Paid with Models: Optimal Contract Design for Collaborative Machine Learning

Bingchen Wang[1*], **Zhaoxuan Wu**[1,2], **Fusheng Liu**[1], **Bryan Kian Hsiang Low**[3]

[1]Institute of Data Science, National University of Singapore
[2]Singapore-MIT Alliance for Research and Technology, Republic of Singapore
[3]Department of Computer Science, National University of Singapore
bingchen@nus.edu.sg, {wu.zhaoxuan, fusheng}@u.nus.edu, lowkh@comp.nus.edu.sg

Abstract

Collaborative machine learning (CML) provides a promising paradigm for democratizing advanced technologies by enabling cost-sharing among participants. However, the potential for rent-seeking behaviors among parties can undermine such collaborations. Contract theory presents a viable solution by rewarding participants with models of varying accuracy based on their contributions. However, unlike monetary compensation, using models as rewards introduces unique challenges, particularly due to the stochastic nature of these rewards when contribution costs are privately held information. This paper formalizes the optimal contracting problem within CML and proposes a transformation that simplifies the non-convex optimization problem into one that can be solved through convex optimization algorithms. We conduct a detailed analysis of the properties that an optimal contract must satisfy when models serve as the rewards, and we explore the potential benefits and welfare implications of these contract-driven CML schemes through numerical experiments.

1 Introduction

Training a state-of-the-art machine learning (ML) model is a Herculean task due to the requirement of an enormous amount of data and computational resources. The exorbitant cost often precludes budget-constrained small parties from training a model on their own, resulting in a high industrial concentration where top-performing models are owned by big firms (AI Index Steering Committee 2024). In this regard, collaborative machine learning (CML) provides a promising crowdsourcing paradigm. The advent of CML schemes like federated learning (McMahan et al. 2017; Kairouz et al. 2021; Sheller et al. 2020; Nguyen et al. 2022) allows participants to join their resources for model training and share the training cost that would otherwise be insurmountable at an individual level. Despite their great potential, such schemes might not make economic sense. As is shown by Karimireddy, Guo, and Jordan (2022), catastrophic freeriding can occur when profit-maximizing parties in a collaboration have the ability to observe each other's data collection costs. This issue can be mitigated through the role of a scheme coordinator who conducts model training on the parties' behalf and rewards models with modified accuracy levels based on the parties' contributions. Practically, this could be achieved through the *design of contracts*, where the scheme coordinator acts as the *principal* and each participating party of the scheme acts as the *agent*.

Prior to our work, there has been a line of research that resorts to contract theory to address the incentive issue in collaborative machine learning (Kang et al. 2019; Ding, Fang, and Huang 2020; Karimireddy, Guo, and Jordan 2022; Liu et al. 2023), but most of them focus on using money as the reward for the collaboration. Karimireddy, Guo, and Jordan (2022) attends to the administration of models with different accuracy levels as rewards, while their primary focus is on the case where the scheme coordinator can directly observe each party's data collection costs. However, in reality, the cost of contribution is typically *private information* known only to the contributing party. For instance, consider a CML scheme where private computing firms pull together their GPUs for the training of a language model for code generation. Each firm could face a different vendor price and incur dissimilar maintenance cost of the chips. As another example, consider the CML scheme where investment firms join their privately curated data for the training of an investment model. To gather the data, each firm needs to recruit analysts, the overheads of which are usually determined by conditions of the local labor market and the firm's own incentive policies. The differences in the operating environments cause the parties of a CML scheme to have a wide range of per-unit contribution costs. While the scheme coordinator can be an expert in the domain field, thereby possessing some general information about the process, it remains challenging for them to gauge the exact costs borne by the parties. Even if the parties willingly inform the coordinator of their costs, the coordinator cannot verify the truthfulness of these reports without incurring significant auditing expenses. Worse still, a rent-seeking party may cheat by misreporting their cost if it leads to higher profits being gained from the scheme. This information asymmetry results in what is known as a *principal-agency problem* in economic literature (see Mas-Colell, Whinston, and Green 1995; Laffont and Martimort 2002; Bolton and Dewatripont 2004 for a comprehensive treatment of the subject).

In the presence of private information, optimal contract design with models as the rewards poses unique challenges that distinguish it from its economic counterparts. For one,

*Corresponding Author.

unlike money, models are a non-rivalrous and non-exclusive good, and can be replicated and offered to the participants at a nominal cost if not free of charge. Therefore, the scheme coordinator would find it tempting to offer less capable parties a good-performing model as long as it does not cause the more capable parties to cheat. For another, the administrable model rewards are constrained by the accuracy level of the model trained using all parties' data or computational resources. Due to *incomplete information*, the coordinator cannot observe the exact numbers of parties with different contribution costs in the CML scheme, and consequently cannot determine the exact accuracy level of the collectively trained model before the training completes. This makes the rewards of the contract stochastic ex-ante. The optimal contracting problem for CML needs to accommodate these challenges, whilst heeding the classical requirements of individual rationality and incentive compatibility. To this end, our paper makes the following contributions:

- We provide a *coherent formalization* of the optimal contracting problem in CML with models as the rewards, casting it as a constrained optimization problem.
- We *simplify* the original non-convex constrained optimization problem into one that can be solved using numerical optimization algorithms.
- We conduct *theoretical analysis of the constraints*, delineating the properties optimally designed contracts should obey for both scenarios when the coordinator can and cannot observe parties' contribution costs.
- We *illustrate the potential and the welfare implications* of optimally designed contracts through numerical experiments, showing that, inter alia, it could help small parties surmount the cost barrier of model training and reaping the reward of emergent technologies.

