The Force of Compensation, A Multi-stage Incentive Mechanism Model for Federated Learning*

Han XU[1][0000−0002−3580−9403], Priyadarsi Nanda[1][0000−0002−5748−155X], Jie Liang[1][0000−0001−7179−5208], and Xiangjian He[2][0000−0001−8962−540X]

[1] University of Technology Sydney, 15 Broadway, Ultimo, NSW, Australia, 2007
han.xu@student.uts.edu.au {priyadarsi.nanda, jie.liang}@uts.edu.au
[2] University of Nottingham Ningbo China, 199 Taikang East Road, Yinzhou District Ningbo, Zhejiang Province, China, 315104 Sean.He@nottingham.edu.cn

Abstract. In federated learning, data owners 'provide' their local data to model owners to train a mature model in a privacy-preserving way. A critical factor in the success of a federated learning scheme is an optimal incentive mechanism that motivates all participants to fully contribute. However, the privacy protection inherent to federated learning creates a dual ethical risk problem in that there is information asymmetry between the two parties, so neither side's effort is observable. Additionally, there is often an implicit cost associated with the effort contributed to training a model, which may lead to self-interested, opportunistic behaviour on both sides. Existing incentive mechanisms have not addressed this issue. Hence, in this paper, we analyse how dual ethical risk affects the performance of federated learning schemes. We also derive an optimal multi-stage contract-theoretic incentive mechanism that minimises this risk, and experiment with calculating an optimal incentive contract for all participants. To our best knowledge, this is the first time that dual ethical risk for federated learning participants has been discussed. It is also the first time that an optimal incentive mechanism to overcome this issue has been developed.

Keywords: federated learning · ethical risk · incentive mechanism

1 Introduction

In this era of AI, more and more complex applications based on machine learning are being introduced into our daily lives. It is now possible to train a highly accurate machine learning model by feeding it vast amounts of real-world data. However, we are also in an era with an emphasis on privacy protection, and various privacy protection regulations around the world, such as the GDPR in the EU [3], restrict data sharing. This creates a significant problem for machine learning where training a well-performing model invariably means accessing private data – and lots of it. In this context, federated learning, an inherently

* This research is supported by Australian Research Council's Linkage grant

private learning scheme introduced by Google in 2016 [6–8], has received much attention. With federated learning, participants train a model collaboratively without ever needing to expose their sensitive raw data. An initialised global model is distributed to the data owners (clients) via a federated learning server, and each client trains the model locally using its own private data. Only the updated parameters of the model are then uploaded to the server for aggregation. After the uploaded parameters have been integrated, the server sends the updated model back to the clients for further training. This process is repeated until the accuracy of the model reaches its target.

In the years since 2016, the concept of federated learning has been expanded to include horizontal federated learning, vertical federated learning, and federated migrated learning [13], while the participants fall into two groups: the model owners and the data owners. The architecture of a simple federated learning scheme is shown in Fig 1. The data owners consume their resources to collect, clean and process large quantities of qualified training data. They also provide the computational and communication resources required for local training. The model owners consume resources throughout the training process for parameter integration, model tuning, optimisation, and more. Thus, an incentive mechanism is needed to compensate both parties for the resources consumed and to motivate them to collaborate. To maximise performance, both parties need to contribute their resources to the fullest degree. However, the privacy-preserving mechanisms within the federated learning paradigm creates information asymmetry between the participants causing a double ethical risk problem where neither side's effort is observable. In addition, there is often a cost associated with the effort, which may lead to self-interested, opportunistic behaviour on both sides given the disparity of interests and the information asymmetry. Consider a practical example: a medical association with several hospital members wants to work with a company that specialises in image recognition to build an automated CT image recognition model that can label suspected lung cancer nodules in CT images. The medical association cannot observe how much effort the model provider puts into the training, and nor can the model provider observe whether the healthcare association is putting enough effort into collecting and processing high-quality/quantity training data. Both parties can only directly observe the training results at particular stages. As such, there is a double ethical risk in this kind of federated learning case.

Much work has been done on different aspects of incentive mechanisms for federated learning – work that can be found in some of the recently published state-of-art surveys [16, 17]. Currently, most reward-based incentive mechanisms focus on model owner-led reward schemes. These are typically designed to maximise federated learning outcomes for model owners, while minimising the incentives offered to the data owners. However, to the best of our knowledge, the issue of dual ethical risk in federated learning has not been addressed. Hence, in this paper, we propose an incentive mechanism that differs from the status quo. In our mechanism, the data owners are Stackelberg game leaders, which address the above dual ethical risk. Our focus is on the ethical risk problems

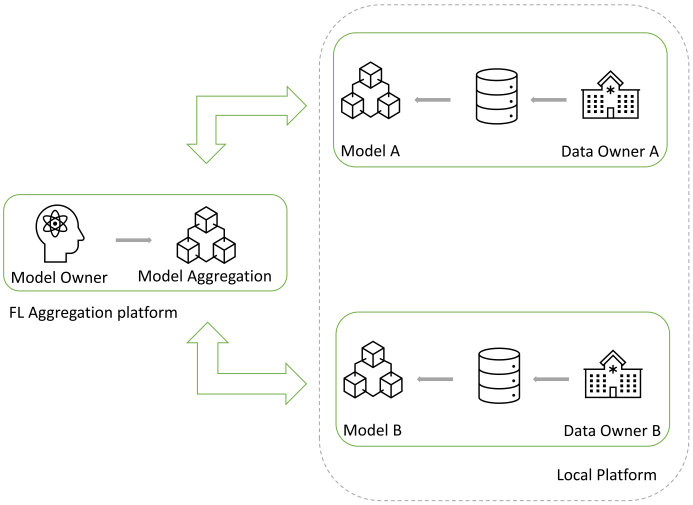

Fig. 1. Client-server architecture FL Model

with federated learning, i.e., how to gauge the implicit efforts of both groups of participants and how this problem might be countered using multi-stage game theory. This is our focus because the implicit efforts targeted by our incentive mechanism are highly significant to the success of federated learning schemes.

