A Mechanism Design Approach for Multi-Party Machine Learning

Mengjing Chen[1], Yang Liu[2], Weiran Shen[3], Yiheng Shen[4], Pingzhong Tang[1], and Qiang Yang[2,5]

[1] Tsinghua University
ccchmj@qq.com, kenshinping@gmail.com
[2] WeBank Co., Ltd.
yangliu@webank.com
[3] Renmin University
shenweiran@ruc.edu.cn
[4] Duke University
ys341@duke.edu
[5] Hong Kong University of Science and Technology
qyang@cse.ust.hk

Abstract. In a multi-party machine learning system, different parties cooperate on optimizing towards better models by sharing data in a privacy-preserving way. A major challenge in learning is the incentive issue. For example, if there is competition among the parties, one may strategically hide his data to prevent other parties from getting better models.

In this paper, we study the problem through the lens of mechanism design and incorporate the features of multi-party learning in our setting. First, each agent's valuation has externalities that depend on others' types and actions. Second, each agent can only misreport a type lower than his true type, but not the other way round. We call this setting *interdependent value with type-dependent action spaces*. We provide the optimal truthful mechanism in the quasi-monotone utility setting. We also provide necessary and sufficient conditions for truthful mechanisms in the most general case. We show the existence of such mechanisms is highly affected by the market growth rate. Finally, we devise an algorithm to find the desirable mechanism that is truthful, individually rational, efficient and weakly budget-balance.

Keywords: Mechanism design · Federated learning · Incentive design.

1 Introduction

In multi-party machine learning, a group of parties cooperates on optimizing towards better models. This concept has attracted much attention recently [12, 22, 23]. The advantage of this approach is that, it can make use of the distributed datasets and computational power to learn a powerful model that anyone in the group cannot achieve alone.

To make multi-party machine learning practical, a large body of works focus on preserving data privacy in the learning process [1, 26, 22]. However, the incentive issues in the multi-party learning have largely been ignored in most previous studies, which results in a significant reduction in the effectiveness when putting their techniques into practice. Previous works usually let all the parties share the same global model with the best quality regardless of their contributions. This allocation works well when there are no conflicts of interest among the parties. For example, an app developer wants to use the users' usage data to improve the user experience. All users are happy to contribute data since they can all benefit from such improvements [16].

When the parties are competing with one another, they may be unwilling to participate in the learning process since their competitors can also benefit from their contributions. Consider the case where companies from the same industry are trying to adopt federated learning to level up the industry's service qualities. Improving other companies' services can possibly harm their own market share, especially when there are several monopolists that own most of the data.

Such a cooperative and competitive relation poses an interesting challenge that prevents the multi-party learning approach from being applied to a wider range of environments. In this paper, we view this problem from the multi-agent system perspective, and address the incentive issues mentioned above with the mechanism design theory.

Our setting is a variant of the so-called interdependent value setting [18]. A key difference between our setting and the standard interdependent value setting is that each agent cannot "make up" a dataset that is of higher quality than his actual one. Thus the reported type of an agent is capped by his true type. We call our setting *interdependent value with type-dependent action spaces*. The setting that agents can never over-report is common in practice. One straightforward example is that the sports competitions where athletes can show lower performance than their actual abilities but not over-perform. The restriction on the action space poses more constraints on agents' behaviors, and allows more flexibility in the design space.

We first formulate the problem mathematically, and then apply techniques from the mechanism design theory to analyze it. Our model is more general than the standard mechanism design framework, and is also able to describe other similar problems involving both cooperation and competition.

We make the following contributions in this paper:

- We model and formulate the mechanism design problem in multi-party machine learning, and identify the differences between our setting and the other mechanism design settings.

– For the quasi-monotone externalities setting, we provide the revenue-optimal and truthful mechanism. For the general valuation functions, we provide both the necessary and the sufficient conditions for all truthful and individually rational mechanisms.
– We analyze the influence of the market size on mechanisms. When the market grows slowly, there may not exist a mechanism that achieves all the desirable properties we focus on.
– We design an algorithm to find the mechanisms that guarantee individual rationality, truthfulness, efficiency and weak budget balance simultaneously when the valuation functions are given.

1.1 Related Works

A large body of literature studies mechanisms with interdependent values [18], where agents' valuations depend on the types of all agents and the intrinsic qualities of the allocated object. Roughgarden et al. [21] extend Myerson's auction to specific interdependent value settings and characterize truthful and rational mechanisms. They consider a bayesian setting while we do not know any prior information. Chawla et al. [5] propose a variant of the VCG auction with reserve prices that can achieve high revenues. They consider value functions that are single-crossing and concave while we consider environments with more general value functions. Mezzetti [17] gives a two-stage Groves mechanism that guarantees truthfulness and efficiency. He requires agents to report their types and valuations before the final monetary transfer are made while in our model, agents can only report their types.

In our setting, agents have restricted action spaces, i.e., they can never report types exceeding their actual types. There is a series of works that focus on mechanism design with a restricted action space [3, 4, 2]. The discrete-bid ascending auctions [7, 6, 2] specify that all bidders' action spaces are the same bid level set. Several works restrict the number of actions, such as bounded communications [4]. Previous works focus on mechanisms with independent values and discrete restricted action spaces, while we study the interdependent values and continuous restricted action spaces setting.

The learned model can be copied and distributed to as many agents as possible, so the supply is unlimited. A line of literature focuses on selling items in unlimited supply such as digital goods [11, 10, 9]. However, the seller sells the same item to buyers while in our setting we can allocate models with different qualities to different agents.

Redko et al. [20] study the optimal strategies of agents for collaborative machine learning problems. Both their work and ours capture the cooperation and competition among the agents, but they only consider the case where agents reveal their total datasets to participate while agents can choose to contribute only a fraction in our setting. Kang et al. [14] study the incentive design problem for federated learning, but all their results are about a non-competitive environment, which may not hold in real-world applications.

Our work contributes to the growing body of literature on incentive mechanism design for federated learning [27, 15]. Jia et al. and Song et al. [13, 24] design mechanisms based on the Shapley value and Ding et al. [8] apply the contract theory. However, the existing works do not consider the interdependent values of participants and type-dependent action space as our model does.

2 Preliminaries

In this section, we introduce the general concepts of mechanism design and formulate the multi-party machine learning as a mechanism design problem. A multi-party learning consists of a central platform and several parties (called agents hereafter). The agents serve their customers with their models trained using their private data. Each agent can choose whether to enter the platform. If an agent does not participate, then he trains his model with only his own data. The platform requires all the participating agents to contribute their data in a privacy-preserving way and trains a model for each participant using a (weighted) combination of all the contributions. Then the platform returns the trained models to the agents.

We assume that all agents use the same model structure. Therefore, each participating agent may be able to train a better model by making use of his private data and the model allocated to him. One important problem in this process is the incentive issue. For example, if the participants have conflicts of interest, then they may only want to make use of others' contributions but are not willing to contribute with all their own data. To align their incentives, we allow the platform to charge the participants according to some predefined rules.

Our goal is to design allocation and payment rules that encourage all agents to join the multi-party learning as well as to contribute all their data.