2 Problem Setup

Collaborative Machine Learning. We consider a typical setting of collaborative machine learning where budget-constrained parties contribute their resources, such as data and GPUs, to collectively train a model through the mediation of a scheme coordinator. We capture the dependency between contributed resources and model accuracy through the function $a(\cdot)$, and make the standard assumptions that more contribution leads to better model performance, $a'(\cdot) > 0$, and that the same amount of contribution has diminishing marginal returns as the total contribution increases, $a''(\cdot) < 0$. To ensure that the framework is well-grounded economically, we additionally represent the dependency between accuracy of a model and the economic profit gained by the participant through a weakly concave and increasing valuation function $v(\cdot)$, with $v'(\cdot) > 0$ and $v''(\cdot) \leq 0$. In the computing firm example mentioned in the introduction, the accuracy function captures the fact that the use of more GPUs enables more iterations of model training to be undertaken within the given time span; the valuation function reflects the fact that improved performance of the language model reduces human overheads—enabling a firm to hire fewer software engineers for the same amount of work. In the investment firm example, $a(\cdot)$ captures the fact that more data leads to a fuller picture of the market landscape, which in turn boosts the predictive accuracy of the trained model; $v(\cdot)$ embodies the fact that a model's predictive accuracy directly affects the quality of a firm's investment decisions and better decisions lead to more revenues being earned by the firm. To consolidate the intuition, we henceforth use data consistently throughout the paper as the resources contributed by parties in the collaboration, whereas the analysis can also apply readily to other resource types, like GPUs or computation devices.

Principal-Agent Problem. There are two distinct features of the CML setting mentioned above: Firstly, there is a *conflict of interests* between the scheme coordinator, who aims to maximize the performance of the collectively trained model by encouraging data contribution, and the scheme participant, who hopes to receive a good-performing model by contributing as little data as possible since collection is costly. Secondly, the scheme participant possesses *information advantage* over the scheme coordinator, as the per-unit data collection cost is directly known by the party themself but may not be observed by the coordinator likely due to costly verification or auditing process. This opens up the possibility that a party might cheat by misreporting their cost so that they can be compensated better by the coordinator. To facilitate the formalization of the model, we introduce the following notation: Let I denote the total number of distinct types of private per-unit data collection costs in the population—henceforth called *private types* or *types*; we denote the per-unit data costs by $c_i, i = 1, \ldots, I$, and order them decreasingly, $c_1 > \ldots > c_I$, for analytical tractability. Hence, a type-1 party incurs the highest per-unit cost while contributing data to the CML scheme. We let $n = (n_1, \ldots, n_I)$ be a vector that counts the numbers of parties of each type in the collaboration and $N = \sum_i^I n_i$ denote the total number of participants of the CML scheme. Here, we assume the coordinator observes N but not n. She nevertheless has the general knowledge that n follows a multinomial distribution $\mathrm{Mul}(N, p)$ with probabilities $p = (p_1, \ldots, p_I)$ for the types. We denote the model reward received by a type-i party by r_i measured in model accuracy, their data contribution by m_i measured in amount, and their reservation utility by f_i measured in monetary terms and defined as the highest profits they could achieve by not participating in the CML scheme.

Model as the Reward. When we use models as the rewards for the CML scheme, a natural *budget constraint* we need to abide by is

$$\|r(n)\|_\infty \leq a\left(\sum_{i=1}^{I} n_i m_i\right), \forall n \in \mathrm{Mul}(N, p)\ ,$$

where $r(n) = \big(r_1(n), \ldots, r_I(n)\big)$ is a vector with elements specifying the model rewards received by different types of parties under the realized combination n. In words, for every possible combination of types, the maximal model accuracy reward assigned to the parties cannot surpass the accuracy of the collectively trained model. Fixing the contributions $m_i, i = 1, \ldots, I$ from each type of parties, we see that this upper bound is dependent on n. Consequently, the model

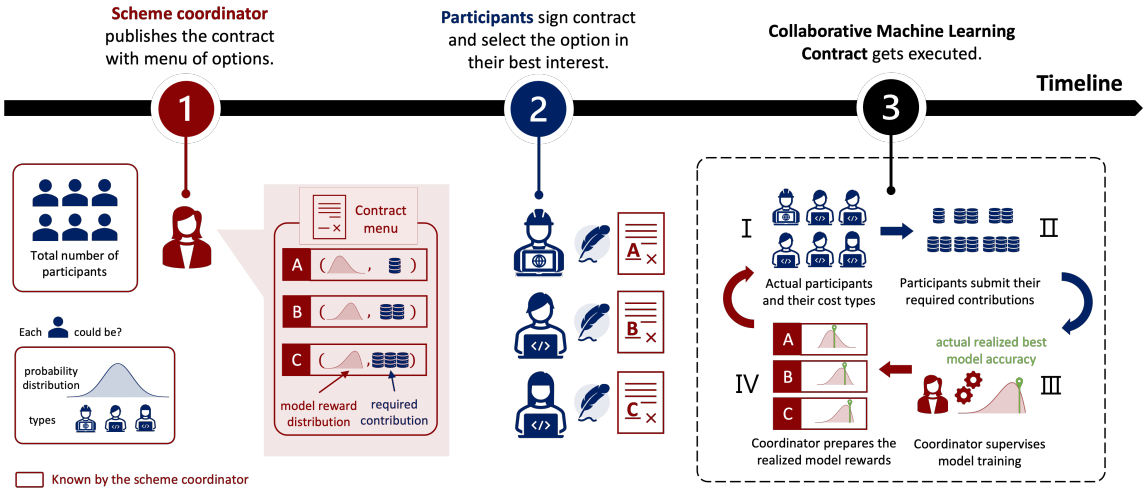

Figure 1: Optimal Contract Design for Collaborative Machine Learning: The Timeline.

