PABP: An Approximately Truthful Auction Mechanism for Federated Learning with Bid Privacy Protection

Anonymous authors
Paper under double-blind review

Abstract

Auction mechanisms incentivize device participation in federated learning (FL), but existing designs typically require devices to disclose their private bids to the server. Protecting these bids is challenging because perturbed bids can distort allocation and payment decisions, thereby undermining economic incentives. We propose PABP, an approximately truthful auction mechanism that jointly designs bid perturbation, device selection, resource allocation, and payment determination. Each device locally perturbs its bid with Gaussian noise, and the server makes decisions using only the perturbed bids. PABP further approximates the theoretical Myerson payment under noisy bids, which provides ε-approximate incentive compatibility and individual rationality. Experiments across standard benchmarks and diverse data settings show that PABP protects bid privacy while consistently outperforming representative baselines.

1 Introduction

Auction mechanisms are widely used in federated learning (FL) to incentivize device participation in model training (Li et al., 2025). Existing studies, such as (Pang et al., 2023; Li et al., 2025; 2026a), mainly pursue efficient allocation while ensuring desirable economic properties. However, they typically require devices to disclose their private bids to a central server. These bids may be exposed by an untrusted server or malicious devices, which is unacceptable to privacy-sensitive participants. Such devices may therefore decline to participate even when offered high rewards, which reduces participant diversity and harms model generalization. This motivates a fundamental question: ***Can an auction mechanism preserve bid privacy without sacrificing its effectiveness?*** Addressing this question is challenging for two reasons.

Challenge 1: Privacy conflicts with effective decision making. The server normally compares bids to determine which devices should participate and how limited training resources should be allocated. When exact bids are hidden, distinguishing devices with similar bids becomes difficult. The first challenge is therefore to protect individual bids while retaining enough information to make effective allocation decisions.

Challenge 2: Bid protection complicates economic incentives. Truthful reporting and individual rationality require payments to be consistent with how a bid affects selection. Without access to exact bids, traditional payment rules are not directly applicable, and payments based on distorted information may reward misreporting. Bid privacy, device selection, and payment determination must therefore be jointly designed.

To address these challenges, we propose **PABP**, a Privacy-Aware Bid Protection auction mechanism for resource-constrained FL. Each device locally randomizes its original bid and reports only the protected bid, and the original bid and the locally sampled randomness never leave devices. With only protected bids, PABP jointly determines device selection, computing and communication resources, and payments. Because directly applying a conventional payment rule to protected bids may undermine truthful incentives, PABP derives

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

an ideal expected payment under randomized allocation and calibrates a payment toward this target before bidding. Our main contributions are summarized as follows:

- **Bid-private FL auction formulation.** We formulate a resource-aware FL auction in which original bids remain private while the server must jointly determine device selection, computing resources, communication resources, and payments. The formulation captures the fundamental interaction among bid privacy, allocation efficiency, and economic incentives;
- **Joint privacy-aware allocation.** We develop PABP, in which devices locally protect their bids and the server makes all decisions using only the protected bids. Moreover, PABP obtains the globally optimal device selection and resource allocation for the corresponding executable welfare objective;
- **Payment with theoretical guarantees.** We design a Myerson-calibrated payment rule that approximates the ideal expected payment while depending only on protected bid. We prove that the expected allocation is monotone in the original bid. If the maximum conditional expected payment error is bounded by δ_i, PABP satisfies $2\delta_i$-approximate expected incentive compatibility and δ_i-approximate expected individual rationality;
- **Extensive empirical evaluation.** Experiments on MNIST, Fashion-MNIST, and CIFAR-10 show that PABP achieves model accuracy close to that of an auction without bid protection, while consistently outperforming baselines. It retains 97.28%–97.82% of the full-information optimal social welfare.

2 BACKGROUND AND RELATED WORK

2.1 Auction Design in FL

FL services let users benefit from collaborative training without sharing raw data(Kourtellis et al., 2020). In our setting, selected users both train the model and pay for the FL service. Prior auctions address related transactions: Li et al. (2024) price model access based on performance, whereas Pang et al. (2023) use a procurement auction to select and pay clients. Li et al. (2025) jointly determine user selection, computing and communication resources, and payments. These studies motivate joint selection and resource allocation, but do not jointly address bid privacy and service payments computed from perturbed bids.

2.2 Incentive Design with Privacy in FL

Sun et al. (2022) account for data owners' privacy costs in a differentially private model marketplace, while Yuan & Wang (2024) jointly design client selection, rewards, and privacy budgets. Both protect information exposed during training, not the bids used for allocation. PABP instead protects bids. Since bid perturbation changes allocation, a payment based on original bids cannot be executed directly. PABP calibrates a payment computed from perturbed bids toward the expected Myerson target, yielding approximate expected IC and IR when payment error is bounded.

3 System Model and Problem Formulation

The system consists of one FL server and N devices, indexed by $i \in \mathcal{N} = \{1, \ldots, N\}$. Device i holds a local dataset $\mathcal{D}_i$ with size $D_i = |\mathcal{D}_i|$, and the total amount of data is $D_{\text{tot}} = \sum_{i=1}^{N} D_i$. FL continues for L global rounds, indexed by $\ell \in \mathcal{L} = \{1, \ldots, L\}$, each lasting T^d. In each round, the server broadcasts the model and receives local updates (McMahan et al., 2017; Nishio & Yonetani, 2019). Then, it evaluates device i's uploaded model $\boldsymbol{w}_i^{(\ell)}$ on a validation set $\mathcal{D}^{\text{val}}$ to obtain accuracy $\text{Acc}(\boldsymbol{w}_i^{(\ell)}; \mathcal{D}^{\text{val}})$. Let $x_i^{(\ell)} = 1$ indicate that device i participates in round ℓ and passes the server's model validation. The model-quality

contribution coefficient of device i is defined as

$$\alpha_i^{(\ell)} = \begin{cases} \dfrac{D_i}{D_{\text{tot}}} \operatorname{Acc}\Big(\boldsymbol{w}_i^{(\ell)}; \mathcal{D}^{\text{val}}\Big), & x_i^{(\ell)} = 1, \\ \alpha_i^{(\ell-1)}, & x_i^{(\ell)} = 0. \end{cases} \tag{1}$$

The server then aggregates the validated devices' models as:

$$\boldsymbol{w}^{(\ell+1)} = \frac{\sum_{i=1}^{N} x_i^{(\ell)} D_i \boldsymbol{w}_i^{(\ell)}}{\sum_{i=1}^{N} x_i^{(\ell)} D_i}, \tag{2}$$

We define refresh rounds to occur every H rounds, with $\mathcal{R} = \{1, 1+H, 1+2H, \ldots\} \cap \mathcal{L}$. In refresh rounds, the server updates clients' model-quality contribution coefficients. Rounds with $\ell \notin \mathcal{R}$ are regular rounds, in which PABP selects devices, and only selected devices update their contribution coefficients by (1). Figure 1 shows the round-wise procedure. The complete pseudocode is given in Algorithm 2 in Appendix A.4.