2.1 Valid Data Size (Type)

Suppose there are n agents, denoted by $N = (1, 2, \dots, n)$, and each of them has a private dataset D_i where $D_i \cap D_j = \emptyset, \forall i \neq j$. For ease of presentation, we assume that a model is fully characterized by its quality Q (e.g., the prediction accuracy), and the quality only depends on the data used to train it. We have the following observation:

Observation 1 *If the agents could fake a dataset with a higher quality, any truthful mechanism would make agents gain equal final utility.*

Suppose that two agents have different true datasets D_1 and D_2. We assume that all other agents truthfully report datasets D_{-i}. If truthfully reporting dataset D_1 and D_2 finally leads to different utility, w.l.o.g, we let $u(D_1, D_{-i}) < u(D_2, D_{-i})$, then if an agent has true dataset D_1, he would report D_2 and use the dataset allocated by the platform in the market. All his behavior is the same as that of an agent with real dataset D_2. Thus if a mechanism is truthful, any reported dataset would lead to the same final utility and it is pointless to discuss the problem. Hence, we make the assumption that the mechanism is able to identify the quality of any dataset. All agents can only report a dataset with a lower quality.

For simplicity, we measure the contribution of a dataset to a trained model by its *valid data size*. Thus we have the following assumption:

Assumption 1 *The model quality Q is bounded and monotone increasing with respect to the valid data size $s \geq 0$ of the training data:*

1. $Q(0) = 0$ *and* $Q(s) \leq 1$, $\forall s$;
2. $Q(s') > Q(s)$, $\forall s' > s$.

The valid data size of every contributor's data is validated by the platform in a secure protocol (which we propose in the full version). Let $t_i \in \mathbb{R}_+$ be the valid data size of agent i 's private dataset D_i. We call t_i the agent's *type*. The agent can only falsify his type by using a dataset of lower quality (for example, using a subset of D_i, or adding fake data), which decreases the contribution to the trained model as well as the size of valid data. As a result, the agent with type t_i cannot contribute to the platform with a dataset higher than his type:

Assumption 2 *Each agent i can only report a type lower than his true type t_i, i.e., the action space of agent i is $[0, t_i]$.*

2.2 Learning Protocol

In this section, we describe the learning protocol that could enable the implementation of our mechanism. We assume that the platform has a validation dataset. The platform requires the agents to report their valid data size t_i. This could be done by asking each agent to submit the best model that he can possibly obtain by using his own dataset. Then the platform computes the model quality q_i using the validation dataset and get the agent's valid data size t_i by $t_i = Q^{-1}(q_i)$. The agent type t_i will be used in the training process (e.g., aggregate weighted model updates), as well as to determine the the final allocation, which is a model with quality x_i.

The platform should also guarantee to deliver to each agent the promised model. However, it is possible that an agent reports t_i in the beginning but only contribute $t'_i < t_i$ in the actual training process. In the extreme case where all agents contribute nothing to the training process, the platform will fail to allocate a model to each agent with the quality determined by the mechanism. To address this issue, the platform can train n additional models simultaneously, with the i-th model trained only using the data from agent i. During the training process, the platform can apply secure multi-party computation techniques, such as homomorphic encryption [26, 25], to prevent the agent from knowing which model is sent to him to compute the update. And after the training, the platform can compute the quality t'_i of the i-th model again using the validation dataset. If the qualities t'_i and t_i match, we know with high probability that the dataset contributed by the agent is consistent with the type he reports. Otherwise, the platform can just exclude the agent and start over the training process again.

The above protocol only ensures that the type reported by each agent is the same as the type he uses in the actual training process with. To encourage all agents to join and contribute all their data, we still need to design mechanisms with desirable properties, to which we devote the rest of the paper.

2.3 Mechanism

Let $t = (t_1, t_2, \dots, t_n)$ and $t_{-i} = (t_1, \dots, t_{i-1}, t_{i+1}, \dots, t_n)$ be the type profile of all agents and all agents without i, respectively. Given the reported types of agents, a mechanism specifies a numerical allocation and payment for each agent, where the allocation is a model in the multi-party learning. Formally, we have:

Definition 1 (Mechanism). *A mechanism $\mathcal{M} = (x, p)$ is a tuple, where*

- $x = (x_1, x_2, \cdots, x_n)$*, where* $x_i \colon \mathbb{R}^n_+ \mapsto \mathbb{R}$ *is the allocation function for agent* i*, which takes the agents' reported types as input and decides the model quality for agent* i *as output;*
- $p = (p_1, p_2, \cdots, p_n)$*, where* $p_i \colon \mathbb{R}^n_+ \mapsto \mathbb{R}$ *is the payment function for agent* i*, which takes the agents' reported types as input and specifies how much agent* i *should pay to the mechanism.*

In a competitive environment, a strategic agent may hide some of data and does not use the model he receives from the platform. Thus the final model quality depends on both the allocation and his actual type. We use valuation function $v_i(x(t'), t)$ to measure the profit of agent i.

Definition 2 (Valuation). *We consider valuation functions* $v_i(x(t'), t)$ *that depend not only on the allocation outcome* $x(t')$ *where* t' *is the reported type profile, but also on the actual type profile* t*.*

We assume the model agent i uses to serve customers is:

$$q_i = \max\{x_i(t'), Q(t_i)\},$$

where $Q(t_i)$ is the model trained with his own data. The valuation of agent i depends on the final model qualities of all agents due to their competition. Hence v_i can also be expressed as $v_i(q_1, \dots, q_n)$.

We make the following assumption on agent i's valuation:

Assumption 3 *Agent* i*'s valuation is monotone increasing with respect to true type* t_i *when the outcome* x *is fixed.*

$$v_i(x, t_i, t_{-i}) \geq v_i(x, \hat{t}_i, t_{-i}), \forall x, \forall t_i \geq \hat{t}_i, \forall t_{-i}, \forall i.$$

This is because possessing more valid data will not lower one's valuation. Otherwise, an agent is always able to discard part of his dataset to make his true type t'_i. Suppose that each agent i's utility $u_i(t, t')$ has the form:

$$u_i(t, t') = v_i(x(t'), t) - p_i(t'),$$

where t and t' are true types and reported types of all agents respectively. As we mentioned above, an agent may lie about his type in order to benefit from the mechanism. The mechanism should incentivize truthful reports to keep agents from lying.

Definition 3 (Incentive Compatibility (IC)). *A mechanism is said to be incentive compatible, or truthful, if reporting truthfully is always the best response for each agent when the other agents report truthfully:*

$$u_i(x(t_i, t_{-i}), t) \geq u_i(x_i(t'_i, t_{-i}), t), \forall t_i \geq t'_i, \forall t_{-i}, \forall i.$$

For ease of presentation, we say agent i reports $\emptyset$ if he chooses not to participate (so we have $x_i(\emptyset, t_{-i}) = 0$ and $p_i(\emptyset, t_{-i}) = 0$). To encourage the agents to participate in the mechanism, the following property should be satisfied:

Definition 4 (Individual Rationality (IR)). *A mechanism is said to be individually rational, if no agent loses by participation when the other agents report truthfully:*

$$u_i(x(t_i, t_{-i}), t) \geq u_i(x(\emptyset, t_{-i}), t), \forall t_i, t_{-i}, \forall i.$$

The revenue and welfare of a mechanism are defined to be all the payments collected from the agents and all the valuations of the agents.

