

Coopetition in Heterogeneous Cross-Silo Federated Learning

Chao Huang, Justin Dachille, and Xin Liu

Abstract. In cross-silo federated learning (FL), companies collaboratively train a shared global model without sharing heterogeneous data. Prior related work focused on algorithm development to tackle data heterogeneity. However, the dual problem of *coopetition*, i.e., FL collaboration and market competition, remains under-explored. This paper studies the FL coopetition using a dynamic two-period game model. In period 1, an incumbent company trains a local model and provides model-based services at a chosen price to users. In period 2, an entrant company enters, and both companies decide whether to engage in FL collaboration and then compete in selling model-based services at different prices to users. Analyzing the two-period game is challenging due to data heterogeneity, and that the incumbent's period one pricing has a temporal impact on coopetition in period 2, resulting in a non-concave problem. To address this issue, we decompose the problem into several concave sub-problems and develop an algorithm that achieves a global optimum. Numerical results on three public datasets show two interesting insights. First, FL training brings model performance gain as well as competition loss, and collaboration occurs only when the performance gain outweighs the loss. Second, data heterogeneity can incentivize the incumbent to limit market penetration in period 1 and promote price competition in period 2.

1 Introduction

Cross-silo federated learning (FL) is a distributed machine learning paradigm where multiple companies or organizations train a shared model collaboratively without directly exchanging local data [12]. Typically, the process involves each participant training a local model on their dataset and then sharing model updates with a coordinating server that aggregates these updates to improve a global model. This local training and aggregation iteration continues until the global model converges [11]. Potential application scenarios for cross-silo FL are abundant [7]. For example, in healthcare, hospitals can collaborate on medical research (e.g., disease diagnosis) without sharing sensitive patient data. In finance, different banks can use FL for improved fraud detection without exposing customer data. In smart manufacturing, companies can enhance predictive maintenance without revealing proprietary operational data.

* Chao Huang (corresponding author) is with the School of Computing, Montclair State University. Email: felixchaohuang@gmail.com. Justin Dachille and Xin Liu are with the Department of Computer Science, University of California, Davis. Email: {jtdachille,xinliu}@ucdavis.edu. This work was partially supported through USDA-020-67021-32855 and NSF OIA-2134901.

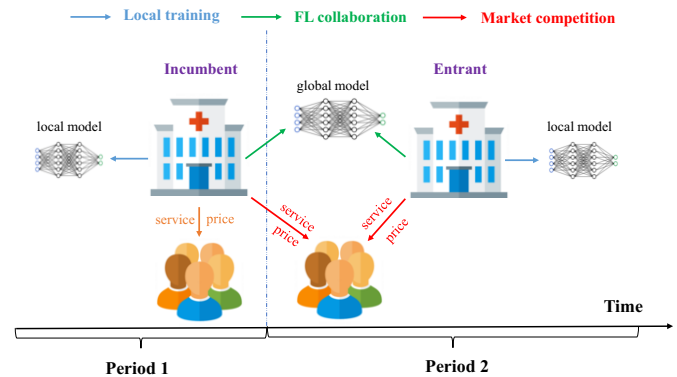


Figure 1: The two-period system model.

While a significant volume of recent research has focused on improving the model performance in cross-silo FL (e.g., [43, 20, 22]), the critical aspect of *market competition* remains less explored. In practice, companies can utilize the shared global model developed through FL to offer model-based services, hence competing for the same pool of potential users [41]. This competition can manifest across various industries. In healthcare, hospitals may compete to attract patients by offering diagnostic services or personalized treatment plans. In finance, banks might leverage an enhanced fraud detection system to attract customers by offering greater security and reduced risk of fraud. In smart manufacturing, companies could use the shared global model to optimize operations and attract clients/partnerships via more efficient production services. The dual focus of FL collaboration and market competition is termed as **FL coopetition**, which is the focus of this paper.

While a few recent studies looked at FL coopetition [10, 35, 39, 41], they focused on static competition and overlooked the important aspect of dynamic market entry. Market entrant usually serves as a catalyst for technology innovation (e.g., model and algorithm development in cross-silo FL) and helps provide insights into the interactions between incumbents and entrants [2]. This also impacts the socioeconomic landscape by promoting coopetition and increasing the quality of services (via FL) available to users in the market.

In this paper, we study FL coopetition using a dynamic market entry model that spans two time periods involving three entities: users, an incumbent company, and an entrant company (see Fig. 1):

- Period 1: Only an incumbent company exists in the market, providing users with the model-based service. For example, consider a healthcare provider that has developed a new AI-based tool for early diagnosis of a specific disease and is the only service

provider in the market. The company trains a local model on its dataset and offers services to patients at a chosen price.

- Period 2: An entrant company enters the market and coexists with the incumbent in providing user services. The incumbent and the entrant need to decide strategically on whether to engage in FL collaboration. If they collaborate, they will use a jointly developed global model to enhance and offer their diagnostic services. If they opt against collaboration, each will continue using individually trained models. In addition, since users can now choose between different service providers, there is a price competition between the incumbent and the entrant.

Given this dynamic model of FL cooperation, our first question is:

Question 1. *Will the incumbent collaborate with the entrant via FL in the presence of competition?*

To answer Question 1, we model the FL collaboration and market competition as two (intertwined) games between the incumbent and the entrant. Each company aims to maximize its profit (from selling model-based services to users). We are interested in solving the game-theoretical equilibrium, which is highly non-trivial due to two reasons. First, companies in practice are heterogeneous regarding their data distributions and quantities [29]. Second, users may exhibit heterogeneous preferences towards companies' services (even if the services have similar qualities). This can happen when some users find that a particular company's service matches more closely with their individual characteristics, e.g., demographic profiles and personal attributes [26]. It is challenging to analyze (or even model) the game. To address this issue, we resort to the renowned Hotelling model in economics [3] and generalize it to model various types of heterogeneity. We characterized the equilibrium with arbitrary data distributions/quantities of companies and under minor assumptions about users' preferences.

Among aforementioned heterogeneity types, data heterogeneity receives the most research attention from the FL community and continues to be the major bottleneck [42]. There are many excellent studies on algorithm development to mitigate the client drift issue caused by data heterogeneity, e.g., [16, 17, 32, 18]. We will show that our model and analysis are orthogonal to any FL algorithms. Since we focus on FL cooperation, our second key question is:

Question 2. *How does data heterogeneity affect FL collaboration and market competition?*

To answer Question 2, we solve the equilibrium of the two-period game and conduct numerical experiments on three public datasets under various levels of data heterogeneity. Notice that the decision-making of the incumbent's pricing in period 1 presents challenges due to its temporal impacts. That is, it not only influences user decisions in period 1, but also affects the decisions of both companies regarding FL collaboration (with heterogeneous data) and price competition, as well as user decisions in period 2. We will show that this interdependence makes the equilibrium analysis a challenging non-concave problem.