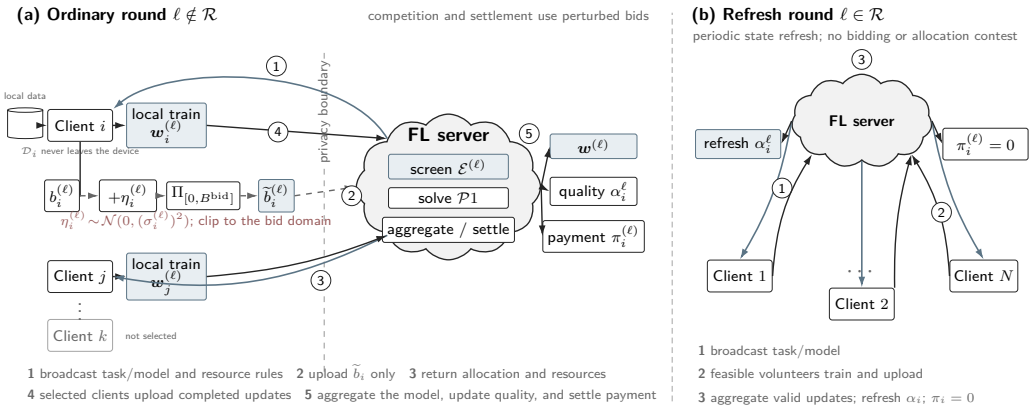

Figure 1: Round-wise system framework of PABP. Regular rounds perform resource-aware allocation under bid protection. Refresh rounds periodically update devices' model-quality contribution coefficients. Original bids and local data remain with devices.

3.1 Incentive Mechanism Model for FL

In regular round ℓ, the server sends the global model to selected devices. Since this model incorporates information from other devices, it can improve local training accuracy. Thus, devices are both local training participants and buyers in FL. They compete for a trained model by submitting bids (Li et al., 2024). For example, in an image-classification task, a device with limited data may not achieve the required classification performance using only local training. A model that incorporates other devices' updates can contribute to future local training and reduce training costs (Li et al., 2026b). Let the private information of device i in regular round ℓ be $b_i^{(\ell)} \in \mathbf{B}$, where $\mathbf{B} = [0, B^{\text{bid}}]$ is the range of $b_i^{(\ell)}$ with upper bound B^{bid}. The bid $b_i^{(\ell)}$ denotes the monetary benefit that one unit of model-quality contribution brings to device i. Under truthful bidding $b_i^{(\ell)} = v_i^{(\ell)}$. Hence, $b_i^{(\ell)} \alpha_i^{\ell-1}$ is the value of device i when its current model-quality contribution coefficient is $\alpha_i^{\ell-1}$. Let the bid vector of all devices be $\boldsymbol{b}^{(\ell)} = \big(b_i^{(\ell)}\big)_{i \in \mathcal{N}} = \big(b_1^{(\ell)}, \ldots, b_N^{(\ell)}\big)$.

To protect private information, each device adds noise to its original bid $b_i^{(\ell)}$. The server publicly specifies the Gaussian noise distribution and its standard deviation $\sigma_i^{(\ell)}$. It does not observe the device's original bid $b_i^{(\ell)}$ or locally generated noise sample $\eta_i^{(\ell)} \sim \mathcal{N}\Big(0, (\sigma_i^{(\ell)})^2\Big)$. Device i submits the perturbed bid as

$$\widetilde{b}_i^{(\ell)} = \min\Big\{B^{\text{bid}}, \max\Big\{0, b_i^{(\ell)} + \eta_i^{(\ell)}\Big\}\Big\}. \tag{3}$$

We assume that noise samples are independent across devices. The perturbed bids form the vector $\widetilde{\boldsymbol{b}}^{(\ell)} = (\widetilde{b}_i^{(\ell)})_{i\in\mathcal{N}} \in \mathbf{B}^N$. The detailed procedure is given in Appendix A.6. After receiving $\widetilde{\boldsymbol{b}}^{(\ell)}$, the server obtains the model-quality contribution coefficient $\alpha_i^{\ell-1}$, the training parameters, and the current channel state. It then determines the device selection $x_i^{(\ell)}$. Moreover, since available computing and communication resources are also limited, at most K devices can be admitted for model training in each round. Thus, we have

$$\sum_{i=1}^{N} x_i^{(\ell)} \le K. \tag{4}$$

Local Training Model. Let $\epsilon_i^{(\ell)} \in (0,1)$ be the local training accuracy parameter. A smaller value requires higher local training accuracy. Following (Li et al., 2025; 2026a), the number of local iterations required to achieve this accuracy is $n_i^{(\ell)} = \xi_i \ln\left(\frac{1}{\epsilon_i^{(\ell)}}\right)$, where $\xi_i > 0$ is a constant determined by the parameters of device i's local loss function(Li et al., 2025; 2026a). Training one sample requires g_i CPU cycles. Thus, one local iteration requires $g_i D_i$ cycles, and round ℓ requires $A_i^{(\ell)} = n_i^{(\ell)} g_i D_i$ cycles. The local training time is $T_i^{\mathrm{tr},(\ell)}(f_i^{(\ell)}) = \frac{A_i^{(\ell)}}{f_i^{(\ell)}}$, where $f_i^{(\ell)}$ is the CPU frequency used for model training. The CPU constraint is

$$0 \le f_i^{(\ell)} \le F_i,\ \forall i, \tag{5}$$

Let κ_i denote the computing-energy coefficient of device i, and let $c_i > 0$ be the monetary cost per unit of energy in the current round. The training resource cost is $\mathcal{C}_i^{\mathrm{tr},(\ell)} = c_i \kappa_i A_i^{(\ell)} \left(f_i^{(\ell)}\right)^2$ (Li et al., 2025; 2026a). The coefficient c_i is not private information reported by the device.

