CARE: Compatibility-Aware Incentive Mechanisms for Federated Learning with Budgeted Requesters

Xiang Liu[1,2], Hau Chan[3], Minming Li[4], Xianlong Zeng[1], Chenchen Fu[1], Weiwei Wu[1,*]

[1]Southeast University, [2]The Chinese University of Hong Kong,

[3]University of Nebraska-Lincoln, [4]City University of Hong Kong

xiangliu@seu.edu.cn, hchan3@unl.edu, minming.li@cityu.edu.hk, {xianlong_zeng, chenchen_fu, weiweiwu}@seu.edu.cn

***Abstract*—Federated learning (FL) is a promising approach that allows requesters (*e.g.*, servers) to obtain local training models from workers (*e.g.*, clients). Since workers are typically unwilling to provide training services/models freely and voluntarily, many incentive mechanisms in FL are designed to incentivize participation by offering monetary rewards from requesters. However, existing studies neglect two crucial aspects of real-world FL scenarios. First, workers can possess inherent incompatibility characteristics (*e.g.*, communication channels and data sources), which can lead to degradation of FL efficiency (*e.g.*, low communication efficiency and poor model generalization). Second, the requesters are budgeted, which limits the amount of workers they can hire for their tasks. In this paper, we investigate the scenario in FL where multiple budgeted requesters seek training services from incompatible workers with private training costs. We consider two settings: the cooperative budget setting where requesters cooperate to pool their budgets to improve their overall utility and the non-cooperative budget setting where each requester optimizes their utility within their own budgets. To address efficiency degradation caused by worker incompatibility, we develop novel compatibility-aware incentive mechanisms, CARE-CO and CARE-NO, for both settings to elicit true private costs and determine workers to hire for requesters and their rewards while satisfying requester budget constraints. Our mechanisms guarantee individual rationality, truthfulness, budget feasibility, and approximation performance. We conduct extensive experiments using real-world datasets to show that the proposed mechanisms significantly outperform existing baselines.**

I. Introduction

Federated learning (FL) [1], [2] is a decentralized machine learning paradigm that enables collaborative model training across a group of workers (*e.g.*, clients, mobile devices and data owners) without directly sharing or revealing workers' raw data openly. Recently, FL has gained significant attention and has been applied to various applications in domains such as edge computing [3], healthcare [4], and finance [5].

In FL, requesters (*e.g.*, servers, and model owners) publish their training tasks and workers participate in the training tasks by using their local data to train local models [6]. It has been observed that workers are commonly unwilling to freely contribute to training due to the costs of using their own data and computational resources [7]. In addition, workers' costs are naturally private and unknown to the platform. Therefore, previous works [8]–[10] have designed (truthful) incentive mechanisms in FL to elicit workers' true private costs, select workers to hire for requesters/tasks, and determine workers' rewards/payments for training and providing local models. As the requesters often have limits on how much they can pay the workers, recent studies have focused on designing incentive mechanisms that ensure that the total payment to workers does not exceed the requester's budget [9], [11]–[13].

However, existing incentive mechanism design studies in FL do not consider two crucial aspects: (1) *multiple budgeted requesters* and (2) the *compatibility of workers* that are prevalent in real-world FL. Regarding multiple budgeted requesters, existing settings primarily concentrate on designing incentive mechanisms for a single budgeted requester [14]–[16]. However, these approaches do not provide reasonable incentive mechanisms for multiple requesters who simultaneously seek to hire workers for their respective tasks, particularly in natural situations where each requester has limited ability to hire workers. Regarding the compatibility of workers, existing incentive mechanisms in FL ignore the fact that workers can be categorized into different groups based on their inherent incompatibility characteristics. Those within the same group are incompatible in jointly performing tasks.

In general, compatibility issues are common in practical FL. For instance, in the FL that utilizes congested wireless communication channels (*e.g.*, the wireless spectrum) to update global model parameters, workers (*e.g.*, mobile devices) can experience congestion or disconnections due to bandwidth limitations [17]. To enhance the stability and efficiency of communication, it is often advisable to restrict the number of workers using the same channel [18]. In such scenarios, workers are often grouped based on their specific communication channels, leading to incompatibilities between workers within the same group [17]. In addition, workers can also be grouped according to the similarity of their datasets and data sources (*e.g.*, data collected from populations with varying demographics, ages or levels of education). To improve the performance (*e.g.*, generalization or robustness [19]) of the global model, each requester prefers a broader selection of workers from various data sources, especially within their budgets [20]. Consequently, workers with the same data source become incompatible when selected simultaneously. There-

*Weiwei Wu is the corresponding author. This work was supported in part by the National Natural Science Foundation of China under grant 62402102, 61972086, 62272099, the Natural Science Foundation of Jiangsu Province under Grant No. BK20241275, BK20230024, BK20231543, and the Key Laboratory of Computer Network and Information Integration (Southeast University), Ministry of Education.

fore, disregarding the compatibility of workers can lead to low efficiency in FL, such as prolonged communication times during the model update process and poor generalization of the trained global model.

Our Goal and Contributions. Motivated by the above real-world scenarios, we investigate the problem of designing **C**ompatibility-**A**ware incentive mechanisms in fede**R**ated l**E**arning (CARE) with multiple budgeted requesters and incompatible workers. Specifically, workers are classified into groups based on their inherent incompatibility characteristics (*e.g.*, communication channels, and data sources). In addition, there are compatibility constraints among workers, *i.e.*, the number of workers assigned to the same requester/task from each group should not exceed a predefined threshold. Our goal is to design incentive mechanisms under the CARE problem to elicit workers' true private costs, select workers to hire for requesters/tasks that optimize the overall reputation (*i.e.*, a common objective in FL [9], [21]), and determine payments to the selected workers subject to the compatibility constraints and requesters' budgets. Moreover, the designed mechanisms should satisfy desirable properties, including individual rationality, truthfulness, budget feasibility, and approximation guarantees. We refer readers to Section III-B for justifications.

Because of multiple budgeted requesters and incompatible workers, designing incentive mechanisms under the CARE problem faces three main new challenges compared to existing studies. *1) Cost-effective worker selection*: It is efficient to prioritize workers with low-cost and high-reputation during the worker selection. However, despite their cost-effectiveness, these workers can violate compatibility constraints, making it difficult to find workers that satisfy both cost-effectiveness and compatibility. *2) Stronger strategic manipulation:* With multiple requesters, it is essential to adaptively match workers to requesters. However, this also creates more opportunities for workers to engage in strategic manipulation, *e.g.*, a worker may bid a false cost to be matched with a different requester and thereby obtain higher utility. *3) Unpredictable payments:* To satisfy requesters' budgets, we should evaluate both the reputation and the payments that each requester can obtain and must pay when selecting workers. However, this process is intractable because workers' payments remain uncertain until the final outcome of worker selection is determined, which is necessary to ensure truthfulness.

We consider the CARE problem under two realistic budget settings: *(i) Cooperative budget setting:* Requesters collaborate by pooling their budgets (*e.g.*, hospitals integrate healthcare resources such as public funds applied from the organization to train a disease recognition model [4], [22]), enabling them to hire more workers and thereby enhance their overall utility (*e.g.*, improving the model accuracy by collaboratively sharing and aggregating trained models between requesters [15], [23]). *(ii) Non-cooperative budget setting:* Each requester hires workers within their own budgets. Our main contributions are summarized as follows:

- To the best of our knowledge, we are the first to design compatibility-aware incentive mechanisms in FL that capture workers' inherent incompatibilities and requesters' limited hiring abilities, thereby preventing efficiency degradation and improving budget utilization.
- We first propose CARE-CO mechanism for cooperative budget setting. Particularly, CARE-CO transforms the selection of workers within the compatibility constraint into a Max-Flow problem, allowing us to explore different potential prices while simultaneously ensuring efficiency. We then propose CARE-NO mechanism for non-cooperative budget setting, which first divides all workers into multiple sets so that each set of workers have similar reputations. Additionally, it introduces a virtual-price based sub-mechanism, named PEA, to address each worker set independently. Specially, PEA utilizes the concept of virtual prices to evaluate requesters' ability to obtain reputation and determines the critical price that aligns with this ability, thereby ensuring both budget feasibility and truthfulness.
- Our mechanisms are proved to guarantee individual rationality, truthfulness, budget feasibility, and computational efficiency. Moreover, our mechanisms achieve approximation guarantees in comparison to the optimal solution that has prior knowledge of workers' private costs.
- Finally, we conduct experiments on two commonly adopted datasets in FL, *i.e.*, Fashion MNIST (FMNIST) and CIFAR-10. Evaluation results show that our mechanisms improve overall reputation of selected workers by about 824% and the global model accuracy by about 57% on average compared to baselines.