rewards the coordinator can assign also depend on n. At the stage of contract design, unless the coordinator can perfectly observe the private types of the parties, she does not know n and therefore the model rewards are stochastic. We accommodate this fact with the notation $\mathbf{r}_i, \forall i = 1, \ldots, I$, which embodies the reward distributions for the types. Figure 1 shows the complete timeline of the contracting process.

3 Optimal Contracting in CML

With the notation in place, we are now ready to formalize the optimal contracting problem in CML.

Coordinator's Objective. The objective of the coordinator reflects the goal of the CML scheme. Here we assume that she aims to maximize the expected model accuracy,

$$\mathbb{E}_{n\sim\text{Mul}(N,p)}\left[a\left(\sum_{i=1}^{I} n_i m_i\right)\right].$$

For readability, we henceforth abbreviate $\mathbb{E}_{n\sim\text{Mul}(N,p)}$ to $\mathbb{E}$ unless otherwise stated.

Participant's Utility Function. Each participant is assumed to be a von Neumann–Morgenstern utility maximizer (von Neumann and Morgenstern 1944), who has the following utility function:

$$u_i(\mathbf{r}_i, m_i) = \mathbb{E}_{n_i\geq 1}[v(r_i)] - c_i m_i \,. \tag{1}$$

Recall that r_i denotes the stochastic model reward (measured in accuracy), which takes on different values for different n; m_i is the amount of data contribution; and c_i denotes the private per-unit data cost. $v(\cdot)$ reflects the profit mechanism that maps a model accuracy level to a pecuniary amount. The expectation operator $\mathbb{E}_{n_i\geq 1}$ codifies the information advantage possessed by the type-i participant—when they decide to participate, they know in advance that there is at least one party making the type-i commitment. The aim of each participant is to maximize their expected net profit, represented by the utility function in (1).

Individual Rationality. To ensure the contract is well designed, it must pass the first test that it gives parties of different types enough incentives to join the CML scheme. Formally, this requires that a party upon choosing the contract option designed for their type cannot be made worse off than them not participating in the CML scheme. This is known as the *individual rationality* (IR) conditions. To formalize the idea, we should specify the reservation utility (a.k.a. opportunity cost) for each type of parties. Here we define a party's reservation utility to be the utility level they achieves by training a model on their own, which amounts to solving the following optimization problem:

$$f_i \triangleq \max_{m\geq 0}\ \tilde{u}_i(m) = v(a(m)) - c_i m \,. \tag{2}$$

We make the following observation upon solving this problem, the proof of which can be found in Appendix A.

Proposition 1. *Let $\bar{m}_i$ denote the data contribution a type-i party is willing to commit to when training a model on their own. If $c_i \leq c_j$, then $\bar{m}_i \geq \bar{m}_j$, and $f_i \geq f_j$.*

The proposition states that when training a model alone, a party with lower data cost is willing to utilize more data and will end up with a better model and generate higher profits. We can write the IR conditions as follows:

$$\mathbb{E}_{n_i\geq 1}[v(r_i)] - c_i m_i \geq f_i \,, \forall i \,.$$

Using this formulation, readers familiar with the economic literature could also interpret f_i as the fixed cost and c_i the variable cost of joining the CML scheme for type-i parties.

Incentive Compatibility. Designing a contract $\mathcal{C}$ in the presence of hidden information seems highly complicated, as participants could lie about their private costs. Luckily, with the aid of the revelation principle (Myerson 1981, 1982; see Mas-Colell, Whinston, and Green 1995 for a lucid explanation on the concept), we could confine the design space to one contract option per private type, with each party choosing the option designed for their type upon signing the contract. In the CML setting, this translates to designing a contract option that stipulates a required data contribution m_i

and the corresponding reward distribution $\mathbf{r}_i$ for each type of parties. The veracity of parties is formally ensured through the *incentive compatibility* (IC) constraints:

$$\mathbb{E}_{n_i \geq 1}[v(r_i)] - c_i m_i \geq \mathbb{E}_{n_j \geq 1}[v(r_j)] - c_i m_j, \forall i, j \ .$$

In words, for each party, choosing the option designed for their type yields a weakly higher expected profit than choosing any other option on offer. We assume a tie is broken in favour of veracity. Note the change from $n_i \geq 1$ to $n_j \geq 1$ in the expectation operator in defining the IC constraints, which again reflects participant's information advantage—they know once they choose another type's option, there is at least one party making such a commitment.