Wireless Communication Model. Similar to (Li et al., 2025; 2026a), an orthogonal transmission scheme is used. Let B_i be the bandwidth allocated to device i, $h_i^{(\ell)}$ its channel gain in round ℓ, and N_0 the noise power. The Shannon formula gives the upload rate $r_i^{\mathrm{up},(\ell)}(p_i^{(\ell)}) = B_i \log_2\left(1 + \frac{p_i^{(\ell)} h_i^{(\ell)}}{N_0}\right)$, where $p_i^{(\ell)}$ is the uplink transmission power. The power constraint is

$$0 \le p_i^{(\ell)} \le P_i,\ \forall i. \tag{6}$$

Let S_i be the local model size. Then, the uploading time and cost are $T_i^{\mathrm{up},(\ell)}(p_i^{(\ell)}) = \frac{S_i}{r_i^{\mathrm{up},(\ell)}(p_i^{(\ell)})}$ and $\mathcal{C}_i^{\mathrm{up},(\ell)} = c_i p_i^{(\ell)} T_i^{\mathrm{up},(\ell)}$, respectively. Therefore, the training and uploading time must not exceed the round duration, i.e.,

$$T_i^{\mathrm{tr},(\ell)}(f_i^{(\ell)}) + T_i^{\mathrm{up},(\ell)}(p_i^{(\ell)}) \le T^d, \quad \forall i. \tag{7}$$

3.2 Device and Server Utilities

In regular round ℓ, a selected device receives a true benefit of $v_i^{(\ell)} \alpha_i^{\ell-1}$. Local training and model upload incur costs. The device also pays the service payment $\pi_i^{(\ell)}$ to the server. Thus, the utility of device i is

$$U_i^{(\ell)} = x_i^{(\ell)} \left[v_i^{(\ell)} \alpha_i^{\ell-1} - \mathcal{C}_i^{\mathrm{tr},(\ell)} - \mathcal{C}_i^{\mathrm{up},(\ell)} \right] - \pi_i^{(\ell)}. \tag{8}$$

Then, the server's utility is $U_s^{(\ell)} = \sum_{i=1}^{N} \pi_i^{(\ell)} x_i^{(\ell)}$. The true social welfare in round ℓ is the sum of the server utility and all device utilities:

$$\begin{aligned} SW^{(\ell)} &= U_s^{(\ell)} + \sum_{i=1}^{N} U_i^{(\ell)} \\ &= \sum_{i=1}^{N} x_i^{(\ell)} \left[v_i^{(\ell)} \alpha_i^{\ell-1} - \mathcal{C}_i^{\mathrm{tr},(\ell)} - \mathcal{C}_i^{\mathrm{up},(\ell)} \right]. \end{aligned} \tag{9}$$

Bid perturbation makes the allocation and payment in each round depend on the noise $\boldsymbol{\eta}^{(\ell)}$. Fix the pre-perturbation bids of all other devices, $\boldsymbol{b}_{-i}^{(\ell)}$. The expected utility of device i is

$$U_i^{\mathrm{E},(\ell)}(v_i^{(\ell)}, b_i^{(\ell)}; \boldsymbol{b}_{-i}^{(\ell)}) = \mathbb{E}_{\boldsymbol{\eta}^{(\ell)}}\big[x_i^{(\ell)}\big(v_i^{(\ell)}\alpha_i^{\ell-1} - \mathcal{C}_i^{\mathrm{tr},(\ell)} - \mathcal{C}_i^{\mathrm{up},(\ell)}\big) - \pi_i^{(\ell)}\big]. \tag{10}$$

where $b_i^{(\ell)}$ is the private information for which truthfulness is required. Resource-cost parameters are fixed before bidding or are verifiable. The expectation is taken over the bid noise in the mechanism. We allow a gap of at most ε between the expected utility under truthful bidding. This condition must hold for any fixed bids of other devices.

Definition 1 (Approximate Expected Incentive Compatibility) *For any true valuation $v_i^{(\ell)}$, reported bid $b_i'^{(\ell)} \in [0, B^{\mathrm{bid}}]$, and $\boldsymbol{b}_{-i}^{(\ell)}$, if*

$$U_i^{\mathrm{E},(\ell)}(v_i^{(\ell)}, v_i^{(\ell)}; \boldsymbol{b}_{-i}^{(\ell)}) \geq U_i^{\mathrm{E},(\ell)}(v_i^{(\ell)}, b_i'^{(\ell)}; \boldsymbol{b}_{-i}^{(\ell)}) - \varepsilon_{\mathrm{IC,i}}^{(\ell)}, \tag{11}$$

then the mechanism is $\varepsilon_{\mathrm{IC,i}}^{(\ell)}$-approximate expected IC.

Definition 2 (Approximate Expected Individual Rationality) *For any $v_i^{(\ell)}$ and $\boldsymbol{b}_{-i}^{(\ell)}$, if*

$$U_i^{\mathrm{E},(\ell)}(v_i^{(\ell)}, v_i^{(\ell)}; \boldsymbol{b}_{-i}^{(\ell)}) \geq -\varepsilon_{\mathrm{IR,i}}^{(\ell)}, \tag{12}$$

then the mechanism is $\varepsilon_{\mathrm{IR,i}}^{(\ell)}$-approximate expected IR.

Since the server observes only the perturbed bid $\widetilde{b}_i^{(\ell)}$, it uses the following welfare objective:

$$SW_{\mathrm{E}}^{(\ell)} = \sum_{i=1}^{N} x_i^{(\ell)} \left[\widetilde{b}_i^{(\ell)}\alpha_i^{\ell-1} - \mathcal{C}_i^{\mathrm{tr},(\ell)} - \mathcal{C}_i^{\mathrm{up},(\ell)}\right]. \tag{13}$$

Thus, the server solves the following joint allocation problem:

$$[\mathcal{P}1]: \quad \max_{\boldsymbol{x}^{(\ell)}, \boldsymbol{f}^{(\ell)}, \boldsymbol{p}^{(\ell)}} \quad SW_{\mathrm{E}}^{(\ell)} \quad \text{s.t. } (4), (6), (5), (7),\ x_i^{(\ell)} \in \{0,1\},\ \forall i.$$

4 Design of PABP

4.1 Overview of PABP

PABP consists of two steps, in which the server selects devices and allocates computation and communication resources based on perturbed bids by solving $[\mathcal{P}1]$. It also determines payment without observing original bids, while meeting the approximate expected incentive compatibility (IC) and individual rationality (IR) requirements in **Definition**s 1 and 2 for each device.

4.2 Solution of Problem $[\mathcal{P}1]$

Screening Devices That Meet the Deadline. Device i needs the least time at its maximum CPU frequency F_i and upload power P_i. It can meet the round deadline only if $T_i^{\mathrm{tr},(\ell)}(F_i) + T_i^{\mathrm{up},(\ell)}(P_i) \leq T^d$. The candidate set $\mathcal{E}^{(\ell)}$ contains devices with a known $\alpha_i^{\ell-1}$ that satisfy this condition.