To this end, our research examines the game between data owners and model owners within a federated learning process, where the efforts of neither party are directly observable. The solution involves a multi-stage incentive mechanism designed for two parties, where the incentive contract is defined before the start of training.

Our contribution to the literature is insight into an optimal multi-stage incentive contract and an endogenous optimal payoff point description. More specifically, this article shows that the optimal scheme for the data owner who leads the incentive contract should, to the extent possible, return all later stages incentive payments to the model owner.

The remainder of the paper is structured as follows. Section 2 reviews the existing incentive mechanisms for federated learning. Section 3 presents the incentive mechanism model used in our research and the results, and Section 4 provides a simulation example to validate the model. Finally, conclusions and future work are drawn in Section 5.

2 Related works

This section positions our research within the existing literature by reviewing relevant studies on incentive mechanisms for federated learning.

When implementing federated learning, participants are often reluctant to participate in training a model unless they receive some benefit for doing so. This is because contributing to a model's training can be a highly resource-intensive undertaking. In addition, an information asymmetry exists between the data owners and the model owners. Thus, a well-designed incentive mechanism can be crucial to the success of federated learning. Such a mechanism is needed to encourage collaboration between all participants and reduces the potentially damaging effects of information asymmetry. For the best possible outcome, the incentive mechanism needs to determine the optimal level of participation and rewards for all parties to keep everyone involved motivated and engaged. Optimisation problems, such as utility maximisation, are all about deriving the best strategy.

Incentive mechanisms typically consist of two phases: contribution assessment and reward allocation. [12] The main contribution assessment strategies are:

- Self-declared contribution assessment. A self-declared contribution assessment is a direct way for data owners to report their contributions to the model owner. Data size and computational resource capacity are among the many metrics used to evaluate the self-declared contribution of a data owner [4].
- Shapley value contribution assessment. Shapley value [9] is a method of utility assessment based on marginal contributions. One of the advantages of this strategy is that it eliminates the effect of the order in which the participants joined the 'collective' in order to calculate a fairer estimate of their marginal contribution. Therefore, payoffs are calculated purely from the contributions provided regardless of sequencing for a fairer distribution of rewards. It is most common in cooperative games. Many recent studies have discussed the assessment of data owner contributions based on Shapley values and its refinement [5, 10].

After assessing the contribution of the data owners, the model owners should allocate the two types of rewards to data owners to maintain and/or increase their participation level.

- Offer rewards. Model owners can reward data owners before training. The payoffs can be determined by the quality of the resources provided [18] or the outcome of a vote [11].
- Share profits. In this scenario, the model owner shares the profits that the model has generated with the data owners after the model has been trained. In these situations, payoff delays may affect the participant's likelihood to contribute. However, a reward-sharing scheme [14, 15] allows for a given budget to be divided dynamically.

3 The multi-stage incentive mechanism model

To ensure the success in federated learning and allow for the best training result, it is crucial to implement an effective incentive mechanism that minimises the

possibility of dual ethical risk. Based on the discussion in the previous section, no existing incentive mechanism has suitably addressed this issue. This section introduces a multi-stage incentive mechanism model based on contract theory. It addresses the dual ethical risks associated with federated learning while incentivising both parties to cooperate successfully. Note that, for simplicity, the game assumes one data owner and one model owner. A contract-theoretic solution for federated learning scenarios with more than one data owner is left to future work.

3.1 The model

The two participants in our model, the data owner and the model owner, are risk-neutral. Both parties agree that the entire training process will be conducted in K stages, with both parties jointly checking the training results at the end of each stage to confirm that the training was successful. Additionally, both parties agree that the contract cannot be ended earlier than these K stages unless the training fails. We assume that the effort value committed by the data owner at stage k is De_k, and the effort value committed by the model owner at stage k is Me_k. De_k and Me_k are both uncorrelated variables. Furthermore, $De_k \geq 0, Me_k \geq 0$.

Table 1 lists the notations commonly used in this paper for ease of reference.

Table 1. Commonly Used Notations

Notation	Description
k	Training stages, $k = 1, \cdots K$.
Me_k	The effort committed by the model owner at stage k
De_k	The effort committed by the data owner at stage k
$P_k(Me_k, De_k)$	The probability of successful training at stage k
$C(Me_k)$	The effort cost of the model owner at stage k
$C(De_k)$	The effort cost of the data owner at stage k
V_k	The incremental value of the model after stage k
M_k	The market value of the model at stage k
I_k	The data owner's costs at stage k
DR_k	Total expected revenue of the data owner from stage k to K
MR_k	Total expected revenue of the model owner from stage k to K
R_k	The reward received by the model owner if training success at stage k
$X_k(Me_k, De_k)$	The model's performance at stage k
ϕ, ν	The weight parameters of the model at stage k

Naturally, the performance of a model, e.g., the accuracy of its inferences, will be higher if the data owner contributes more effort to providing more and higher

quality data. Similarly, if the model owner puts in more effort, such as improving the algorithm, model performance will also increase. The model's performance is assumed to be

$$X_k(Me_k, De_k) = 1 - e^{-\phi(Me_k, De_k)^{\nu}},$$

where ϕ and ν are the weight parameters.

Fig. 2 shows the relationship between the performance of a typical federation learning model and the effort values Me and De of the training participants.