Our key contributions are summarized as follows:

- To our best knowledge, this is the first work to study FL cooperation from a temporal market entry perspective. We formulate a generic two-period model that accommodates heterogeneity including companies' data distributions and quantities as well as users' service preferences. Our model is also orthogonal to any FL algorithms.

- We provide game-theoretical solutions to the two-period model. The analysis involves solving a non-concave problem. To tackle this challenge, we decompose the problem into several manageable concave sub-problems and further develop an algorithm that achieves the global optimum.
- We conduct numerical experiments on three public datasets and show two interesting results. First, even under highly heterogeneous data, FL training improves model performance (compared to local learning). However, companies may avoid collaborating in the presence of competition, as FL collaboration also benefits the competitor. Second, data heterogeneity can incentivize the incumbent to limit market penetration in period 1 and promote price competition in period 2.

1.1 Related Work

Heterogeneous FL. Data heterogeneity, also commonly referred to as non-IID data,¹ is known to cause significant performance loss in FL due to the client drift issue. Many studies have focused on algorithm development to tackle client drift. One commonly used technique is regularization by adding proximal terms to restrain local updates with respect to the global model, e.g., FedProx [17], SCAFFOLD [13], MOON [16], FedUV [32]. Other works have approached the problem using personalization (e.g., [25, 28]) and local learning generality (e.g., [24]).

This paper focuses on how data heterogeneity affects FL cooperation. Our model and analysis are orthogonal to any FL algorithms mentioned above and beyond.

Incentives for cross-silo FL. Our work is related to game theoretical and, in particular, incentive studies for cross-silo FL, e.g., [36, 44, 6, 19, 7]. For example, Tang and Wong [36] proposed an auction-like mechanism to encourage organizations' FL training participation. Zhang et al. [44] studied how organizations participate in long-term collaboration. However, these studies overlooked the important aspect of market competition.

Only until very recently, a few papers have looked at market competition in cross-silo FL [10, 35, 41, 39]. Huang et al. [10] studied an oligopoly market. Wu and Yu [41] focused on a fully competitive market. Tan [35] developed an algorithm to find stable collaboration structures among companies. Tsoy and Konstantinov [39] studied both the price and quantity competitions in cross-silo FL. However, these studies did not study how data heterogeneity affects competition. More importantly, they looked at a static competition model while we study a dynamic model where the the incumbent and entrant have temporal interactions. The dynamic model renders the problem a challenging non-concave one.

Market competition in economics. Market competition is an extensively studied topic in economics, e.g., see [37, 34] for comprehensive discussions on theories of competition. Most pertaining to our work is the renowned Hotelling model (e.g., [26, 4]), where companies at different spatial locations compete for users. Our model and analysis differ from the conventional Hotelling due to the unique features of data heterogeneity and collaborative training in FL. That is, prior studies considered that each company can independently decide their service qualities to attract users. In cross-silo FL, however, the qualities of companies' model-based services are dependent on their heterogeneous data distributions and the FL collaboration strategies. This makes our analysis more challenging than and conclusions different from prior literature.

¹ Non-IID means not identically and independently distributed.

In summary, to our best knowledge, this paper is the first to study FL coopetition using a dynamic market entry model and also considering data heterogeneity.

The rest of the paper is organized as follows. Sec. 2 introduces the system model, and Sec. 3 provides both analytical and algorithmic solutions. Sec. 4 presents the numerical results. We discuss the extensions of our model in Sec. 5 and conclude in Sec. 6.

2 Model

Sec. 2.1 introduces the companies and users. Sec. 2.2 models their two-period game-theoretical interactions.

2.1 Companies and Users

Companies. We consider a two-period interaction with time index $t \in \mathcal{T} = \{1, 2\}$. In period $t = 1$, there is an incumbent company I who holds a private dataset \mathcal{D}_I with size $D_I = |\mathcal{D}_I|$. It trains a machine learning model using \mathcal{D}_I and provides model-based services (e.g., disease diagnosis) to users at a price $p_{I,1}$. In period $t = 2$, an entrant company E who has a local private dataset \mathcal{D}_E with size $D_E = |\mathcal{D}_E|$ enters the market.² Since both companies I and E are in the market, they can choose to collaborate via FL to improve the performance of their models and the quality of model-based services. We use $r_i \in \{1, 0\}$, $i \in \{I, E\}$ to denote the collaboration decision, where $r_i = 1$ means FL collaboration and $r_i = 0$ means no collaboration. More specifically,

- If $r_I \cdot r_E = 1$, both companies train a shared global model using \mathcal{D}_I and \mathcal{D}_E till convergence. They will use the converged model to generate model-based services.
- If $r_I \cdot r_E = 0$, each company uses its own data to train a local model, based on which services are generated.

We use $q_{i,t}$ to denote company i 's quality of service in period t . In particular, $q_{i,2}$ is a function of $\mathbf{r} = \{r_I, r_E\}$ that depends on both companies' decisions on FL collaboration (and data heterogeneity), which we refer to as $q_{i,2}(\mathbf{r})$ hereafter.

Besides the potential FL collaboration, both companies compete for the same users at prices $p_{i,2}$. For simplicity, let $\mathbf{p} = \{p_{I,1}, p_{I,2}, p_{E,2}\}$. Next, we introduce how users respond to companies' model-based services and prices.

Users. The market consists of a large number of users. To model this, we consider a continuum of users and normalize the total mass to be one. Each user decides whether to buy the service, and if so, from which company. Similar to [21], we assume that each user buys at most once, i.e., if a user buys from company I in period 1, it will not buy from company I or E in period 2 due to them offering substitutable services. Let $d_{n,t} \in \{\emptyset, i\}$ denote user n 's purchasing decision in period t , where $d_{n,t} = \emptyset$ means no purchasing and $d_{n,t} = i$ means purchasing from company i .

A user's payoff consists of three parts discussed below.

- First, a user obtains a utility from enjoying the model-based service, where the utility increases in the quality of service $q_{i,t}$. This paper considers a linear function (e.g., [21]). It is easy to extend the model to an arbitrary non-decreasing function (similar to [10]).
- Second, users have *heterogeneous preferences* towards companies' services. This can arise due to the traveling cost or brand loyalty over a particular company [26]. In the context of cross-silo




Figure 2: The two-period game.