Minimum-Cost Resource Allocation for Each Candidate Device. For feasible training and upload durations, the minimum cost of selecting device i is

$$\mathrm{RC}_i^{(\ell)} = \min_{t_i^{\mathrm{tr}}, t_i^{\mathrm{up}}} \mathcal{C}_i^{(\ell)}(t_i^{\mathrm{tr}}, t_i^{\mathrm{up}}). \tag{15}$$

The total cost $\mathcal{C}_i^{(\ell)}$ is convex for positive durations (Boyd & Vandenberghe, 2004). Its expression and the feasibility constraints for (15) are given in Appendix A.2. Let $\bar{t}_i^{\mathrm{tr},(\ell)}$ and $\bar{t}_i^{\mathrm{up},(\ell)}$ attain this minimum. Inverting the training- and upload-time formulas gives the corresponding CPU frequency $\bar{f}_i^{(\ell)}$ and upload power $\bar{p}_i^{(\ell)}$. For each device, the minimum resource cost and its resource configuration conditional on selection depend on its workload, channel state, resource limits, and deadline, not on its bid.

Resource Costs in Device Selection. Because the server observes only perturbed bids, it estimates the value of device i's model-quality contribution as $\phi_i^{(\ell)} = \alpha_i^{\ell-1}\widetilde{b}_i^{(\ell)}$, Selecting device i thus increases noisy executable welfare by the net score $s_i^{(\ell)} = \phi_i^{(\ell)} - \mathrm{RC}_i^{(\ell)} = \alpha_i^{\ell-1}\widetilde{b}_i^{(\ell)} - \mathrm{RC}_i^{(\ell)}$. Therefore, device selection in a regular round reduces to

$$[\mathcal{P}2]: \quad \max_{\boldsymbol{x}^{(\ell)}} \quad \sum_{i\in\mathcal{E}^{(\ell)}} x_i^{(\ell)} s_i^{(\ell)} \\ \text{s.t.} \quad x_i^{(\ell)} \in \{0,1\}, \quad \forall i \in \mathcal{E}^{(\ell)}, \\ \sum_{i\in\mathcal{E}^{(\ell)}} x_i^{(\ell)} \le K. \tag{16}$$

We solve (16) by branch-and-bound. Algorithm 1 in Appendix A.3 gives the bounds, branching, and pruning rules. The solution derivations are given in Appendix A.2.

Proposition 1 (Optimality of Joint Allocation for Fixed Perturbed Bids). For any perturbed bids, Algorithm 1 obtains the globally optimal results for $[\mathcal{P}2]$, which is denoted as $\left(\boldsymbol{x}^{*(\ell)}, \boldsymbol{f}^{*(\ell)}, \boldsymbol{p}^{*(\ell)}\right)$. The complete proof is provided in Appendix B.1.

4.3 Myerson-Calibrated Payment (MCP)

When device i reports z, its expected winning probability is

$$q_i(z; b_{-i}) = \mathbb{E}_{\boldsymbol{\eta}}\left[x_i^*(\widetilde{\boldsymbol{b}}) \mid b_i = z, b_{-i}\right]. \tag{17}$$

Then, the expected quality-weighted allocation quantity is defined as

$$Q_i(z; b_{-i}) = \alpha_i q_i(z; b_{-i}). \tag{18}$$

Myerson Target Service Payment. Since model training and upload incur costs, based on Myerson's theorem, we define the target service payment π_i^{ref} that exactly satisfies expected IC/IR as

$$\pi_i^{\mathrm{ref}}(b_i; b_{-i}) = b_i Q_i(b_i; b_{-i}) - \int_0^{b_i} Q_i(z; b_{-i})\,\mathrm{d}z - \mathrm{RC}_i q_i(b_i; b_{-i}), \tag{19}$$

where $\mathrm{RC}_i q_i(b_i; b_{-i})$ is the expected resource cost. The IC/IR derivation is provided in Appendix B.4.

From the Expected Target Payment to the Actual Payment. The target payment in (19) cannot be used for online settlement: the server observes only the perturbed bid vector $\widetilde{\boldsymbol{b}}$ and the current allocation, not the original bid or the expectations q_i and Q_i. For a single perturbed bid, define the payment before calibration as

$$\pi_i^{\mathrm{pre}}(\widetilde{\boldsymbol{b}}) = \widetilde{b}_i \alpha_i x_i^*(\widetilde{\boldsymbol{b}}) - \int_0^{\widetilde{b}_i} \alpha_i x_i^*(u, \widetilde{b}_{-i})\,\mathrm{d}u - x_i^*(\widetilde{\boldsymbol{b}})\mathrm{RC}_i, \tag{20}$$

where u replaces device i's perturbed bid, and $x_i^*(u, \widetilde{b}_{-i})$ is the optimal allocation under that bid, obtained from (16). Because the original and perturbed bids differ, the noise expectation of π_i^{pre} need not equal the target payment. We therefore add a correction based on the server-visible $\alpha_i\widetilde{b}_i$ for selected devices, i.e.,

$$\pi_i(\widetilde{\boldsymbol{b}}; \boldsymbol{\theta}_i) = \pi_i^{\mathrm{pre}}(\widetilde{\boldsymbol{b}}) + x_i^*(\widetilde{\boldsymbol{b}})\left(\theta_{i1} + \theta_{i2}\alpha_i\widetilde{b}_i\right). \tag{21}$$

The correction parameter is $\boldsymbol{\theta}_i = (\theta_{i1}, \theta_{i2})^{\mathsf{T}}$. The constant θ_{i1} adjusts the charge after selection, while θ_{i2} controls its dependence on $\alpha_i\widetilde{b}_i$. Appendix C.5 compares the payment errors with and without this correction.

Pre-Bid Payment Calibration and Actual Service Payment. To estimate the target payment and the conditional expected payment under (21), sample bids and noises are randomly drawn from the distributions $[0, B^{\text{bid}}]$ and $\mathcal{N}(0, (\sigma_i^{(\ell)})^2)$, respectively. The correction $\theta_{i1}^{(\ell)} + \theta_{i2}^{(\ell)}\phi$ must remain nonnegative for every $0 \le \phi \le \alpha_i^{\ell-1} B^{\text{bid}}$, so a selected device's service payment is not reduced by calibration. This is a linear program, and its formulation and estimations are in Appendix A.5. With perturbed bids and allocation results, the server calculates device i's actual service payment as

$$\pi_i^{(\ell)}(\widetilde{\boldsymbol{b}}^{(\ell)}) = \pi_i^{\text{pre},(\ell)}(\widetilde{\boldsymbol{b}}^{(\ell)}) + x_i^{*(\ell)}(\widetilde{\boldsymbol{b}}^{(\ell)}) \left(\theta_{i1}^{*(\ell)} + \theta_{i2}^{*(\ell)} \phi_i^{(\ell)}\right). \tag{22}$$

The actual service payment in (22) is nonnegative for selected devices and zero for unselected devices. The complete proof is provided in Appendix B.2.

4.4 Theoretical Properties and Performance Guarantees

Lemma 1 (Allocation Monotonicity in a Realized State). For fixed $\widetilde{b}_{-i}$, the exact solution of (16) selects device i monotonically in its perturbed bid $\widetilde{b}_i$. Proof in Appendix B.3.

