mathcal{X}_n$.
- Each client n sends the model updates ω_n^r to the server for aggregation, which generates a new global model ω^r to be trained in the next round. The celebrated aggregation algorithm FedAvg [24], [25] works as follows: $\omega^r = \sum_{n\in\mathcal{N}} \frac{x_n}{\sum_{n'\in\mathcal{N}} x_{n'}} \omega_n^r$.

The above iterative process terminates when the global model converges. We use $A^{\mathrm{g}}(x) \in [0, 1]$ to denote the global model accuracy, and it depends on the data $x \triangleq (x_1, x_2)$ chosen by the two clients.

2) Local Fine-Tuning: In practical cross-silo FL, after the global model converges, the clients further fine-tune the global model. That is, each client retrains part (or all) parameters of a converged global model on its local training data [26].[1] Many empirical studies have shown that fine-tuning enhances the global model performance [28]. This is mainly due to two reasons below.

- First, when different clients have non-i.i.d. data, the converged global model can have a poor performance on each client's local data distribution. This is known as client drift [29], and a promising remedy is to fine-tune the global model to improve the performance on local data.
- Second, even if clients have i.i.d. data, they may not have an incentive to contribute many data in global iteration. It is because this would give other clients (who are their business competitors) a better global model, which can damage the competitiveness of the contributing client. This leads to a bad global model, and clients find it beneficial to generate a better local model via fine-tuning.

After fine-tuning, each client obtains a final local model, and we use $A_n(x) \in [0, 1]$ to denote client n's local model accuracy. Each client uses its fine-tuned local model to generate model-related services (e.g., loan predictions by banks, disease

[1] While there are other methods (e.g., multi-task learning and knowledge distillation [27]), a detailed discussion is out of the scope of this article.

Authorized licensed use limited to: Univ of Calif Davis. Downloaded on January 20,2024 at 01:48:26 UTC from IEEE Xplore. Restrictions apply.

diagnosis by hospitals [30]). Then, the clients compete in selling the services to customers in the market. Next, we model the decisions and payoffs of customers and clients.

B. Customer's Decision and Payoff

We first introduce the customer characteristic, and then define each customer's decision and payoff function.

1) Customer Characteristic: Consider a continuum of customers with normalized population size one [19]. We denote θ as a customer's valuation toward the model-related service. Intuitively, a larger θ means that a customer draws a higher utility from the service. Customers are heterogeneous and each is characterized by a random variable θ, where the distribution of θ on the entire population has a PDF $h(\theta)$ and CDF $H(\theta)$ on support $[0, \theta_{\max}]$. We consider that each customer's individual valuation is unknown but the valuation distribution is known to the clients (e.g., due to market research [31]).

2) Customer Decision: The two competing clients may offer services with different qualities (due to having different local models) at different prices. Given the service qualities and prices from the two clients, a customer decides whether to purchase the service, and if so, which one client to purchase service from. Let $d_\theta \in \{\emptyset, 1, 2\}$ denote a type-θ customer's decision, where $d_\theta = \emptyset$ means no purchasing, and $d_\theta = n$ means purchasing from client n, for both $n \in \{1, 2\}$. In this article, we consider that the competing clients offer substitutable services (e.g., loan and insurance) to customers and it is natural to consider that a customer only purchases from one of the two clients.

3) Customer Payoff: A customer's payoff is the difference between his utility and the payment to the client. If a customer does not purchase the service, he achieves a zero payoff. If a customer purchases the service from client n, he draws a utility from enjoying the model-related service and needs to pay a non-negative amount $p_n \geq 0$ to client n. More specifically, a customer with valuation θ has the following payoff function:

$$u_\theta(d_\theta; \boldsymbol{x}, \boldsymbol{p}) = \begin{cases} 0, & \text{if} \quad d_\theta = \emptyset, \\ \theta \cdot A_n(\boldsymbol{x}) - p_n, & \text{if} \quad d_\theta = n \in \{1, 2\}, \end{cases} \quad (1)$$

where $\boldsymbol{p} \triangleq (p_1, p_2)$. Intuitively, a higher model accuracy $A_n(\boldsymbol{x})$ corresponds to a service of higher quality, and hence the customer gains a higher utility $\theta \cdot A_n(\boldsymbol{x})$.[2] Note that the customers do not need to know the clients' data contribution decisions $\boldsymbol{x}$ but the model accuracy $A_n(\boldsymbol{x})$. In practice, the customers can obtain information on different clients' service quality by exploring their reputation and feedback systems [32].

C. Clients' Decisions and Profits

In this subsection, we introduce the cross-silo clients' possible decisions and profit functions.

1) Client Data Contribution and Pricing Strategies: Each client first needs to decide how much data to use in cross-silo FL, which is related to the two phases below.

[2]Note that our model and analysis can be easily extend to a more general case where a customer has a valuation function $V(A_n(\boldsymbol{x}))$, which is concave and strictly increasing in $A_n(\boldsymbol{x})$.

- *Global iteration:* each client n decides the date contribution $x_n \in [0, D_n]$ for global model iteration. More specifically, a client chooses a subset $\mathcal{X}_n$ of its local data set to iteratively train models until the global model converges.
- *Local fine-tuning:* after the global model converges, each client uses local data to fine-tune the global model and generates a final local model. Similar to practical implementations, we consider that each client uses *all* of its local data $\mathcal{D}_n$ for fine-tuning, as it has a superior empirical performance in improving the global model [26], [27], [28].

Remark: (i) Since local fine-tuning uses all data, we do not model it as an active decision. Our model can be easily extended to the case where a client further optimizes the amount of data for fine-tuning, as the privacy cost does not depend on the use of fine-tuning data.

(ii) If a client does not contribute any data to FL in global iteration, then local fine-tuning degenerates to the case where a client trains a local model with all of its own data.

Besides deciding the data contribution x_n for FL (in global iteration), each client also needs to decide the pricing strategy $p_n \geq 0$. A client n charges a customer p_n if the customer purchases service from client n.

2) Client Revenue: A client n's revenue from a single customer is the price p_n it charges. Hence, the total revenue client n obtains from the entire customer population is

$$R_n(\boldsymbol{x}, \boldsymbol{p}) = \int_0^{\theta_{\max}} p_n \cdot \mathbb{1}_{d_\theta = n}(\boldsymbol{x}, \boldsymbol{p}) \cdot h(\theta) d\theta, \quad (2)$$

where $\mathbb{1}$ is an indicator function, i.e., $\mathbb{1}_{d_\theta = n} = 1$ iff $d_\theta = n$.

3) Client Costs: Clients participating in cross-silo FL incur three main types of costs: privacy cost, computation cost, and communication cost, which we elaborate on as follows:

Privacy cost: Training with local data during global iteration incur privacy costs. For example, an adversary can infer information of clients' private training data through uploaded model updates by launching inference attacks and model inversion attacks [33]. Although recent studies applied methods such as differential privacy to mitigate this issue, they may not fully address the privacy concern [34]. In fact, cross-silo clients usually have stringent privacy requirements, e.g., hospitals can be highly sensitive to their patients' medical data [2]. We use $C_n^{\text{pri}}(x_n) = \mu_n \cdot f_n(x_n)$ to denote client n's privacy cost, where $\mu_n \geq 0$ represents a client's privacy sensitivity and $f_n(x_n)$ increases in x_n. Intuitively, if a privacy attack succeeds, a client experiences more data leakage (i.e., larger privacy cost) when it uses more data for FL training.