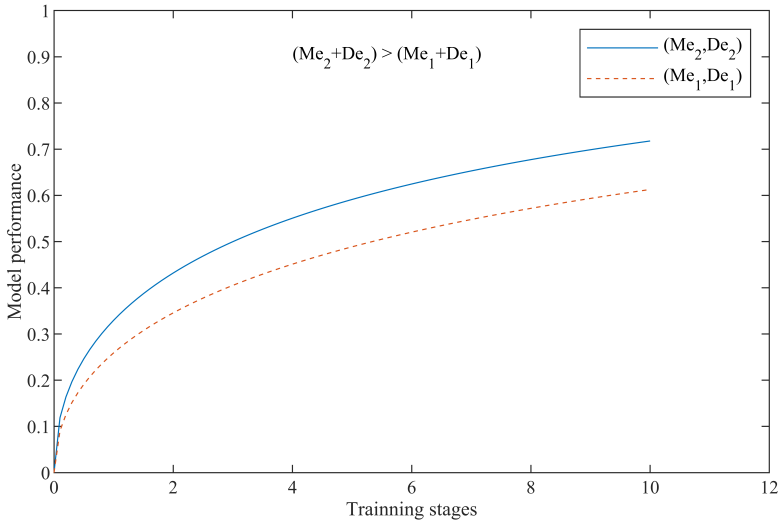

Fig. 2. Federated Machine Learning Performance

The following assumptions are made over the probability that training at stage k will be successful:

$$P_k(Me_k, De_k),$$

and

$$1 \geq P_k(Me_k, De_k) \geq 0, \frac{\partial P_k(Me_k, De_k)}{\partial Me_k} > 0, \frac{\partial P_k(Me_k, De_k)}{\partial De_k} > 0,$$

$$\frac{\partial^2 P_k(Me_k, De_k)}{\partial {Me_k}^2} < 0, \frac{\partial^2 P_k(Me_k, De_k)}{\partial {De_k}^2} < 0, (k = 1, \cdots, K).$$

Thus, there is a positive correlation between the probability of successful training and the efforts contributed by the data and model owners. The probability of success increases as Me_k and De_k increase with diminishing marginal returns.

The cost of inputting effort by the two parties in the training stage k are $C(Me_k)$ and $C(De_k)$. Obviously, these costs increase as the effort increases, i.e., $C'(Me_k) > 0$, $C'(De_k) > 0$. Similarly, the marginal cost of effort increases as well, i.e., $C''(Me_k) > 0$, $C''(De_k) > 0$.

Suppose that federated learning is successful in stage k. In that case, the data owner receives the incremental value of the upgraded model as V_k (V_k is a constant agreed upon by both participants before the contract), and the training continues into stage $k+1$. Assuming that the model's market value at the end of stage k is M_k and the data owner's cost at stage k is I_k, we have $V_k = M_k - I_k$. After all K stages of training have been completed, the data owner receives the final value of the model as $\sum_{k=1}^{K} V_k = \sum_{k=1}^{K}(M_k - I_k)$.

DR_k and MR_k are defined as the total expected revenues of the data owner and model owner from stage k to K. Logically, the data owner will only participate in training if they believe that the total expected revenue will be positive. If the total expected revenue in stages k to K turns out to be a loss, the data owner will drop out at any stage from $k+1$ to K and terminate the contract. Therefor, we can assume that $V_k + DR_{k+1} > 0$ and $DR_k \geq 0$. This assumption is reasonable because it assumes that the parties have some opportunity to argue success or failure at each stage. If the data owner expects a negative payoff, they will claim failure to get out of the contract. It is assumed that before a particular point in the training $V_k < 0$, i.e., the data owner's contribution is more significant than the benefit. After that point, the data owner's payoff becomes positive. This assumption ensures that the data owner agrees to cooperate with the model owner for the purposes of training the model. R_k represents the reward given by the data owner to the model owner if the training is successful at stage k. The event sequence in the contract is shown in Fig. 3.

Before entering the federated learning scheme, the data owner and the model owner need to agree on the reward $R_k > 0 \quad (k = 1, \cdots, K)$ and set up the contract. The model owner receives R_k from the data owner after training is confirmed to be successful in stage k. According to the contract, the model owner commits the optimal level of effort ${Me_k}^*$ to maximise their expected return MR_k. At the same time, the data owner also to commit the optimal level of effort ${De_k}^*$ to maximise DR_k. If the training result is successful at the end of stage k, the value of the updated model held by the data owner increases by V_k, and the model owner receives the reward R_k from the data owner. Training then proceeds to the next stage. If stage k training fails, both the model owner

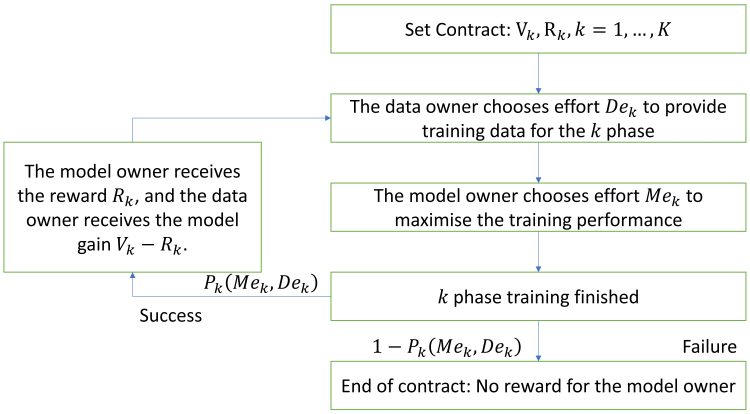

Fig. 3. Contract Execution Stages

and the data owner gain nothing for that stage. Note that the optimal strategy for the Stackelberg game leader is to not reward the follower for failure at each stage of the game [1, 2]. Both parties will pay $C(Me_k)$ and $C(De_k)$ regardless of success or failure. Thus, the following recursive equation describes the profit of the data owner and the model owner,

$$MR_k = P_k(Me_k, De_k)[R_k + MR_{k+1}] - C(Me_k) \tag{1}$$

and

$$DR_k = P_k(Me_k, De_k)[V_k - R_k + DR_{k+1}] - C(De_k), k = 1, \cdots, K. \tag{2}$$