Corollary 1 (Monotonicity of the Expected Effective Allocation). For fixed b_{-i}, couple two reports with the same bid-independent noise. Then Lemma 1 implies that the expected winning probability $q_i(z; b_{-i})$ and expected quality-weighted allocation quantity $Q_i(z; b_{-i})$ are nondecreasing in the pre-perturbation bid z. Proof in Appendix B.3.

Theorem 1 (Approximate Expected IC and IR). For fixed b_{-i}, Corollary 4.4 and the target payment π_i^{ref} yield exact expected IC and IR. If the conditional expected deviation of the actual payment from this target is at most δ_i over the full bid domain, as defined in (33), then it satisfies $2\delta_i$-approximate expected IC and δ_i-approximate expected IR. The proof is in Appendix B.4.

PABP's Expected Social Welfare Ratio. Let $SW^{(\ell)}$ be the social welfare obtained by PABP in round ℓ, and let $SW_{\text{FI}}^{*,(\ell)}$ be the maximum social welfare when the true values are available. Under truthful bidding, define $\gamma_{\text{SW}}^{(\ell)} = \mathbb{E}_{\boldsymbol{\eta}^{(\ell)}}[SW^{(\ell)}]/SW_{\text{FI}}^{*,(\ell)}$ as the approximation ratio. The explicit formula and numerical estimation are given in Appendix B.5.

5 EXPERIMENTAL EVALUATIONS

5.1 EXPERIMENTAL SETUP

Datasets & Settings. We conduct experiments on three commonly used image classification datasets in FL: MNIST(LeCun et al., 1998), Fashion-MNIST(Xiao et al., 2017), and CIFAR-10(Krizhevsky & Hinton, 2009). The global model is the VGG9. We use truthful bids sampled uniformly from $[0, 1]$ and add Gaussian noise with a standard deviation of 0.12 before bid submission.

Baselines & Metrics. We compare PABP with Non-private auction, IAFL-adapted based on IAFL(Wu et al., 2024), and QI-DPFL-adapted based on QI-DPFL(Yuan & Wang, 2024). We use test accuracy to measure learning performance, true-bid inference error to measure bid protection, and devices' expected utilities under different bids to evaluate approximate IC and approximate IR.

5.2 Main Comparison Experiments

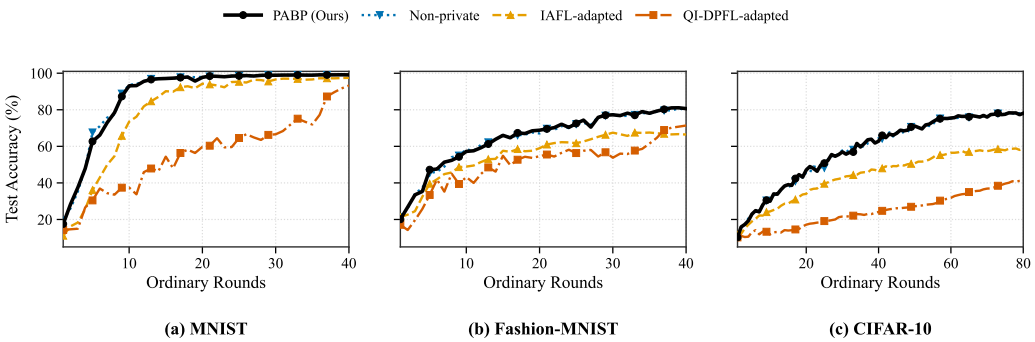

Figure 2: Learning curves under the common capacity-constrained protocol.

Table 1: Test accuracy and server-visible bids.

Method	Bid visible to server	Test accuracy (%)		
		MNIST	Fashion-MNIST	CIFAR-10
PABP (Ours)	**Noise-perturbed bid**	**99.12**	**82.11**	**79.75**
Non-private auction	True bid	99.09	82.04	79.34
IAFL-adapted	No monetary bid	97.53	69.76	59.81
QI-DPFL-adapted	True bid; perturbed model update	93.26	71.34	41.37

Figure 2 and Table 1 show the global-model test-accuracy curves and final values of PABP, Non-private auction, IAFL-adapted, and QI-DPFL-adapted on MNIST, Fashion-MNIST, and CIFAR-10, respectively. PABP outperforms both adapted methods on all three datasets, with only a small difference from Non-private auction. This is because QI-DPFL-adapted perturbs the model updates used for aggregation, while IAFL-adapted does not use monetary bids. In contrast, PABP perturbs only the bids and does not directly modify the model updates. These results indicate that PABP preserves learning performance while protecting bids and outperforms both adapted methods under the same capacity and resource constraints.

5.3 Server Attack for Inferring True Bids

Figure 3 shows a white-box attack experiment, where the server uses known information to infer each device's true fixed bid. As shown, PABP remains in the range 0.1040–0.1059 on the three datasets, whereas Non-private auction remains at zero. These results show that PABP substantially reduces the server's ability to infer true bids, confirming the effectiveness of bid perturbation.

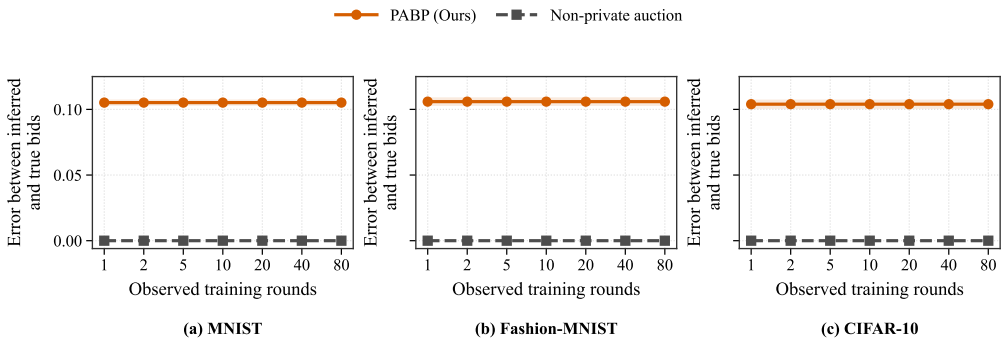

Figure 3: Server inference error for true bids.

5.4 Approximate Incentive Compatibility and Individual Rationality

Approximate Incentive Compatibility. Figure 4 compares users' expected utilities across reported bids under the theoretical target payment and the MCP payment actually executed by PABP. The target payment serves as the exact expected-IC benchmark, under which expected utility is maximized at the truthful bid $b_i = v_i$. MCP is the payment rule that PABP executes under a single perturbed bid realization; its utility curve remains nearly maximized at the truthful bid, and the maximum additional utility attainable through misreporting is small. This shows that the pre-bid calibration preserves the incentive structure of the target payment, yielding the approximate expected-IC guarantee in Definition 1. The complete audit results are provided in Appendix Table 7.