FL, some users may find that a company's service matches more closely with their individual characteristics (e.g., demographic profiles and personal attributes). To model users' heterogeneous preferences, we use the renowned Hotelling model [3, 37], in which users are located on a line $[0, 1]$ and companies I and E are at the ending points $\phi = 0$ and $\phi = 1$, respectively. The users' locations ϕ_n follow a distribution with a PDF $h(\phi)$ and a CDF $H(\phi)$, which are known to both companies due to market research [21, 10]. Here, we model a user's *preference misalignment* as the distance to the company where the service is purchased.³ One can understand such preference misalignment as the dissimilarity between users' characteristics and the companies' data distributions.

- Third, a user needs to pay a price for the service.

Next, we define a user n 's payoff as follows:

$$u_{n,t}(d_{n,t}; \mathbf{r}, \mathbf{p}) = \begin{cases} 0, & \text{if } d_{n,t} = \emptyset, \\ w_q q_{I,t}(\mathbf{r}) - w_\phi \phi_n - w_p p_{I,t}, & \text{if } d_{n,t} = I, \\ \underbrace{w_q q_{E,t}(\mathbf{r})}_{\text{utility}} - \underbrace{w_\phi (1 - \phi_n)}_{\text{preference misalignment}} - \underbrace{w_p p_{E,t}}_{\text{price}}, & \text{if } d_{n,t} = E, \end{cases} \quad (1)$$

where w_q, w_ϕ, w_p are positive constants. Without loss of generality, we normalize w_ϕ to 1. Note that in period 1, only company I is in the market, so a user can only choose to buy from I or not buy.

2.2 The Two-Period Game

We model the interactions among the two companies and users as a two-period game (see Fig. 2). At the beginning of period 1, company I provides service using locally trained model with quality $q_{I,1}$ at price $p_{I,1}$. Then, users make purchasing decisions. In period 2, company E enters. Both companies decide whether to collaborate via FL \mathbf{r} , based on which they generate model-based services with quality $q_{i,2}(\mathbf{r})$. Then, they initiate price competition by selling services at prices $p_{i,2}$.

Users' decision problem. We consider that users are *myopic* [14, 21]. That is, in period 1, if a user finds purchasing from company I is better off than not purchasing, it will purchase without anticipating the potential entry of company E and its service in period 2. Hence, each user's decision problem is as follows.

$$\mathbf{P1}: \quad \max_{d_{n,t}} u_{n,t}(d_{n,t}; \mathbf{r}, \mathbf{p}), \quad \forall t. \quad (2)$$

Companies' decision problem. Both companies aim to maximize their own profit. For company E , its goal is to maximize the profit obtained in period 2, i.e.,

$$W_E(\mathbf{r}, \mathbf{p}) = \int_0^1 (p_{E,2} - c_E) \cdot \mathbb{1}_{d_{n,2}=E}(\mathbf{r}, \mathbf{p}) \cdot h(\phi) d\phi, \quad (3)$$

² See Sec. 5 for discussions on how to extend our model and analysis 1) with more than two companies, and 2) across $t \geq 3$ periods.

³ The Hotelling model and its variants (e.g., [30]) have been widely used to model user preferences in the economics literature.

where $\mathbb{1}$ is an indicator function, meaning $\mathbb{1}_{d_{n,2}=E} = 1$ if and only if $d_{n,2} = E$. Here, c_E is the marginal service cost.⁴ Hence, company E 's problem is as follows.

$$\mathbf{P2} : \max_{\mathbf{r}_E, \mathbf{p}_{E,2}} W_E(\mathbf{r}, \mathbf{p}). \quad (4)$$

Unlike users who are myopic, company I is assumed to be forward-looking [33].⁵ That is, it will strategically set the price in period 1, anticipating the reactions of users and company E in period 2. Here, we consider that company I knows the information about $q_{i,2}(\mathbf{r}), \forall \mathbf{r}$, which can be acquired via market research [21, 10]. In particular, company I 's aims to maximize its two-period profits:

$$W_{I,1}(p_{I,1}) = \int_0^1 (p_{I,1} - c_I) \cdot \mathbb{1}_{d_{n,1}=I}(p_{I,1}) \cdot h(\phi) d\phi, \quad (5)$$

$$W_{I,2}(\mathbf{r}, \mathbf{p}) = \int_0^1 (p_{I,2} - c_I) \cdot \mathbb{1}_{d_{n,2}=I}(\mathbf{r}, \mathbf{p}) \cdot h(\phi) d\phi, \quad (6)$$

where c_I is the marginal service cost. Hence, we formulate company I 's problem below.

$$\mathbf{P3} : \max_{p_{I,1}, p_{I,2}} W_I(\mathbf{r}, \mathbf{p}) = W_{I,1}(p_{I,1}) + W_{I,2}(\mathbf{r}, \mathbf{p}). \quad (7)$$

Next, we solve the two-period game.

3 Solving the Two-Period Game

We first solve the user decision in Sec. 3.1. Then, we solve the companies' period 2 decisions in Sec. 3.2. Sec. 3.3 presents an algorithm to solve company I 's pricing in period 1. Due to space limits, we leave all proofs to the appendix in the online technical report [9].

3.1 Users' Optimal Purchasing in Two Periods

Since users are myopic, we can solve users' optimal purchasing in the two periods separately.

Lemma 1. *Each user n 's optimal purchasing in period 1 is*

$$d_{n,1}^*(p_{I,1}) = \begin{cases} I, & \text{if } \phi_n \leq w_q q_{I,1} - w_p p_{I,1}, \\ \emptyset, & \text{else.} \end{cases} \quad (8)$$

Lemma 1 shows that a user is more likely to purchase from company I if the model-based service has a better quality $q_{I,1}$, or the price $p_{I,1}$ is lower, or the preference misalignment ϕ_n is smaller.

Lemma 2. *Assume $w_q q_{I,2}(\mathbf{r}) - w_p p_{I,2} \geq w_q q_{E,2}(\mathbf{r}) - w_p p_{E,2}$. Each user n 's optimal purchasing in period 2 is given in (9).*

The assumption in Lemma 2 is without loss of generality. One can similarly analyze the case where $w_q q_{I,2}(\mathbf{r}) - w_p p_{I,2} < w_q q_{E,2}(\mathbf{r}) - w_p p_{E,2}$. Importantly, Lemma 2 shows that users' purchasing in period 2 depends on companies' FL collaboration \mathbf{r} and their prices \mathbf{p} , including company I 's period 1 price $p_{I,1}$. As will be shown, such temporal dependence will render company I 's decision problem challenging.

⁴ Similar to [7], this paper normalizes the FL collaboration cost, in terms of computation and communication, to be zero. Our analysis and conclusions will not change even if we consider a non-zero cost.

⁵ This is reasonable due to companies usually having more resources and information than individual users when making decisions.