Definition 5. *The revenue and welfare of a mechanism* (x, p) *are:*

$$\textsc{Rev}(x, p) = \textstyle\sum_{i=1}^n p_i(t'), \ \ \textsc{Wel}(x, p) = \textstyle\sum_{i=1}^n v_i(x, t).$$

We say that a mechanism is efficient *if*

$$(x, p) = \arg\max_{(x,p)} \textsc{Wel}(x, p),$$

A mechanism is *weakly budget-balance* if it never loses money.

Definition 6 (Weak Budget Balance). *A mechanism is weakly budget-balance if:*

$$\textsc{Rev}(x, p) \geq 0, \forall t.$$

Definition 7 (Desirable Mechanism). *We say a mechanism is desirable if it is IC, IR, efficient and weakly budget-balance.*

2.4 Comparison with the Standard Interdependent Value Setting

Although each agent's valuation depend on both the outcome of the mechanism and all agent's true types, our interdependent value with type-dependent action spaces setting, however, fundamentally different from standard interdependent value settings:

- In our setting, the type of each agent is the "quality" of his dataset, thus each agent cannot report a higher type than his true type. While in the standard interdependent value setting, an agent can possibly report any type.
- In our setting, the agents do not have the "exit choice" (not participating in the mechanism and getting 0 utility) as they do in the standard setting. This is due to the motivation of this paper: companies from the same industry are trying to improve their service quality, and they are always in the game regardless of their choices. A non-participating company may even have a negative utility if all other companies improved their services.
- To capture the cooperation among the agents, the item being sold, i.e., the learned model, also depends on all agents types. The best model learned by the multi-party learning platform will have high quality if all agents contribute high-quality datasets. However, the objects for allocation are usually fixed in standard mechanisms instead.

3 Quasi-Monotone Externality Setting

In the interdependent value with type-dependent action spaces setting, each agent's utility may also depend on the models that other agents actually use. Such externalities lead to interesting and complicated interactions between the agents. For example, by contributing more data, one may improve the others' model quality, and end up harming his own market share. In this section, we study the setting where agents have *quasi-monotone* externalities.

Definition 8 (Quasi-Monotone Valuation). *Let q_i be the final selected model quality of the agent and q_{-i} be the profile of model qualities of all the agents except i. A valuation function is quasi-monotone if it is in the form:*

$$v_i(q_i, q_{-i}) = F_i(q_i) + \theta_i(q_{-i}),$$

where F_i is monotone and θ_i is an arbitrary function.

Example 1. Let's consider a special quasi-monotone valuation: the linear externality setting, where the valuation for each agent is defined as $v_i = \sum_j \alpha_{ij} q_j$ with q_j being the model that agent j uses. The externality coefficient α_{ij} means the influence of agent j to agent i and captures either the competitive or cooperative relations among agents. If the increase of agent j's model quality imposes a negative (positive) effect on agent i's utility (e.g. major opponents or collaborators in the market), α_{ij} would be negative (positive). Additionally, α_{ii} should always be positive, naturally.

In the linear externality setting, the efficient allocation is straightforward. For each agent i, we give i the training model with best possible quality if $\sum_j \alpha_{ji} \geq 0$. Otherwise, agent i are not allocated any model if $\sum_j \alpha_{ji} < 0$.

We introduce a payment function called *maximal exploitation payment*, and show that the mechanism with efficient allocation and the maximal exploitation payment guarantees IR, IC, efficiency and revenue optimum.

Definition 9 (Maximal Exploitation Payment (MEP)). *For a given allocation function x, if the agent i reports a type t_i' and the other agents report t_{-i}', the maximal exploitation payment is to charge agent i*

$$p_i(t_i', t_{-i}') = v_i(x(t_i', t_{-i}'), t_i', t_{-i}') - v_i(x(\emptyset, t_{-i}'), t_i', t_{-i}').$$

We emphasize that our MEP mechanism and the VCG are quite different. The VCG charges each agent for the harm he causes to others due to his participation while the MEP charges each agent the profit he gets from the mechanism due to his participation. We will show that the MEP is truthful in the quasi-monotone valuation setting in the following theorem, while it is already known that VCG cannot guarantee truthfulness in the interdependent setting [17].

Theorem 1. *Under the quasi-monotone valuation setting, any mechanism with MEP is the mechanism with the maximal revenue among all IR mechanisms, and it is IC.*

Corollary 1. *Any efficient allocation mechanism with MEP under the linear externality setting with all the linear coefficients $\alpha_{ji} \geq 0$ should be IR, IC, weakly budget-balance and efficient.*

In the standard mechanism design setting, the Myerson-Satterthwaite Theorem [19] is a well-known classic result, which says that no mechanism is simultaneously IC, IR, efficient and weakly budget-balance. The above Corollary 1 shows that in our setting, the Myerson-Satterthwaite Theorem fails to hold.

4 General Externality Setting

In this section, we consider the general externality setting where the valuations of agents can have any forms of externalities. The restrictions on the action space and the value functions make the IC and IR mechanisms hard to characterize. It is possible that given a allocation rule, there exist several mechanisms with different payments that satisfy both IC and IR constraints. To understand what makes a mechanism IC and IR, we analyze some properties of truthful mechanisms in this section. For ease of presentation, we assume that the functions $v(\cdot)$, $x(\cdot)$ and $p(\cdot)$ are differentiable.

Theorem 2 (Necessary Condition). *If a mechanism (x, p) is both IR and IC, for all possible valuation functions satisfying Assumption 3, then the payment function satisfies $\forall t_i \geq t_i', \forall t_i, \forall t_{-i}, \forall i$,*

$$p_i(0, t_{-i}) \leq v_i(x(0, t_{-i}), 0, t_{-i}) - v_i(x(\emptyset, t_{-i}), 0, t_{-i}), \tag{1}$$

$$p_i(t_i, t_{-i}) - p_i(t'_i, t_{-i}) \leq \int_{t'_i}^{t_i} \frac{\partial v_i(x(s', t_{-i}), s, t_{-i})}{\partial s'}\bigg|_{s=s'} \mathrm{d}s', \tag{2}$$

where we view $v_i(x(t'_i, t_{-i}), t_i, t_{-i})$ as a function of t_i, t'_i and t_{-i} for simplicity. The partial derivative in Equation (2) is computed using the chain rule, i.e.,

$$\frac{\partial v_i(x(t'_i, t'_{-i}), t_i, t_{-i})}{\partial t'_i} = \sum_{j=1}^{n} \frac{\partial v_i(x(t'_i, t'_{-i}), t_i, t_{-i})}{\partial x_j(t'_i, t'_{-i})} \frac{\partial x_j(t'_i, t'_{-i})}{\partial t'_i}$$

Theorem 2 describes what the payment p is like in all IC and IR mechanisms. In fact, the conditions in Theorem 2 are also crucial in making a mechanism truthful. However, to ensure IC and IR, we still need to restrict the allocation.