This paper is structured as follows. Section II reviews the related works. The system model and the definition of the problem are given in Section III. We propose CARE-CO and CARE-NO in Section IV and V, respectively. Section VI presents the experimental results. Finally, the conclusion is given in Section VII.

II. Related Work

Recently, a large body of literature has investigated incentive mechanisms in FL. We refer the reader to the comprehensive survey [7], [10]. In the following, we focus on discussing related works on (budgeted) reverse auction-based or procurement mechanisms in FL and general settings.

Reverse Auction or Procurement Based Mechanisms for Federated Learning. Reverse auction or procurement based mechanisms have been extensively used in various FL scenarios, effectively guiding the requesters in selecting high-quality workers to participate in training tasks and maximize the objective such as the social welfare [17], [24], model accuracy [9] and the requester's utility [25]. While previous works have overlooked the budget constraint of the requester, Fan *et al.* [11] address this issue by considering the requester's budget and introducing a data quality-driven reverse auction. This approach aims to maximize the requester's global model accuracy. In a similar vein, Zhang *et al.* [9] design a reputation calculation method to indirectly capture the data quality of workers when designing the reverse auction-based incentive

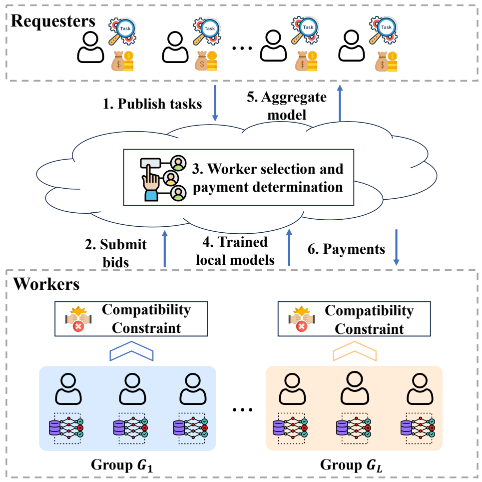

Fig. 1: FL with incompatible workers and budgeted requesters.

mechanism under the budget constraint. Building upon this, Zhang *et al.* [13] focus on the online setting of the aforementioned problems, where workers arrive in a specific order. In addition, there are other works that examine budgeted incentive mechanisms in FL. These works explore various aspects such as differential private noises [26], [27], trustworthy data acquisition [28], and competition among federations [29]. However, all these works do not explicitly consider worker compatibility and typically assume a single requester in FL.

Budget Feasible Mechanisms. Regarding designing budgeted reverse-auction or procurement mechanisms in general settings, our work is also related to the settings of budget feasible mechanism design (see *e.g.*, [30]–[33]). In a typical setting of budget feasible mechanism design, a single buyer (*i.e.*, requester) with a budget wants to procure services from sellers (*i.e.*, workers) with private costs. However, all these settings assume a single buyer (requester) and disregard entirely the inherent compatibility issues among workers. We note that there are few works in budget feasible mechanisms consider multiple buyers [34], [35]. However, they ignore the compatibility among workers. Therefore, existing mechanisms do not apply to our problem.

In contrast to the aforementioned previous works, our study focuses on incentive mechanisms for multiple budgeted requesters and incompatible workers in FL.

III. System Model and Problem Definition

A. System Model

As shown in Fig. 1, we consider the FL system consisting of incompatible workers and multiple budgeted requesters. Requesters first publish their training tasks and workers submit their bids (*e.g.*, each worker's bid corresponds to the cost of training local models). Subsequently, we implement the designed incentive mechanism, which takes into account workers' bids, reputations (*e.g.*, calculated by historical task performance [9]), compatibility constraints and requesters' budgets, to assign workers to requesters. The selected workers are provided with initial global models from their corresponding requesters and train their local models, which they then upload to requesters for aggregation. This iterative process continues until requesters' models converge. Finally, requesters compensate workers with monetary rewards for their services.

B. Problem Definition

Let $S = \{s_1, s_2, \ldots, s_n\}$ denote the set of workers who can be recruited to conduct model training tasks locally. Each worker s_i has a private raw dataset and a cost c_i to participate in a task, *e.g.*, the consumption of energy and computational resource. There are m requesters $A = \{a_1, a_2, \ldots, a_m\}$, and each requester a_j holds a training task and seeking to hire a subset of workers $S_j \subseteq S$ to carry out the training using their datasets[1]. Each requester a_j has a budget B_j for hiring workers, and we use $\mathcal{B} = \{B_1, B_2, \ldots, B_m\}$ to denote requesters' budgets. Suppose that workers' costs are no higher than requesters' budgets [34], [35]. We consider two different budget settings: *1) Cooperative budget setting:* Requesters are willing to collaborate by pooling their budgets, and denoted by $B = \sum_{j \le m} B_j$ the total budget of all requesters. *2) Non-cooperative budget setting:* Each requester hire workers within their individual budget.

Compatibility Constraints: Due to the inherent incompatibility characteristics of workers, workers are categorized into L different groups, *i.e.*, $\mathcal{G} = \{G_1, G_2, \cdots, G_l, \cdots, G_L\}$. We define τ_{lj} as the compatibility level of group G_l for requester a_j, which indicates the maximum number of workers in G_l can be selected for a_j (*e.g.*, $\tau_{lj} = 1$ means that only one worker in G_l can be selected for a_j). Then, we define the *compatibility constraint* such that $|S_j \cap G_l| \le \tau_{lj}, \forall j \le m, l \le L$.

Incomplete Information: We consider the incomplete information scenario where each worker's cost is *private* (known by themselves). Thus, each worker can behave strategically to misreport their private cost to improve their utility (defined below). Let b_i denote the cost reported by worker s_i, which may not equal (or potentially much higher than) the true cost c_i. Denote by $\mathbf{b} = \{b_1, b_2, ..., b_n\}$ and $\mathbf{b}_{-i}$ the set of workers' bids and the set of workers' bids except b_i, respectively.

Incentive Mechanism: Let S_w be the winner (or selected) worker set, *i.e.*, $S_w = \cup_{j \le m} S_j$. The incentive mechanism $\mathbb{M} = (X, P)$ consists of the allocation rule X which maps the bid profile $\mathbf{b}$ to S_w and the payment rule P which decides the payment for each winner. Let $x_{ij} \in \{0, 1\}$ indicate whether worker s_i is allocated to requester a_j, and $x_i := \sum_{j \le m} x_{ij} \le 1$. In particular, $x_i = 1$ implies $s_i \in S_w$. Let p_{ij} be the payment paid to worker s_i from requester a_j and $p_i := \sum_{j \le m} p_{ij}$. Specially, if $x_{ij} = 0$, then $p_{ij} = 0$. Given the mechanism $\mathbb{M}$, the utility of worker s_i is the difference between the true cost and the received payment, *i.e.*, $u_s^i(\mathbf{b}, \mathbb{M}) = p_i - x_i \cdot c_i$.

[1] We focus on scenarios where workers have the same type of datasets, *e.g.*, image classification or NLP datasets, while requesters seek their own models using data from these workers. Our mechanisms can also be readily generalized to scenarios where workers have different dataset types by applying them separately to groups of workers with the same dataset type.

Algorithm 1: CARE-CO($B, \mathbf{b}, A, S, \mathcal{G}$)

Output: P, S_w

1 $P \leftarrow 0, S_w \leftarrow \emptyset$;
2 Sort all workers as $\frac{b_1}{v_1} \leq \frac{b_2}{v_2} \leq \cdots \frac{b_n}{v_n}$;
3 // **Worker selection**;
4 **for** $1 \leq i \leq n$ **do**
5   Compute $\mathcal{M}(\mathcal{S}_i)$ by Eq. (1) - (4);
6   If $\frac{b_i}{v_i} \cdot \mathcal{M}(\mathcal{S}_i) \leq B$, $i \leftarrow i+1$; otherwise, break;
7 **end**
8 $k \leftarrow i-1$;
9 Decide the allocation X by the solution of $\mathcal{M}(\mathcal{S}_k)$;
10 Selected workers perform training tasks for requesters;
11 // **Payment scheme**;
12 **for** $i \leq n, j \leq m$ **do**
13     $p_{ij} = v_i \cdot x_{ij} \min\{\frac{b_{k+1}}{v_{k+1}}, \frac{B}{\mathcal{M}(\mathcal{S}_k)}\}$;
14 **end**

The Objective: *Reputation* is commonly employed as a metric to reflect the quality of workers in FL [36], *e.g.*, their exerted efforts and the quality of their datasets. The reputation of workers can be calculated through historical task performance [9], [21]. Given that v_i is the reputation of worker s_i and $\mathbf{v} = \{v_1, v_2, \cdots, v_n\}$ is the workers' reputation profile. Thus, the objective of the designed mechanism is to maximize the overall reputation of selected workers $\max \sum_{i \leq n, j \leq m} x_{ij} v_i$, which is a standard objective in FL [9], [13], [21]. Moreover, the proposed mechanism should satisfy the following properties:

(1) Individual Rationality: The utility of each worker s_i is non-negative, *i.e.*, $u_s^i(\mathbf{b}, \mathbb{M}) \geq 0$ for any $\mathbf{b}$. **(2) Truthfulness:** Bidding the true cost, each worker gets the maximum utility, *i.e.*, $\forall b_i, u_s^i(c_i, \mathbf{b}_{-i}, \mathbb{M}) \geq u_s^i(b_i, \mathbf{b}_{-i}, \mathbb{M})$. **(3) Budget Feasibility:** 1) For the cooperative budget setting, the total payment of requesters does not exceed the total budget B, *i.e.*, $\sum_{1 \leq i \leq n} p_i \leq B$, and 2) for the non-cooperative budget setting, the total payment of requester $a_j, \forall j \leq m$, does not exceed budget B_j, *i.e.*, $\sum_{1 \leq i \leq n} p_{ij} \leq B_j$. **(4) Computational Efficiency:** The output of the mechanism can be computed in polynomial time. **(5) Approximation:** Let $ALG(I)$ be the obtained reputation of mechanism $\mathbb{M}$ with instance I. We compare the output of the mechanism with the optimal achievable overall reputation when workers' costs are known in advance. A mechanism is β-approximate if $\forall I, ALG(I) \geq \frac{1}{\beta} OPT(I)$.

IV. The Cooperative Budget Setting

In this section, we propose CARE-CO for the cooperative budget setting. To address the challenge posed by the compatibility constraint, we first disregard the requesters' budgets and assess the optimal overall reputation for a specific worker set with bids not exceeding a particular price. Subsequently, by incrementally increasing the given price, we can identify a price that maximizes the overall reputation while ensuring budget feasibility. The detail of CARE-CO is in Algorithm 1.

As a crucial component of CARE-CO, we first introduce the method to calculate the maximum overall reputation of selected workers under the compatibility constraint, ignoring costs and budgets. We sort workers in non-decreasing order of their bids relative to their reputations, *i.e.*, $\frac{b_1}{v_1} \leq \frac{b_2}{v_2} \leq \cdots \leq \frac{b_n}{v_n}$. Let $\mathcal{S}_i = \{s_1, s_2, \ldots, s_i\}$ be the worker set containing the first i workers. Next, we introduce the optimal overall reputation computation problem (ORP) for the worker set $\mathcal{S}_i$.

Definition 1 (Sub-problem ORP)**.** *Given a worker set $\mathcal{S}_i$, then*

$$\textbf{ORP}(\mathcal{S}_i): \quad \max \sum_{t \leq i, j \leq m} v_t x_{tj} \tag{1}$$

$$s.t., \sum_{s_t \in G_l} x_{tj} \leq \tau_{lj}, \forall j \leq m, l \leq L; \tag{2}$$

$$\sum_{j \leq m} x_{tj} \leq 1, \forall t \leq i; \tag{3}$$

$$x_{tj} \in \{0, 1\}, \forall t \leq i, j \leq m; \tag{4}$$

where (2) indicates compatibility constraints and (3) means that each worker can only be assigned to at most one requester[2]*. Let $\mathcal{M}(\mathcal{S}_i)$ denote the optimal value of ORP($\mathcal{S}_i$).*

Then, we are ready to introduce CARE-CO. We start from the first worker and find the *key worker* s_k such that $\frac{b_k}{v_k} \cdot \mathcal{M}(\mathcal{S}_k) \leq B$ and $\frac{b_{k+1}}{v_{k+1}} \cdot \mathcal{M}(\mathcal{S}_{k+1}) > B$. In the optimal solution of ORP($\mathcal{S}_k$), if $x_i = \sum_{j \leq m} x_{ij} = 1$, then $s_i \in S_w$. The payment of each worker is $p_{ij} = v_i \cdot x_{ij} \min\{\frac{b_{k+1}}{v_{k+1}}, \frac{B}{\mathcal{M}(\mathcal{S}_k)}\}$.

Next, we prove the theoretical guarantees of CARE-CO.

Theorem 1. *CARE-CO guarantees individual rationality, truthfulness, budget feasibility and computational efficiency, and achieves a $(2 + \frac{v_{max}}{v_{min}})$-approximation where $v_{max} := \max_{i \leq n} v_i$ and $v_{min} := \min_{i \leq n} v_i$.*

Proof Sketch. (1) Individual rationality: For each winner $s_i \in S_w$, we have $p_i = v_i \cdot \min\{\frac{b_{k+1}}{v_{k+1}}, \frac{B}{\mathcal{M}(\mathcal{S}_k)}\} \geq v_i \cdot \frac{b_i}{v_i}$, which indicates that s_i' utility is $p_i - b_i \geq 0$.

(2) Budget feasibility: The total payment to the winner is $\min\{\frac{b_{k+1}}{v_{k+1}}, \frac{B}{\mathcal{M}(\mathcal{S}_k)}\} \cdot \sum_{s_i \in S_w} v_i \leq \frac{B}{\mathcal{M}(\mathcal{S}_k)} \cdot \mathcal{M}(\mathcal{S}_k) = B$.

(3) Computational efficiency: The running time of CARE-CO is dominated by the loop in computing the optimal reputation ORP($\mathcal{S}_i$) in the given worker set (line 4-7) As ORP problem can be converted to the Max-Flow problem, the final total complexity of $O(MN(N+L)(ML+2N))$.

(4) Truthfulness: We leverage the famous Monotone Theorem [37] to prove truthfulness, which shows that truthful mechanisms satisfy monotonicity and workers are paid threshold payments. Monotonicity means that when the selected worker reports a lower cost, the worker remains selected. Threshold payments guarantee that if a worker reports a cost higher than the threshold payment, this worker will not be selected.

i) Monotonicity: For any worker $s_i \in S_w$, if s_i decreases their bid to $b'_i < c_i$, we prove that worker s_i will still be selected. Thus, CARE-CO satisfies monotonicity. *ii) Threshold payments:* According to the relationship between $\frac{B}{\mathcal{M}(\mathcal{S}_k)}$ and $\frac{b_{k+1}}{v_{k+1}}$ in the payment p_i, we consider two cases: $\frac{B}{\mathcal{M}(\mathcal{S}_k)} \leq \frac{b_{k+1}}{v_{k+1}}$ and $\frac{b_{k+1}}{v_{k+1}} < \frac{B}{\mathcal{M}(\mathcal{S}_k)}$. Then, we prove that any winner bidding a

[2] Note that all integer programs introduced in this paper can be solved in polynomial time by converting it to the Max-Flow problem.

cost higher than the threshold payment in both two cases will not obtain a higher utility. Therefore, CARE-CO guarantees truthfulness.

(5) Approximation ratio: Let ALG_g, OPT_g denote the procured reputation of CARE-CO and the optimal solution, respectively. We divide workers into two groups: workers before s_{k+1} and workers after s_k. For the workers before s_{k+1}, the optimal solution can achieve at most $\mathcal{M}(\mathcal{S}_k) = ALG_g$ reputation with cost zero. For the workers after s_k, the optimal solution can obtain at most $\mathcal{M}(\mathcal{S}_{k+1})$ reputation under the budget constraint. Then, we prove that $\mathcal{M}(\mathcal{S}_{k+1}) - \mathcal{M}(\mathcal{S}_k) \leq v_{k+1}$. Thus, we have $OPT_g \leq \mathcal{M}(\mathcal{S}_k) + \mathcal{M}(\mathcal{S}_{k+1}) \leq 2\mathcal{M}(\mathcal{S}_k) + v_{k+1}$, which implies $\frac{OPT_g}{ALG_g} \leq \frac{2\mathcal{M}(\mathcal{S}_k)+v_{k+1}}{\mathcal{M}(\mathcal{S}_k)} \leq 2 + \frac{v_{max}}{v_{min}}$. □

V. THE NON-COOPERATIVE BUDGET SETTING

In this section, we further propose CARE-NO to address the non-cooperative budget setting. Two key questions arise: (1) How can we measure requesters' employability (*i.e.*, the reputation they can obtain from workers) on varying budgets to ensure budget feasibility? (2) Given the varying employability of requesters, how can we efficiently assign workers to requesters under compatibility constraints and determine appropriate payments while ensuring truthfulness.