Computation and communication costs: Cross-silo FL also incurs computation costs (e.g., for model training) and communication costs (e.g., for model uploading and downloading). However, cross-silo clients are companies or organizations who are likely to have strong computational resources (e.g., powerful servers) and reliable communication channels (e.g., high-speed wired connections) [2]. Hence, we normalize the computation and communication costs to be zero, as they are less major concerns in cross-silo FL. We concur that in cross-device FL, however, computation and communication costs can be the bottleneck since edge devices are usually resource constrained.

Authorized licensed use limited to: Univ of Calif Davis. Downloaded on January 20,2024 at 01:48:26 UTC from IEEE Xplore. Restrictions apply.

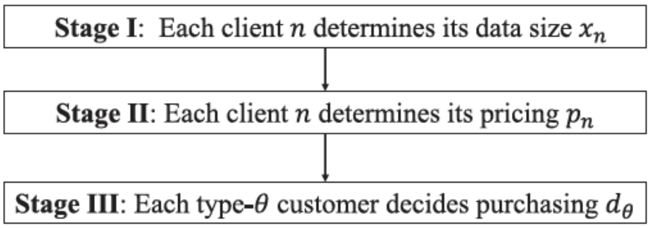

Fig. 2. Three-stage stackelberg game.

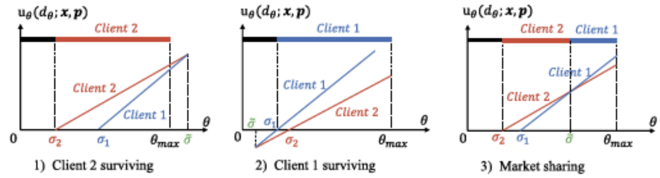

Fig. 3. Three market partition modes in duopoly market.

TABLE I
KEY NOTATIONS

Decision Variables	
d_θ	purchasing strategy of type-θ customer
$\boldsymbol{p} = (p_1, p_2)$	pricing strategies of two clients
$\boldsymbol{x} = (x_1, x_2)$	data contributions of two clients used for FL
Parameters	
θ	customer type (i.e., valuation)
$h(\cdot)$, $H(\cdot)$	PDF and CDF of customer type
μ_n	privacy sensitivity of client n
Functions	
$u_\theta(d_\theta; \boldsymbol{x}, \boldsymbol{p})$	payoff of type-θ customer
$A^{\mathrm{g}}(\boldsymbol{x})$	global model accuracy
$A_n(\boldsymbol{x})$	client n's local model accuracy after fine-tuning
$R_n(\boldsymbol{x}, \boldsymbol{p})$	revenue of client n
$C_n^{\mathrm{pri}}(x_n)$	privacy cost of client n
$W_n(\boldsymbol{x}, \boldsymbol{p})$	profit of client n

4) Client Profit: A client n's profit (i.e., objective function) is defined as the difference between the revenue from the customers and the privacy cost:

$$W_n(\boldsymbol{x}, \boldsymbol{p}) = R_n(\boldsymbol{x}, \boldsymbol{p}) - C_n^{\mathrm{pri}}(x_n). \quad (3)$$

D. Three-Stage Game Formulation

We model the interactions among cross-silo clients and customers as a three-stage game (see Fig. 2). The clients decide the data contribution $\boldsymbol{x}$ used for FL training in Stage I and the pricing strategies $\boldsymbol{p}$ in Stage II, with each client aiming to maximize its own profit in (3). In Stage III, each customer decides the purchasing strategies d_θ to maximize his own payoff in (1). We analyze the game via backward induction.

We further summarize the key notations in Table I.

IV. Solving Duopoly Competition

In this section, we solve the three-stage game. We will solve the customers' purchasing, the clients' pricing, and the clients' data contributions in Sections IV-A, IV-B, IV-C, respectively.

A. Stage III: Customer Purchasing in Duopoly

Given $\boldsymbol{x}$ and $\boldsymbol{p}$, each customer chooses his purchasing strategy to maximize his own payoff in (1). Proposition 1 computes the customer's optimal purchasing strategies.

Proposition 1: Given $\boldsymbol{x}$ and $\boldsymbol{p}$, a type-θ customer's optimal purchasing strategy $d_\theta^*(\boldsymbol{x}, \boldsymbol{p})$ is

$$d_\theta^*(\boldsymbol{x}, \boldsymbol{p}) = \begin{cases} 1, & \text{if } \theta A_1(\boldsymbol{x}) - p_1 \geq \max\{\theta A_2(\boldsymbol{x}) - p_2, 0\}, \\ 2, & \text{if } \theta A_2(\boldsymbol{x}) - p_2 \geq \max\{\theta A_1(\boldsymbol{x}) - p_1, 0\}, \\ \emptyset, & \text{else}. \end{cases} \quad (4)$$

Due to space limitations, we defer all the technical proofs to the online appendix (available online).

Proposition 1 shows that in a duopoly market, a type-θ customer will buy service from client n if client n can bring him a larger (among the two clients) and non-negative payoff.

To better illustrate the insights, we present the market partition [19], which gives a holistic picture of customers' purchasing decisions among two competing clients. To facilitate the analysis in Stage III (and Stage II), we make Assumption 1, i.e., given $\boldsymbol{x}$, client 1 has a better local model than client 2.

Assumption 1: Without loss of generality, we assume $A_1(\boldsymbol{x}) > A_2(\boldsymbol{x})$, where $\boldsymbol{x}$ is clients' data contributions.

Assumption 1 is without loss of generality. We can similarly analyze the case where $A_1(\boldsymbol{x}) < A_2(\boldsymbol{x})$ by switching clients' indices. Also, the case where $A_1(\boldsymbol{x}) = A_2(\boldsymbol{x})$ corresponds to the renowned Bertrand competition [35], which is a degeneration of the case $A_1(\boldsymbol{x}) > A_2(\boldsymbol{x})$ or $A_1(\boldsymbol{x}) < A_2(\boldsymbol{x})$.

Before we present the market partition, we introduce two definitions for notation simplicity.

Definition 1 (Threshold Customer Type): The client n's threshold customer type $\sigma_n \in [0, \theta_{\max}]$ corresponds to the customer who obtains a zero payoff, i.e., $u_{\sigma_n}(n; \boldsymbol{x}, p_n) = 0$. Hence,

$$\sigma_n(\boldsymbol{x}, p_n) = \frac{p_n}{A_n(\boldsymbol{x})}. \quad (5)$$

Based on Definition 1, any customer with valuation $\theta \geq \sigma_n$ achieves a non-negative payoff if purchasing from client n.

Definition 2 (Neutral Customer Type): Denote $\tilde{\sigma}$ as the neutral customer type. A type-$\tilde{\sigma}$ customer obtains the same payoff by purchasing service from either of the clients, i.e., $u_{\tilde{\sigma}}(1; \boldsymbol{x}, \boldsymbol{p}) = u_{\tilde{\sigma}}(2; \boldsymbol{x}, \boldsymbol{p})$. Under Assumption 1, we can derive $\tilde{\sigma}$ as follows:

$$\tilde{\sigma}(\boldsymbol{x}, \boldsymbol{p}) = \frac{\sigma_1(\boldsymbol{x}, p_1) A_1(\boldsymbol{x}) - \sigma_2(\boldsymbol{x}, p_2) A_2(\boldsymbol{x})}{A_1(\boldsymbol{x}) - A_2(\boldsymbol{x})}, \quad (6)$$

where $\sigma_1(\boldsymbol{x}, p_1)$ and $\sigma_2(\boldsymbol{x}, p_2)$ are two clients' threshold customer types defined in Definition 1.

Based on Definitions 1 and 2, we summarize the market partition.