Theorem 3 (Sufficient Condition). *A mechanism (x, p) satisfies both IR and IC, for all possible valuation functions satisfying Assumption 3, if for each agent i, for all $t_i \geq t'_i$, and all t_{-i}, Equations (1) and the following two hold*

$$t'_i = \arg\min_{t_i : t_i > t'_i} \frac{\partial v_i(x(t'_i, t_{-i}), t_i, t_{-i})}{\partial t'_i} \tag{3}$$

$$\begin{aligned} &p_i(t_i, t_{-i}) - p_i(t'_i, t_{-i}) \\ &\leq \int_{t'_i}^{t_i} \frac{\partial v_i(x(s', t_{-i}), s, t_{-i})}{\partial s'}\bigg|_{s=s'} \mathrm{d}s' - \int_{t'_i}^{t_i} \frac{\partial v_i(x(\emptyset, t_{-i}), s, t_{-i})}{\partial s} \mathrm{d}s. \end{aligned} \tag{4}$$

5 Market Growth Rate

In this section, we will analyze a factor, the market growth rate, for the existence of the desirable mechanism. Expanding the market size would reduce competition among the agents, meaning that the damage to an agent's existing market caused by joining the mechanism is more likely to be covered by the market growth. Thus our intuition is that if the market grows quickly, a desirable mechanism is more likely to exist.

As mentioned above, each agent's valuation is the profit made from the market, so formally we define the market size to be the sum of the valuations of all the agents. Let $M(q)$ be the agents' total valuations where $q = (q_1, q_2, \ldots, q_n)$ is the set of actual model qualities they use. We have:

$$M(q) = \textstyle\sum_{i=1}^{n} v_i(x, t).$$

In general, the multi-party learning process improves all agents' models. So we do not consider the case where the market shrinks due to the agents' participation, and assume that the market is growing.

Assumption 4 (Growing Market) *$q \succeq q'$ implies $M(q) \geq M(q')$.*

A special case of the growing market is the non-competitive market where agent's values are not affected by others' model qualities, formally:

Definition 10 (Non-competitive Market). *A market is non-competitive iff* $\dfrac{\partial v_i(q)}{\partial q_j} \geq 0, \forall i, j.$

Theorem 4. *In a non-competitive market, there always exists a desirable mechanism, that gives the best possible model to each agent and charges nothing.*

Since the efficient mechanism both redistributes existing markets and enlarges the market size by giving the best learned model when the market is growing, it is difficult to determine whether a desirable mechanism exists if the competition exists. We will give the empirical analysis of the influence of the growth rate of competitive growing markets for desirable mechanisms in Appendix G.

6 Finding a Desirable Mechanism

In the linear externality setting, we provide a mechanism that satisfies all the desirable properties. But this mechanism is not applicable to all valuation functions in the general setting, since the existence of a desirable mechanism depends on the agents' actual valuation functions. We provide an algorithm, that given the agents' valuations, computes whether such a mechanism exists, and outputs the one that optimizes revenue, if any.

Since each agent can only under-report, according to the IR property, we must have:

$$u_i(x(t_i, t_{-i}), t) \geq u_i(x(\emptyset, t_{-i}), t), \forall t, \forall i.$$

Equivalently, we get $\forall t, \forall i$,

$$\begin{aligned} u_i(x(\emptyset, t_{-i}), t) &\leq v_i(x(t_i, t_{-i}), t) - p_i(t_i, t_{-i}), \\ p_i(t_i, t_{-i}) &\leq v_i(x(t_i, t_{-i}), t) - u_i(x(\emptyset, t_{-i}), t), \\ p_i(t) &\leq v_i(x(t_i, t_{-i}), t) - u_i(x(\emptyset, t_{-i}), t). \end{aligned}$$

For simplicity, we define the upper bound of $p(t')$ as

$$\overline{p(t)} \triangleq \{v_i(x(t_i, t_{-i}), t) - u_i(x(\emptyset, t_{-i}), t).$$

The IC property requires that $\forall t_i \geq t_i', \forall t_{-i}, \forall i$,

$$u_i(x(t_i, t_{-i}), t) \geq u_i(x(t_i', t_{-i}), t).$$

A little rearrangement gives:

$$p_i(t_i, t_{-i}) - p_i(t_i', t_{-i}) \leq v_i(x(t_i, t_{-i}), t) - v_i(x(t_i', t_{-i}), t) \triangleq Gap_i(t_i', t_i, t_{-i}).$$

Note that the inequality correlations between the payments form a system of difference constraints. The form of update of the payments is almost identical to that of the shortest path problem. Therefore, we make use of this observation to design the algorithm.

We assume that all the value functions are common knowledge, the efficient allocation is then determined because the mechanism always chooses the one that maximizes the social welfare. Thus it suffices to figure out whether there is a payment rule $p(t')$ which makes the mechanism IR, IC and weakly budget-balance. Since the valid data size for each agent is bounded in practice, we assume the mechanism only decides the payment functions on the data range $[0, D]$, and discretize the type space into intervals of length ϵ, which is also the minimal size of the data. Thus each agent's type is a multiple of ϵ. Note that since the utility function is general, all the points in the action space would influence the properties and existence of the mechanism, thus it is necessary to enumerate all the points in the space. The exponential value function space, i.e., the exponential input space, determines that the complexity of our algorithm is exponential in D.

We give the following algorithmic characterization for the existence of a desirable mechanism.

Algorithm 1: Finding desirable mechanisms

```
input: Agents' valuation functions v.
Use the function v_i to calculate all the Gap_i(t'_i, t_i, t_{-i}) and \overline{p_i(t_i, t_{-i})} for each i;
Initialize all p_i^{max}(t_i, t_{-i}) to be \overline{p_i(t_i, t_{-i})} for each i;
for i = 1 to n do
    for t_{-i} = (∅, ∅, ··· , ∅) to (D, D, ··· , D) (increment = ε on each
      dimension) do
        Build an empty graph;
        For each p_i(t_i, t_{-i}), construct a vertex V_{t_i t_{-i}} and insert it into the
          graph;
        Construct a base vertex VB_{t_{-i}} which denotes the payment zero into
          the graph;
        for t_i = 0 to D (increment = ε) do
            Add an edge from VB_{t_{-i}} to V_{t_i t_{-i}} with weight \overline{p(t_i, t_{-i})};
            for t'_i = 0 to t_i (increment = ε) do
                Add an edge with weight Gap_i(t'_i, t_i, t_{-i}) from V_{t'_i t_{-i}} to V_{t_i t_{-i}};
        Use the Single-Source Shortest-Path algorithm to find the shortest
          path from VB_{t_{-i}} to all the other vertices. These are the maximum
          solutions p_i^{max}(t_i, t_{-i}) for each payment case;
        if \sum_{j=1}^{n} p_j^{max}(t) < 0 then
            return There is no desirable mechanism.
return p_i^{max} as the payment functions.
```

The following theorem proves the correctness of Algorithm 1.