To address these two critical questions, we first introduce a virtual-price based sub-mechanism called PEA, which serves as a core component of CARE-NO. PEA treats all workers as having the same reputation. Specially, PEA introduces a non-trivial concept of *virtual price*, which helps to understand each requester's employability. By Utilizing the virtual price and the requester's employability, PEA employs an integer program to assign workers under the compatibility constraint and identifies a critical price that ensures both efficiency and truthfulness. Subsequently, we introduce CARE-NO, which divides all workers into multiple sets, ensuring that each set of workers has similar reputations, and applies PEA to address each worker set separately. Detailed explanations of PEA and CARE-NO are provided in Section V-A and V-B, respectively.

A. Design of PEA

We first consider the scenario that all workers have the same reputations and propose PEA (detailed in Algorithm 2) to hire as many as possible number of workers for requesters.

a) Virtual price set: We first introduce the concept of *virtual price* which determines the maximum employable number of workers for requesters under budget constraints. We sort all workers in the non-decreasing order of their bids, *i.e.*, $b_1 \leq b_2 \leq \cdots \leq b_n$, and assign a weight $w_i = 2^i$ to worker s_i. Let $W = \{w_1, ..., w_n\}$ denote the weight profile of the workers. For each requester a_j, the employable number of workers falls within the range $[1, n]$. Denote by $\frac{B_j}{t}$ the maximum price at which requester a_j can hire t workers. Then, we can use the set of prices $\{\frac{B_j}{t}\}_{t \leq n}$ to differentiate the employability of the requester a_j. We define *the virtual price set* $R_b = \{\frac{B_j}{t}\}_{\forall j \leq m, t \leq n}$ to save these prices from all requesters. Specially, we define $\mathcal{E}(r) = \sum_{j \leq m} \lfloor \frac{B_j}{r} \rfloor$ as the *requesters' employability* under price $r \in R_b$. Let $\mathcal{S}(r) = \{s_i | b_i \leq r\}$ represent the *available worker set* with bids no higher than r.

Algorithm 2: PEA $(\mathcal{B}, \mathbf{b}, A, \mathcal{G})$

Output: P, S_w
$P \leftarrow 0, S_w \leftarrow \emptyset$;
Sort all workers as $b_1 \leq b_2 \leq \cdots b_n$;
$w_i \leftarrow 2^i, \forall i \leq n$;
Generate the virtual price R_b;
for $r \in R_b$ **do**
  Compute $\mathcal{M}_f(r)$ by Eq. (5) - (9);
end
Find $r^* \in \arg\min_{r \in \mathcal{R}_b} \{\mathcal{E}(r) = \mathcal{M}_f(r)\}$;
Candidate worker set is $\mathcal{S}(r^*)$;
// **Winner selection**;
Compute the allocation by Eq. (10) - (15) with worker set $\mathcal{S}(r^*)$;
If $x_i = \sum_{j \leq m} x_{ij} = 1$, $s_i \in S_w$;
Selected workers perform training tasks for requesters;
// **Payment Scheme**;
for $s_i \in S_w$ **do**
  for $b_l \geq b_i$ **do**
    Run PEA$(\mathcal{B}, \mathbf{b}'_{b_l}, S, G)$;
    if *s_i is still selected as a winner* **then**
      $P_i \leftarrow P_i \cup b_{l+1}$;
    end
  end
  $p_i = \min\{r^*, \max_{b \in P_i}\{b\}\}$;
end

b) Optimal worker selection problem (OSP): Utilizing virtual prices and the requester's employability, we next define the problem of computing the maximum number of selected workers under compatibility constraints at a given price.

Definition 2 (Sub-problem OSP)**.** *Given $\forall r \in R_b$, requesters' employability $\mathcal{E}(r)$ and the available worker set $\mathcal{S}(r)$, then*

$$\textbf{\textit{OSP(r):}} \quad \max \sum_{s_t \in \mathcal{S}(r), j \leq m} x_{tj} \tag{5}$$

$$s.t. \sum_{s_t \in G_l \cup \mathcal{S}(r)} x_{tj} \leq \tau_{lj}, \forall j \leq m, l \leq L; \tag{6}$$

$$\sum_{j \leq m} x_{tj} \leq 1, \forall s_t \in \mathcal{S}(r); \tag{7}$$

$$x_{tj} \in \{0, 1\}, \forall s_t \in \mathcal{S}(r), j \leq m; \tag{8}$$

$$\sum_{s_t \in \mathcal{S}(r)} x_{tj} \leq \lfloor \frac{B_j}{r} \rfloor, \forall j \leq m; \tag{9}$$

where (9) indicates that the number of workers allocated to each requester cannot exceed their employment ability at price r. Denote by $\mathcal{M}_f(r)$ the maximum value of OSP(r). Fig. 2 illustrates the employability under price set R_b and the corresponding values of OSP(r).

c) Winner selection and payment scheme: Given the solution of OSP problem, we can find the minimum price $r^* \in R_b$, namely the *critical price*, such that the maximum number of allocated workers equals requesters' employability, *i.e.*, $r^* \in \arg\min_{r \in R_b} \{\mathcal{E}(r) = \mathcal{M}_f(r)\}$. We use $r^*_{\prec}$ and $r^*_{\succ}$ to denote the left and right adjacent prices of r^* in R_b. Suppose that s_k is the last worker with a cost no higher than r^*.

Winner Selection: To ensure that any winner remains selected after bidding a lower cost (as shown in Lemma 1),

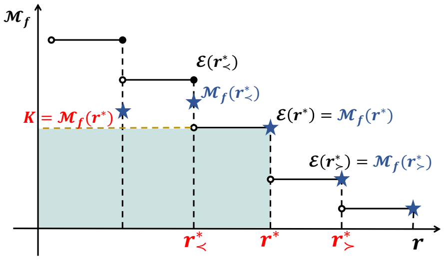

Fig. 2: The maximum employable worker curves under the price set R_b, where the black lines represent requesters' employability and the blue star represents the value of $\mathcal{M}_f(r)$.

which is crucial to guarantee the truthfulness of workers, we should choose a specific set of winners from worker set $\mathcal{S}(r^*)$. Particularly, we use the following integer program to select winners with the minimum sum of weights such that the total number of selected workers is still $\mathcal{M}_f(r^*)$, *i.e.*,

$$\min \sum_{i\le k, j\le m} w_i x_{ij} \tag{10}$$

$$s.t. \sum_{s_i \in G_l} x_{ij} \le \tau_{lj}, \forall j \le m, l \le L \tag{11}$$

$$\sum_{j\le m} x_{ij} \le 1, \forall i \le n \tag{12}$$

$$\sum_{i\le n} x_{ij} \le \lfloor \frac{B_j}{r^*} \rfloor, \forall j \le m \tag{13}$$

$$\sum_{i\le k, j\le m} x_{ij} = \mathcal{M}_f(r^*) \tag{14}$$

$$x_{ij} \in \{0,1\}, \forall i \le k, j \le m \tag{15}$$

where (13) indicates that the number of workers selected for each requester cannot exceed their employment ability at the price r^* and (14) means that we choose $\mathcal{M}_f(r^*)$ items. Note that there is only one optimal solution to Eq. (10) since the weight of s_i is 2^i. If $x_i = \sum_{j\le m} x_{ij} = 1, \forall i \le k$, then $s_i \in S_w$.

Payment Scheme: The intuition behind the payment scheme is to determine the maximum bid that the winner s_i can report while ensuring their status as a winner. For every winner $s_i \in S_w$, assume that s_i bids to the l-th position in the worker order, *i.e.*, bidding a higher cost $b_i' = b_l > b_i, \forall 1 \le l < n$, resulting in s_i becoming the new l-th worker in the new worker order $\mathbf{b}'_{b_l}$, *i.e.*, $\mathbf{b}'_{b_l} = \{b_1, \cdots, b_{i-1}, b_{i+1}, \cdots, b_l, b_i', b_{l+1}, \cdots, b_{n_h}\}$. Then, we run **PEA**$(\mathcal{B}, \mathbf{b}'_{b_l}, S, \mathcal{G})$ with input $\mathbf{b}'_{b_l}$ for every $b_l \ge b_i$. If s_i is still selected under the false cost $b_i' = b_l$, then add b_{l+1} to the candidate bid set, denoted by P_i. Finally, the payment of $s_i \in S_w$ is $p_i = \min\{r^*, \max_{b\in P_i}\{b\}\}$.

Next, we analyze the theoretical performance of PEA.

Theorem 2. *PEA guarantees individual rationality, budget feasibility, and computational efficiency, and achieves a* $(2\alpha + 1)$*-approximation where* $\alpha = \min\{m, \max_{l\le L, j\le m}\{\lceil |G_l|/\tau_{lj}\rceil\}\}$.