In our model, the contract is set before the first phase. The relevant payoffs in the first phase are DR_1 for the data owner and MR_1 for the model owner. Note that the payoff for stage k is directly effected by the payoffs for stage $k+1$. Expanding the above recursive equations, we have:

$$\begin{aligned} MR_m = &\sum_{k=m}^{K} \left\{ \prod_{j=m}^{k} P_j(Me_j, De_j) R_k \right\} \\ &- \sum_{k=m}^{K} \left\{ \prod_{j=m}^{k-1} P_j(Me_j, De_j) C(Me_k) \right\} \end{aligned} \tag{3}$$

and

$$
\begin{aligned}
DR_m = &\sum_{k=m}^{K} \left\{ \prod_{j=m}^{k} P_j(Me_j, De_j)(V_k - R_k) \right\} \\
&- \sum_{k=m}^{K} \left\{ \prod_{j=m}^{k-1} P_j(Me_j, De_j) C(De_k) \right\}.
\end{aligned} \tag{4}
$$

3.2 Research Findings

In this section, we outline the findings of the above model, beginning with the optimal effort ${De_k}^*$ of the data owner.

The derivative of the data owner's payoff with respect to their effort De_k from Equation 2 is

$$
\begin{aligned}
\frac{dDR_k}{dDe_k} =& \frac{dP_k(Me_k, De_k)}{dDe_k}(V_k - R_k + DR_{k+1}) - \frac{dC(De_k)}{dDe_k} \\
=& 0 \quad (k = 1, \cdots, K),
\end{aligned} \tag{5}
$$

where

$$
\frac{dP_k(Me_k, De_k)}{dDe_k}(V_k - R_k + DR_{k+1}) = \frac{dC(De_k)}{dDe_k} \quad (k = 1, \cdots, K). \tag{6}
$$

Thus, the optimal effort ${De_k}^*$ of the data owner is:

$${De_k}^* = {De_k}^*(V_k - R_k + DR_{k+1}). \tag{7}$$

Corollary 1. *The optimal effort of the data owner is a function of the incremental value of the model, the reward to the model owner, and the data owner's expectation of future payoffs. Reducing the reward to the model owner and increasing the incremental value of the model and the data owner's expectations for the future should motivate the data owner to put in more effort and reduce their ethical risk.*

In the same way, we can solve the optimal effort ${Me_k}^*$ of the model owner. The derivative of the model owner's payoff with respect to it's effort Me_k from equation 1 is

$$
\begin{aligned}
\frac{dMR_k}{dMe_k} =& \frac{dP_k(Me_k, {De_k}^*)}{dMe_k}(R_k + MR_{k+1}) - \frac{dC(Me_k)}{dMe_k} \\
=& 0 \quad (k = 1, \cdots, K).
\end{aligned} \tag{8}
$$

Thus, the optimal effort ${Me_k}^*$ of the model owner is:

$${Me_k}^* = {Me_k}^*(R_k + MR_{k+1}). \tag{9}$$

Corollary 2. *The optimal effort level of the model owner is positively correlated with the reward and their expected future payoff. Higher rewards from the data owner and increasing the model owner's future expectations should motivate the model owner to work harder and reduce any ethical risks.*

Based on Corollaries 1 and 2, we have the following conditions:

$$\begin{cases} \frac{dP_k(Me_k,De_k)}{dDe_k}(V_k - R_k + DR_{k+1}) = \frac{dC(De_k)}{dDe_k}; \\ \frac{dP_k(Me_k,De_k)}{dMe_k}(R_k + MR_{k+1}) = \frac{dC(Me_k)}{dMe_k} \quad (k = 1, \cdots, K). \end{cases} \tag{10}$$

Corollary 3. *An optimal incentive mechanism should be such that the marginal benefit of each participant's effort equals their marginal cost.*

Given the optimal level of effort ${Me_k}^*$ and ${De_k}^*$ for the model owner and data owner, MR_m in Equation 3 satisfies the following conditions:

$$\frac{\partial MR_m}{\partial R_k} = \prod_{j=m}^{k} P_j({Me_j}^*, {De_j}^*) \quad (k = 1, \cdots, K; m \leq k). \tag{11}$$

From Equation 11,

$$\frac{\frac{\partial MR_1}{\partial R_k}}{\frac{\partial MR_1}{\partial R_{k+1}}} = \frac{\prod_{j=1}^{k} P_j({Me_j}^*, {De_j}^*)}{\prod_{j=1}^{k+1} P_j({Me_j}^*, {De_j}^*)} = \frac{1}{P_{k+1}({Me_{k+1}}^*, {De_{k+1}}^*)} > 1 \quad (k = 1, \cdots, K-1). \tag{12}$$

Then

$$\left.\frac{\partial MR_1}{\partial R_k}\right|_{{Me_j}^*, {De_j}^*} > \left.\frac{\partial MR_1}{\partial R_{k+1}}\right|_{{Me_j}^*, {De_j}^*} \quad (k = 1, \cdots, K-1). \tag{13}$$

Corollary 4. *The marginal utility of the rewards diminishes for the model owner over time. Therefore, to encourage the model owner to increase their effort, the rewards for the model owner in the incentive mechanism should be gradually increased as training continues. This should mean the incentive mechanism stays effective in motivating the model owner to work hard.*