Theorem 5. *Taking agents' valuation functions as input, Algorithm 1 outputs the answer of the decision problem of whether there exists a mechanism that guar-*

antees IR, IC, efficiency and weak budget balance simultaneously, and specifies the payments that achieve maximal revenue if the answer is yes.

7 Conclusion

In this paper, we study the mechanism design problem for multi-party machine learning. We restrict the action space of each agent where he can only misreport a lower type than his actual type and consider the valuation function that is about the allocation outcome and the true types of all agents. The VCG mechanism does not guarantee IR and DSIC and the Myerson-Satterthwaite Theorem in the standard mechanism design setting does not hold in our setting, implying the desirable mechanisms that are both IR, DSIC, efficient and weakly budget-balance exist in our setting. We propose a maximal exploitation payment mechanism and show that this mechanism is truthful and revenue-optimal in the quasi-monotone externalities setting. Then we give sufficient and necessary conditions for designing a truthful mechanism for the general setting. These conditions restrict both the allocation function and the payment function. We show that the data size disparity between agents and the market growth rate highly affect the existence of the desirable mechanism. If the market grows fast and the disparity is small, a desirable mechanism is more likely to exist. Finally, we devise an algorithm to find desirable mechanisms that are truthful, individually rational, efficient and weakly budget-balance simultaneously.

References

1. Abadi, M., Chu, A., Goodfellow, I., McMahan, H.B., Mironov, I., Talwar, K., Zhang, L.: Deep learning with differential privacy. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. pp. 308–318. ACM (2016)
2. Ausubel, L.M.: An efficient ascending-bid auction for multiple objects. American Economic Review **94**(5), 1452–1475 (2004)
3. Blumrosen, L., Feldman, M.: Implementation with a bounded action space. In: Proceedings of the 7th ACM conference on Electronic commerce. pp. 62–71. ACM (2006)
4. Blumrosen, L., Nisan, N.: Auctions with severely bounded communication. In: The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings. pp. 406–415. IEEE (2002)
5. Chawla, S., Fu, H., Karlin, A.R.: Approximate revenue maximization in interdependent value settings. In: Babaioff, M., Conitzer, V., Easley, D.A. (eds.) ACM Conference on Economics and Computation, EC '14, Stanford , CA, USA, June 8-12, 2014. pp. 277–294. ACM (2014)
6. Chwe, M.S.Y.: The discrete bid first auction. Economics Letters **31**(4), 303–306 (1989)
7. David, E., Rogers, A., Jennings, N.R., Schiff, J., Kraus, S., Rothkopf, M.H.: Optimal design of english auctions with discrete bid levels. ACM Transactions on Internet Technology (TOIT) **7**(2), 12 (2007)
8. Ding, N., Fang, Z., Huang, J.: Optimal contract design for efficient federated learning with multi-dimensional private information. IEEE Journal on Selected Areas in Communications **39**(1), 186–200 (2020)
9. Fang, W., Tang, P., Zuo, S.: Digital good exchange. In: Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems. pp. 1277–1278. International Foundation for Autonomous Agents and Multiagent Systems (2016)
10. Goldberg, A.V., Hartline, J.D.: Envy-free auctions for digital goods. In: Proceedings of the 4th ACM conference on Electronic commerce. pp. 29–35. ACM (2003)
11. Goldberg, A.V., Hartline, J.D., Wright, A.: Competitive auctions and digital goods. In: Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms. pp. 735–744. Society for Industrial and Applied Mathematics (2001)
12. Hu, Y., Niu, D., Yang, J., Zhou, S.: Fdml: A collaborative machine learning framework for distributed features. In: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. pp. 2232–2240. ACM (2019)
13. Jia, R., Dao, D., Wang, B., Hubis, F.A., Hynes, N., Gürel, N.M., Li, B., Zhang, C., Song, D., Spanos, C.J.: Towards efficient data valuation based on the shapley value. In: The 22nd International Conference on Artificial Intelligence and Statistics. pp. 1167–1176. PMLR (2019)
14. Kang, J., Xiong, Z., Niyato, D., Yu, H., Liang, Y.C., Kim, D.I.: Incentive design for efficient federated learning in mobile networks: A contract theory approach. arXiv preprint arXiv:1905.07479 (2019)
15. Lim, W.Y.B., Luong, N.C., Hoang, D.T., Jiao, Y., Liang, Y.C., Yang, Q., Niyato, D., Miao, C.: Federated learning in mobile edge networks: A comprehensive survey. IEEE Communications Surveys & Tutorials **22**(3), 2031–2063 (2020)

16. McMahan, B., Moore, E., Ramage, D., Hampson, S., y Arcas, B.A.: Communication-efficient learning of deep networks from decentralized data. In: Artificial Intelligence and Statistics. pp. 1273–1282 (2017)
17. Mezzetti, C.: Mechanism design with interdependent valuations: Efficiency. Econometrica **72**(5), 1617–1626 (2004)
18. Milgrom, P.R., Weber, R.J.: A theory of auctions and competitive bidding. Econometrica **50**(5), 1089–1122 (1982)
19. Myerson, R.B., Satterthwaite, M.A.: Efficient mechanisms for bilateral trading. Journal of economic theory **29**(2), 265–281 (1983)
20. Redko, I., Laclau, C.: On fair cost sharing games in machine learning. In: Thirty-Third AAAI Conference on Artificial Intelligence (2019)
21. Roughgarden, T., Talgam-Cohen, I.: Optimal and robust mechanism design with interdependent values. ACM Trans. Economics and Comput. **4**(3), 18:1–18:34 (2016)
22. Shokri, R., Shmatikov, V.: Privacy-preserving deep learning. In: Proceedings of the 22nd ACM SIGSAC conference on computer and communications security. pp. 1310–1321. ACM (2015)
23. Smith, V., Chiang, C.K., Sanjabi, M., Talwalkar, A.S.: Federated multi-task learning. In: Advances in Neural Information Processing Systems. pp. 4424–4434 (2017)
24. Song, T., Tong, Y., Wei, S.: Profit allocation for federated learning. In: 2019 IEEE International Conference on Big Data (Big Data). pp. 2577–2586. IEEE (2019)
25. Takabi, H., Hesamifard, E., Ghasemi, M.: Privacy preserving multi-party machine learning with homomorphic encryption. In: 29th Annual Conference on Neural Information Processing Systems (NIPS) (2016)
26. Yonetani, R., Naresh Boddeti, V., Kitani, K.M., Sato, Y.: Privacy-preserving visual learning using doubly permuted homomorphic encryption. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2040–2050 (2017)
27. Zhan, Y., Zhang, J., Hong, Z., Wu, L., Li, P., Guo, S.: A survey of incentive mechanism design for federated learning. IEEE Transactions on Emerging Topics in Computing (2021)

Appendix

A Proof of Theorem 1

Proof. Intuitively, the MEP rule charges agent i the profit he gets from an model that the mechanism allocates to him. If the mechanism charges higher than the MEP, an agent would have negative utility after taking part in. The IR constraint would then be violated. So it's easy to see that the MEP is the maximal payment among all IR mechanisms.