Order Index	1	⋯	l	$l+1$	$l+2$	$l+3$	⋯	$i-1$	i	$i+1$	⋯	k
weights	2^1	⋯	2^l	2^{l+1}	2^{l+2}	2^{l+3}	⋯	2^{i-1}	2^i	2^{i+1}	⋯	2^k
Bids	b_1	⋯	b_l	b_{l+1}	b_{l+2}	b_{l+3}	⋯	b_{i-1}	b_i	b_{i+1}	⋯	b_k
Bids (s_i bids lower)	b_1	⋯	b_l	b_i'	b_{l+1}	b_{l+2}	⋯	b_{i-2}	b_{i-1}	b_{i+1}	⋯	b_k

$S_{<1,l>}$: Workers' weights remain unchanged. $S_{<l+1,i-1>}$: Workers' weights double. $S_{<i+1,k>}$: Workers' weights remain unchanged.

Fig. 3: k workers' bids and weights when s_i bids a false cost.

Before proving truthfulness of PEA, we prove the following useful lemmas.

Lemma 1. *Assume that the winner* $\forall s_i \in S_w$ *bids a lower cost* $b_i' < c_i \le r^*$. *If the new critical price is still* r^*, *i.e., PEA chooses* $\mathcal{E}(r^*)$ *winners from worker set* $\mathcal{S}(r^*)$, *then* s_i *is still a winner.*

Proof. Note that s_k is the last worker with the bid no higher than r^*. When s_i bids b_i, the order of the first k workers' bids should be $b_1 \le b_2 \le b_3 \le \cdots b_k$, and the corresponding weight of the worker s_t for every $t \le k$ is $w_t = 2^t$. Based on the value of b_i', we consider two cases:

Case 1: If $b_{i-1} < b_i' \le b_{i+1}, \forall 2 \le i \le k$ or $b_i' \le b_2, i = 1$, then s_i remains the i-th worker and the weight is still 2^i. According to the computation of Eq. (10) with the input worker set $\mathcal{S}(r^*) = \{s_1, s_2, ..., s_k\}$, s_i is still a winner.

Case 2: If $b_l < b_i' \le b_{l+1}, \forall l \le i-2, i \le k$, then s_i will be the $(l+1)$-th worker with a new weight $2^{l+1} < 2^i$. Let $w_t', \forall t \le k$, and S_w' be the new weight of worker s_t and the new winner set after s_i bids a lower bid, respectively. As shown in Fig. 3, we can divide the workers except s_i into three groups: $S_{<1,l>} = \{s_1, ..., s_l\}, S_{<l+1,i-1>} = \{s_{l+1}, ..., s_{i-1}\}, S_{<i+1,k>} = \{s_{i+1}, ..., s_k\}$. We can find that the weights of the workers in $S_{<1,l>}$ and $S_{<i+1,k>}$ remain unchanged, whereas the weights of the workers in $S_{<l+1,i-1>}$ double, *i.e.*, $w_t' = 2w_t, \forall l+1 \le t \le i-1$. Let $\mathbb{S}_w(r^*)$ denote the set that contains all possible sets of winners computed by Eq. (5) with price r^* and $S_w \in \mathbb{S}_w(r^*)$. Thus, $\sum_{s_t\in S_w} w_t = \min_{S\in\mathbb{S}_w(r^*)} \sum_{s_t\in S} w_t$ since S_w is the final winner set.

Assume that s_i is not selected as a winner after bidding at a lower cost, *i.e.*, $S_w' \cap \{s_i\} = \emptyset$. Then, we have $\sum_{s_t\in S_w'} w_t' \ge \sum_{s_t\in S_w'} w_t$ since $w_t' \ge w_t, \forall t \le k, t \ne i$. Thus,

$$\sum_{s_t\in S_w'} w_t' \ge \sum_{s_t\in S_w'} w_t \ge \sum_{s_t\in S_w} w_t, \tag{16}$$

where the last inequality holds because S_w is the winning set with the smallest sum of weights. Let S_{l+1}^{i-1} and $< i_1, i_2, ..., i_{|S_{l+1}^{i-1}|} >$ denote the set of winners selected from set $S_{<l+1,i-1>}$ when s_i bids the true cost and the sequence of indexes of these workers, respectively. Then, we have

$$\begin{aligned}\sum_{s_t\in S_{l+1}^{i-1}\cup\{s_i\}} w_t' &= 2^{l+1} + 2(w_{i_1} + w_{i_2} + \cdots + w_{i_{|S_{l+1}^{i-1}|}}) \\ &\le 2^{l+1} + (w_{i_2} + \cdots + w_{i_{|S_{l+1}^{i-1}|}} + 2^i) \\ &\le w_i + w_{i_1} + \cdots + w_{i_{|S_{l+1}^{i-1}|}} = \sum_{s_t\in S_{l+1}^{i-1}\cup\{s_i\}} w_t.\end{aligned}$$

Consequently, we can have $\sum_{s_t \in S_w} w_t \geq \sum_{s_t \in S_w} w'_t$. If $\sum_{s_t \in S_w} w_t = \sum_{s_t \in S_w} w'_t$, then S_w remains the winner set and s_i is a winner, leading to the contradiction that s_i is not selected. If $\sum_{s_t \in S_w} w_t > \sum_{s_t \in S_w} w'_t$, by combining Eq. (16), we can conclude that $\sum_{s_t \in S'_w} w'_t > \sum_{s_t \in S_w} w'_t$. This suggests that, when worker s_i bids a lower cost, we should choose set S_w in order to obtain a smaller sum of weights compared with S'_w, which leads to the contradiction. Therefore, combining the above cases, this lemma holds. □

Lemma 2. *$\forall r_1, r_2 \in R_b, r^* \leq r_1 < r_2$, we have $\mathcal{M}_f(r) = \mathcal{E}(r), \forall r \in \{r_1, r_2\}$ and $\mathcal{M}_f(r_1) > \mathcal{M}_f(r_2)$.*

Proof. Based on the definition of price set R_b, $\mathcal{E}(r_1) > \mathcal{E}(r_2)$. As r^* is the critical price, we have $\mathcal{M}_f(r^*) = \mathcal{E}(r^*)$. For every price $r \in R_b$ and $r > r^*$, we have $\mathcal{S}(r^*) \subseteq \mathcal{S}(r)$, and we can therefore assign $\mathcal{E}(r)$ workers to requesters as the mechanism can assign $\mathcal{M}_f(r^*)$ workers to requesters where $\mathcal{M}_f(r^*) = \mathcal{E}(r^*) > \mathcal{E}(r)$. Thus, we have $\mathcal{M}_f(r_1) = \mathcal{E}(r_1) > \mathcal{M}_f(r_2) = \mathcal{E}(r_2)$. □

Lemma 3. *Given $\forall r \in R_b$ and the corresponding set of workers $\mathcal{S}(r)$, if we remove any worker $s_t \in \mathcal{S}(r)$, then the new value $\mathcal{M}_f(r)'$ satisfies $\mathcal{M}_f(r)' \geq \mathcal{M}_f(r) - 1$.*

Proof. Let $S_w(r)$ denote the selected workers of $\mathcal{M}_f(r)$, and we have $\mathcal{S}(r)' = \mathcal{S}(r) \setminus \{s_t\}$. If $s_t \in S_w(r)$, then we can still choose $S_w(r) \setminus \{s_t\}$ that satisfies the compatibility constraints in Eq. (5), which implies $\mathcal{M}_f(r)' = |S_w(r) \setminus \{s_t\}| = \mathcal{M}_f(r) - 1$. If $s_t \notin S_w(r)$, then $\mathcal{M}_f(r)' = \mathcal{M}_f(r)$. □

Theorem 3. *PEA guarantees truthfulness.*

Proof. We still leverage the famous Monotone Theorem [37] so that we can prove the mechanisms to be truthful. Monotone Theorem shows that truthful mechanisms satisfy monotonicity and workers are paid threshold payments.

Monotonicity: Based on the definition of the price set R_b, the value of $\mathcal{E}(r), \forall r \in R_b$, remains unchanged if any worker reports a false cost. If worker s_i reports a lower cost $b'_i \leq r' \in R_b$ and $b'_i < c_i$ where $r' \leq r^*$, we consider two cases: **(1)** If the new critical price remains r^*, s_i maintains its status as the winner according to Lemma 1. **(2)** If the new critical price decreases to $r' \in R_b$ where $r' < r^*$: Let $\mathcal{S}(r')'$ denote the new set of workers with costs no higher than r' after s_i reports a false cost. Note that $\mathcal{S}(r')' = \mathcal{S}(r') \cup \{s_i\}$. We have $\mathcal{E}(r') = \mathcal{M}_f(r')'$ as r' is the new critical price, and $\mathcal{M}_f(r') < \mathcal{E}(r')$ as the critical price is r^* when s_i does not report a false cost. Thus, s_i will also be selected as a winner. Otherwise, we must have $\mathcal{M}_f(r') = \mathcal{M}_f(r')'$ due to $\mathcal{S}(r')' = \mathcal{S}(r') \cup \{s_i\}$. Thus, $\mathcal{M}_f(r')' < \mathcal{E}(r')$ since $\mathcal{M}_f(r') < \mathcal{E}(r')$.