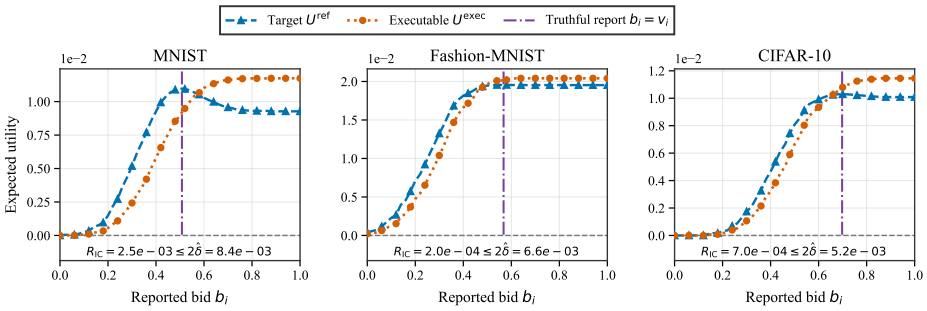

Figure 4: Expected utility versus reported bid under target and executable payments. The purple dashed line marks truthful bidding ($b_i = v_i$), at which the target-payment utility under exact IC attains its maximum.

Approximate Individual Rationality. Figure 5 compares the expected utility under truthful bidding with the approximate IR lower bound $-\widehat{\delta}_i^{\mathrm{val}}$. In all audit cases across the three datasets, the expected utility under truthful bidding is no lower than the minimum utility permitted by approximate individual rationality, $-\widehat{\delta}_i^{\mathrm{val}}$. This result is consistent with the approximate expected IR guarantee in Definition 2 and indicates that the experiments found no utility loss from truthful participation beyond the allowed range. The numbers of cases and the audit results are provided in Appendix Table 7.

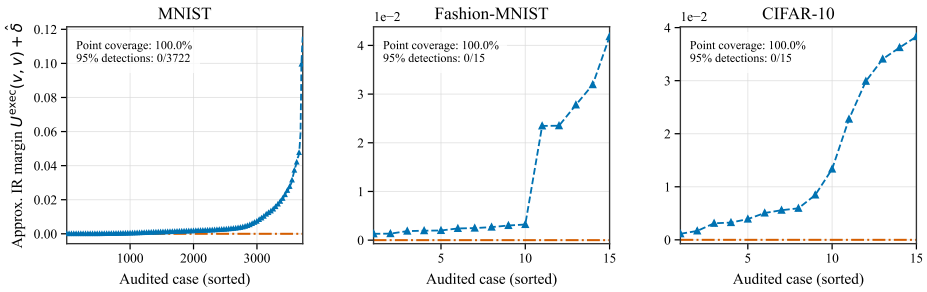

Figure 5: Expected utility under truthful bidding and the approximate IR lower bound.

6 CONCLUSION

This paper proposed PABP, an auction mechanism for bid privacy protection in FL. It combines perturbed-bid allocation with MCP payment and provides approximate expected IC, IR, and social-welfare guarantees. Experiments show that PABP protects true bids while preserving model accuracy.

7 AI USE STATEMENT

In this work, we used generative AI tools to assist with paper polish and code writing. We did not use generative AI tools for paper content generation. Additionally, we have reviewed all AI-assisted work. For example, LLM-generated code was verified and tested for correctness by at least 2 authors. We take responsibility for the final content of this work, including text, claims, or content produced with the aid of generative AI.

8 ETHICS STATEMENT

This work adheres to the ICLR Code of Ethics. Our research does not involve human subjects, animal experiments, or the collection of sensitive personal data. The datasets used are publicly available benchmarks. We have considered the potential societal impacts and do not foresee risks of misuse, privacy violations, or discrimination.

9 REPRODUCIBILITY STATEMENT

All datasets used in this work are publicly available. The datasets, model hyperparameters, computational resources, and other experimental settings are described in the Experimental Setup section. We believe that these materials are sufficient to reproduce the main findings of this paper.

REFERENCES

Pierre Bonami, Lorenz T. Biegler, Andrew R. Conn, Gerard Cornuejols, Ignacio E. Grossmann, Carl D. Laird, Jon Lee, Andrea Lodi, Francois Margot, Nicolas Sawaya, and Andreas Wachter. An algorithmic framework for convex mixed integer nonlinear programs. *Discrete Optimization*, 5(2):186–204, 2008. doi: 10.1016/j.disopt.2006.10.011.

Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography Conference*, pp. 635–658, 2016.

Nicolas Kourtellis, Kleomenis Katevas, and Diego Perino. FLaaS: Federated learning as a service. In *Proceedings of the 1st Workshop on Distributed Machine Learning*, pp. 7–13. ACM, 2020. doi: 10.1145/3426745.3431337.

Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.

Ailsa H. Land and Alison G. Doig. An automatic method of solving discrete programming problems. *Econometrica*, 28(3):497–520, 1960.

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

G. Li, J. Cai, J. Lu, and H. Chen. Incentive mechanism design for cross-device federated learning: A reinforcement auction approach. *IEEE Transactions on Mobile Computing*, 24(4):3059–3075, 2025.

G. Li, J. Cai, L. You, and X. Zhang. A sustainable incentive mechanism for long-term cross-device federated learning with energy limits and fairness. *IEEE Transactions on Mobile Computing*, 25(7):9487–9504, 2026a.

Xiang Li, Bing Luo, Jianwei Huang, and Yuan Luo. Mechanism design for federated learning with non-monotonic network effects. *IEEE Transactions on Mobile Computing*, pp. 1–13, 2026b. doi: 10.1109/TMC.2026.3717853.

Zitao Li, Bolin Ding, Liuyi Yao, Yaliang Li, Xiaokui Xiao, and Jingren Zhou. Performance-based pricing for federated learning via auction. *Proceedings of the VLDB Endowment*, 17(6):1269–1282, 2024.

Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pp. 1273–1282. PMLR, 2017. URL `https://proceedings.mlr.press/v54/mcmahan17a.html`.

Takayuki Nishio and Ryo Yonetani. Client selection for federated learning with heterogeneous resources in mobile edge. In *2019 IEEE International Conference on Communications (ICC)*, pp. 1–7. IEEE, 2019. doi: 10.1109/ICC.2019.8761315.

Jinlong Pang, Jieling Yu, Ruiting Zhou, and John C. S. Lui. An incentive auction for heterogeneous client selection in federated learning. *IEEE Transactions on Mobile Computing*, 22(10):5733–5750, 2023. doi: 10.1109/TMC.2022.3182876.