The optimal incentive ${R_k}^* > 0 \quad (k = 1, \cdots, K)$ for the model owner is determined before starting the first stage of training. Therefore, the optimal payoff ${R_k}^*$ of the data owner can also be solved. The first-order condition of data owner with respect to payoff R_k from Equation 2 is

$$\begin{aligned} \left.\frac{\partial DR_1}{\partial R_k}\right|_{{R_i}^*, i=1,\cdots,K} = & \left[{P_1}'({Me_1}^*, {De_1}^*) {Me_1}^{*\prime} \frac{\partial MR_2}{\partial R_k} \right. \\ & \left. + {P_1}'({Me_1}^*, {De_1}^*) {De_1}^{*\prime} \frac{\partial DR_2}{\partial R_k} \right] (V_1 - R_1 + DR_2) \\ & + P_1({Me_1}^*, {De_1}^*) \frac{\partial DR_2}{\partial R_k} - C'(De_1^*) De_1^{*\prime} \frac{\partial DR_2}{\partial R_k} = 0. \end{aligned} \tag{14}$$

From Corollary 4, we can derive $\frac{\partial MR_2}{\partial R_k} = \prod_{j=2}^{k} P_j(Me_j{}^*, De_j{}^*)$ and from Corollary 1, we can derive $P_1'(Me_1{}^*, De_1{}^*)(V_1 - R_1 + DR_2) - C'(De_1{}^*) = 0, V_1 - R_1 + DR_2 > 0$. Substituting both of these into Equation 14 and rearranging the terms yield:

$$\begin{aligned} &\left\{De_1{}^{*\prime}[P_1{}'(Me_1{}^*, De_1{}^*)(V_1 - R_1 + DR_2) - C'(De_1{}^*)] + P_1(Me_1{}^*, De_1{}^*)\right\} \\ &\frac{\partial DR_2}{\partial R_k} + P_1{}'(Me_1{}^*, De_1{}^*)Me_1{}^* \left[\prod_{j=2}^{k} P_j(Me_j{}^*, De_j{}^*)\right](V_1 - R_1 + DR_2) = 0. \end{aligned} \tag{15}$$

Then,

$$\begin{aligned} \left.\frac{\partial DR_2}{\partial R_k}\right|_{R_i{}^*, i=1,\cdots,K} = &-\frac{1}{P_1(Me_1{}^*, De_1{}^*)} P_1{}'(Me_1{}^*, De_1{}^*)Me_1{}^{*\prime} \\ &\left[\prod_{j=2}^{k} P_j(Me_j{}^*, De_j{}^*)\right](V_1 - R_1 + DR_2) < 0. \end{aligned} \tag{16}$$

Thus, if $R_k{}^* > 0$ and $R_j{}^* > 0, j > k$, then

$$\begin{aligned} \left.\frac{\partial DR_2}{\partial R_j}\right|_{R_i{}^*, i=1,\cdots,K} &= \left(\prod_{i=k+1}^{j} P_i(Me_i{}^*, De_i{}^*)\right) \left.\frac{\partial DR_2}{\partial R_k}\right|_{R_i{}^*, i=1,\cdots,K} \\ &> \left.\frac{\partial DR_2}{\partial R_k}\right|_{R_i{}^*, i=1,\cdots,K}. \end{aligned} \tag{17}$$

Corollary 5. *The expected payoff to the model owner increases marginal utility for the data owner over time. Intuitively, the data owner always wants to delay the reward to the model owner, while the model owner wants to receive the reward as early as possible. For the data owner, the later the reward is given to the model owner, the more likely it is for ethical risk to be avoided.*

From Corollary 5, for $k > 1$,

$$\begin{aligned} \frac{\partial DR_1}{\partial R_k} = &\left[P_1(Me_1{}^*, De_1*)'Me_1{}^{*\prime}\frac{\partial MR_2}{\partial R_k} + P_1(Me_1{}^*, De_1*)'De_1{}^{*\prime}\frac{\partial DR_2}{\partial R_k}\right] \\ &(V_1 - R_1 + DR_2) + P_1(Me_1{}^*, De_1*)\frac{\partial DR_2}{\partial R_k} - C(De_1{}^*)'De_1{}^{*\prime}\frac{\partial DR_2}{\partial R_k}. \end{aligned} \tag{18}$$

For every $m < k$,

$$\begin{aligned} \frac{\partial DR_m}{\partial R_k} = &\left[P_m(Me_m{}^*, De_m*)'Me_m{}^{*\prime}\frac{\partial MR_{m+1}}{\partial R_k}\right. \\ &\left. + P_m(Me_m{}^*, De_m*)'De_m{}^{*\prime}\frac{\partial DR_{m+1}}{\partial R_k}\right](V_m - R_m + DR_{m+1}) \\ &+ P_m(Me_m{}^*, De_m*)\frac{\partial DR_{m+1}}{\partial R_k} - C(De_m{}^*)'De_m{}^{*\prime}\frac{\partial DR_{m+1}}{\partial R_k}, \end{aligned} \tag{19}$$

and for every k,

$$\begin{aligned}\frac{\partial DR_k}{\partial R_k} =&[P_k({Me_k}^*, De_k*)' {Me_k^*}' - P_k({Me_k}^*, De_k*)' {De_k^*}'](V_k - R_k + DR_{k+1}) \\ &- P_k({Me_k}^*, De_k*) + C({De_m}^*)' {De_k^*}'.\end{aligned} \tag{20}$$