Optimal Contract Design in CML. Amalgamating the constraints and the coordinator's objective, we obtain the constrained optimization problem for optimal contract design in CML using models as the rewards:

$$\max_{\mathcal{C} \triangleq (\mathbf{r}_i, m_i)_{i=1}^I} \ \mathbb{E}_{n \sim \text{Multi}(N,p)} \left[a \left(\sum_{i=1}^I n_i m_i \right) \right] \ \text{s.t.} \tag{3}$$

$$\begin{cases} \mathbb{E}_{n_i \geq 1}[v(r_i)] - c_i m_i \geq f_i, \forall i \ ; \\ \mathbb{E}_{n_i \geq 1}[v(r_i)] - c_i m_i \geq \mathbb{E}_{n_j \geq 1}[v(r_j)] - c_i m_j, \forall i, j \ ; \\ \|r(n)\|_\infty \leq a \left(\sum_{i=1}^I n_i m_i \right), \forall n \in \text{Multi}(N,p) \ . \end{cases} \tag{4}$$

Before solving this problem, we detour to analyze a special benchmark where the coordinator has the ability to observe the private information possessed by the participants, which is conventionally called the *complete information* scenario.

4 Contracting with Observable Costs

We conduct an analysis of the complete information scenario, as it offers insights on two fronts. Firstly, it helps establish a welfare benchmark, against which we could evaluate the welfare loss incurred by the existence of hidden information. Secondly, it models particular CML settings, in which it is relatively cheap to obtain or elicit the private cost information. When the coordinator can observe a party's type, the IC conditions become redundant, as the coordinator can directly contract parties based on their costs. The problem thus reduces significantly to

$$\max_{\{r_i, m_i\}_{i=1}^I} \ a \left(\sum_{i=1}^I n_i m_i \right)$$
$$\text{s.t.} \begin{cases} v(r_i) - c_i m_i - f_i \geq 0 \ , \forall i \ ; \\ r_i \leq a \left(\sum_{i=1}^I n_i m_i \right), \forall i \ . \end{cases}$$

Note that we drop the expectation operators because the coordinator now fully knows the number of parties present for each type. Consequently, the rewards also become deterministic. Solving the problem leads to the following proposition.

Proposition 2. *Under the complete information scenario, the optimal strategy for the principal is to offer the best model to all participating parties and require them utilize the amount of data such that an party's IR constraint binds. The parties will in general utilize more data than they would when training a model on their own.*

The detailed proof is deferred to Appendix A, and we provide here a brief intuition on offering the best model to all. The key is that the full observability of a party's cost eliminates the possibility of cheating. Even if a party wishes to choose the option designed for another type that requires less data contribution, they can longer do so as the coordinator can embed the type into the option and easily verify a party's eligibility at the time of contract signing. Since models are freely replicable, granting the highest rewards to parties incentivizes them to make the highest possible level of contribution while satisfying the IR condition.

While the idea of offering the same model to all parties making different contributes may seem counter-intuitive from a fairness perspective, it actually still obeys the principle that a bigger contributor ends up with a higher profit.

Proposition 3. *Under the complete information scenario, a party with lower cost makes more contribution and obtains higher profits, thereby obeying the principle that a bigger contributor ends up with a higher profit.*

The proof is slightly involved and thus deferred to Appendix A. We conclude this section by noting that the results hold with the goal of maximizing the accuracy of the collectively trained model. In practice, if the coordinator wants to enact additional fairness requirements, such as letting bigger contributors gain more through the CML scheme, she can do so by modifying the IR constraints. We provide such a framework in Appendix B for interested readers.

5 Contracting with Private Costs

When data costs are private information of the parties, it could have serious implications on the CML scheme without contract. As we demonstrate in Appendix B, a complete collaboration failure could occur where all parties contribute nothing to the scheme. In other cases, equilibrium does not exist, making the learning outcome unpredictable. Designing a contract helps address these issues but solving the problem defined by (3) and (4) directly is difficult due to the non-convexity of the constraints and the enormous number of choice variables. Luckily, the problem can be simplified on two fronts to improve its tangibility. Firstly, we transform the original problem into a convex constrained optimization problem with respect to first moments (termed the *first-moment problem*), which significantly reduces the number of choice variables we need to optimize with. We then derive a mapping from the solution of the first-moment problem to one that elegantly solves the original problem. Secondly, we conduct constraint analysis of the first-moment problem, further removing redundant constraints and delineating the properties an optimal contract should satisfy.

5.1 First-moment problem

The key to converting the original problem into a first-moment problem lies in the relaxation of the budget constraint. The following provides a necessity result.

Proposition 4. *The budget constraint in the original problem implies the budget constraint in first moments. Namely,*

$$\|r(n)\|_\infty \leq a\left(\sum_{i=1}^{I} n_i m_i\right) \implies$$
$$\mathbb{E}_{n_i\geq 1}[v(r_i)] \leq \mathbb{E}_{n_i\geq 1}\left[v\left(a\left(\sum_{i=1}^{I} n_i m_i\right)\right)\right], \forall i\ .$$

The proof is left to Appendix A. With this relaxation, all constraints in the original optimization problem are related to only the first moments of $v(r_i), \forall i$. For notational brevity, we let $t_i \triangleq \mathbb{E}_{n_i\geq 1}[v(r_i)]$, which denotes the expected revenue gained by a type-i party using the rewarded model from the scheme. The first-moment problem is defined as follows:

$$\max_{\{(t_i,m_i)_{i=1}^I\}} \mathbb{E}_{n\sim\text{Multi}(N,p)}\left[a\left(\sum_{i=1}^{I} n_i m_i\right)\right]$$
$$\text{s.t.} \begin{cases} t_i - c_i m_i \geq f_i\ , \forall i\ ; \\ t_i - c_i m_i \geq t_j - c_i m_j\ , \forall i, j\ ; \\ t_i \leq \mathbb{E}_{n_i\geq 1}\left[v\left(a\left(\sum_{i=1}^{I} n_i m_i\right)\right)\right], \forall i\ . \end{cases}$$

Note that instead of optimizing with respect to the distributions $\mathbf{r}_i, \forall i$, we only need to optimize with respect to scalars $t_i, \forall i$ in the first-moment problem. To complete the transformation, we are left to show that there exists a mapping from the solution of the first-moment problem to that of the original problem. We do this by construction.