Threshold payments: Recall that the payment of s_i is $p_i = \min\{r^*, \max_{b \in P_i}\{b\}\}$. According to the relationship between r^* and $\max_{b \in P_i}\{b\}$, we consider the following two cases:

Case 1 $p_i = r^*$: When s_i bids a higher cost $r^* < b'_i$, we consider two subcase: **Subcase 1** $\mathcal{M}_f(r^*)' < \mathcal{E}(r^*)$: we have $\mathcal{S}(r^*)' = \mathcal{S}(r^*) \setminus \{s_i\}$, and we should choose a higher price $r^{*\prime} > r^*$ as the new critical price. As r^* is the critical price, we have $\mathcal{M}_f(r^*) > \mathcal{M}_f(r^{*\prime})$ according to Lemma 2. If $r^{*\prime} < b'_i$, PEA selects the winners from the workers with a cost at most $r^{*\prime}$, resulting in s_i not being selected as the winner. If $r^{*\prime} \geq b'_i$, then it follows that $\mathcal{S}(r^{*\prime})' = \mathcal{S}(r^{*\prime})$. As proven in Lemma 3, $\mathcal{M}_f(r^*)' \geq \mathcal{M}_f(r^*) - 1$, thus we have $\mathcal{M}_f(r^*)' \geq \mathcal{M}_f(r^*) - 1 \geq \mathcal{M}_f(r^{*\prime}) = \mathcal{M}_f(r^{*\prime})'$ which means that we can choose $\mathcal{M}_f(r^{*\prime})'$ workers from the set $\mathcal{S}(r^*)'$. Let $w_{\mathcal{S}(r^*)'}$ represent the maximum weight among the workers in $\mathcal{S}(r^*)'$. Since s_i bids higher than r^*, the weight of s_i will be at least $2w_{\mathcal{S}(r^*)'}$. Consequently, s_i will not be selected since we can select a minimum of $\mathcal{M}_f(r^{*\prime})'$ workers from the set $\mathcal{S}(r^*)'$ whose total weight is smaller than $2w_{\mathcal{S}(r^*)'}$. **Subcase 2** $\mathcal{M}_f(r^*)' = \mathcal{E}(r^*)$: the threshold price is still r^*, and we will choose winners from workers with costs no higher than r^*, then s_i will not be selected.

Case 2 $p_i = \max_{b \in P_i}\{b\}$: If s_i bids a higher cost $p_i = \max_{b \in P_i}\{b\} < b'_i \leq r^*$, according to the definition of P_i, the worker s_i will not be selected as the winner. When s_i bids a higher cost $b'_i > r^*$, by **Case 1**, worker s_i will also not be selected as a winner.

Therefore, by applying Myerson's theorem, PEA guarantees truthfulness. □

B. Design of CARE-NO

Given the PEA sub-mechanism, we are ready to introduce CARE-NO. Intuitively, we divide all workers into multiple sets so that each set of workers has similar reputations. This way, we can treat each set of workers as having the same reputation and call PEA to address each worker set. In detail, let $\rho_i = \frac{v_i}{v_{min}} \geq 1$ represent the virtual reputation of s_i, and $\rho_{max} := \max_{i \leq n} \rho_i$. We divide all workers in S into $\gamma = \lceil \log_\varepsilon \rho_{max} \rceil$ sets $\mathcal{D} = \{D_1, ..., D_\gamma\}$ by their virtual reputations and $\varepsilon > 1$ is an appropriate predetermined parameter, *i.e.*,

$$\mathbb{D}(s_i) = \begin{cases} D_h, & \text{if } v_i \in (\varepsilon^{h-1}, \varepsilon^h], 1 \leq h \leq \gamma \\ D_1, & \rho_i = 1, \end{cases} \tag{17}$$

where $\mathbb{D}(s_i)$ refers to the set to which worker s_i is selected. Specifically, workers with virtual reputation 1 are assigned to the set D_1. Then, we view that each worker in the same set owns the same reputation and call PEA to deal with the workers in the same set. Denote by S_w^h and $\mathcal{P}_h$ the winners and the corresponding payment returned by PEA on set D_h, *i.e.*, $(\mathcal{P}_h, S_w^h) = \textbf{PEA}(\mathcal{B}, \mathbf{b}, D_h, \mathcal{G})$. Finally, we sample one of the outputs from all sets with probability $\frac{1}{\gamma}$ as the final solution.

Theorem 4. *CARE-NO guarantees individual rationality, truthfulness, budget feasibility, and computational efficiency, and achieves $(2\alpha + 1)\varepsilon\gamma$-approximation in expectation where $\alpha = \min\{m, \max_{l \leq L, j \leq m}\{\lceil |G_l| / \tau_{lj} \rceil\}$ and $\gamma = \lceil \log_\varepsilon \rho_{max} \rceil$.*

VI. Experiment

A. Experimental Settings

1) Setup: To vary the quality of different workers and subsequently their reputation, we consider the *noise label datasets* [38]: part of the worker data samples are incorrectly labeled, and the data accuracy rate and cost range for workers are presented in Table I. Before starting the experiment,

TABLE I: Bid Ranges w.r.t. Data Accuracy Rates.

Data accuracy rate	[0.4, 0.6)	[0.6, 0.8)	[0.8, 1.0]
Bid range	[2, 4]	[3, 5]	[4, 6]

the costs are randomly generated within their corresponding cost range. In addition, we first conduct 10 training tasks to calculate the reputations of workers using the method in [9]. We randomly assign the budget of each requester within the range [40, 80]. Furthermore, τ_{lj} is randomly assigned within the range $[1, |G_l|]$, and we set $\varepsilon = 10$ for CARE-NO. A total of 120 workers are established in FL. To demonstrate the influence of the number of requesters m, we fix the number of groups at 10 and vary the number of requesters from 2 to 12 in increments of 2. Similarly, to evaluate the impact of the number of groups L, we fix the number of requesters at 5 and vary the number of groups from 4 to 24 in increments of 4. All workers are randomly assigned to groups.

2) Datasets and Models: To validate the performance of the proposed mechanisms, we consider the task on two commonly adopted datasets: Fashion MNIST (FMNIST) and CIFAR-10. For FMNIST, we adopt a three-layer neural network [9], while for CIFAR-10, we use a CNN with three convolutional layers, followed by a maximum-pooling layer and two fully connected layers [39]. For each dataset, each worker is provided with a training set of size 2000, while the requesters have test and validation datasets of size 2000 each. Individual data is randomly drawn from the corresponding dataset.

3) Benchmarks: Because no prior work has considered our challenging settings, we compare the proposed mechanisms with the following two reasonable benchmarks. (1) **RRAFL**: The most relevant mechanism from [9], which focuses only on a single requester. We extend it to handle multiple requesters and groups of workers by assuming a virtual sum of requester budgets and randomly assigning winners to requesters without compatibility violation. (2) **RanPri**: A simple pricing mechanism that sets a random price within the cost range for each worker. If the price is at least equal to the worker's cost, the worker is selected as a winner and assigned to a random requester with sufficient remaining budget, ensuring no compatibility violations.

4) Metrics: We evaluate these mechanisms using the following two metrics: *(1) Overall Reputation*: $\sum_{i \le n} \sum_{j \le m} x_{ij} v_i$, which is the objective of our proposed mechanisms; *(2) Average Global Accuracy*: the average global model accuracy of requesters, *i.e.*, $\frac{1}{m} \sum_{j \le m} q_i$ where q_i is the global model accuracy of requester a_j.

B. Experimental Results

1) Overall Reputation of Selected Workers: Fig. 4 presents the overall reputation of the proposed mechanisms for different numbers of requesters and groups. Specifically, Figs. 4a, 4b and Figs. 4c, 4d show the results on FMNIST and CIFAR-10, respectively. **Firstly**, we observe that our proposed mechanisms, CARE-CO and CARE-NO, consistently outperform RRAFL and RanPri. In particular, the overall reputation of CARE-CO is significantly higher, with average improvements of 3 and 15 times compared to RRAFL and RanPri, respectively. Similarly, the overall reputation of CARE-NO has improved by factors of 2 and 12, respectively, compared to RRAFL and RanPri. The superior performance of our proposed mechanisms can be attributed to their consideration of worker cost-efficiency (*i.e.*, bids relative to reputations) and their ability to efficiently allocate the budget to select more efficient workers while accommodating worker compatibility. In contrast, RRAFL and RanPri struggle to allocate workers effectively in the presence of multiple budgets and worker compatibility. **Secondly**, we consistently find that the overall reputation of CARE-CO exceeds that of CARE-NO. This aligns with the intuitive understanding that sharing the budget can facilitate the hiring of more high-quality workers, thereby enhancing the overall reputation attained by requesters. **Lastly**, our proposed mechanisms demonstrate a stable increase in overall reputation as the number of groups increases. This can be attributed to the fact that a higher number of groups indicates a lower level of incompatibilities among workers, which in turn contributes to achieving a higher overall reputation.