Theorem 1 (Market Partition Equilibrium in Stage III): Under Assumption 1, given $\boldsymbol{x}$ and $\boldsymbol{p}$, the market partition equilibrium denoted by $\Phi_1^*(\boldsymbol{x}, \boldsymbol{p})$ and $\Phi_2^*(\boldsymbol{x}, \boldsymbol{p})$ is (see also Fig. 3):

1) If $\sigma_1(\boldsymbol{x}, p_1) > \sigma_2(\boldsymbol{x}, p_2)$ and $\tilde{\sigma}(\boldsymbol{x}, \boldsymbol{p}) \geq \theta_{\max}$, then

$$\Phi_1^*(\boldsymbol{x}, \boldsymbol{p}) = \emptyset, \quad \Phi_2^*(\boldsymbol{x}, \boldsymbol{p}) = [\sigma_2(\boldsymbol{x}, p_2), \theta_{\max}]. \quad (7)$$

2) If $\sigma_1(\boldsymbol{x}, p_1) \leq \sigma_2(\boldsymbol{x}, p_2)$, then

$$\Phi_1^*(\boldsymbol{x}, \boldsymbol{p}) = [\sigma_1(\boldsymbol{x}, p_1), \theta_{\max}], \quad \Phi_2^*(\boldsymbol{x}, \boldsymbol{p}) = \emptyset. \quad (8)$$

Authorized licensed use limited to: Univ of Calif Davis. Downloaded on January 20,2024 at 01:48:26 UTC from IEEE Xplore. Restrictions apply.

3) If $\sigma_1(x,p_1) > \sigma_2(x,p_2)$ and $\tilde{\sigma}(x,p) < \theta_{\max}$, then

$$\Phi_1^*(x,p) = [\tilde{\sigma}(x,p), \theta_{\max}],$$
$$\Phi_2^*(x,p) = [\sigma_2(x,p_2), \tilde{\sigma}(x,p)]. \quad (9)$$

We discuss the intuitions corresponding to the three cases in Theorem 1 as follows:

1) Client 1 has a much larger threshold type than client 2. This means that client 1 sets a much larger price (see (5)), and no customer purchases from client 1. Hence, only client 2 survives and has a non-empty market share.
2) Client 1 has a smaller threshold customer type than client 2. This implies that client 1 sets a smaller price (or slightly larger price) than client 2, which together with a better model is more attractive to all potential customers. Hence, only client 1 survives and has a non-empty market share.
3) Client 1 has a reasonably larger threshold type (and price) than client 2, and hence both clients share the market.

B. Stage II: Client Pricing in Duopoly

In Stage II, given the data contributions x in Stage I, the clients decide the pricing strategies p to maximize their own profits in (3), anticipating the market partition equilibrium in Stage III.

By substituting $\Phi_1^*(x,p)$ and $\Phi_2^*(x,p)$ in Theorem 1 into (3), we derive the clients' profit functions below:

$$W_1(x,p) = A_1(x)\sigma_1(x,p_1)[1 - H(\tilde{\sigma}(x,p))] - C_1^{\text{pri}}(x_1), \quad (10)$$

$$W_2(x,p) = A_2(x)\sigma_2(x,p_2)[H(\tilde{\sigma}(x,p)) - H(\sigma_2(x,p_2))] - C_2^{\text{pri}}(x_2), \quad (11)$$

where $H(\cdot)$ is the CDF of customers' valuation distribution.

According to (6), (10), and (11), we can see that the clients' profits are uniquely determined by the data contributions x and the threshold customer types $\boldsymbol{\sigma} = (\sigma_1, \sigma_2)$. Moreover, based on (5), client n can achieve an arbitrary σ_n by adjusting p_n. Hence, the clients' price competition in Stage II can be equivalently formulated as the following threshold competition game.

Game 1 (Threshold Competition Game in Stage II): Given the data contributions x, the two clients' threshold competition in Stage II can be modeled as the following game:

- *Players:* client n for both $n = \{1, 2\}$.
- *Strategies:* each client n decides $\sigma_n \in [0, \theta_{\max}]$.
- *Objectives:* each clients n obtains a profit $W_n(x, \boldsymbol{\sigma})$.

We are interested in solving the Nash equilibrium (NE) of Game 1, which is defined below.

Definition 3 (Nash Equilibrium in Stage II): Given x, a strategy profile $\boldsymbol{\sigma}^*(x) \triangleq (\sigma_n^*(x), \sigma_j^*(x))$ constitutes a Nash equilibrium of Game 1 if for all $n \in \{1, 2\}$,

$$W_n(\sigma_n^*(x), \sigma_j^*(x)) \geq W_n(\sigma_n'(x), \sigma_j^*(x)),\ \forall \sigma_n'(x) \neq \sigma_n^*(x), \quad (12)$$

where $j \neq n$ and $j \in \{1, 2\}$.

At a NE, each client's strategy is a best response to the strategy played by the other client, i.e., the equilibrium is the fixed point of all clients' best response strategies [36].

To present closed-form solutions, we make Assumption 2.

Assumption 2: The customers' valuation θ follows a uniform distribution, i.e., $h(\theta) = 1/\theta_{\max},\ \forall\, \theta \in [0, \theta_{\max}]$.

Uniform distribution is commonly adopted in business competition [19] and marketing literature [37]. There are also empirical studies showing that customers' valuations can be approximated by uniform distribution [38]. It can also model where there is a lack of market research. Please refer to Appendix (available online) where we provide a detailed discussion on how to generalize to other distributions.

Next, we present the NE of Game 1 in Theorem 2.

Theorem 2 (Threshold Competition Equilibrium in Stage II): Under Assumptions 1, 2, given any x, the clients' threshold competition equilibrium $(\sigma_1^*(x), \sigma_2^*(x))$ is given by

$$\sigma_1^*(x) = \frac{2(A_1(x) - A_2(x))\theta_{\max}}{4A_1(x) - A_2(x)}, \quad \sigma_2^*(x) = \frac{\sigma_1^*(x)}{2}. \quad (13)$$

Moreover, the resulting market partition in Stage III is

$$\Phi_1^*(x) = [\tilde{\sigma}^*(x), \theta_{\max}], \quad \Phi_2^*(x) = [\sigma_2^*(x), \tilde{\sigma}^*(x)], \quad (14)$$

where $\tilde{\sigma}^*(x) = \frac{\sigma_1^*(x)A_1(x) - \sigma_2^*(x)A_2(x)}{A_1(x) - A_2(x)}$.

Insights: Theorem 2 sheds light on how cross-silo clients compete with each other, which we discuss as follows.

- Each client's equilibrium threshold type σ_n^* (and price p_n^* based on (5)) increases in its own model accuracy A_n but decreases in its competitor's model accuracy A_j, $j \neq n$. As client n has a better model, it will set a higher price to obtain more profits. However, if its competitor has a better model, it will lower the price to be more competitive.
- The two clients coexist in the duopoly market at equilibrium, i.e., both have a non-empty market share (see (14)). More specifically, client 1 with a better model obtains high-valuation customers at a higher price, while client 2 attracts low-valuation customers at a lower price. This result is consistent with many real life practices, e.g., hospitals with more accurate disease-diagnosis-service usually charge a higher price than those with less accurate services.

So far, we have solved the equilibrium in Stage II. Next, we solve the clients' data contribution design in Stage I.

C. Stage I: Client Data Contribution in Duopoly

In Stage I, the clients decide their data contributions $x = (x_1, x_2)$ for FL, considering the responses from Stages II and III. Note that we no longer need Assumption 1, which is only used without loss of generality to facilitate the analysis in Stages II and III. Note that the analysis in Stage I is very different from and more challenging than previous vertical differentiation studies (e.g., [39], [40]). Prior related studied typically assumed that clients independently decide their own qualities. In our work, however, the clients' qualities (i.e., model accuracy) are jointly decided by their data contribution levels through an indirect fashion.