Then we prove that this payment rule also guarantees the IC condition. It suffices to show that if an agent hides some data, no matter which model he chooses to use, he would never get more utility than that of truthful reporting. We suppose that agent i's type is t_i' and he untruthfully reports t_i'.

Suppose that the agent i truthfully reports the type $t_i' = t_i$, since the payment function is defined to charge this agent until he reaches the valuation when he does not take part in the mechanism, the utility of this honest agent would be

$$u_i^0(t') = F_i(Q(t_i)) + \theta_i(q_{-i}(\emptyset, t'_{-i})).$$

If the agent does not report truthfully, we suppose that the agent reports t_i' where $t_i' \leq t_i$. According to the MEP, the payment function for agent i would be

$$p_i(t_i', t'_{-i}) = F_i(q_i(t_i', t'_{-i})) + \theta_i(q_{-i}(t_i', t'_{-i})) - F_i(Q(t_i')) - \theta_i(q_{-i}(\emptyset, t'_{-i})).$$

It can be seen that the mechanism would never give an agent a worse model than the model trained by its reported data, otherwise the agents would surely select their private data to train models. Hence it is without loss of generality to assume that the allocation $x_i(t_i', t'_{-i}) \geq Q(t_i')$, $\forall t_i', t'_{-i}, \forall i$. Thus we have $q_{-i}(t_i', t'_{-i}) = x_{-i}(t_i', t'_{-i})$. We discuss the utility of agent i by two cases of choosing models.

Case 1: the agent chooses the allocation x_i. Since agent i selects the allocated model, we have $q_i = x_i(t_i', t'_{-i})$. Then the utility of agent i would be

$$\begin{aligned} u_i^1 =& v_i(t_i', t'_{-i}) - p_i(t_i', t'_{-i}) \\ =& F_i(x_i(t_i', t'_{-i})) + \theta_i(x_{-i}(t_i', t'_{-i})) + F_i(Q(t_i')) \\ &+ \theta_i(x_{-i}(\emptyset, t'_{-i})) - F_i(x_i(t_i', t'_{-i})) - \theta_i(x_{-i}(t_i', t'_{-i})) \\ =& F_i(Q(t_i')) + \theta_i(x_{-i}(\emptyset, t'_{-i})). \end{aligned}$$

Because both F_i and Q are monotone increasing functions and $t_i \geq t_i'$, we have

$$u_i^1 \leq F_i(Q(t_i)) + \theta_i(x_{-i}(\emptyset, t'_{-i})) = u_i^0.$$

Case2: the agent chooses $Q(t_i)$. Since agent i selects the model trained by his private data, we have $q_i = Q(t_i)$. The final utility of agent i would be

$$\begin{aligned} u_i^2 =& v_i(t_i', t'_{-i}) - p_i(t_i', t'_{-i}) \\ =& F_i(Q(t_i)) + \theta_i(x_{-i}(t_i', t'_{-i})) + F_i(Q(t_i')) \\ &+ \theta_i(x_{-i}(\emptyset, t'_{-i})) - F_i(x_i(t_i', t'_{-i})) - \theta_i(x_{-i}(t_i', t'_{-i})) \\ =& F_i(Q(t_i)) + F_i(Q(t_i')) + \theta_i(x_{-i}(\emptyset, t'_{-i})) - F_i(x_i(t_i', t'_{-i})). \end{aligned}$$

Subtract the original utility from the both sides, then we have

$$\begin{aligned}u_i^2 - u_i^0 =&F_i(Q(t_i)) + F_i(Q(t_i')) + \theta_i(x_{-i}(\emptyset, t_{-i}'))\\ &- F_i(x_i(t_i', t_{-i}')) - F_i(Q(t_i)) - \theta_i(x_{-i}(\emptyset, t_{-i}'))\\ =&F_i(Q(t_i')) - F_i(x_i(t_i', t_{-i}')).\end{aligned}$$

Because $x_i(t_i', t_{-i}') \geq Q(t_i')$, $\forall t_i', t_{-i}', \forall i$ and because F_i is a monotonically increasing function, we can get $u_i^2 - u_i^0 \leq 0$. Therefore $\max\{u_i^1, u_i^2\} \leq u_i^0$, lying would not bring more benefits to any agent, and the mechanism is IC.

B Proof of Corollary 1

Proof. In Theorem 1 we know that the MEP mechanism is IR and IC. Since the linear coefficients are all positive and the externality setting is linear, any efficient mechanism would allocate the best model to all the agents. Since each agent gets a model with no less quality than his reported one and the payment is equal to the value difference between the case an agent truthfully report and the case he exit the mechanism. The agent's value is always larger than the value when he exits the mechanism. Then the payment is always positive and the mechanism should satisfy all of the four properties.

C Proof of Theorem 2

Proof. We first prove that Equation (1) holds. Observe that

$$\begin{aligned}&u_i(x(t_i, t_{-i}'), t_i, t_{-i}) - u_i(x(t_i', t_{-i}'), t_i', t_{-i})\\ =&[v_i(x(t_i, t_{-i}'), t_i, t_{-i}) - p_i(t_i, t_{-i}')] - [v_i(x(t_i', t_{-i}'), t_i', t_{-i}) - p_i(t_i', t_{-i}')]\\ \geq&[v_i(x(t_i, t_{-i}'), t_i, t_{-i}) - p_i(t_i, t_{-i}')] - [v_i(x(t_i', t_{-i}'), t_i, t_{-i}) - p_i(t_i', t_{-i}')]\\ =&u_i(x(t_i, t_{-i}'), t_i, t_{-i}) - u_i(x(t_i', t_{-i}'), t_i, t_{-i})\\ \geq&0, \end{aligned} \tag{5}$$

where the first inequality is because of Assumption 3, and the last inequality is because of the DSIC property.

Let $t_i' = 0$ in Equation (5). We have

$$u_i(x(t_i, t_{-i}'), t_i, t_{-i}) \geq u_i(x(0, t_{-i}'), 0, t_{-i}).$$

The IR property further requires that $u_i(x(0, t_{-i}'), 0, t_{-i}) \geq u_i(x(\emptyset, t_{-i}'), 0, t_{-i})$, which Equation (1) follows.