Proposition 5. *Let $(t_i^*, m_i^*)_{i=1}^I$ denote the solution to the first-moment problem, and $\bar{t}_i \triangleq \mathbb{E}_{n_i\geq 1}\left[v\left(a\left(\sum_{i=1}^{I} n_i m_i\right)\right)\right]$. Then, the following mapping (**proportional assignment**) maximizes the original problem defined by (3) and (4):*

$$r_i(n) = v^{-1}\left(\frac{t_i^*}{\bar{t}_i} v\left(a\left(\sum_{i=1}^{I} n_i m_i\right)\right)\right). \tag{5}$$

The proof is deferred to Appendix A, while the idea is intuitive: Regardless of the actual realization of the type distribution, we will assign the realized model rewards according to the ratios defined by the first moments. Note that there exist other assignment policies that solve the original problem while aligning with the first-moment problem solution, while the proportional assignment rule is one that is intuitive and easy to implement in reality.

5.2 Constraint analysis

The first-moment problem has in total $(I^2 + I)$ inequality constraints, which gets unwieldy as I rises. As we will see later, many of these constraints turn out to be superfluous, thus inviting a more precise description of the feasible set.

Proportional Fairness. We start with a useful observation about parties' levels of contribution and rewards in an optimal contract under incomplete information.

Proposition 6. *With incomplete information, a party with lower cost should contribute weakly more data and receive weakly better rewards, i.e., given $c_i > c_{i+1}$,*

$$m_{i+1} \geq m_i \quad \textit{and} \quad t_{i+1} \geq t_i\ .$$

The proof makes use of the IC constraints and is postponed to Appendix A. Importantly, Proposition 6 establishes a notion of fairness as a necessary condition of an optimal contract, suggesting that a bigger contributor must receive a better reward. This differs from incentive mechanism designs using cooperative game theory (CGT) where this property is treated as a desideratum based on which one devises a reward scheme (Sim et al. 2020).

Adjacent Comparisons Constraints. As each participant needs to compare their tailored option with the rest $(I-1)$ option, we have a copious set of $I(I-1)$ incentive compatibility constraints. The following result shows that the actual key constraints is much smaller in number. The proof is involved and is left to Appendix A for interested readers.

Theorem 1. *With incomplete information, the following two sets of constraints are equivalent:*

(a) $t_i - c_i m_i \geq t_j - c_j m_j, \forall i, j\ ;$

(b) $\begin{cases} m_{i+1} \geq m_i, \forall i \in \{1, \ldots, I-1\}\ ; \\ t_i - c_i m_i = t_{i-1} - c_i m_{i-1}, \forall i \in \{2, \ldots, I\}\ . \end{cases}$

Weak Efficiency. Proposition 6 establishes one connection with incentive mechanism designs in CGT. Here, we show another shared feature between the paradigms.

Proposition 7. *With incomplete information, the most cost-efficient type (party with the lowest variable cost) must be offered the best model, i.e.,*

$$t_I = \mathbb{E}_{n_I\geq 1}\left[v\left(a\left(\sum_{i=1}^{I} n_i m_i\right)\right)\right].$$

This result aligns with the concept of weak efficiency proposed by mechanism design studies using CGT (Sim et al. 2020). Namely, when employing only model rewards to incentivize collaborative machine learning, the best model must be awarded to one of the participants. Like Proposition 6, weak efficiency in this context is attained as a necessary condition rather than being treated as a desideratum.

Highest-Cost Type Break Even Condition. Finally, we show that the IR condition must bind for participants with the highest per-unit cost, leaving them no rent to be gained.

Proposition 8. *With incomplete information, a type-1 party would obtain utilities no greater than their reservation level if choosing options designed for the other types. In addition, the individual rationality constraint must bind for parties of type 1. Namely,*

$$t_j - c_1 m_j \leq f_1, \forall j \in \{2, \ldots, I\}\ ,$$
$$t_1 - c_1 m_1 = f_1\ .$$

In other words, the optimal contract that maximizes the expected accuracy of the collectively trained model should make participants with the highest per-unit cost of contribution indifferent between participating and opting out.