We first derive the clients' profit functions by plugging in the equilibrium results from Stages II and III, i.e., for all n,

$$W_n(x) = R_n(x) - C_n^{\text{pri}}(x_n), \quad (15)$$

Authorized licensed use limited to: Univ of Calif Davis. Downloaded on January 20,2024 at 01:48:26 UTC from IEEE Xplore. Restrictions apply.

where

$$R_n(\boldsymbol{x}) = \begin{cases} \frac{4A_n^2(\boldsymbol{x})(A_n(\boldsymbol{x})-A_j(\boldsymbol{x}))\theta_{\max}}{(4A_n(\boldsymbol{x})-A_j(\boldsymbol{x}))^2}, & \text{if } A_n(x) > A_j(x), \\ \frac{A_j(\boldsymbol{x})A_n(\boldsymbol{x})(A_j(\boldsymbol{x})-A_n(\boldsymbol{x}))\theta_{\max}}{(4A_j(\boldsymbol{x})-A_n(\boldsymbol{x}))^2}, & \text{if } A_n(x) < A_j(x), \\ 0, & \text{if } A_n(x) = A_j(x), \end{cases} \tag{16}$$

for any $j \in \{1,2\}$ that satisfies $j \neq n$.

Next, we model the interactions between two clients in Stage I as a data contribution game.

Game 2 (Data Contribution Game in Stage I): The data contribution game in Stage I is defined as follows:

- *Players:* client n for both $n \in \{1,2\}$.
- *Strategies:* each client n chooses $x_n \in [0, D_n]$.
- *Objectives:* each client n's profit function $W_n(x)$ in (15).

We aim to solve Game 2's Nash equilibrium defined below.

Definition 4 (Nash Equilibrium in Stage I): A strategy profile $x^* \triangleq (x_n^*, x_j^*)$ is a Nash equilibrium of Game 2 if $\forall n \in \{1,2\}$,

$$W_n(x_n^*, x_j^*) \geq W_n(x_n', x_j^*),\ \forall x_n' \neq x_n^*, \tag{17}$$

where $j \neq n$ and $j \in \{1,2\}$.

With the general function forms of local model accuracy $A_n(\cdot)$ and privacy cost $C_n^{\text{pri}}(\cdot)$, there may not exist a Nash equilibrium of Game 2. Even if one exists, providing a closed-form equilibrium characterization is difficult, as it is challenging to characterize the impact of clients' data contribution strategies on the local model accuracy as well as the privacy cost. However, under some minor assumptions, we can guarantee the existence of a Nash equilibrium and analyze the equilibrium properties.

Assumption 3: For both $n \in \{1,2\}$, both the local model accuracy function $A_n(\cdot)$ and the global model accuracy function $A^g(\cdot)$ are concave; the privacy cost function $C_n^{\text{pri}}(\cdot)$ is convex.

Assumption 3 means that as a client contributes more data in FL training, it will experience more significant increments in privacy cost (i.e., loss aversion [41]), and a smaller marginal accuracy improvement. Nevertheless, our numerical experiments using the MNIST dataset showcase that the local model accuracy indeed exhibits a concave increasing trend in the data size for both i.i.d. and non-i.i.d. data (see Fig. 4 in Section VI).

We first show that a Nash equilibrium exists.

Lemma 1: Under Assumptions 2, 3, there exists a Nash equilibrium of Game 2.

Lemma 1 is based on the conclusion in [42] where there exists a NE if the game has a finite set of players, and the players' profit functions are quasi-concave in their strategies.

Next, we show the monotonicity property of the NE.

Theorem 3: Under Assumptions 2, 3, each client's equilibrium data contribution x_n^* weakly decreases in μ_n.

Theorem 3 implies that as a client is more privacy sensitive, it will use a smaller amount of data for FL training.

Before we end this section, we try to answer *Key Question 1*, i.e., *how will business competition affect cross-silo FL in terms of the global model accuracy and the clients' profits?* To this end, we compare duopoly with monopoly, which is a benchmark case where there are two markets with each having a monopoly client. Specifically, each market has the same population with size a half and is characterized by $h(\theta)$. A customer within market n can only purchase service from client n. The two clients do not have business competitions but can train FL models together. Due to space limitations, we defer the theoretical details of monopoly to Appendix (available online), and provide more numerical results in Section VI.

We summarize the answers to *Key Question 1* in Theorem 4.

Theorem 4: Under Assumptions 2, 3, compared to monopoly, duopoly competition leads to (i) a lower total client profit.

(iia) a better global model if $\max\{\mu_1, \mu_2\} \leq \underline{\mu}$ for some $\underline{\mu} > 0$.

(iib) a worse global model if $\min\{\mu_1, \mu_2\} \geq \bar{\mu}$ for some $\bar{\mu} > 0$.

Insights: Theorem 4 answers Key Question 1, which implies that duopoly competition hurts the clients' profits (Case (i)). Clients need to compete for customers via selling services at a lower price (than in monopoly). This leads to a lower profit. Following this result, one may guess that a lower profit decentivizes clients to contribute to FL, leading to a worse global model. However, it is surprising that duopoly can achieve a better global model when clients' privacy costs are low (Case (iia)). This is because even if competition limits profits, clients strive to survive in the market by improving its competitiveness. To this end, a client may not only lower the price but also contribute more data to obtain a better global and local model.

On the other hand, when clients' privacy costs are large, duopoly leads to a worse global model than monopoly (Case (iib)). This undesirable outcome is likely to occur in practice, as hospitals or banks usually have stringent privacy requirements (i.e., large privacy cost) [16]. To avoid such undesirable outcomes, we propose a revenue-sharing framework among the clients, which helps alleviate the competition and achieve a better global model and larger profit. We present the revenue-sharing framework and analyze its impact in Section V.

V. Duopoly With Revenue Sharing

In this section, we present the revenue-sharing framework to mitigate competition and enhance client collaboration. In Section V-A, we propose a general revenue-sharing framework. Then, we solve the three-stage game considering revenue-sharing. More specifically, we solve the customers' purchasing strategies, the clients' pricing design, the clients' data contribution design in Sections V-B, V-C, and V-D, respectively.

A. General Revenue-Sharing Framework

To mitigate the competition, we propose a framework where clients can share the revenues obtained from all customers. Our design is inspired by *contribution evaluation* (i.e., evaluating clients' contribution to the global model) [43], [44], [45] in FL. However, the existing contribution evaluation methods cannot be directly applied to our setting, since clients are strategic in both FL training strategies and pricing decisions.

Authorized licensed use limited to: Univ of Calif Davis. Downloaded on January 20,2024 at 01:48:26 UTC from IEEE Xplore. Restrictions apply.

We use the following mapping to denote a revenue-sharing mechanism:

$$\mathcal{G}: \boldsymbol{x} \mapsto \{I_n(\boldsymbol{x})\}_{n\in\{1,2\}}, \tag{18}$$

where $I_n(\boldsymbol{x}) \geq 0$ is client n's contribution index. Either the clients themselves or a trusted third party (e.g., the government) can perform the evaluation via running a sandbox simulation to estimate the effect of a client's contribution on the global model accuracy [46]. Contribution index $I_n(\boldsymbol{x})$ serves as a metric to evaluate how much a client contributes to FL, and it is a function of clients' data contributions used for FL training.