From Corollary 1, we can derive $P_k({Me_k}^*, De_k*)'(V_k - R_k + DR_{k+1}) - C(De_k)^{*'} = 0$, and substituting this into the three equations above, we have

$$\begin{aligned}\frac{\partial DR_1}{\partial R_k} = &\left(\prod_{j=1}^{k} P_j({Me_j}^*, {De_j}^*)\right) \sum_{i=1}^{k} \frac{1}{P_i({Me_i}^*, {De_i}^*)} \\ &{P_i}'({Me_i}^*, {De_i}^*){Me_i}^*[V_i - R_i + DR_{i+1}] - \prod_{j=1}^{k} P_j({Me_j}^*, {De_j}^*),\end{aligned} \tag{21}$$

and

$$\begin{aligned}\frac{\partial DR_1}{\partial R_{k+1}} =&\frac{\partial DR_1}{\partial R_k} P_{k+1}({Me_{k+1}}^*, {De_{k+1}}^*) \\ &+ \left(\prod_{j=1}^{k} P_j({Me_j}^*, {De_j}^*)\right) {P_{k+1}}'({Me_{k+1}}^*, {De_{k+1}}^*) \\ &{Me_{k+1}^*}'[V_{k+1} - R_{k+1} + DR_{k+2}] = 0 \quad (k = 1, \cdots, K-1).\end{aligned} \tag{22}$$

Since $\left.\frac{\partial DR_1}{\partial R_k}\right|_{{R_i}^*, i=1,\cdots,K} = 0$, from Equation 22, we can derive $V_{k+1} - R_{k+1} + DR_{k+2} = 0$. It is known that $DR_{K+1} = 0$, so it follows that ${R_K}^* = V_K$, so $DR_K = 0$. Similarly, for any δ, there is $1 \le \delta \le K-1$. If ${Me_\delta}^* > 0$ and ${R_\delta}^* > 0$, then:

$$\begin{cases} {R_k}^* = V_k \quad (k = \delta + 1, \cdots, K). \\ DR_k = 0 \quad (k = \delta + 1, \cdots, K). \end{cases} \tag{23}$$

Then,

$$\begin{aligned}DR_1 =&\sum_{j=1}^{\delta-1} \left(\prod_{i=1}^{j} P_i({Me_i}^*, {De_i}^*)(V_j - C({De_i}^*))\right) \\ &+ \left(\prod_{i=1}^{\delta-1} P_i({Me_i}^*, {De_i}^*)\right) P_\delta({Me_\delta}^*, {De_\delta}^*)[V_\delta - R_\delta].\end{aligned} \tag{24}$$

Theorem 1. *The data owner can receive their optimal payoff at point δ during training such that*

$$\begin{cases} {R_k}^* = 0 \quad (k < \delta), \\ {R_k}^* = {V_k}^*, {DR_k}^* = 0 \quad (k > \delta), \end{cases} \tag{25}$$

and

$$\begin{cases} DR_1 \geq \sum_{j=1}^{\delta-1} \left(\prod_{i=1}^{j} P_i({Me_i}^*, {De_i}^*) \left(V_j - C({De_j}^*) \right) \right), \\ DR_1 \leq \sum_{j=1}^{\delta} \left(\prod_{i=1}^{j} P_i({Me_i}^*, {De_i}^*) \left(V_j - C({De_j}^*) \right) \right). \end{cases} \tag{26}$$

Theorem 1 shows an optimal payoff point for the data owner, where the data owner receives the total payoff from the federated learning process and the reward given to the model owner is zero in phases $1 - \delta$. However, after that point, the data owner does not have any profit, the expected future payoffs are zero, and the benefit goes entirely to the model owner. Thus, point δ is the optimal payoff point for the data owner. Essentially, what Theorem 1 indicates is that, for a federated learning scenario initiated by the data owner, the optimal incentive scheme is one where as much of the incremental value of the model as possible is paid to the model owner. Therefore, success in the later stages of training is based on the success in the earlier stages and, in turn, rewards in the later stages incentivise effort in the earlier stages. Overall, giving back as much of the value created by the model owner's efforts as possible in the later stages is the least costly incentive scheme for the data owner.

4 Experimental Evaluation

To complement the analytical findings and evaluate the performance of our incentive mechanism for federated learning, we create a multi-stage contract simulator for the data and model owners. The simulator evaluates the impact of different reward settings on the level of effort contributed by each participant and gives the total payoff for both parties.

4.1 Experiment Settings

Assume that the incremental model values are $V_1 = 1, V_2 = 2$ and $V_3 = 3$, where federated learning is carried out in 3 stages (i.e., $K = 3$) and the functional expression for the probability of success at each stage is $P_k(Me_k, De_k) = MIN(0.6(Me_k + De_k), 1)$. As we will see later, the equilibrium effort satisfies $0.6({Me_k}^* + {De_k}^*) < 1$, so we can count $P_k(Me_k, De_k) = 0.6(Me_k + De_k)$. We also assume that the effort cost of the model owner's function is $C(Me_k) = {Me_k}^2$, and the effort cost of the data owner's function is $C(De_k) = {De_k}^2$, such that the utility function of the model owner is

$$mr_k = 0.6(Me_k + De_k)(R_k + mr_{k+1}) - {Me_k}^2, \quad k = 1, 2, 3,$$
$$mr_4 = 0.$$

Taking the utility function for each stage and deriving it to its effort level determines the optimal effort yield for the model owner:

$$Me_k{}^* = \frac{\partial mr_k}{\partial Me_k} = 0.3(R_k + mr_{k+1}) \quad k = 1, 2, 3.$$

Repeating the same approach, and its based on Equation 24, we can derive the utility function of the data owner and their optimal effort:

$$dr_k = 0.6(Me_k + De_k)(V_k - R_k + dr_{k+1}) - De_k{}^2, \quad k = 1, 2, 3,$$
$$dr_3 = 0, dr_4 = 0.$$
$$De_k{}^* = \frac{\partial dr_k}{\partial De_k} = 0.3(V_k - R_k + dr_{k+1}) \quad k = 1, 2, 3,$$
$$De_3 = 0.$$

The utility functions and the optimal efforts of the two parties in different stages are listed in Table 2.