2) Average Global Accuracy of Requesters: To validate the effectiveness of our proposed mechanisms, which not only exhibit high overall reputation efficiency but also improve model accuracy, we present the performance of average global accuracy for requesters as follows. Table II shows the impact of the number of requesters and groups on the average global accuracy for both FMNIST and CIFAR-10. **Firstly**, it is evident that the proposed mechanisms, CARE-CO and CARE-NO, consistently outperform RRAFL and RanPri. Specifically, the average accuracy of CARE-CO is 32.21% and 84.49% higher than that of RRAFL and RanPri on average, respectively. Similarly, the average accuracy of CARE-NO is 28.70% and 83.89% higher than that of RRAFL and RanPri on average, respectively. **Secondly**, it is worth noting that the improvement in accuracy achieved by our proposed mechanisms is not as significant as the improvement in overall reputation. This is because achieving higher accuracy requires the participation of more high-quality workers. Nevertheless, both CARE-CO and CARE-NO are capable of achieving significantly better average accuracy due to their considerably higher reputation. **Lastly**, we observe a slight decrease in average accuracy with an increasing number of requesters. This can be attributed to the fact that as the number of requesters increases, the number of workers allocated to each requester may slightly decrease, resulting in a decrease in average accuracy. Moreover, the overall average accuracy of our proposed mechanisms increases as the number of groups increases. The reason is that our proposed mechanisms can attain higher reputation in the presence of a lower level of incompatibilities among workers (*i.e.*, a larger number of groups), which in turn contributes to achieving higher model accuracy for the requesters.

3) Accuracy on Non-IID Label Distribution: In the previous experiments, we have considered the performance of the

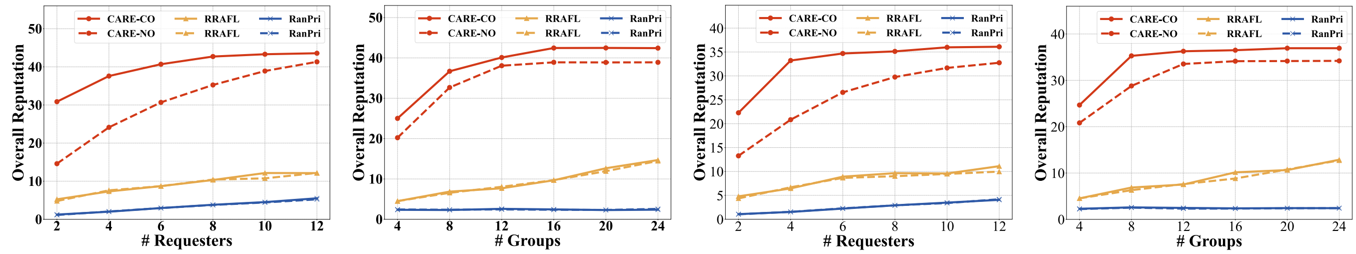

(a) #requesters for FMNIST (b) #groups for FMNIST (c) #requesters for CIFAR-10 (d) #groups for CIFAR-10

Fig. 4: The overall reputation of selected workers, with solid lines representing the reputation in the cooperative budget setting and dashed lines representing the reputation in the non-cooperative budget setting.

TABLE II: The Performance of the Average Global Accuracy

		Cooperative			Non-cooperative				Cooperative			Non-cooperative		
	# Req.	CARE-CO	RPAFL	RanPri	CARE-NO	RPAFL	RanPri	# Group	CARE-CO	RPAFL	RanPri	CARE-NO	RPAFL	RanPri
FMN-IST	2	**0.797**	0.796	0.588	**0.775**	0.759	0.653	4	**0.784**	0.571	0.573	**0.779**	0.664	0.570
	4	**0.795**	0.694	0.542	**0.787**	0.742	0.551	8	**0.781**	0.612	0.513	**0.782**	0.599	0.494
	6	**0.786**	0.654	0.519	**0.779**	0.653	0.531	12	**0.787**	0.648	0.509	**0.785**	0.621	0.507
	8	**0.790**	0.504	0.516	**0.782**	0.657	0.488	16	**0.789**	0.620	0.547	**0.782**	0.627	0.493
	10	**0.782**	0.691	0.559	**0.781**	0.621	0.501	20	**0.798**	0.683	0.564	**0.785**	0.638	0.553
	12	**0.789**	0.560	0.507	**0.781**	0.568	0.507	24	**0.797**	0.741	0.537	**0.793**	0.670	0.535
CIFA-R10	2	**0.546**	0.422	0.296	**0.537**	0.477	0.243	4	**0.549**	0.340	0.246	**0.533**	0.340	0.249
	4	**0.553**	0.431	0.211	**0.518**	0.427	0.235	8	**0.548**	0.355	0.263	**0.536**	0.392	0.258
	6	**0.543**	0.398	0.227	**0.508**	0.375	0.219	12	**0.559**	0.391	0.250	**0.555**	0.357	0.251
	8	**0.539**	0.375	0.354	**0.511**	0.368	0.249	16	**0.555**	0.378	0.221	**0.558**	0.444	0.240
	10	**0.533**	0.373	0.233	**0.538**	0.364	0.237	20	**0.551**	0.460	0.239	**0.553**	0.421	0.276
	12	**0.523**	0.316	0.232	**0.514**	0.341	0.238	24	**0.556**	0.473	0.233	**0.559**	0.444	0.234

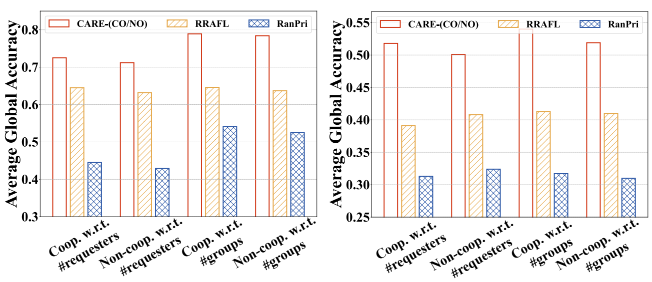

(a) FMNIST (b) CIFAR-10

Fig. 5: The average global accuracy of the proposed mechanisms under non-IID label distribution datasets.

proposed mechanisms in noisy label datasets where noise was added to create incorrect samples. In order to further validate the robustness and efficiency of our proposed mechanisms, we also compare the average global model accuracy with baselines on *non-IID label distribution* datasets [40], *i.e.*, the sample labels in the dataset of the workers are non-uniformly distributed. Fig. 5 illustrates the average accuracies of the proposed mechanisms in the FMNIST and CIFAR-10 datasets. We observe that the average global model accuracies of the proposed mechanisms decrease by only 4.25% and 3.75% on average compared to the accuracy obtained in Section VI-B2 for the FMNIST and CIFAR-10 datasets, respectively. This further validates the robustness of our proposed mechanisms on non-IID label distribution datasets. Furthermore, our proposed mechanisms continue to significantly outperform the baseline mechanisms in non-IID label distribution datasets. Specifically, CARE-CO shows an improvement of 24.01% and 61.13% compared to RRAFL and RanPri, respectively, in terms of average accuracy. Similarly, CARE-NO demonstrates an improvement of 21.28% and 59.33% compared to RRAFL and RanPri, respectively.

VII. Conclusion

In this paper, we consider compatibility-aware incentive mechanisms for incompatible workers in FL, where multiple requesters with budgets want to procure training services from groups of workers. For the cooperative budget setting, we propose CARE-CO, which leverages the Max-Flow solution to find feasible allocations under the compatibility constraint. Then, we propose CARE-NO for non-cooperative budget setting, which divides all workers into multiple sets and introduces a sub-mechanism PEA to address each worker set separately. In particular, PEA uses virtual prices to evaluate requesters' ability to obtain reputation and determines the critical price that aligns with their ability to ensure both budget feasibility and truthfulness. The proposed mechanisms can ensure individual rationality, budget feasibility, truthfulness, and approximation guarantee. Experimental results in real-world datasets, FMNIST and CIFAR-10, validate that our proposed mechanisms significantly outperform baseline mechanisms in terms of the overall reputation of selected workers and the average global accuracy.

References

[1] L. Li, Y. Fan, M. Tse, and K.-Y. Lin, "A review of applications in federated learning," *Computers & Industrial Engineering*, vol. 149, p. 106854, 2020.

[2] 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*, vol. 14, no. 1–2, pp. 1–210, 2021.