Remark: This article does not consider malicious clients who try to exaggerate their contributions (e.g., using fabricated or duplicated data in FL), since such behaviors can be detected by a trusted execution environment [47]. Also, we consider that the information of clients' contributed data $\boldsymbol{x}$ is shared among clients due to binding contracts or strict market regulations. It is a widely adopted assumption in literature (e.g., FedAvg [1] and FedProx [48]) where the server uses x_n to weight the clients' model updates during global aggregation.

Given $I_n(\boldsymbol{x})$, client n obtains a proportion $g_n(\boldsymbol{x})$ of the total revenues where

$$g_n(\boldsymbol{x}) = \frac{I_n(\boldsymbol{x})}{\sum_{n'\in\{1,2\}} I_{n'}(\boldsymbol{x})} \in [0,1], \quad \forall n \in \{1,2\}. \tag{19}$$

In practice, the clients themselves can negotiate a mutually agreed revenue-sharing mechanism $\mathcal{G}$, or a market regulator (e.g., the government) can design and implement $\mathcal{G}$.

Next, we solve the three-stage game taking into account the clients' revenue sharing.

B. Stage III: Customer Purchasing in Duopoly With Sharing

Note that the clients' revenue sharing does not change the decision space or the payoff function of customers. Hence, given $\boldsymbol{x}$, any sharing mechanism $\mathcal{G}$, and pricing $\boldsymbol{p}$, the equilibrium in Stage III can still be characterized by Theorem 1.

C. Stage II: Client Pricing in Duopoly With Sharing

With revenue sharing, each client n's new profit functions W_n^{share} are given as follows: for all $n \neq j \in \{1,2\}$,

$$W_n^{\text{share}}(\boldsymbol{x},\boldsymbol{p}) = g_n(\boldsymbol{x}) \cdot [R_n(\boldsymbol{x},\boldsymbol{p}) + R_j(\boldsymbol{x},\boldsymbol{p})] - C_n^{\text{pri}}(x_n), \tag{20}$$

where $g_n(\boldsymbol{x})$ is in (19), $R_n(\boldsymbol{x},\boldsymbol{p}) = W_n(\boldsymbol{x},\boldsymbol{p}) + C_n^{\text{pri}}(x_n)$, and $W_n(\boldsymbol{x},\boldsymbol{p})$ is given in (10), (11).

Similarly, we can formulate the pricing competition in Stage II as a threshold competition game in Game 1, except that the clients' profit functions are updated by the ones in (20). Next, we show the Stage II equilibrium in Theorem 5.

Theorem 5 (Threshold Competition Equilibrium in Stage II with Revenue Sharing): Under Assumptions 1, 2, given any $\boldsymbol{x}$ and any revenue-sharing mechanism $\mathcal{G}$, the clients' threshold competition equilibrium $(\sigma_1^{\text{share}*}(\boldsymbol{x}), \sigma_2^{\text{share}*}(\boldsymbol{x}))$ is given by

$$\sigma_1^{\text{share}*}(\boldsymbol{x}) = \sigma_2^{\text{share}*}(\boldsymbol{x}) = \theta_{\max}/2. \tag{21}$$

Moreover, the resulting market partition in Stage III is

$$\Phi_1^{\text{share}*}(\boldsymbol{x}) = [\theta_{\max}/2, \theta_{\max}], \quad \Phi_2^{\text{share}*}(\boldsymbol{x}) = \emptyset. \tag{22}$$

Note that Theorem 5 is a general result that applies to any revenue-sharing mechanism $\mathcal{G}$. Interestingly, only client 1 survives while client 2 has zero market share (see (22)). Revenue sharing incentivizes the clients to collaboratively maximize the total revenue, and the best strategy to this end is for the "stronger" client 1 to obtain all potential customers. Even if client 2 has zero market share, it can still benefit from sharing a proportion of the total revenue.

However, the above pricing strategy may lead to collusion issues and violate the antitrust laws. In what follows, we restrict our analysis to where clients cannot conduct price collusion.[3] In this case, the Stage II equilibrium can still be characterized by Theorem 2.

Next, we solve the clients' data contribution in Stage I.

D. Stage I: Client Data Contribution in Duopoly With Sharing

In Stage I, the two clients decide their data contributions. We derive the clients' profit functions considering the responses from Stages II and III, i.e., for all $n \in \{1,2\}, j \neq n$,

$$W_n^{\text{share}}(\boldsymbol{x}) = g_n(\boldsymbol{x}) \cdot [R_n(\boldsymbol{x}) + R_j(\boldsymbol{x})] - C_n^{\text{pri}}(x_n), \tag{23}$$

where $R_n(\boldsymbol{x})$ and $R_j(\boldsymbol{x})$ are given in (16).

We model the interactions between clients in Stage I as a data contribution game similar to Game 2, except that the profit functions are replaced by $W_n^{\text{share}}(\boldsymbol{x})$ in (23).

Next, we try to characterize the NE. With the general form of $A_n(\cdot)$, $C_n^{\text{pri}}(\cdot)$, and $g_n(\cdot)$, it is highly challenging to analyze the existence or properties of NE. Nevertheless, we are able to show the existence of NE under a mild assumption below.

Assumption 4: $\frac{\partial^2(g_n \cdot R_n(A_n))}{\partial x_n \partial x_j} \leq 0, \quad \text{for all} \quad j \neq n.$

Assumption 4 is not restrictive, as it applies to revenue sharing mechanisms such as the *egalitarian* mechanism, where two clients equally share the revenue, i.e., $g_n = 0.5, \forall n$.[4]

Next, we discuss the existence of NE in Lemma 2.

Lemma 2: Under Assumptions 2–4, there exists an NE of the data contribution game considering revenue sharing.

The key to Lemma 2 is that we can equivalently formulate the game as a supermodular game, in which according to [49], a best response algorithm converges to a NE.

We further characterize the properties of NE in Theorem 6.

Theorem 6: Under Assumptions 2–4, Theorem 3 continues to hold when we consider revenue sharing among clients.

With revenue sharing, each client' equilibrium data contribution is still non-increasing in μ_n. A more privacy sensitive client will use a smaller amount of data for FL training.

An important question is how revenue sharing affects the global model accuracy and the client profit at equilibrium. We

[3] For completeness of our analysis, we have also included the analysis where collusion is possible. This can provide insights on how to detect and prevent price collusion in practical cross-silo federated learning systems. We leave the details to Appendix (available online).

[4] We have tested various other revenue-sharing mechanisms in the numerical experiments, and find that the best response always converges to a NE.

Authorized licensed use limited to: Univ of Calif Davis. Downloaded on January 20,2024 at 01:48:26 UTC from IEEE Xplore. Restrictions apply.

note that a general analytical characterization is challenging, as there is a lack of the sharing mechanism $g_n(\cdot)$. Also, revenue sharing in the context of cross-silo federated learning is further complicated by the contribution evaluation process (see Section V-A). We will numerically study several practical sharing mechanisms and their impact on the equilibrium in Section VI-C.

VI. Numerical Results

We conduct extensive numerical experiments to gain more useful insights. Specifically, in Section VI-A, we introduce the experimental setup. In Section VI-B, we investigate how the clients' data contribution strategies affect cross-silo FL. In Section VI-C, we compare the equilibrium results of monopoly, duopoly, and duopoly with revenue sharing.

A. Experimental Setup

We conduct numerical experiments for both the i.i.d. and non-i.i.d. cases. We introduce the experimental setup for the two cases separately as follows:

1) Experimental Setup for I.i.d. Case: We use MNIST to train FL models with FedAvg. The MNIST dataset contains 10^6 handwritten digits, and it has been used in many existing FL studies [6]. To simulate the i.i.d. data partition, we assign each client 2600 data points randomly sampled from MNIST.