Table 2. the utility functions and the optimal efforts, $K = 3$

K	**Data Owner**	**Model Owner**
1	$dr_1 = 0.6(Me_1 + De_1)(V_1 - R_1 + dr_2) - De_1{}^2$ $De_1 = 0.3(V_1 - R_1 + dr_2)$	$mr_1 = 0.6(Me_1 + De_1)(R_1 + mr_2)$ $Me_1 = 0.3(R_1 + mr_2)$
2	$dr_2 = 0.6(Me_2 + De_2)(V_2 - R_2 + dr_3) - De_2{}^2$ $De_2 = 0.3(V_2 - R_2 + dr_3)$	$mr_2 = 0.6(Me_2 + De_2)(R_2 + mr_3)$ $Me_2 = 0.3(R_2 + mr_2)$
3	$dr_3 = 0$ $De_3 = 0$	$mr_3 = 0.6(Me_3 + De_3)(R_3 + mr_4)$ $mr_4 = 0, Me_3 = 0.3(R_3 + mr_3)$

4.2 Experimental Result and Discussion

Fig. 4 shows the optimal rewards yielded for the model owner, calculated by recurring the above equations in Table 2 and the derivative of the data owner's payoff dr_1 with respect to the reward $R_2{}^*$:

$$R_1{}^* = 0, R_2{}^* = 0.4085, R_3{}^* = 3,$$

where the probabilities of successes are $P_1(Me_1{}^*, De_1{}^*) = 0.3707, P_2(Me_2{}^*, De_2{}^*) = 0.5058, P_3(Me_3{}^*, De_3{}^*) = 0.54$. As predicted by Theorem 1, the optimal payoff point for the data owner is $\delta = 2$, and $R_1{}^* = 0, R_3{*} = V_3$, and $0 < R_2 < V_2$. The data owner's expected payoff is $dr_1 = 0.3608$, which is consistent with Theorem 1,

$$\begin{cases} dr_1 \geq P_1(Me_1{}^*, De_1{}^*)(V_1 - De_1{}^2) = 0.2878, \\ dr_1 \leq P_1(Me_1{}^*, De_1{}^*)(V_1 - De_1{}^2) + P_2(Me_2{}^*, De_2{}^*)(V_2 - De_2{}^2) = 1.1841. \end{cases}$$

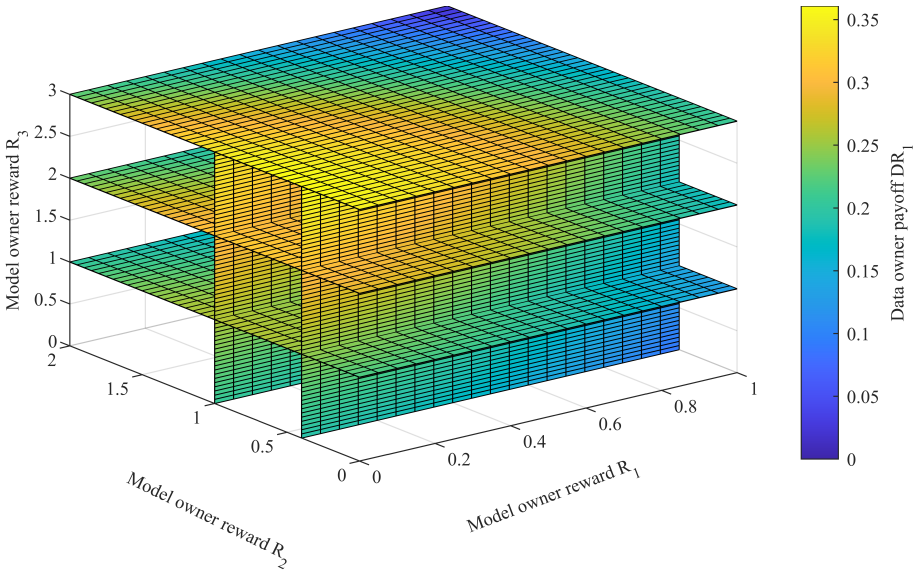

Fig. 4. The optimal rewards yielded for the model owner

Table 3. the reward settings vs the best efforts of each stage

Reward settings	**DO expected payoff** dr_1	**Stg1 BEs** $Me_1 + De_1$	**Stg2 BEs** $Me_2 + De_2$	**Stg3 BEs** $Me_3 + De_3$
$R_1 = 0, R_2 = 0.2, R_3 = 3$	0.3508	0.6114	0.843	0.9
${R_1}^* = 0, {R_2}^* = 0.4085, {R_3}^* = 3$	***0.3608***	***0.6179***	***0.843***	***0.9***
$R_1 = 0, R_2 = 1, R_3 = 3$	0.3386	0.6109	0.843	0.9
$R_1 = 0.5, R_2 = 0.4085, R_3 = 3$	0.2949	0.6179	0.843	0.9

We have taken some relevant data from the simulator to make it easier to understand, as shown in Table 3. This table shows the effects of the reward value settings at different stages on the efforts of the participants and the expected payoff for the data owner in the incentive contract. Some settings around the optimal one have been selected as comparisons: ${R_1}^* = 0, {R_2}^* = 0.4085, {R_3}^* = 3$. From the results, we can see that:

1. Any deviation from the optimal value of ${R_2}^* = 0.4085$ negatively impacts the efforts of both participants and the expected training payoff for the data owner. This means that any reward setting that deviates from the optimal value ${R_2}^*$ will increase the ethical risk of the participants.
2. If the data owner keeps $R_2 = {R_2}^*$ and increases the reward R_1 for stage 1, this scenario is identical to the optimal incentive scenario in terms of the effort values at each stage. However, the data owner's expected training payoffs will be significantly lower. From a self-interested perspective by the data owner, as the leader of the incentive contract, there is no incentive to increase the reward given to the model owner at Stage 1.