[3] S. Wang, T. Tuor, T. Salonidis, K. K. Leung, C. Makaya, T. He, and K. Chan, "Adaptive federated learning in resource constrained edge computing systems," *IEEE journal on selected areas in communications*, vol. 37, no. 6, pp. 1205–1221, 2019.

[4] J. Xu, B. S. Glicksberg, C. Su, P. Walker, J. Bian, and F. Wang, "Federated learning for healthcare informatics," *Journal of Healthcare Informatics Research*, vol. 5, pp. 1–19, 2021.

[5] G. Long, Y. Tan, J. Jiang, and C. Zhang, "Federated learning for open banking," in *Federated Learning: Privacy and Incentive*. Springer, 2020, pp. 240–254.

[6] C. Zhang, Y. Xie, H. Bai, B. Yu, W. Li, and Y. Gao, "A survey on federated learning," *Knowledge-Based Systems*, vol. 216, p. 106775, 2021.

[7] X. Tu, K. Zhu, N. C. Luong, D. Niyato, Y. Zhang, and J. Li, "Incentive mechanisms for federated learning: From economic and game theoretic perspective," *IEEE transactions on cognitive communications and networking*, vol. 8, no. 3, pp. 1566–1593, 2022.

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

[9] J. Zhang, Y. Wu, and R. Pan, "Incentive mechanism for horizontal federated learning based on reputation and reverse auction," in *Proceedings of the Web Conference 2021*, 2021, pp. 947–956.

[10] Y. Zhan, J. Zhang, Z. Hong, L. Wu, P. Li, and S. Guo, "A survey of incentive mechanism design for federated learning," *IEEE Transactions on Emerging Topics in Computing*, vol. 10, no. 2, pp. 1035–1044, 2021.

[11] S. Fan, H. Zhang, Y. Zeng, and W. Cai, "Hybrid blockchain-based resource trading system for federated learning in edge computing," *IEEE Internet of Things Journal*, vol. 8, no. 4, pp. 2252–2264, 2020.

[12] P. Sun, H. Che, Z. Wang, Y. Wang, T. Wang, L. Wu, and H. Shao, "Pain-fl: Personalized privacy-preserving incentive for federated learning," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 12, pp. 3805–3820, 2021.

[13] J. Zhang, Y. Wu, and R. Pan, "Online auction-based incentive mechanism design for horizontal federated learning with budget constraint," *arXiv preprint arXiv:2201.09047*, 2022.

[14] P. Roy, S. Sarker, M. A. Razzaque, M. Mamun-or Rashid, M. M. Hassan, and G. Fortino, "Distributed task allocation in mobile device cloud exploiting federated learning and subjective logic," *Journal of Systems Architecture*, vol. 113, p. 101972, 2021.

[15] J. Xu, H. Wang, and L. Chen, "Bandwidth allocation for multiple federated learning services in wireless edge networks," *IEEE transactions on wireless communications*, vol. 21, no. 4, pp. 2534–2546, 2021.

[16] T. Mai, H. Yao, J. Xu, N. Zhang, Q. Liu, and S. Guo, "Automatic double-auction mechanism for federated learning service market in internet of things," *IEEE Transactions on Network Science and Engineering*, vol. 9, no. 5, pp. 3123–3135, 2022.

[17] Y. Jiao, P. Wang, D. Niyato, B. Lin, and D. I. Kim, "Toward an automated auction framework for wireless federated learning services market," *IEEE Transactions on Mobile Computing*, vol. 20, no. 10, pp. 3034–3048, 2020.

[18] M. M. Amiri and D. Gündüz, "Federated learning over wireless fading channels," *IEEE Transactions on Wireless Communications*, vol. 19, no. 5, pp. 3546–3557, 2020.

[19] W. Huang, M. Ye, Z. Shi, G. Wan, H. Li, B. Du, and Q. Yang, "Federated learning for generalization, robustness, fairness: A survey and benchmark," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.

[20] A. Doan, P. Domingos, and A. Y. Halevy, "Reconciling schemas of disparate data sources: A machine-learning approach," in *Proceedings of the 2001 ACM SIGMOD international conference on Management of data*, 2001, pp. 509–520.

[21] J. Kang, Z. Xiong, D. Niyato, S. Xie, and J. Zhang, "Incentive mechanism for reliable federated learning: A joint optimization approach to combining reputation and contract theory," *IEEE Internet of Things Journal*, vol. 6, no. 6, pp. 10 700–10 714, 2019.

[22] E. L. Hultberg, C. Glendinning, P. Allebeck, and K. Lönnroth, "Using pooled budgets to integrate health and welfare services: a comparison of experiments in england and sweden." *Heart Lung & Circulation*, vol. 13, no. 6, pp. 531–541, 2010.

[23] Y. Qin and M. Kondo, "Mlmg: Multi-local and multi-global model aggregation for federated learning," in *2021 IEEE international conference on pervasive computing and communications workshops and other affiliated events (PerCom Workshops)*. IEEE, 2021, pp. 565–571.

[24] T. H. T. Le, N. H. Tran, Y. K. Tun, M. N. Nguyen, S. R. Pandey, Z. Han, and C. S. Hong, "An incentive mechanism for federated learning in wireless cellular networks: An auction approach," *IEEE Transactions on Wireless Communications*, vol. 20, no. 8, pp. 4874–4887, 2021.

[25] R. Zeng, S. Zhang, J. Wang, and X. Chu, "Fmore: An incentive scheme of multi-dimensional auction for federated learning in mec," in *2020 IEEE 40th international conference on distributed computing systems (ICDCS)*. IEEE, 2020, pp. 278–288.

[26] Y. Yang, M. Hu, Y. Zhou, X. Liu, and D. Wu, "Csra: Robust incentive mechanism design for differentially private federated learning," *IEEE Transactions on Information Forensics and Security*, 2023.

[27] K. Ren, G. Liao, Q. Ma, and X. Chen, "Differentially private auction design for federated learning with non-iid data," *IEEE Transactions on Services Computing*, 2023.

[28] S. Zheng, Y. Cao, M. Yoshikawa, H. Li, and Q. Yan, "Fl-market: Trading private models in federated learning," in *2022 IEEE International Conference on Big Data (Big Data)*. IEEE, 2022, pp. 1525–1534.

[29] X. Tan, W. Y. B. Lim, D. Niyato, and H. Yu, "Reputation-aware opportunistic budget optimization for auction-based federation learning," in *2023 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2023, pp. 1–8.

[30] Y. Singer, "Budget feasible mechanisms," in *2010 IEEE 51st Annual Symposium on foundations of computer science*. IEEE, 2010, pp. 765–774.

[31] N. Chen, N. Gravin, and P. Lu, "On the approximability of budget feasible mechanisms," in *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*. SIAM, 2011, pp. 685–699.

[32] N. Gravin, Y. Jin, P. Lu, and C. Zhang, "Optimal budget-feasible mechanisms for additive valuations," *ACM Transactions on Economics and Computation (TEAC)*, vol. 8, no. 4, pp. 1–15, 2020.

[33] G. Amanatidis, G. Birmpas, and E. Markakis, "On budget-feasible mechanism design for symmetric submodular objectives," in *International Conference on Web and Internet Economics*. Springer, 2017, pp. 1–15.

[34] W. Wu, X. Liu, and M. Li, "Budget-feasible procurement mechanisms in two-sided markets." in *IJCAI*, 2018, pp. 548–554.

[35] X. Liu, C. Fu, W. Wu, M. Li, W. Wang, V. Chau, and J. Luo, "Budget-feasible mechanisms in two-sided crowdsensing markets: Truthfulness, fairness, and efficiency," *IEEE Transactions on Mobile Computing*, 2022.

[36] J. Kang, Z. Xiong, D. Niyato, Y. Zou, Y. Zhang, and M. Guizani, "Reliable federated learning for mobile networks," *IEEE Wireless Communications*, vol. 27, no. 2, pp. 72–80, 2020.

[37] R. B. Myerson, "Optimal auction design," *Mathematics of operations research*, vol. 6, no. 1, pp. 58–73, 1981.

[38] Y. Li, J. Yang, Y. Song, L. Cao, J. Luo, and L.-J. Li, "Learning from noisy labels with distillation," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 1910–1918.

[39] Q. Li, B. He, and D. Song, "Model-contrastive federated learning," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 10 713–10 722.

[40] H. Wang, Z. Kaplan, D. Niu, and B. Li, "Optimizing federated learning on non-iid data with reinforcement learning," in *IEEE INFOCOM 2020-IEEE conference on computer communications*. IEEE, 2020, pp. 1698–1707.

[41] A. V. Goldberg and T. Radzik, *A heuristic improvement of the Bellman-Ford algorithm*. Citeseer, 1993.