Simplified First-Moment Problem. With the above simplifications, the first-moment problem becomes

$$\max_{\{(t_i,m_i)_{i=1}^I\}} \mathbb{E}_{n\sim\mathrm{Mul}(N,p)}\left[a\left(\sum_{i=1}^I n_i m_i\right)\right] \tag{6}$$

$$\text{s.t.} \begin{cases} t_1 - c_1 m_1 - f_1 = 0\,; \\ t_i \le \mathbb{E}_{n_i\ge 1}\left[v\left(a\left(\sum_{i=1}^I n_i m_i\right)\right)\right], \forall i \in \mathcal{I}\,; \\ t_i - c_i m_i = t_{i-1} - c_i m_{i-1}, \forall i \in \{2,\dots,I\}; \\ m_i \ge m_{i-1}, \forall i \in \{2,\dots,I\}\,; \\ t_i - c_i m_i - f_i \ge 0, \forall i \in \{2,\dots,I\}\,. \end{cases} \tag{7}$$

We keep the inequality of the budget constraint for type-I parties so that the resulting problem is still convex (cf. Appendix B). The set of constraints in (7) fully specifies model rewards as a function of contributions:

$$t_i = \begin{cases} f_1 + c_1 m_1 & \text{if } i = 1\,; \\ f_1 + c_1 m_1 + \sum_{k=2}^i c_k (m_k - m_{k-1}) & \text{if } i > 1\,. \end{cases}$$

The problem defined by (6) and (7) is convex and can be solved by numerical optimization methods, such as the trust-region interior-point algorithm, which works by staying away from the boundary of the feasible region defined by the inequality constraints and weakening the barrier effects as the estimate of the solution gets increasingly accurate (Nocedal and Wright 2006).

6 Experiments

To gain numerical insights into optimal contract design, we conduct a series of experiments with specified forms of the accuracy function and the valuation function. Following Karimireddy, Guo, and Jordan (2022), we adopt the standard generalization bound (Mohri, Rostamizadeh, and Talwalkar 2018) as the accuracy function, expressed as follows:

$$a(m) := \max\left\{0, a_{opt} - \frac{\sqrt{2k(2+\log(m/k))+4}}{\sqrt{m}}\right\},$$

where m measures the quantity of data used for model training; a_{opt} is the optimal accuracy achievable by the model, and k captures the difficulty of the learning task. Arguably, this choice of $a(m)$ itself is not concave due to its piecewise nature, but this non-concavity can be addressed via the concavity of each piece of the function. We set $k = 1$ and $a_{opt} = 1$ for the experiments. We assume a constant return to the model accuracy, $v(x) := 100x$, so a model with perfect accuracy is worth 100 in monetary terms.

6.1 Two-type Case

We first focus on the case where there are two private types of parties, $\mathcal{I} = \{1, 2\}$, to analyze how contextual factors such as the total number of participants N and the type distribution p affect the design of optimal contracts. We specify the per-unit data cost for the high-cost type as $c_1 = 0.02$ and that for the low-cost type as $c_2 = 0.01$. In this setting, both types would have initial incentives to train a model on their own without the CML scheme—the high-cost type would use 715.6 units of data and obtain a model valued at 69.6 in monetary terms, and the low-cost type would use 1148.5 units of data to achieve a model reward worth 75.6. The top panel of Figure 2 depicts the optimal contract for the CML scheme under incomplete information for varied probability of high-cost type $p_1 \in [0, 1]$ and total number of participants $N \in [2, 100]$, the solutions of which are obtained by solving the corresponding first-moment problems. Under the incentivized CML scheme, both types contribute more data and obtain better models than their reservation levels. For contract design, ceteris paribus, a higher probability of the high-cost type in the population makes it less favorable to create distinct contract options for the types—making a pooling contract more likely to be optimal from the coordinator's perspective. In contrast, the total number of participants has a relatively marginal effect. When all other factors are held constant, a larger participant pool leads to greater differentiation between the options in a separating contract.

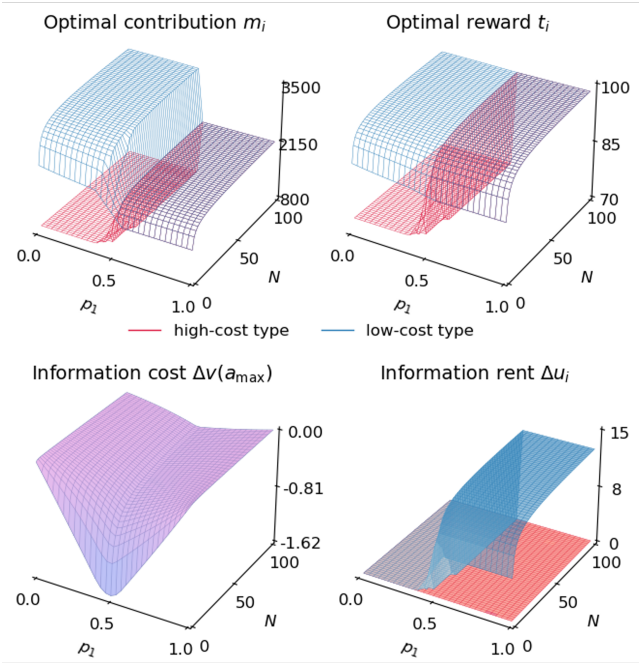

Figure 2: **Top**: Optimal contracts under incomplete information for varied probability of high-cost type $p_1 \in (0, 1)$ and total number of participants $N \in [2, 100]$, with $c = \{0.02, 0.01\}$. **Bottom**: Information costs for the coordinator and information rents for the parties under incomplete information vis-à-vis complete information.