Peng Sun, Xu Chen, Guocheng Liao, and Jianwei Huang. A profit-maximizing model marketplace with differentially private federated learning. In *IEEE INFOCOM 2022 – IEEE Conference on Computer Communications*, pp. 1439–1448, 2022. doi: 10.1109/INFOCOM48880.2022.9796833.

Zhaoxuan Wu, Mohammad Mohammadi Amiri, Ramesh Raskar, and Bryan Kian Hsiang Low. Incentive-aware federated learning with training-time model rewards. In *The Twelfth International Conference on Learning Representations*, 2024.

Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-MNIST: A novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.

Wenhao Yuan and Xuehe Wang. QI-DPFL: Quality-aware and incentive-boosted federated learning with differential privacy. *arXiv preprint arXiv:2404.08261*, 2024. To appear in the 2024 International Joint Conference on Neural Networks.

Appendix

Contents

1 Introduction **1**

2 BACKGROUND AND RELATED WORK **2**
2.1 Auction Design in FL 2
2.2 Incentive Design with Privacy in FL 2

3 System Model and Problem Formulation **2**
3.1 Incentive Mechanism Model for FL 3
3.2 Device and Server Utilities 4

4 Design of PABP **5**
4.1 Overview of PABP 5
4.2 Solution of Problem [$\mathcal{P}1$] 5
4.3 Myerson-Calibrated Payment (MCP) 6
4.4 Theoretical Properties and Performance Guarantees 7

5 EXPERIMENTAL EVALUATIONS **7**
5.1 EXPERIMENTAL SETUP 7
5.2 Main Comparison Experiments 8
5.3 Server Attack for Inferring True Bids 8
5.4 Approximate Incentive Compatibility and Individual Rationality 9

6 CONCLUSION **9**

7 AI use statement **10**

8 Ethics statement **10**

9 Reproducibility statement **10**

A Additional Details of PABP **14**
A.1 Major Notations 14
A.2 Minimum-Cost Resource Allocation and Recovery 14
A.3 Branch-and-Bound Device Selection and Resource Allocation 15
A.4 Complete Round-Wise PABP Procedure 17
A.5 Pre-Bid Payment Calibration 17
A.6 Formal Privacy Analysis of Bid Randomization 18

B Theoretical Proofs and Welfare Evaluation **19**

B.1 Optimality of Joint Allocation for Fixed Perturbed Bids 19
B.2 Nonnegativity of the Executable Service Payment 19
B.3 Monotonicity of Allocation 19
B.4 Approximate Expected Incentive Compatibility and Individual Rationality . . 20
B.5 Numerical Estimation and Cross-Round Aggregation of the True Welfare Ratio 21

C Additional Experiments and Results **22**
C.1 Experimental Setup 22
C.2 Model Performance under Bid Perturbation 22
C.3 Main Comparisons and Stability 23
C.4 Welfare Retention 24
C.5 Payment Calibration and Approximate IC/IR Audit 24
C.6 Label Heterogeneity and IAFL Adaptation 25
C.7 Ablation Study 26
C.8 Solver Correctness Checks 27

A Additional Details of PABP

This appendix collects the notation, resource-allocation details, algorithms, payment calibration, and bid-privacy analysis used in the main paper.

A.1 Major Notations

Table 2 summarizes the major notations used in this paper.

Table 2: Major Notations

Notation	Interpretation
$\mathcal{N} = \{1, \dots, N\}$	Set of devices, $N = \lvert\mathcal{N}\rvert$, indexed by $i \in \mathcal{N}$
$\mathcal{L} = \{1, \dots, L\}$	Set of global training rounds, indexed by $\ell \in \mathcal{L}$
i, ℓ	Device and training-round indices
$H, \mathcal{R}$	Refresh interval and set of refresh rounds
$x_i^{(\ell)}$	Device selection variable in a regular round
$f_i^{(\ell)}, p_i^{(\ell)}$	CPU frequency and upload power
$g_i, A_i^{(\ell)}$	CPU cycles per sample and total CPU cycles in round ℓ
$n_i^{(\ell)}, \epsilon_i^{(\ell)}$	Number of local iterations and local training accuracy parameter
$t_i^{\mathrm{tr}}, t_i^{\mathrm{up}}$	Training and upload times in a fixed round
$B_i, h_i^{(\ell)}$	Uplink bandwidth and channel gain
F_i, P_i, T^d	Maximum CPU frequency, maximum upload power, and round deadline
$\pi_i^{(\ell)}$	Service payment from device i to the server
c_i	Monetary cost per unit of energy
$v_i^{(\ell)}$	True monetary value per unit of model-quality contribution
$\boldsymbol{b}^{(\ell)} = (b_i^{(\ell)})_{i \in \mathcal{N}}$	Vector of original bids, $\boldsymbol{b}^{(\ell)} \in \mathbf{B}^N$
$b_i^{(\ell)}, \widetilde{b}_i^{(\ell)}$	Original bid and perturbed bid of device i
$\eta_i^{(\ell)}, \sigma_i^{(\ell)}$	Gaussian noise sample and its standard deviation
$\mathbf{B} = [0, B^{\mathrm{bid}}]$	Bid range with upper bound B^{bid}
$\mathcal{D}^{\mathrm{val}}$	Server validation set
α_i^{ℓ}	Model-quality contribution coefficient of device i in round ℓ
$x_i^{(\ell)} \alpha_i^{\ell-1}$	Quality-weighted allocation in a regular round
$K, \mathcal{E}^{(\ell)}$	Maximum number of selected devices and set of eligible devices
$\mathbf{v} = (\mathcal{I}_{\mathbf{v}}, \mathcal{O}_{\mathbf{v}}, \mathcal{U}_{\mathbf{v}})$	Branch-and-bound node and its selected, excluded, and undecided sets
$\phi_i^{(\ell)}$	Estimated value based on the perturbed bid
$\mathrm{RC}_i^{(\ell)}, s_i^{(\ell)}$	Minimum resource cost and net score
q_i, Q_i	Expected winning probability and quality-weighted allocation
$\pi_i^{\mathrm{ref}}, \pi_i^{\mathrm{pre}}, \pi_i$	Target service payment, payment before calibration, and actual payment
$\boldsymbol{\theta}_i$	Payment parameters fixed before bidding
$\widehat{\delta}_i, \delta_i, \widehat{\delta}_i^{\mathrm{val}}$	Calibration, theoretical, and validation payment errors
S, G, M	Numbers of bid vectors, integration intervals, and noise samples
M_w, M_u	Numbers of noise samples for welfare and utility estimates
$T_i^{\mathrm{tr},(\ell)}, T_i^{\mathrm{up},(\ell)}$	Training and upload times