The key hyper-parameters are as follows [50]. We use $LeNet$-5 as the model structure, set local epoch number to 5, local batch size to 256, local learning rate to 0.1, global learning rate to 0.5. Furthermore, since the MNIST dataset is relatively easy to train and we consider two clients, the global model converges quickly in dozens of rounds, and hence we set the communication round to 50. After 50 rounds, each client uses all local data to fine-tune the global model for 5 epochs, and each obtains a final local model.

2) Experimental Setup for Non-I.i.d. Case: We use CIFAR-10 [51] to train FL models using FedAvg. The CIFAR-10 dataset contains 10 classes with $6 * 10^4$ data points. To simulate the non-i.i.d. case, we apply the widely used Dirichlet distribution with a controlling parameter $\alpha > 0$ [52]. Here, a smaller α means that clients' label distributions are less similar, i.e., a higher degree of non-i.i.d. In the experiment, we assign each client 6000 data points with $\alpha = 10$.

The key hyper-parameters are as follows. We use $ResNet-18$ as the model structure. We set the local epoch number as 5 and batch size as 64. We set the local and global learning rates to be 0.1 and 1, respectively. We set the communication round to be 50. After 50 rounds, each client uses all local data to fine-tune the global model for 5 epochs, and each obtains a final local model.

B. Impact of Data Contribution on Cross-Silo FL

We first investigate how clients' data contributions affect cross-silo FL in terms of the global model and the local fine-tuned models.

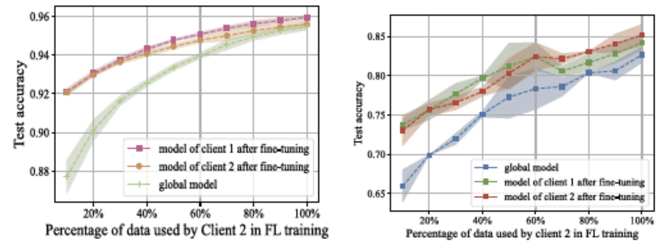

(a) Model accuracies on i.i.d. data. (b) Model accuracies on non-i.i.d. data.

Fig. 4. Impact of clients' data contributions on model accuracies.

In the i.i.d. case, client 1 uses 10% of local data for FL, while we change the client 2's used data size as $x_2 \in 2600 \cdot \{10\%, 20\%, \ldots, 100\%\}$, and for each x_2 we repeat the entire training process for 5 times. In the non-i.i.d. case, client 1 uses 60% of local data for FL, while we change client 2's used data size as $x_2 \in 6000 \cdot \{10\%, 20\%, \ldots, 100\%\}$. For each x_2 we repeat the entire training process for 5 times.[5] Fig. 4(a) and (b) plot the i.i.d. and non-i.i.d. cases respectively, how the global model accuracy and local model accuracy (after fine-tuning) depend on x_2. In Fig. 4(a)–(b), the lines present the average accuracies over 5 runs, and the shaded areas reflect the variances.

First, in both Fig. 4(a) and (b), we observe that the average accuracies of local and global models generally concavely increase in x_2. This observation is consistent with Assumption 3. Second, we observe that compared to i.i.d. data, non-i.i.d. data introduces larger variance (i.e., larger shaded areas). This is mainly due to the dissimilar data distributions among the two clients [10]. Nevertheless, we see that both clients' local models have a higher average accuracy than the global model (e.g., Fig. 4(a)). This demonstrates the usefulness of local fine-tuning in terms of improving the global model performance.

We summarize the key observations in Fig. 4 below.

Observation 1: (i) Both the local and global model accuracies concavely increase in clients' data contributions.

(ii) Local fine-tuning helps improve the global model performance.

C. Monopoly Vs. Duopoly Vs. Duopoly With Revenue Sharing

Next, we compare the equilibrium results of monopoly, duopoly, and Duopoly with Revenue Sharing (D-w-S). To evaluate our D-w-S framework, we use four benchmarks below, where the superscript of the contribution index represents the corresponding benchmark.

- *Egalitarian (EG)* [46]: clients equally share the revenues:

$$I_n^{\mathrm{EG}} = 0.5, \quad \forall n. \tag{24}$$

[5]Please refer to Appendix (available online) where we report all the training results.

Authorized licensed use limited to: Univ of Calif Davis. Downloaded on January 20,2024 at 01:48:26 UTC from IEEE Xplore. Restrictions apply.

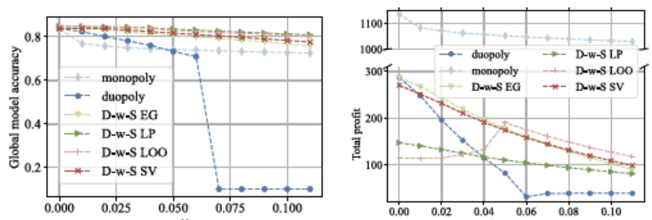

(a) Global model accuracy vs. μ. (b) Clients' total profit vs. μ.

Fig. 5. Impact of clients' privacy sensitivity on i.i.d. data.

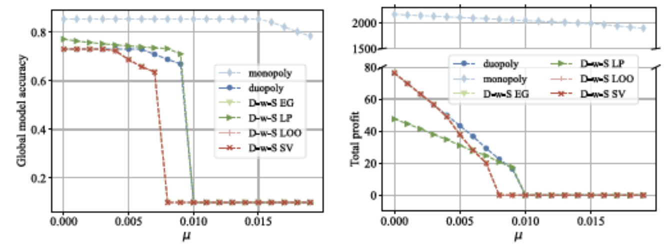

(a) Global model accuracy vs. μ. (b) Clients' total profit vs. μ.

Fig. 7. Impact of clients' privacy sensitivity on non-i.i.d. data.

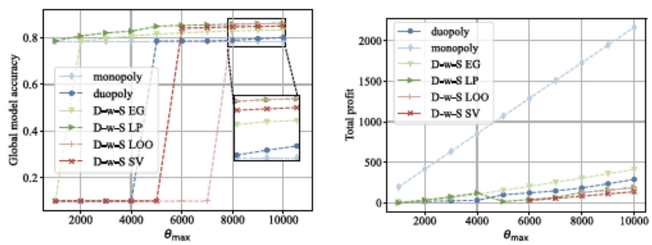

(a) Global model accuracy vs. $\theta_{\max}$. (b) Clients' total profit vs. $\theta_{\max}$.

Fig. 6. Impact of customers' valuation on i.i.d. data.

- *Linearly Proportional (LP)* [53]: a client's contribution index is linearly proportional to its data contribution:

$$I_n^{\text{LP}}(\boldsymbol{x}) = x_n, \quad \forall n. \tag{25}$$

- *Leave-One-Out (LOO)* [54]: the index is calculated by the marginal contribution to the global model accuracy:

$$I_n^{\text{LOO}}(\boldsymbol{x}) = A^g(\{x_i\}_{i\in\mathcal{N}}) - A^g(\{x_j\}_{j\in\mathcal{N}\setminus\{n\}}). \tag{26}$$

- *Shapley Value (SV)* [55]: $I_n(\boldsymbol{x})$ is the average marginal contribution to FL calculated by the Shapley value method:

$$I_n^{\text{SV}}(\boldsymbol{x}) = \sum_{\mathcal{S}\subseteq\mathcal{N}\setminus\{n\}} \frac{A^g(\{x_i\}_{i\in\mathcal{S}\cup\{n\}}) - A^g(\{x_j\}_{j\in\mathcal{S}})}{N\binom{N-1}{|\mathcal{S}|}}. \tag{27}$$

We use a linear cost function, i.e., $C_n^{\text{pri}}(x_n) = \mu x_n$. It can characterize the risk associated with the privacy leakage in revealing x_n data points [17].