To gauge the welfare implications of the information asymmetry, we calculate the information cost and information rent under incomplete information. The *information cost*, $\Delta v(a_{\max})$, is defined as the expected difference between the value of the collectively trained model under incomplete information and that under complete information:

$$\Delta v(a_{\max}) = \mathbb{E}\left[t_C - v\left(a\left(\sum_{i=1}^I n_i m_i^{\text{complete}}\right)\right)\right],$$

where $t_C = v(a(\sum_{i=1}^I n_i m_i))$ is the value of the col-

lectively trained model under incomplete information, and m_i^{complete} is the required data contribution from a type-i party under complete information, which varies with different realizations of n. Similarly, the *information rent*, Δu_i, is defined as the expected utility surplus of the party under incomplete information vis-à-vis complete information:

$$\Delta u_i = \mathbb{E}\left[u_i(t_i, m_i) - u_i\left(v(r_i^{\text{complete}}), m_i^{\text{complete}}\right)\right],$$

where (t_i, m_i) is the contract option designed for a type-i party under incomplete information and $v(r_i^{\text{complete}})$ is the model reward given to type-i under complete information, which depends on the realization of n. The bottom panel of Figure 2 shows the information costs and information rents for the two-type setting for varied p_1 and N. Aligned with general intuition, information asymmetry affects the collaboratively trained model most conspicuously when the number of participants is low and the probabilities of different types are similar. When one private type becomes dominant in the population, the optimal contract design accommodates by asking the corresponding type to contribute more at the sacrifice of the contribution from the other type, narrowing the gap with the solution under complete information. The low-cost type earns information rent under pooling contracts due to its ability to satisfy the stipulated contribution requirements at a lower cost than its high-cost counterparts.

6.2 Multi-type Case

Next, we consider more general cases where there are more than two private types in the population. For the ease of comparisons, we set $N = 10$ and $I = 5$, $p_i = 0.2, \forall i \in \mathcal{I}$ but vary the private costs c. We consider three different scenarios: 1) all types find it in their interest to train a model on their own, with $c = \{0.2, 0.16, 0.12, 0.08, 0.04\}$; 2) all types would not train a model on their own due to high per-unit costs, with $c = \{1, 0.85, 0.7, 0.55, 0.4\}$; 3) some types would train the model on their own and others would not, with $c = \{0.5, 0.4, 0.03, 0.02, 0.001\}$. Figure 3 shows the simulation results for the three scenarios. We highlight some of the key observations below.

A party may be incentivized to contribute less than their reservation level. This happens with the lowest-cost type (Type 5) in the first simulated scenario. This intriguing result is partially driven by the incentive compatibility constraint—as a higher contribution requirement would cause the type to deviate to the contract option designed for the adjacent type (cf. Appendix B for a graphical illustration).

Incentivized collaborative scheme can democratize machine learning. Row 2 of Figure 3 considers the case when the difficulty of the learning task prohibits all types from training a model on their own. With an effectively designed contract, this hurtle is overcome, with the cost of data collection shared among the participants and each of them receiving a model with decent accuracy as the reward.

In the presence of dominant players, small players can still gain from collaboration. The last scenario illustrated in Figure 3 conveys the idea that incentivized collaboration

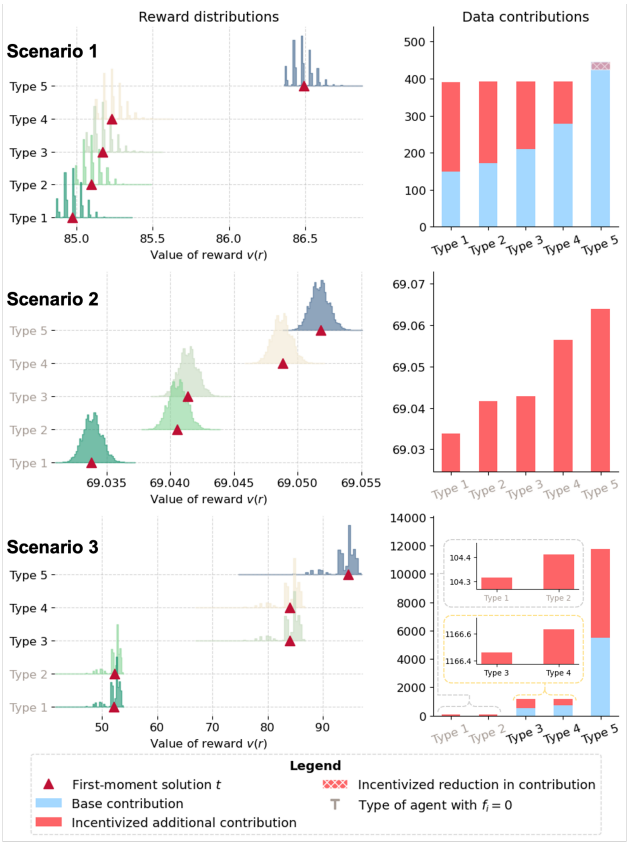

Figure 3: Optimal contract designs for multi-type scenarios. **Scenario 1**: All types would train a model on their own. **Scenario 2**: All types would not train a model on their own due to prohibitive costs. **Scenario 3**: Some types would train the model on their own and others would not.

is still possible despite significant differences in cost structures among the participants. In the simulated scenario, the per-unit data contribution cost of a type-1 party is 500 times the cost of a type-5 party. Yet, by contributing 104.3 units of data, a type-1 party can receive a model reward valued at 52.2, showing the trickle-down effect of the collaboration.