3.2 Companies' Decisions in Period 2

In period 2, given $p_{I,1}$ in period 1, both companies decide the FL collaboration \mathbf{r} and prices $\{p_{I,2}, p_{E,2}\}$ to maximize profits. We solve the pricing in Sec. 3.2.1 and FL collaboration in Sec. 3.2.2.

3.2.1 Companies' Pricing in Period 2

We first solve companies' optimal pricing in period 2, given $p_{I,1}$ and \mathbf{r} . Since both companies' prices will affect users' purchasing decisions and hence the company profits, the companies are playing a price competition game, which we model as follows.

Game 1. (Price Competition in Period 2) *The price competition game in period 2 is defined as a tuple $\langle \{I, E\}, \mathcal{P} = \prod p_{i,2}, \mathcal{W} = \prod W_{i,2} \rangle$, where each company i in $\{I, E\}$ decides the pricing $p_{i,2}$ to maximize its own profit $W_{i,2}$ in (3) and (6).*

We aim to find the Nash equilibrium (NE) of Game 1.

Definition 1. *A profile $\mathbf{p}_2^*(\mathbf{r}, p_{I,1}) = (p_{I,2}^*(\mathbf{r}, p_{I,1}), p_{E,2}^*(\mathbf{r}, p_{I,1}))$ is an NE of Game 1 if for all $i \in \{I, E\}$, $p_{i,2}^*(\mathbf{r}, p_{I,1}) \neq p_{i,2}^*(\mathbf{r}, p_{I,1})$,*

$$W_{i,2}(p_{i,2}^*(\mathbf{r}, p_{I,1}), p_{j,2}^*(\mathbf{r}, p_{I,1})) \geq W_{i,2}(p_{i,2}'(\mathbf{r}, p_{I,1}), p_{j,2}^*(\mathbf{r}, p_{I,1})), \quad (10)$$

where $j \neq i$ and $j \in \{I, E\}$.

NE is considered as a stable strategy profile as no company can achieve a higher profit via unilaterally changing its pricing. Notice that an arbitrary choice of $H(\cdot)$ can easily render the equilibrium analysis intractable. To facilitate theoretical analysis, we make some minor assumptions on $H(\cdot)$ (and $h(\cdot)$).

Assumption 1. *Users' preference distribution satisfies: (i) $h(\theta) > 0$ and is continuous. (ii) $h(\theta)/[1 - H(\theta)]$ is increasing in θ .*

Assumption 1 holds for many widely used distributions, e.g., Gaussian, uniform, and gamma distributions. Now, we discuss the equilibrium existence below.

Proposition 1. *Under Assumption 1, Game 1's NE exists.*

Proposition 1 can be proved by showing that the price competition game is a concave game, and in particular, each organization's profit is quasi-concave in its price. Refer to the online appendix [9] where we have developed a best response algorithm to compute the price equilibrium.

Next, we provide a result that establishes the relationship between the optimal pricing and companies' quality of service.

Corollary 1. *Each company i 's optimal pricing in period 2 is a non-decreasing function in $q_{i,2}(\mathbf{r})$ but a non-increasing function in $q_{j,2}(\mathbf{r})$, where $i, j \in \{I, E\}$, $i \neq j$.*

Corollary 1 can be proved via contradiction. It means that each company will set a higher price if it has a better quality of service, while it will lower the price if its competitor has a better quality.

3.2.2 Companies' FL Collaboration in Period 2

Now, we solve companies' optimal FL collaboration in period 2, given $p_{I,1}$ in period 1. The companies decide collaboration strategy \mathbf{r} for FL, anticipating the price equilibrium of Game 1. We model the two companies' interactions as a collaboration game.

$$d_{n,2}^*(\mathbf{r}, \mathbf{p}) = \begin{cases} I, & \text{if } d_{n,1}^*(p_{I,1}) = \emptyset, \phi_n \in \left[0, \max \left(\min \left(w_q q_{I,2}(\mathbf{r}) - w_p p_{I,2}, \frac{1+w_q(q_{I,2}(\mathbf{r})-q_{E,2}(\mathbf{r}))-w_p(p_{I,2}-p_{E,2})}{2}, 1 \right), 0 \right) \right], \\ E, & \text{if } d_{n,1}^*(p_{I,1}) = \emptyset, \phi_n \in \left[\min \left(\max \left(1 - w_q q_{E,2}(\mathbf{r}) + w_p p_{E,2}, \frac{1+w_q(q_{I,2}(\mathbf{r})-q_{E,2}(\mathbf{r}))-w_p(p_{I,2}-p_{E,2})}{2}, 0 \right), 1 \right), 1 \right], \\ \emptyset, & \text{else.} \end{cases} \quad (9)$$

Game 2. (FL Collaboration Game in Period 2) The FL collaboration game is a tuple $\langle \{I, E\}, \mathcal{R} = \prod r_i, \mathcal{W} = \prod W_{i,2} \rangle$, where each company i decides its FL collaboration r_i to maximize $W_{i,2}(\mathbf{r}, \mathbf{p}_2^*(\mathbf{r}, p_{I,1}))$, where $\mathbf{p}_2^*(\mathbf{r}, p_{I,1})$ is the NE of Game 1.

We aim to solve Game 2's NE defined below.

Definition 2. A strategy profile $\mathbf{r}^* = (r_i^*, r_j^*)$ is an NE of Game 2 if $\forall i \in \{I, E\}, \forall r_i' \neq r_i^*$,

$$W_{i,2}(\mathbf{r}^*, \mathbf{p}_2^*(\mathbf{r}^*, p_{I,1})) \geq W_{i,2}((r_i', r_j^*), \mathbf{p}_2^*((r_i', r_j^*), p_{I,1})), \quad (11)$$

where $j \neq i$ and $j \in \{I, E\}$.

Next, we characterize the NE in Proposition 2.

Proposition 2. The profile $(1, 1)$ is the NE of Game 2 if and only if $\forall i \in \{I, E\}, W_{i,2}((1, 1), \mathbf{p}_2^*((1, 1), p_{I,1})) \geq W_{i,2}((0, 0), \mathbf{p}_2^*((0, 0), p_{I,1}))$. Otherwise, $(0, 0)$ is the NE.

Proposition 2 means that the two companies will collaborate in FL training if and only if both achieves a higher profit than no collaboration. Next, we characterize a somewhat counter-intuitive result.

Theorem 1. Under Assumption 1, there exists a $p_{I,1}$ such that $(0, 0)$ is the NE of Game 2 even if $q_{i,2}((1, 1)) > q_{i,2}((0, 0)), \forall i$.

Theorem 1 shows that even if FL collaboration improves both companies' quality of service, it may not be the NE. This is because FL brings model performance gain and competition loss at the same time. That is, the shared FL model also improves the competitor's quality of service. According to Corollary 1, a company needs to set a lower price to attract users, which can lead to a lower profit. We will show in Sec. 4 that FL is the NE only when the model performance gain is significant and outweighs the competition loss.