To show Equation (2) must hold, we rewrite Equation (5):

$$\begin{aligned}p_i(t_i, t_{-i}') - p_i(t_i', t_{-i}') &\leq v_i(x(t_i, t_{-i}'), t_i, t_{-i}) - v_i(x(t_i', t_{-i}'), t_i', t_{-i})\\ &= \int_{t_i'}^{t_i} \frac{\mathrm{d}v_i(x(s', t_{-i}'), s(s'), t_{-i})}{\mathrm{d}s'} \,\mathrm{d}s'.\end{aligned} \tag{6}$$

Fixing t_{-i} and t'_{-i}, the total derivative of $v_i(x(s', t'_{-i}), s, t_{-i})$ is:

$$\begin{aligned}&\mathrm{d}v_i(x(s', t'_{-i}), s, t_{-i})\\=&\frac{\partial v_i(x(s', t'_{-i}), s, t_{-i})}{\partial s'}\,\mathrm{d}s' + \frac{\partial v_i(x(s', t'_{-i}), s, t_{-i})}{\partial s}\,\mathrm{d}s.\end{aligned}$$

View s as a function of s' and let $s(s') = s'$:

$$\begin{aligned}&\frac{\mathrm{d}v_i(x(s', t'_{-i}), s(s'), t_{-i})}{\mathrm{d}s'}\\=&\left.\frac{\partial v_i(x(s', t'_{-i}), s, t_{-i})}{\partial s'}\right|_{s=s'} + \frac{\partial v_i(x(s', t'_{-i}), s(s'), t_{-i})}{\partial s(s')}\frac{\mathrm{d}s(s')}{\mathrm{d}s'}.\end{aligned}$$

Plug into Equation (6), and we obtain:

$$\begin{aligned}&p_i(t_i, t'_{-i}) - p_i(t'_i, t'_{-i})\\\le&\int_{t'_i}^{t_i}\left.\frac{\partial v_i(x(s', t'_{-i}), s, t_{-i})}{\partial s'}\right|_{s=s'} + \int_{t'_i}^{t_i}\frac{\partial v_i(x(s', t'_{-i}), s(s'), t_{-i})}{\partial s(s')}\,\mathrm{d}s'.\end{aligned}$$

Since the above inequality holds for any valuation function with $v_i(x, t_i, t_{-i}) \ge v_i(x, t'_i, t_{-i}), \forall x, \forall t_{-i}, \forall t_i \ge t'_i$, we have:

$$p_i(t_i, t'_{-i}) - p_i(t'_i, t'_{-i}) \le \int_{t'_i}^{t_i}\left.\frac{\partial v_i(x(s', t'_{-i}), s, t_{-i})}{\partial s'}\right|_{s=s'}\,\mathrm{d}s'.$$

D Proof of Theorem 3

Proof. Equation (3) indicates that the function $\frac{\partial v_i(x(t'_i, t'_{-i}), t_i, t_{-i})}{\partial t'_i}$ is minimized at t'_i:

$$\left.\frac{\partial v_i(x(t'_i, t'_{-i}), s, t_{-i})}{\partial t'_i}\right|_{s=t'_i} \le \frac{\partial v_i(x(t'_i, t'_{-i}), t_i, t_{-i})}{\partial t'_i}. \tag{7}$$

Therefore, we have

$$\begin{aligned}&u_i(x(t_i, t'_{-i}), t_i, t_{-i}) - u_i(x(t'_i, t'_{-i}), t_i, t_{-i})\\=&\int_{t'_i}^{t_i}\frac{\partial v_i(x(s', t'_{-i}), t_i, t_{-i})}{\partial s'}\,\mathrm{d}s' - p_i(t_i, t'_{-i}) + p_i(t'_i, t'_{-i})\\\ge&\int_{t'_i}^{t_i}\left.\frac{\partial v_i(x(s', t'_{-i}), s, t_{-i})}{\partial s'}\right|_{s=s'}\,\mathrm{d}s' - p_i(t_i, t'_{-i}) + p_i(t'_i, t'_{-i})\\\ge&\int_{t'_i}^{t_i}\frac{\partial v_i(x(\emptyset, t'_{-i}), s, t_{-i})}{\partial s}\mathrm{d}s,\end{aligned} \tag{8}$$

where the two inequalities are due to Equation (7) and (4), respectively. Since $v_i(x, t_i, t_{-i}) \ge v_i(x, t'_i, t_{-i})$, $\forall x, \forall t_{-i}, \forall t_i \ge t'_i$ indicates $\frac{\partial v_i(x(\emptyset, t'_{-i}), s, t_{-i})}{\partial s} \ge 0$, the above inequality shows that the mechanism guarantees the DSIC property.

To prove that the mechanism is IR, we first observe that

$$\begin{aligned}
&[u_i(x(t_i,t'_{-i}),t_i,t_{-i}) - v_i(x(\emptyset,t'_{-i}),t_i,t_{-i})] - [u_i(x(t'_i,t'_{-i}),t'_i,t_{-i}) \\
&\quad - v_i(\emptyset,x(t_{-i}),t'_i,t_{-i})] \\
=&u_i(x(t_i,t'_{-i}),t_i,t_{-i}) - u_i(x(t'_i,t'_{-i}),t'_i,t_{-i}) - \int_{t'_i}^{t_i} \frac{\partial v_i(x(\emptyset,t'_{-i}),s,t_{-i})}{\partial s}\mathrm{d}s \\
\geq&u_i(x(t_i,t'_{-i}),t_i,t_{-i}) - u_i(x(t'_i,t'_{-i}),t_i,t_{-i}) - \int_{t'_i}^{t_i} \frac{\partial v_i(x(\emptyset,t'_{-i}),s,t_{-i})}{\partial s}\mathrm{d}s \\
\geq&0,
\end{aligned}$$

where the two inequalities are Assumption 3 and Equation (8). Letting $t'_i = 0$ using Equation (2), we get:

$$\begin{aligned}
&u_i(x(t_i,t'_{-i}),t_i,t_{-i}) - v_i(x(\emptyset,t'_{-i}),t_i,t_{-i}) \\
\geq&u_i(x(0,t'_{-i}),0,t_{-i}) - v_i(x(\emptyset,t'_{-i}),0,t_{-i}) \\
=&v_i(x(0,t'_{-i}),0,t_{-i}) - p_i(0,t'_{-i}) - v_i(x(\emptyset,t'_{-i}),0,t_{-i}) \\
\geq&0.
\end{aligned}$$

E Proof of Theorem 4

Proof. Suppose that the platform uses the mechanism mentioned in the theorem. Then for each agent, contributing with more data increases all participants' model qualities. By definition, in a non-competitive market, improving others' models does not decrease one's profit. Therefore, the optimal strategy for each participant is to contribute with all his valid data, making the mechanism truthful. Also because of the definition, entering the platform always weakly increases one's model quality. Thus the mechanism is IR. With the IC and IR properties, it is easy to see that the mechanism is also efficient and weakly budget-balance.

F Proof of Theorem 5

Proof. Suppose that there is a larger payment for agent i such that $p_i(t') > p_i^{\max}(t')$ where t' is the profile of reported types. In the process of our algorithm, the $p_i^{\max}(t')$ is the minimal path length from VB_{-i} to $V_{t_i}t_{-i}$, denoted by $(VB_{-i}, V_{t_{i1}}t_{-i}, V_{t_{i2}}t_{-i}, \cdots, V_{t_{ik}=t'_i}t_{-i})$. By the definition of edge weight, we have the following inequalities:

$$\begin{gathered}
p_i(t_{i1},t_{-i}) \leq \overline{p_i(t_{i1},t_{-i})}, \\
p_i(t_{i2},t_{-i}) - p_i(t_{i1},t_{-i}) \leq Gap_i(t_{i1},t_{i2},t_{-i}), \\
\vdots \\
p_i(t_{ik},t_{-i}) - p_i(t_{i(k-1)},t_{-i}) \leq Gap_i(t_{i1},t_{i2},t_{-i}).
\end{gathered}$$

Adding these inequalities together, we get

$$p_i(t') \le \overline{p_i(t_{i1}, t_{-i})} + \sum_{j=1}^{k-1} Gap_i(t_{ij}, t_{i(j+1)}, t_{-i}) = p_i^{\max}(t').$$

If $p_i(t') < p_i^{\max}(t')$ holds, this would violate at least 1 of the k inequalities above. If the first inequality is violated, the mechanism would not be IR, by the definition of $\overline{p_i(t_{i1}, t_{-i})}$. If any other inequality is violated, the mechanism would not be IC, by the definition of $Gap_i(t_{ij}, t_{i(j+1)}, t_{-i})$.