1) Equilibrium Results Under I.i.d. Case: We first compare the equilibrium results under the i.i.d. case. Figs. 5 and 6 plot, at equilibrium, how the average global model accuracy and clients' total profit change with clients' privacy sensitivity μ and customers' valuation $\theta_{\max}$, respectively.

Impact of Clients' Privacy Sensitivity: In Fig. 5(a)–(b), we observe that the global model accuracy and clients' total profit generally decrease in μ. Clients with a larger privacy sensitivity will contribute less data for FL, leading to a worse global model (and hence worse fined-tuned local models). The customers will pay less due to receiving service of lower quality, resulting in a lower client profit.

Impact of Customer Valuation: In Fig. 6(a)–(b), we see that both the global model accuracy and the clients' profits generally increase in $\theta_{\max}$. A larger $\theta_{\max}$ means that the customers have a higher average valuation toward the model-related services and hence are more willing to pay. This incentivizes the clients to use more data to train FL models and obtain better fine-tuned models. As a result, the clients obtain higher profits due to selling better services to customers at a higher price.

We summarize the above observations below.

Observation 2: Both the global model accuracy and clients' total profit tend to decrease in clients' privacy sensitivity, but they tend to increase in customers' valuations.

Monopoly vs. Duopoly: In Figs. 5(b) and 6(b), we see that duopoly significantly reduces the clients' profits (by more than 80% at $\mu = 0.04$ in Fig. 5(b)). Despite the profit reduction, it is interesting to see that duopoly achieves a better global model than monopoly when clients' privacy sensitivity is low (e.g., $\mu = 0.02$ in Fig. 5(a)), or when customers' valuations are high (e.g., $\theta_{\max} = 10^4$ in Fig. 6(a)). In such cases, the clients strive to compete for potential customers by offering lower prices and contributing more data in FL, leading to a better global model.

We summarize the above observations below.

Observation 3: (i) Compared to monopoly, duopoly greatly limits the clients' profits.

(ii) Counter-intuitively, despite the profit reduction, it can lead to a better global model when clients' privacy sensitivity is low or when customers' valuations are high.

Duopoly vs. D-w-S: In Figs. 5 and 6, we observe that compared to duopoly, all D-w-S mechanisms can lead to a better global model (e.g., by up to 25.7% at $\mu = 0.06$ in Fig. 5(a)) and a much higher client profit (e.g., by more than 300% at $\mu = 0.8$ in Fig. 5(b)). This shows that revenue-sharing can enhance collaboration between clients.

In addition, among the 4 D-w-S mechanisms, LP achieves the best global model (e.g., $\theta_{\max} = 8000$ in Fig. 6(a)). A client's contribution index under LP is linear in its used data size, while it is sub-linear under SV/LOO [17] or constant under EG. That is, given data sizes, LP provides the largest proportion of shared revenue. Hence, LP incentivizes most data contribution and has the best global model.

We summarize the above key observations below.

Observation 4: (i) Compared to duopoly, D-w-S can greatly improve the global model accuracy and clients' total profit.

(ii) Among 4 benchmark D-w-S mechanisms, LP tends to lead to the best global model.

2) Equilibrium Results Under Non-I.i.d. Case: We now compare the equilibrium results under the non-i.i.d. case. Figs. 7 and 8 plot, at equilibrium, how the average global model accuracy

Authorized licensed use limited to: Univ of Calif Davis. Downloaded on January 20,2024 at 01:48:26 UTC from IEEE Xplore. Restrictions apply.

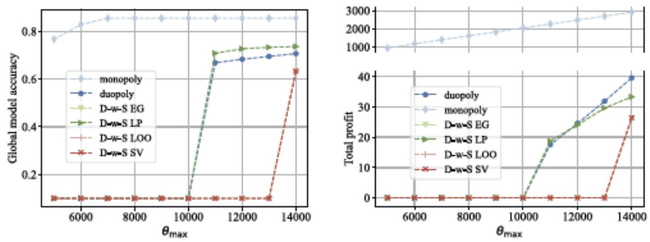

(a) Global model accuracy vs. $\theta_{\max}$. (b) Clients' total profit vs. $\theta_{\max}$.

Fig. 8. Impact of customers' valuation on non-i.i.d. data.

and clients' total profit change with clients' privacy sensitivity μ and customers' valuation $\theta_{\max}$, respectively.

We notice that the key results under the i.i.d. case continue to hold in the non-i.i.d. case. For example, in Fig. 7, we observe that both the global model accuracy and the total profit generally decrease in clients' privacy sensitivity. In addition, in Fig. 8, we observe that the global model accuracy and the total profit generally increase in customers' valuation. These results are consistent with Observation 2. Also, in both Figs. 7(b) and 8(b), we again see that compared to monopoly, duopoly significantly reduces the clients' profits due to fierce market competition, validating Observation 3(i). Nevertheless, with proper revenue sharing (e.g., LP), the clients have better incentives to contribute to FL training. This can lead to a better global model (e.g., $\mu = 0.005$ in Fig. 8(a)), and a larger client profit (e.g., $\theta_{\max} = 11000$ in Fig. 8(b)).

VII. Conclusion

This article studies the overlooked but important problem of business competition in cross-silo FL. The problem is challenging due to the conflicting objectives and complex coupling between clients' FL collaboration and pricing competition. We manage to characterize the equilibrium properties and show that market competition always reduces clients' profits, and it can damage the global model accuracy under certain conditions. To address this issue, we propose a general revenue-sharing framework that is shown to effectively enhance client collaboration. Surprisingly, we show that market competition can induce clients to use more data and generate a better global model.

For the future work, it is important to extend duopoly to a more general oligopoly framework where there are more than two competing clients. Also, it is interesting to study more sophisticated customer behaviors (e.g., multi-homing where a customer can buy services from more than one client).