3.3 Company I's Pricing in Period 1

We solve company I 's optimal pricing in period 1. In particular, company I decides $p_{I,1}$ to maximize its total profit in (7), anticipating the equilibria of the price competition game and FL collaboration game.

Note that solving this problem is challenging due to a few reasons. First, company I 's period-1-price $p_{I,1}$ affects users' decisions in both periods 1 and 2 (see (8)-(9)). This in turn will affect companies' price competition and FL collaboration in period 2, leading to a highly coupled analysis. Second, even if we established equilibrium existence in Propositions 1-2, we still lack closed-form equilibrium solutions, making the characterization of the profit function challenging. Third, it is easy to showcase that company I 's total profit is non-concave in $p_{I,1}$ if we use a Gaussian distribution for ϕ_n .

To address this challenge, we develop an algorithm to compute company I 's optimal pricing in period 1 in Algorithm 1. For ease of presentation, we use $\text{NE}_p(p_{I,1}, l, h)$ and $\text{NE}_r(p_{I,1}, l, h)$ to denote the equilibrium calculated on price $p_{I,1} \in [l, h]$ of Game 1 and Game 2, respectively. The calculation of NE_p is based on the a best response algorithm (details in online appendix [9]). The calculation of NE_r is based on Proposition 2. Note that even if the problem is non-concave, the proposed algorithm can return the global optimal solution.

Algorithm 1 Optimization of Period 1 Pricing

- 1: **initialize** $o_I, o_E, w_q, w_p, w_\phi$. Let $p_{I,1,a} = p_{I,1,b} = p_{I,1,c} = 0$.
- 2: **while** not convergent **do**
- 3: Compute $\text{NE}_p(p_{I,1,a}, 0, (w_q q_{I,1} - 1)/w_p)$ and $\text{NE}_r(p_{I,1,a}, 0, (w_q q_{I,1} - 1)/w_p)$
- 4: Update $p_{I,1,a}$ using gradient ascent
- 5: **end while**
- 6: Return converged solution as $p_{I,1,a}^*$
- 7: **while** not convergent **do**
- 8: Compute $\text{NE}_p(p_{I,1,b}, (w_q q_{I,1} - 1)/w_p, w_q q_{I,1}/w_p)$ and $\text{NE}_r(p_{I,1,b}, (w_q q_{I,1} - 1)/w_p, w_q q_{I,1}/w_p)$
- 9: Update $p_{I,1,b}$ using gradient ascent
- 10: **end while**
- 11: Return converged solution as $p_{I,1,b}^*$
- 12: **while** not convergent **do**
- 13: Compute $\text{NE}_p(p_{I,1,c}, w_q q_{I,1}/w_p, \infty)$ and $\text{NE}_r(p_{I,1,c}, w_q q_{I,1}/w_p, \infty)$
- 14: Update $p_{I,1,c}$ using gradient ascent
- 15: **end while**
- 16: Return converged solution as $p_{I,1,c}^*$
- 17: $p_{I,1}^* \leftarrow \arg \max \{W_I(p_{I,1,a}^*), W_I(p_{I,1,b}^*), W_I(p_{I,1,c}^*)\}$.

Theorem 2. Under Assumption 1, Algorithm 1 converges to company I 's optimal pricing in period 1.

The key rationale is that even if the problem is non-concave, we can decompose the problem into several concave sub-problems (i.e., sub-problem a in lines 2-6, sub-problem b in lines 7-11, and sub-problem c in lines 12-16). Then it suffices to locate the optimal solutions for each sub-problem and then compare those to achieve the global optimum.

4 Numerical Results

We conduct numerical experiments to gain more useful insights. In Sec. 4.1, we discuss the simulation setup. In Sec. 4.2, we train FL models on three different datasets and report the test accuracy. In Sec. 4.3, we use the accuracy results to compute the solutions to the two-period game.

4.1 Simulation Setup

We train ResNet-18 on CIFAR-10 and HAM10000, as well as train ResNet-50 on CIFAR-100 using FedAvg [23].⁶ CIFAR-10 (100) is a balanced dataset on 10(100) classes with 50k training and 10k test data [15]. HAM10000 is an imbalanced medical dataset on 7 classes with 10,015 images from four domains, where each domain represents an institution who collected dermoscopic (skin disease) images [38]. We consider that the two companies have heterogeneous data,

⁶ Our model is compatible with any FL algorithms, and this paper focuses on FedAvg, which till today, remains one of the state-of-the-art FL algorithms.

Table 1: Impact of data heterogeneity β on model accuracy in %.

| CIFAR10 | | | | | | CIFAR100 | | | | | |
|-----------|------------------|------------------|------------------|------------------|------------------|-----------|------------------|------------------|------------------|------------------|------------------|
| | $\beta = \infty$ | $\beta = 1.0$ | $\beta = 0.5$ | $\beta = 0.1$ | $\beta = 0.01$ | | $\beta = \infty$ | $\beta = 1.0$ | $\beta = 0.5$ | $\beta = 0.1$ | $\beta = 0.01$ |
| I local | 60.81 \pm 4.32 | 55.34 \pm 1.82 | 53.42 \pm 5.37 | 39.23 \pm 6.84 | 43.27 \pm 4.03 | I local | 19.00 \pm 4.28 | 18.50 \pm 2.79 | 17.37 \pm 2.11 | 17.17 \pm 3.06 | 16.94 \pm 1.89 |
| E local | 57.56 \pm 3.75 | 54.50 \pm 3.58 | 49.30 \pm 7.41 | 48.17 \pm 5.93 | 33.51 \pm 4.95 | E local | 17.93 \pm 2.78 | 17.46 \pm 2.39 | 17.12 \pm 2.66 | 18.09 \pm 4.49 | 14.26 \pm 1.79 |
| FedAvg | 69.81 \pm 2.26 | 69.01 \pm 2.22 | 66.96 \pm 2.77 | 57.40 \pm 9.02 | 51.89 \pm 6.49 | FedAvg | 32.80 \pm 3.64 | 31.16 \pm 3.46 | 29.92 \pm 3.45 | 26.58 \pm 3.84 | 21.03 \pm 1.63 |

Table 2: Impact of data quantity D_E on model accuracy in %.