7 Conclusion

In this work, we consider optimal contract design for CML with models as the rewards. We convert the original non-convex problem of optimizing with reward distributions into one solvable through convex constrained optimization algorithms. Our constraint analysis establishes the necessary conditions for an optimal contract. We further demonstrate the framework through numerical experiments, showing its ability to overcome high cost of model training and improve participant welfare. Our findings highlight that optimal contract design is a viable tool for democratizing future technology in an incentive-driven economy. Future research could explore relaxing the distribution assumption to improve scalability. For a detailed discussion on future research directions, we refer interested readers to Appendix B.

Acknowledgments

This research is supported by the National Research Foundation Singapore and DSO National Laboratories under the AI Singapore Programme (AISG Award No: AISG2-RP-2020-018). The authors extend their gratitude to Dr. Wenjie Feng, Dr. Wenbo Zhao, Mingzhe Du, and four anonymous reviewers for their invaluable feedback on this work.

References

AI Index Steering Committee, S. U. 2024. AI Index 2024 Annual Report. Technical report, Stanford University.

Bolton, P.; and Dewatripont, M. 2004. *Contract Theory*. Cambridge, MA: MIT Press. ISBN 9780262025768.

Ding, N.; Fang, Z.; and Huang, J. 2020. Incentive Mechanism Design for Federated Learning with Multi-Dimensional Private Information. In *2020 18th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOPT)*, 1–8.

Kairouz, P.; McMahan, H. B.; Avent, B.; Bellet, A.; Bennis, M.; Nitin Bhagoji, A.; Bonawitz, K.; Charles, Z.; Cormode, G.; Cummings, R.; D'Oliveira, R. G. L.; Eichner, H.; El Rouayheb, S.; Evans, D.; Gardner, J.; Garrett, Z.; Gascón, A.; Ghazi, B.; Gibbons, P. B.; Gruteser, M.; Harchaoui, Z.; He, C.; He, L.; Huo, Z.; Hutchinson, B.; Hsu, J.; Jaggi, M.; Javidi, T.; Joshi, G.; Khodak, M.; Konecný, J.; Korolova, A.; Koushanfar, F.; Koyejo, S.; Lepoint, T.; Liu, Y.; Mittal, P.; Mohri, M.; Nock, R.; Özgür, A.; Pagh, R.; Qi, H.; Ramage, D.; Raskar, R.; Raykova, M.; Song, D.; Song, W.; Stich, S. U.; Sun, Z.; Suresh, A. T.; Tramèr, F.; Vepakomma, P.; Wang, J.; Xiong, L.; Xu, Z.; Yang, Q.; Yu, F. X.; Yu, H.; and Zhao, S. 2021. Advances and Open Problems in Federated Learning. *Foundations and Trends ® in Machine Learning*, 14(1-2): 1–210.

Kang, J.; Xiong, Z.; Niyato, D.; Yu, H.; Liang, Y.-C.; and Kim, D. I. 2019. Incentive Design for Efficient Federated Learning in Mobile Networks: A Contract Theory Approach. In *2019 IEEE VTS Asia Pacific Wireless Communications Symposium (APWCS)*, 1–5.

Karimireddy, S. P.; Guo, W.; and Jordan, M. 2022. Mechanisms that Incentivize Data Sharing in Federated Learning. In *International Workshop on Federated Learning: Recent Advances and New Challenges in Conjunction with NeurIPS 2022*.

Laffont, J.-J.; and Martimort, D. 2002. *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press. ISBN 978-0-691-09184-6.

Liu, Y.; Tian, M.; Chen, Y.; Xiong, Z.; Leung, C.; and Miao, C. 2023. *A Contract Theory Based Incentive Mechanism for Federated Learning*, 117–137. Cham: Springer International Publishing. ISBN 978-3-031-11748-0.

Mas-Colell, A.; Whinston, M. D.; and Green, J. R. 1995. *Microeconomic Theory*. New York: Oxford University Press. ISBN 9780195073409.

McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and Arcas, B. A. y. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 1273–1282. PMLR. ISSN: 2640-3498.

Mohri, M.; Rostamizadeh, A.; and Talwalkar, A. 2018. *Foundations of Machine Learning*. The MIT Press, 2nd edition. ISBN 0262039400.

Myerson, R. B. 1981. Optimal Auction Design. *Mathematics of Operations Research*, 6(1): 58–73. Publisher: INFORMS.

Myerson, R. B. 1982. Optimal coordination mechanisms in generalized principal–agent problems. *Journal of Mathematical Economics*, 10(1): 67–81.

Nguyen, A.; Do, T.; Tran, M.; Nguyen, B. X.; Duong, C.; Phan, T.; Tjiputra, E.; and Tran, Q. D. 2022. Deep Federated Learning for Autonomous Driving. In *2022 IEEE Intelligent Vehicles Symposium (IV)*, 1824–1830.

Nocedal, J.; and Wright, S. J. 2006. *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering. Springer New York. ISBN 978-0-387-30303-1.

Sheller, M. J.; Edwards, B.; Reina, G. A.; Martin, J.; Pati, S.; Kotrotsou, A.; Milchenko, M.; Xu, W.; Marcus, D.; Colen, R. R.; and Bakas, S. 2020. Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Scientific Reports*, 10(1): 12598. Publisher: Nature Publishing Group.

Sim, R. H. L.; Zhang, Y.; Chan, M. C.; and Low, B. K. H. 2020. Collaborative Machine Learning with Incentive-Aware Model Rewards. In *Proceedings of the 37th International Conference on Machine Learning*, 8927–8936. PMLR. ISSN: 2640-3498.

von Neumann, J.; and Morgenstern, O. 1944. *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.