References

[1] P. Kairouz et al., "Advances and open problems in federated learning," *Foundations Trends® Mach. Learn.*, vol. 14, no. 1/2, pp. 1–210, 2021.
[2] C. Huang, J. Huang, and X. Liu, "Cross-silo federated learning: Challenges and opportunities," 2022, *arXiv:2206.12949*.
[3] [Online]. Available: https://www.lifeinsuranceinternational.com/news/swiss-re-webank/
[4] [Online]. Available: https://owkin.com/federated-learning/
[5] 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.
[6] M. Tang and V. W. S. Wong, "An incentive mechanism for cross-silo federated learning: A public goods perspective," in *Proc. IEEE Conf. Comput. Commun.*, 2021, pp. 1–10.
[7] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," *ACM Trans. Intell. Syst. Technol.*, vol. 10, no. 2, pp. 1–19, 2019.
[8] 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.
[9] [Online]. Available: https://www.swissre.com/en/china/
[10] Y. Huang et al., "Personalized cross-silo federated learning on non-IID data," in *Proc. AAAI Conf. Artif. Intell.*, 2021, pp. 7865–7873.
[11] X. Zhang, Y. Li, W. Li, K. Guo, and Y. Shao, "Personalized federated learning via variational Bayesian inference," in *Proc. Int. Conf. Mach. Learn.*, 2022, pp. 26293–26310.
[12] A. Xu et al., "Closing the generalization gap of cross-silo federated medical image segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 20866–20875.
[13] O. Marfoq, C. Xu, G. Neglia, and R. Vidal, "Throughput-optimal topology design for cross-silo federated learning," in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, pp. 19478–19487.
[14] M. A. Heikkilä, A. Koskela, K. Shimizu, S. Kaski, and A. Honkela, "Differentially private cross-silo federated learning," 2020, *arXiv:2007.05553*.
[15] C. Zhang, S. Li, J. Xia, W. Wang, F. Yan, and Y. Liu, "BatchCrypt: Efficient homomorphic encryption for Cross-Silo federated learning," in *Proc. USENIX Conf. Annu. Tech. Conf.*, 2020, pp. 493–506.
[16] C. Wang, X. Wu, G. Liu, T. Deng, K. Peng, and S. Wan, "Safeguarding cross-silo federated learning with local differential privacy," *Digit. Commun. Netw.*, vol. 8, pp. 446–454, 2022.
[17] C. Huang, H. Zhang, and X. Liu, "Incentivizing data contribution in cross-silo federated learning," 2022, *arXiv:2203.03885*.
[18] M. Chen, Y. Liu, W. Shen, Y. Shen, P. Tang, and Q. Yang, "Mechanism design for multi-party machine learning," in *Proc. Int. Workshop Front. Algorithmics*, Cham, Aug. 2022, pp. 248–268. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-031-20796-9_18
[19] 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.
[20] R. T. B. Ma, "Usage-based pricing and competition in congestible network service markets," *IEEE/ACM Trans. Netw.*, vol. 24, no. 5, pp. 3084–3097, Oct. 2016.
[21] S. Chowdhury, "Resource allocation in cognitive radio networks using Stackelberg game: A survey," *Wireless Pers. Commun.*, vol. 122, no. 1, pp. 807–824, 2022.
[22] Z. Zhang, C. Guan, H. Chen, X. Yang, W. Gong, and A. Yang, "Adaptive privacy-preserving federated learning for fault diagnosis in internet of ships," *IEEE Internet Things J.*, vol. 9, no. 9, pp. 6844–6854, May 2022.
[23] 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.
[24] 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. Statist.*, 2017, pp. 1273–1282.
[25] S. Ke, C. Huang, and X. Liu, "On the impact of label noise in federated learning," in *Proc. IEEE WiOpt*, 2023.
[26] T. Yu, E. Bagdasaryan, and V. Shmatikov, "Salvaging federated learning by local adaptation," 2020, *arXiv:2002.04758*.
[27] 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 Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 10174–10183.
[28] L. Collins, H. Hassani, A. Mokhtari, and S. Shakkottai, "Exploiting shared representations for personalized federated learning," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 2089–2099.
[29] S. P. Karimireddy, S. Kale, M. Mohri, S. Reddi, S. Stich, and A. T. Suresh, "Scaffold: Stochastic controlled averaging for federated learning," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 5132–5143.
[30] 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.

Authorized licensed use limited to: Univ of Calif Davis. Downloaded on January 20,2024 at 01:48:26 UTC from IEEE Xplore. Restrictions apply.

[31] 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.
[32] 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, 2018.
[33] Y. Huang, S. Gupta, Z. Song, K. Li, and S. Arora, "Evaluating gradient inversion attacks and defenses in federated learning," in *Proc. Adv. Neural Inf. Process. Syst.*, 2021, pp. 7232–7241.
[34] K. Wei et al., "Federated learning with differential privacy: Algorithms and performance analysis," *IEEE Trans. Inf. Forensics Secur.*, vol. 15, pp. 3454–3469, 2020.
[35] R. Deneckere and C. Davidson, "Incentives to form coalitions with Bertrand competition," *RAND J. Econ.*, vol. 16, pp. 473–486, 1985.
[36] M. J. Osborne and A. Rubinstein, *A Course in Game Theory*. Cambridge, MA, USA: MIT Press, 1994.
[37] T. Dai, R. Ke, and C. T. Ryan, "Incentive design for operations-marketing multitasking," *Manage. Sci.*, vol. 67, no. 4, pp. 2211–2230, 2021.
[38] P. Lech and P. Włodarski, "Analysis of the IoT WiFi mesh network," in *Proc. Computer Sci. On-Line Conf.*, 2017, pp. 272–280.
[39] A. Shaked and J. Sutton, "Relaxing price competition through product differentiation," *Rev. Econ. Stud.*, vol. 49, no. 1, pp. 3–13, 1982.
[40] O. Shy, *Industrial Organization: Theory and Applications*. Cambridge, MA, USA: MIT Press, 1996.
[41] N. C. Barberis, "Thirty years of prospect theory in economics: A review and assessment," *J. Econ. Perspectives*, vol. 27, no. 1, pp. 173–96, 2013.
[42] D. Fudenberg and J. Tirole, *Game Theory*. Cambridge, MA, USA: MIT Press, 1991.
[43] A. N. Bhagoji, S. Chakraborty, P. Mittal, and S. Calo, "Analyzing federated learning through an adversarial lens," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 634–643.
[44] A. Ghorbani and J. Zou, "Data shapley: Equitable valuation of data for machine learning," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 2242–2251.
[45] T. Song, Y. Tong, and S. Wei, "Profit allocation for federated learning," in *Proc. IEEE Int. Conf. Big Data*, 2019, pp. 2577–2586.
[46] H. Yu et al., "A fairness-aware incentive scheme for federated learning," in *Proc. AAAI/ACM Conf. AI Ethics Soc.*, 2020, pp. 393–399.
[47] X. Zhang, F. Li, Z. Zhang, Q. Li, C. Wang, and J. Wu, "Enabling execution assurance of federated learning at untrusted participants," in *Proc. IEEE Conf. Comput. Commun.*, 2020, pp. 1877–1886.
[48] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," in *Proc. Mach. Learn. Syst.*, 2020, pp. 429–450.
[49] D. M. Topkis, *Supermodularity and Complementarity*. Princeton, NJ, USA: Princeton Univ. Press, 1998.
[50] J. Wang et al., "A field guide to federated optimization," 2021, *arXiv:2107.06917*.
[51] [Online]. Available: https://www.cs.toronto.edu/~kriz/cifar.html
[52] 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*.
[53] 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.
[54] T. Wang, J. Rausch, C. Zhang, R. Jia, and D. Song, "A principled approach to data valuation for federated learning," in *Federated Learning*. Berlin, Germany: Springer, 2020, pp. 153–167.
[55] R. Jia et al., "Towards efficient data valuation based on the shapley value," in *Proc. 22nd Int. Conf. Artif. Intell. Statist.*, 2019, pp. 1167–1176.

Chao Huang received the Ph.D. degree in information engineering from the Chinese University of Hong Kong, Hong Kong, in 2021. During July–November 2021, he was a Postdoctoral Research Fellow with the Department of Management Sciences, City University of Hong Kong, Hong Kong. He is currently a Postdoctoral Researcher with the Department of Computer Science, University of California, Davis, Davis, CA, USA. His research interests include the spectrum of distributed machine learning, network economics, and low-carbon systems.

Shuqi Ke received the B.Sc. degree in mathematics from the Chinese University of Hong Kong, Shenzhen, China. His research interests include the broad area of differential privacy, reasoning with large language models, distribution shift, and multi-agent reinforcement learning.

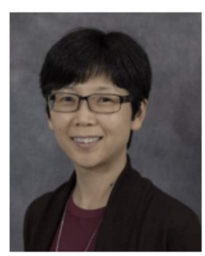

Xin Liu (Fellow, IEEE) received the Ph.D. degree in electrical engineering from Purdue University, West Lafayette, IN, USA, in 2002. She is currently a Professor of computer science with the University of California, Davis, Davis, CA, USA. Her research interests include machine learning algorithm development 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.

Authorized licensed use limited to: Univ of Calif Davis. Downloaded on January 20,2024 at 01:48:26 UTC from IEEE Xplore. Restrictions apply.