| CIFAR10 | | | | | CIFAR100 | | | | | HAM10000 | | | | |
|-----------|------------------|------------------|------------------|------------------|-----------|------------------|------------------|------------------|------------------|-----------|------------------|------------------|------------------|------------------|
| | $D_E = 3k$ | $D_E = 5k$ | $D_E = 8k$ | $D_E = 10k$ | | $D_E = 3k$ | $D_E = 5k$ | $D_E = 8k$ | $D_E = 10k$ | | $D_E = 2k$ | $D_E = 3k$ | $D_E = 4k$ | $D_E = 5k$ |
| I local | 38.47 \pm 9.91 | 38.63 \pm 8.10 | 39.80 \pm 8.08 | 40.56 \pm 8.25 | I local | 17.44 \pm 2.35 | 17.20 \pm 1.61 | 18.64 \pm 3.03 | 19.85 \pm 4.60 | I local | 76.37 \pm 1.65 | 76.28 \pm 0.52 | 75.61 \pm 0.87 | 75.68 \pm 1.24 |
| E local | 48.07 \pm 8.99 | 49.27 \pm 2.84 | 51.84 \pm 0.66 | 53.23 \pm 0.77 | E local | 16.66 \pm 2.17 | 18.30 \pm 1.16 | 21.50 \pm 3.84 | 22.92 \pm 4.87 | E local | 73.67 \pm 1.87 | 76.88 \pm 1.44 | 77.99 \pm 1.24 | 78.11 \pm 1.29 |
| FedAvg | 59.69 \pm 9.55 | 60.66 \pm 1.62 | 62.12 \pm 3.63 | 62.88 \pm 5.50 | FedAvg | 25.40 \pm 2.95 | 26.14 \pm 0.41 | 29.10 \pm 1.15 | 32.08 \pm 1.45 | FedAvg | 79.90 \pm 0.91 | 80.50 \pm 0.72 | 80.89 \pm 0.99 | 81.09 \pm 1.48 |

which are sampled using the Dirichlet distribution with a controlling parameter $\beta > 0$, where a smaller β implies higher data heterogeneity [5]. We study two scenarios:

- *Impact of data heterogeneity:* On CIFAR-10 and CIFAR-100, we use $\beta \in \{\infty, 1, 0.5, 0.1, 0.01\}$. For each β , we sample $5k$ data for company I and $5k$ data for company E . We do not explore the impact of data heterogeneity on HAM10000 as the dataset itself is highly imbalanced.
- *Impact of data quantity:* On CIFAR-10 and CIFAR-100, we fix $\beta = 0.1$, assign company I $5k$ data, and change company E 's data volume $D_E \in \{3k, 5k, 8k, 10k\}$. On HAM10000, we randomly assign company I $2k$ data and company E data of size $D_E \in \{2k, 3k, 4k, 5k\}$.

The key hyper-parameters are summarized as follows. We use 100 rounds for FL training, where each round consists of 5 local epochs. We use SGD as the optimizer, and choose learning rate $l_r = 0.001$ and batch size $B = 64$. We further report the results when each company only uses its own data to train a local model. Each experiment is repeated over 3 runs using different random seeds.

4.2 Training Results

Table 1 reports the training results under different levels of data heterogeneity, where “ I (E) local” means that company I (E) trains a model using its local data without FL. From this table, we make two observations.

First, on both datasets, FedAvg consistently outperforms local training models across all values of β , including the highly heterogeneous case $\beta = 0.01$. This suggests the benefit of FL collaboration in improving local model performances. Second, as β decreases, there is a general trend of performance decline for FedAvg. This is due to the notorious client drift issue in FL, which has been observed in many distributed learning literature (e.g., [40, 17]). Despite that data heterogeneity hurts FL performance, we will show a counter-intuitive result in Sec. 4.3 that data heterogeneity promotes price competition.

Table 2 reports the accuracy results under different data quantities. We observe that FedAvg again outperforms local training across different data quantities on the three datasets. Also, increasing data volume D_E consistently improves model accuracy across all datasets, underscoring the importance of data quantity in model training.

4.3 Equilibrium Results

Now, we use Tables 1-2 to calculate the equilibrium of the two period game. We consider that companies' quality of service is an increasing function in the model accuracy. In particular, we use $q_{i,t}(A) = A$, where A is the mean accuracy value from the tables. We further consider that users' type (preference) ϕ_n is heterogeneous, which follows a uniform distribution on support $[0, 1]$.⁷ We choose coefficients

⁷ We have also tested Gaussian distributions and results carry over.

Table 3: FL collaboration under different β .

| | $\beta = \infty$ | $\beta = 1$ | $\beta = 0.5$ | $\beta = 0.1$ | $\beta = 0.01$ |
|-----------|------------------|-------------|---------------|---------------|----------------|
| CIFAR-10 | ✓ | ✓ | ✓ | ✓ | ✓ |
| CIFAR-100 | ✓ | ✓ | ✓ | ✓ | ✓ |

Table 4: FL collaboration under different D_E .

| | $D_E = 3k$ | $D_E = 5k$ | $D_E = 8k$ | $D_E = 10k$ |
|-----------|------------|------------|------------|-------------|
| CIFAR-10 | ✓ | ✓ | ✓ | ✓ |
| CIFAR-100 | ✓ | ✓ | ✓ | ✓ |
| | $D_E = 2k$ | $D_E = 3k$ | $D_E = 4k$ | $D_E = 5k$ |
| HAM10000 | × | × | × | × |

$w_p = w_\phi = w_p = 1$ and normalize the service costs c_I and c_E to 0. We use Proposition 2 to calculate the equilibrium FL collaboration, in which we use the best response algorithm [27] to compute the equilibrium of Game 1. We further use Algorithm 1 to compute the optimal period 1 pricing of company I .

Optimal FL collaboration. We first investigate the equilibrium FL collaboration and report the results in Tables 3-4, where ✓ means collaboration and × means no collaboration. We observe that while FL enhances model performance across the three datasets (see Table 1), intriguingly, collaboration is observed only on CIFAR datasets but not on HAM10000. This may seem counter-intuitive at first glance because the improved performance through FL suggests a natural incentive to collaborative behavior. However, it is crucial to note that improved performance alone does not necessitate collaboration in a competitive environment. Recall from Corollary 1 that a company's price increases in its own model performance but at the same time decreases in its competitor's model performance. Hence, collaboration is a beneficial strategy only when the performance gains are substantial enough to outweigh the loss from price competition. For instance, CIFAR datasets exhibit significant performance enhancements (e.g., $\beta = 0.1$ on CIFAR-10 in Table 1), which motivates collaboration, while the less significant improvements in HAM10000 (e.g., $D_E = 3k$ in Table 2) are insufficient to motivate the same collaborative effort. We summarize the above observation below.

Observation 1. *FL collaboration is the NE only when the model performance improvement is substantial.*

Optimal Pricing. We now investigate how the optimal pricing changes with data heterogeneity β and data quantity D_E . The results are reported in Fig. 3. Note that in period 2 the two companies' optimal pricing are the same due to 1) both companies having the same (global) model, and 2) users' preference ϕ_n follows a uniform distribution, which is symmetric for both companies.