Thus, we can conclude that our model is able to reduce the dual ethical risk of federated learning due to information asymmetry. It can motivate the participants to exert an optimized effort to training, confirming the intuition behind our model that the success in the later stages is based on success in the earlier stages. Thus, rewards in the later stages incentivise efforts in the earlier stages. Moreover, giving back as much of the value created by the model owner's efforts in the later stages is the least costly incentive scheme for the data owner.

5 Conclusion and Future Works

In this paper, we have used the framework of a dynamic game to investigate the dual ethical risk problem between model owners and data owners in federated learning. The model used is novel and it has derived optimal incentive payoff contracts for the data and model owners through two sets of analyses: one for a multi-stage incentive payoff game and the other for the dual ethical risk affecting the contract design. The output is an optimal payoff point for the data owners. Our approach has provided insights into the characteristics of optimal incentive contracts between data owners and model owners in federated learning schemes, including their endogenous optimality. Specifically, our study has shown that, for a federated learning scenario initiated by the data owner, the optimal incentive scheme is one where as much of the incremental value of the model as possible is paid to the model owner. There could be several possible extensions of this paper, which requires further research in this field. First, we explored the dual ethical risk problem in the data owner-led federated learning scenario using a multi-stage incentive model. Further work will extend this model in other scenarios and can be compared comprehensively with existing incentive mechanisms. Second, we negated the possibility of multiple data owners to treat them as a single entity. It would be interesting to consider multi-data owners joining the game at different

stages as a possible extension to our proposed model. The third extension of this paper would be to investigate how the efforts of model and data owners with fair preferences in the later stages of cooperation (based on fair preference theory) are affected by the value of benefits and new compensation schemes.

References

1. Bergemann, D., Hege, U.: Venture capital financing, moral hazard, and learning. Journal of Banking & Finance **22**(6-8), 703–735 (1998)
2. Elitzur, R., Gavious, A.: A multi-period game theoretic model of venture capitalists and entrepreneurs. European Journal of Operational Research **144**(2), 440–453 (2003)
3. European Parliament, C.o.t.E.U.: Guide to the general data protection regulation (2018), https://www.gov.uk/government/publications/guide-to-the-general-data-protection-regulation
4. Feng, S., Niyato, D., Wang, P., Kim, D.I., Liang, Y.C.: Joint service pricing and cooperative relay communication for federated learning. In: 2019 International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData). pp. 815–820. IEEE (2019)
5. 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)
6. Konečnỳ, J., McMahan, H.B., Ramage, D., Richtárik, P.: Federated optimization: Distributed machine learning for on-device intelligence. arXiv preprint arXiv:1610.02527 (2016)
7. Konečnỳ, J., McMahan, H.B., Yu, F.X., Richtárik, P., Suresh, A.T., Bacon, D.: Federated learning: Strategies for improving communication efficiency. arXiv preprint arXiv:1610.05492 (2016)
8. McMahan, H.B., Moore, E., Ramage, D., y Arcas, B.A.: Federated learning of deep networks using model averaging. arXiv preprint arXiv:1602.05629 (2016)
9. Nowak, A.S., Radzik, T.: The shapley value for n-person games in generalized characteristic function form. Games and Economic Behavior **6**(1), 150–161 (1994)
10. Sim, R.H.L., Zhang, Y., Chan, M.C., Low, B.K.H.: Collaborative machine learning with incentive-aware model rewards. In: International Conference on Machine Learning. pp. 8927–8936. PMLR (2020)
11. Toyoda, K., Zhang, A.N.: Mechanism design for an incentive-aware blockchain-enabled federated learning platform. In: 2019 IEEE International Conference on Big Data (Big Data). pp. 395–403. IEEE (2019)
12. Tu, X., Zhu, K., Luong, N.C., Niyato, D., Zhang, Y., Li, J.: Incentive mechanisms for federated learning: From economic and game theoretic perspective. IEEE Transactions on Cognitive Communications and Networking (2022)
13. Yang, Q., Liu, Y., Chen, T., Tong, Y.: Federated machine learning: Concept and applications. ACM Transactions on Intelligent Systems and Technology (TIST) **10**(2), 1–19 (2019)
14. Yu, H., Liu, Z., Liu, Y., Chen, T., Cong, M., Weng, X., Niyato, D., Yang, Q.: A fairness-aware incentive scheme for federated learning. In: Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society. pp. 393–399 (2020)

15. Yu, H., Liu, Z., Liu, Y., Chen, T., Cong, M., Weng, X., Niyato, D., Yang, Q.: A sustainable incentive scheme for federated learning. IEEE Intelligent Systems **35**(4), 58–69 (2020)
16. Zeng, R., Zeng, C., Wang, X., Li, B., Chu, X.: A comprehensive survey of incentive mechanism for federated learning. arXiv preprint arXiv:2106.15406 (2021)
17. 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)
18. Zhang, W., Lu, Q., Yu, Q., Li, Z., Liu, Y., Lo, S.K., Chen, S., Xu, X., Zhu, L.: Blockchain-based federated learning for device failure detection in industrial iot. IEEE Internet of Things Journal **8**(7), 5926–5937 (2020)