On the other hand, if we select $p_i^{\max}(t')$ to be payment of agent i, all the inequalities should be satisfied, otherwise the shortest path would be updated to a smaller length.

Therefore the $p_i^{\max}(t')$ must be the maximum payment for agent i. If the maximal payment sum up to less than 0, there would obviously be no mechanism that is IR, IC and weakly budget-balance under the efficient allocation function.

G Experiments

We design experiments to demonstrate the performance of our mechanism for practical use. We first show the mechanism with the maximal exploitation payments can guarantee a good quality of trained model and high revenues under the linear externality cases. Then we conduct simulations to exhibit the relation of the market growth of competitive markets to the existence of desirable mechanisms.

G.1 The MEP Mechanism

We consider the valuation with linear externalities setting where α_{ij}'s (defined in Example 1) are generated uniformly in $[-1, 1]$. Each agent's type is drawn uniformly from $[0, 1]$ independently and the $Q(t)$ is $\frac{1-e^{-t}}{1+e^{-t}}$. The performance of a mechanism is measured by the platform's revenue and its best quality of trained model under the mechanism. All the values of each instance are averaged over 50 samples. We both show the performance changes as the number of agents increases and as the agents' type changes.

When the number of agents becomes larger, the platform can obtain more revenues and train better models (see Figure 1). Particularly, the model quality is close to be optimal when the number of agents over 12. An interesting phenomenon is that the revenue may surpass the social welfare. This is because the average external effect of other agents on one agent i tends to be negative when agent i does not join in the mechanism, thus the second term in the MEP payment is averagely negative and revenue is larger than the welfare.

To see the influence of type on performance, we fix one agent's type to be 1 and set the other agent's type from 0 to 10. It can be seen in Figure 2 that the welfare and opponent agent's utility (`uti_2`) increase as the opponent's type increases but the platform's revenue and the utility of the static agent (`uti_1`) are almost not affected by the type. So we draw the conclusion that the most efficient way for the platform to earn more revenue is to attract more small companies to join the mechanism, since in the Figure 1 the revenue obviously increases as the number of agents increases.

G.2 Existence of Desirable Mechanisms

We assume all the agents' types lie in $[0, D]$, and the type space can be discretized into intervals of length ϵ, which can be viewed as the minimal size of a dataset. Thus each agent's type is a multiple of ϵ. The data disparity is defined as the ratio of the largest possible data size to the smallest possible data size, namely, D/ϵ. We measure the condition for existence of desirable mechanisms by the maximal data disparity when the market growth rate is given.

To describe the market growth, we use the following form of valuation function and model quality function:

$$Q(t) = t \text{ and } v_i(q) = \left(\sum_{j=1}^{n} Q(q_j)\right)^{\gamma} \cdot Q(q_i), \forall i,$$

where γ indicates the market growth rate. We consider the competitive growing market case where $-1 \leq \gamma < 0$.[6]

The algorithm we use to find desirable mechanisms under different valuation functions is described in Section 6. We enumerate the value of γ from -1 to -0.668 with step length 0.002 and run the algorithm to figure out the boundary of D/ϵ under different γ in a market with 2 agents. The range of γ is determined by our computing capability, and the disparity boundary has been over 10000 when γ is near -0.66.

Figure 3 shows the boundary of data disparity for existence of desirable mechanisms under different market growth rates. For every fixed γ, there does not exist any desirable mechanism when the data disparity is larger than the point on the red line. It can be seen an obvious trend that when γ becomes larger, the constraint on data size disparity would become looser. A desirable mechanism is more likely to exist in a market that grows faster. When the market is not growing, there would not be such a desirable mechanism at all. On the other hand, if the market grows so fast such that there does not exist any competition between the agents, the desirable mechanism always exists.

[6] When $\gamma < -1$, the market is not a growing market; when $\gamma \geq 0$, the market becomes non-competitive, therefore by Theorem 4, a desirable mechanism trivially exists.

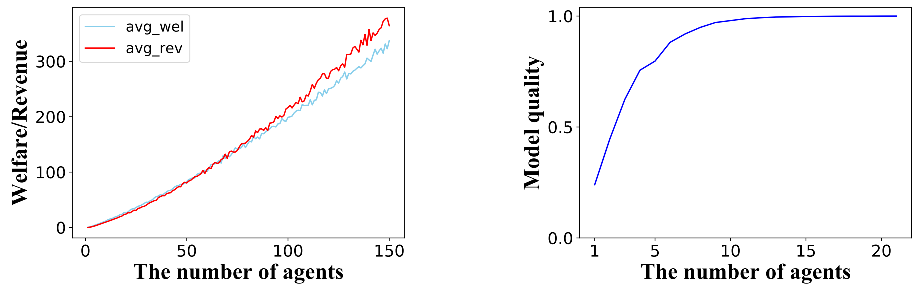

(a) Welfare & revenue

(b) Best quality of trained model

Fig. 1. Performance of MEP under different numbers of agents

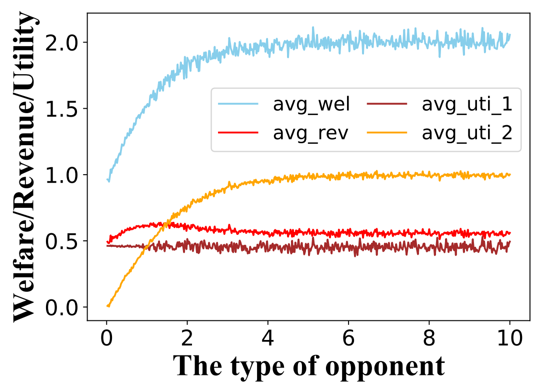

Fig. 2. Performance of MEP under different types

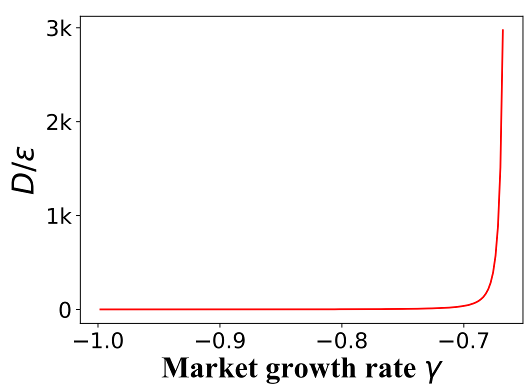

Fig. 3. Data disparity vs. Market growth