From Fig. 3(a)-(b), it is a bit surprising to observe that both companies optimal pricing can increase in data heterogeneity (i.e., a smaller β). While the expectation is that higher heterogeneity leads to worse performance and hence users are willingly to pay, leading to smaller prices. However, this is not always the case. The reason is that when data is highly heterogeneous, even if the model performance declines, the companies still obtain a big performance improvement from FL compared to local learning (e.g., see $\beta = 0.1$ in Table 1).

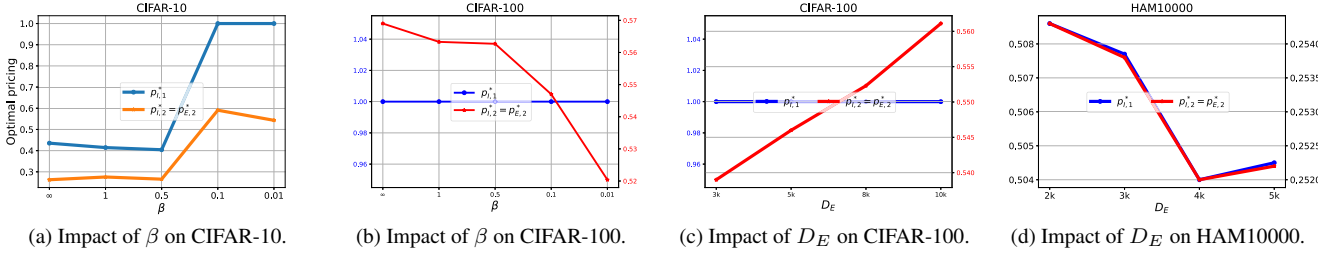


Figure 3: Optimal two-period pricing on different datasets.

As a result, company I sets a high price in period 1 so that no users buy low-quality service in this period. For example, in Fig. 3(a) at $\beta = 0.1$, I sets $p_{i,1}^* = 1$ and based on Lemma 1, none of the users will purchase from I . Instead, both companies attract users to buy high-quality services at higher prices in period 2, which results in higher profits. This observation has an interesting implication, i.e., data heterogeneity does not always prevent but can promote price competition. In particular, the incumbent company I does not harvest most users in the starting period but prefers to compete with the entrant company using better model-based services. We summarize the above observations as follows.

Observation 2. *Data heterogeneity can promote price competition between incumbent and entrant companies.*

From Fig. 3(c)-(d), we observe different trends of pricing as a function of D_E on the two datasets. The reason is that the optimal FL collaboration decisions are different (refer to Table 4). On CIFAR-100, the companies collaborate via FL and the performance improvement tends to be larger as D_E increases. Hence, company I uses a very large price in period 1 so that no users buy in period 1, and instead users will buy higher-quality services in period 2 at a larger price. On HAM10000, however, the companies do not collaborate in FL. As a result, company I is inclined to set a smaller price in period 1 so that most users subscribe to I in the first period. Then in period 2, the companies participate in more fierce price competition due to a smaller user pool, leading to a smaller price.

5 Discussions

We have provided a first study on FL coopetition using a dynamic two-period model with two companies. The analysis is highly non-trivial with this stylized model involving a non-concave problem. Next, we discuss how to extend our model with more than two companies across multiple time periods.

Extension to multiple companies. Our model can be extended to scenarios with two or more companies. In particular, consider $N \geq 3$ companies where $N - 1$ of them coexist in period 1, and an entrant enters in period 2. This can model real-world market conditions where multiple companies often coexist and new entrants periodically disrupt existing equilibriums. In this case, we can similarly analyze the user behavior by comparing the payoffs obtained when purchasing from different companies. For companies' FL collaboration structures, e.g., [1, 35]. For companies' price competitions, prior studies on oligopoly competition [37] can offer theoretical and algorithmic solutions. However, how to optimally set prices in period 1 (which affects period 2 outcomes) remains a challenging but interesting problem to explore in future work.

Extension to multiple time periods. It is possible to extend our model to where there are $t \geq 3$ periods, and in each period a new entrant company enters the market. For user decisions, it suffices to restrict attention in each period and compute the highest payoff achieved among existing companies. However, we note that solving companies' FL collaboration and price competition over $t \geq 3$ periods in closed-form is generally analytically intractable. One particular challenge is to estimate future entrants and their properties (e.g., data distributions). One possible remedy is to resort to reinforcement learning for sequential decision making, and in particular, multi-armed bandit [31, 8]. More specifically, one could model the *agents* as companies, the *arms* as FL collaboration and pricing, and the *reward signals* as profits achieved from users, which takes into account competitions and model improvement from FL training.

6 Conclusion

This work studied the under-explored problem of FL coopetition (FL collaboration and market competition) using a dynamic two-period model. One challenge pertains to multi-dimensional heterogeneity in terms of companies' data distributions and quantities, as well as users' service preferences. Another challenges is associated with solving a non-concave optimization problem. We decomposed the problem into multiple concave sub-problems and managed to characterize the solutions to the two-period game. Theoretical analysis and numerical experiments on three datasets have led to two implications. First, FL brings both model performance gain and competition loss, and collaboration occurs only when the performance gain outweighs the loss. Second, data heterogeneity can incentivize the incumbent to limit market penetration in period 1 and promote price competition in period 2.

For future work, it is interesting to study how data quality (e.g., label correctness and label sparsity) affects FL coopetition and further develop robust algorithms to enhance performance. It is also interesting to incorporate privacy enhancing techniques, which can be of particular interest to companies, into the coopetition analysis.

References

- [1] K. Donahue and J. Kleinberg. Optimality and stability in federated learning: A game-theoretic approach. *Advances in Neural Information Processing Systems*, 34:1287–1298, 2021.
- [2] J. Duffy and E. Hopkins. Learning, information, and sorting in market entry games: theory and evidence. *Games and Economic behavior*, 51(1):31–62, 2005.
- [3] H. Hotelling. Stability in competition. *The economic journal*, 39(153):41–57, 1929.
- [4] H. Hotelling. Stability in competition. *The Economic Journal*, 39(153):41–57, 1929. ISSN 00130133, 14680297. URL <http://www.jstor.org/stable/2224214>.

- [5] T.-M. H. Hsu, H. Qi, and M. Brown. Measuring the effects of non-identical data distribution for federated visual classification. *arXiv preprint arXiv:1909.06335*, 2019.
- [6] C. Huang, S. Ke, C. Kamhoua, P. Mohapatra, and X. Liu. Incentivizing data contribution in cross-silo federated learning. *arXiv preprint arXiv:2203.03885*, 2022.
- [7] C. Huang, M. Tang, Q. Ma, J. Huang, and X. Liu. Promoting collaborations in cross-silo federated learning: Challenges and opportunities. *IEEE Communications Magazine*, 2023.
- [8] C. Huang, H. Yu, J. Huang, and R. Berry. An online inference-aided incentive framework for information elicitation without verification. *IEEE Journal on Selected Areas in Communications*, 41(4):1167–1185, 2023.
- [9] C. Huang, J. Dachille, and X. Liu. Technical report: Competition in heterogeneous cross-silo federated learning. *arXiv preprint arXiv:2408.11355*, 2024.
- [10] C. Huang, S. Ke, and X. Liu. Duopoly business competition in cross-silo federated learning. *IEEE Transactions on Network Science and Engineering*, 11(1):340–351, 2024.
- [11] Y. Huang, L. Chu, Z. Zhou, L. Wang, J. Liu, J. Pei, and Y. Zhang. Personalized cross-silo federated learning on non-iid data. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 7865–7873, 2021.
- [12] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings, et al. Advances and open problems in federated learning. *Foundations and trends® in machine learning*, 14(1–2):1–210, 2021.
- [13] S. P. Karimireddy, S. Kale, M. Mohri, S. Reddi, S. Stich, and A. T. Suresh. Scaffold: Stochastic controlled averaging for federated learning. In *International conference on machine learning*, pages 5132–5143. PMLR, 2020.
- [14] M. Kremer, B. Mantin, and A. Ovchinnikov. Dynamic pricing in the presence of myopic and strategic consumers: Theory and experiment. *Production and Operations Management*, 26(1):116–133, 2017.
- [15] A. Krizhevsky, G. Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [16] Q. Li, B. He, and D. Song. Model-contrastive federated learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10713–10722, 2021.
- [17] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith. Federated optimization in heterogeneous networks. *Proceedings of Machine learning and systems*, 2:429–450, 2020.
- [18] T. Li, S. Hu, A. Beirami, and V. Smith. Ditto: Fair and robust federated learning through personalization. In *International conference on machine learning*, pages 6357–6368. PMLR, 2021.
- [19] Y. Li, X. Wang, R. Zeng, M. Yang, K. Li, M. Huang, and S. Dustdar. Varf: An incentive mechanism of cross-silo federated learning in mec. *IEEE Internet of Things Journal*, 2023.
- [20] K. Liu, S. Hu, S. Z. Wu, and V. Smith. On privacy and personalization in cross-silo federated learning. *Advances in neural information processing systems*, 35:5925–5940, 2022.
- [21] Q. Ma, B. Shou, J. Huang, and T. Başar. Dynamic pricing in the presence of participation-dependent social learning. In *Proceedings of the Eighteenth ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 101–110, 2018.
- [22] O. Marfoq, C. Xu, G. Neglia, and R. Vidal. Throughput-optimal topology design for cross-silo federated learning. *Advances in Neural Information Processing Systems*, 33:19478–19487, 2020.
- [23] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017.
- [24] M. Mendieta, T. Yang, P. Wang, M. Lee, Z. Ding, and C. Chen. Local learning matters: Rethinking data heterogeneity in federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8397–8406, 2022.
- [25] J. Oh, S. Kim, and S.-Y. Yun. Fedbabu: Towards enhanced representation for federated image classification. *Proc. ICLR*, 2022.
- [26] M. J. Osborne and C. Pitchik. Equilibrium in hotelling’s model of spatial competition. *Econometrica: Journal of the Econometric Society*, pages 911–922, 1987.
- [27] M. J. Osborne and A. Rubinstein. *A course in game theory*. MIT press, 1994.
- [28] K. Pillutla, K. Malik, A.-R. Mohamed, M. Rabbat, M. Sanjabi, and L. Xiao. Federated learning with partial model personalization. In *International Conference on Machine Learning*, pages 17716–17758. PMLR, 2022.
- [29] Z. Qin, S. Deng, M. Zhao, and X. Yan. Fedapen: Personalized cross-silo federated learning with adaptability to statistical heterogeneity. In *Proc. ACM SIGKDD*, pages 1954–1964, 2023.
- [30] S. C. Salop. Monopolistic competition with outside goods. *The Bell Journal of Economics*, pages 141–156, 1979.
- [31] A. Slivkins et al. Introduction to multi-armed bandits. *Foundations and Trends® in Machine Learning*, 12(1–2):1–286, 2019.
- [32] H. M. Son, M. H. Kim, T.-M. Chung, C. Huang, and X. Liu. Feduv: Uniformity and variance for heterogeneous federated learning. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2024.
- [33] M. Spann, M. Fischer, and G. J. Tellis. Skimming or penetration? strategic dynamic pricing for new products. *Marketing Science*, 34(2):235–249, 2015.
- [34] G. J. Stigler. A theory of oligopoly. *Journal of political Economy*, 72(1):44–61, 1964.
- [35] S. Tan, H. Cheng, X. Wu, H. Yu, T. He, Y. S. Ong, C. Wang, and X. Tao. Fedcompetitors: Harmonious collaboration in federated learning with competing participants. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 15231–15239, 2024.
- [36] M. Tang and V. W. Wong. An incentive mechanism for cross-silo federated learning: A public goods perspective. In *IEEE INFOCOM 2021-IEEE Conference on Computer Communications*, pages 1–10. IEEE, 2021.
- [37] J. Tirole. *The theory of industrial organization*. MIT press, 1988.
- [38] P. Tschandl, C. Rosendahl, and H. Kittler. The ham10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific data*, 5(1):1–9, 2018.
- [39] N. Tsoy and N. Konstantinov. Strategic data sharing between competitors. *Advances in Neural Information Processing Systems*, 36, 2024.
- [40] B. E. Woodworth, K. K. Patel, and N. Srebro. Minibatch vs local sgd for heterogeneous distributed learning. *Advances in Neural Information Processing Systems*, 33:6281–6292, 2020.
- [41] X. Wu and H. Yu. Mars-fl: Enabling competitors to collaborate in federated learning. *IEEE Transactions on Big Data*, 2022.
- [42] M. Ye, X. Fang, B. Du, P. C. Yuen, and D. Tao. Heterogeneous federated learning: State-of-the-art and research challenges. *ACM Computing Surveys*, 56(3):1–44, 2023.
- [43] C. Zhang, S. Li, J. Xia, W. Wang, F. Yan, and Y. Liu. {BatchCrypt}: Efficient homomorphic encryption for {Cross-Silo} federated learning. In *USENIX annual technical conference*, pages 493–506, 2020.
- [44] N. Zhang, Q. Ma, and X. Chen. Enabling long-term cooperation in cross-silo federated learning: A repeated game perspective. *IEEE Transactions on Mobile Computing*, 2022.