A.2 Minimum-Cost Resource Allocation and Recovery

Combining the training and upload costs in Section 3, the total cost of device i in the time domain is

$$\mathcal{C}_i^{(\ell)}(t_i^{\mathrm{tr}}, t_i^{\mathrm{up}}) = c_i \left[\frac{\kappa_i (A_i^{(\ell)})^3}{(t_i^{\mathrm{tr}})^2} + \frac{N_0 t_i^{\mathrm{up}}}{h_i^{(\ell)}} \left(2^{S_i/(B_i t_i^{\mathrm{up}})} - 1 \right) \right]. \tag{23}$$

The minimum cost over feasible training and upload durations is

$$\mathrm{RC}_i^{(\ell)} = \min_{\substack{t_i^{\mathrm{tr}} \geq T_i^{\mathrm{tr},(\ell)}(F_i),\ t_i^{\mathrm{up}} \geq T_i^{\mathrm{up},(\ell)}(P_i) \\ t_i^{\mathrm{tr}} + t_i^{\mathrm{up}} \leq T^d}} \mathcal{C}_i^{(\ell)}(t_i^{\mathrm{tr}}, t_i^{\mathrm{up}}). \tag{24}$$

For positive durations, the training-cost and upload-cost terms are convex functions of their respective durations. The feasible time constraints in (24) are linear. Therefore, (24) is a convex optimization problem.

If $\bar{t}_i^{\mathrm{tr},(\ell)}$ and $\bar{t}_i^{\mathrm{up},(\ell)}$ attain the minimum in (24), the corresponding CPU frequency and upload power are

$$\bar{f}_i^{(\ell)} = \frac{A_i^{(\ell)}}{\bar{t}_i^{\mathrm{tr},(\ell)}}, \qquad \bar{p}_i^{(\ell)} = \frac{N_0}{h_i^{(\ell)}} \left[2^{S_i/(B_i \bar{t}_i^{\mathrm{up},(\ell)})} - 1 \right]. \tag{25}$$

After solving (16), assign each device's resource variables by

$$\left(t_i^{\mathrm{tr},*(\ell)}, t_i^{\mathrm{up},*(\ell)}, f_i^{*(\ell)}, p_i^{*(\ell)}\right) = x_i^{*(\ell)} \left(\bar{t}_i^{\mathrm{tr},(\ell)}, \bar{t}_i^{\mathrm{up},(\ell)}, \bar{f}_i^{(\ell)}, \bar{p}_i^{(\ell)}\right). \tag{26}$$

A.3 Branch-and-Bound Device Selection and Resource Allocation

Section 4.2 formulates the device-selection problem $[\mathcal{P}2]$ in (16). Here we give the branch-and-bound details for solving it. Each search node is $\mathbf{v} = (\mathcal{I}_\mathbf{v}, \mathcal{O}_\mathbf{v}, \mathcal{U}_\mathbf{v})$, where the three disjoint sets contain the selected, excluded, and undecided devices, respectively. The objective contribution fixed at node $\mathbf{v}$ is

$$G_\mathbf{v} = \sum_{i \in \mathcal{I}_\mathbf{v}} s_i^{(\ell)}. \tag{27}$$

If $|\mathcal{I}_\mathbf{v}| > K$, the node violates the capacity constraint and is discarded. Otherwise, setting the devices in $\mathcal{U}_\mathbf{v}$ to zero gives a complete feasible allocation with value $G_\mathbf{v}$ and may improve the incumbent lower bound LB. An upper bound for the node is

$$UB_\mathbf{v} = \begin{cases} G_\mathbf{v} + \displaystyle\sum_{\substack{i \in \mathcal{U}_\mathbf{v} \\ s_i^{(\ell)} > 0}} s_i^{(\ell)}, & |\mathcal{I}_\mathbf{v}| < K, \\ G_\mathbf{v}, & |\mathcal{I}_\mathbf{v}| = K. \end{cases} \tag{28}$$

The first case includes every positive score that an undecided device could add, so it is no smaller than the value of any feasible descendant. Hence a node with $UB_\mathbf{v} \le LB$ can be pruned. Otherwise, the algorithm selects an undecided device j with the largest score and creates one child that adds j to $\mathcal{I}_\mathbf{v}$ and another that adds j to $\mathcal{O}_\mathbf{v}$(Land & Doig, 1960; Bonami et al., 2008).

Algorithm 1 solves (16). It identifies eligible devices and computes their minimum resource costs before applying branch-and-bound. The resource allocation is then obtained from (26).

Algorithm 1: Branch-and-bound device selection and resource allocation in a regular round

Input: $\alpha_i^{\ell-1}$, $\widetilde{b}_i^{(\ell)}$, K, T^d, and the round-ℓ resource parameters
Output: $\boldsymbol{x}^{*(\ell)}$, $t_i^{\mathrm{tr},*(\ell)}$, $t_i^{\mathrm{up},*(\ell)}$, $f_i^{*(\ell)}$, and $p_i^{*(\ell)}$

Construct $\mathcal{E}^{(\ell)}$ from devices with known $\alpha_i^{\ell-1}$ that pass the maximum-resource deadline test;
foreach $i \in \mathcal{E}^{(\ell)}$ **do**
  Solve (24) and store $\mathrm{RC}_i^{(\ell)}$, $\bar{t}_i^{\mathrm{tr},(\ell)}$, and $\bar{t}_i^{\mathrm{up},(\ell)}$;
  Compute $\bar{f}_i^{(\ell)}$ and $\bar{p}_i^{(\ell)}$ using (25);
  Compute $s_i^{(\ell)} = \alpha_i^{\ell-1}\widetilde{b}_i^{(\ell)} - \mathrm{RC}_i^{(\ell)}$;
end
Initialize the all-zero incumbent allocation and set $LB \leftarrow 0$;
Create $\mathbf{v}_0 = (\varnothing, \varnothing, \mathcal{E}^{(\ell)})$, compute $UB_{\mathbf{v}_0}$ using (28), and set $\mathcal{A} \leftarrow \{\mathbf{v}_0\}$;
while $\mathcal{A} \neq \varnothing$ **do**
  Remove a node $\mathbf{v}$ with the largest upper bound from $\mathcal{A}$; break ties by generation order;
  if $|\mathcal{I}_{\mathbf{v}}| > K$ **then**
    Discard $\mathbf{v}$ and **continue**;
  end
  Form a complete allocation by selecting $\mathcal{I}_{\mathbf{v}}$ and excluding $\mathcal{O}_{\mathbf{v}} \cup \mathcal{U}_{\mathbf{v}}$;
  if $G_{\mathbf{v}} > LB$ **then**
    Update the incumbent allocation and set $LB \leftarrow G_{\mathbf{v}}$;
  end
  Compute $UB_{\mathbf{v}}$ using (28);
  if $UB_{\mathbf{v}} \leq LB$ **then**
    Prune $\mathbf{v}$ and **continue**;
  end
  Select $j \in \mathcal{U}_{\mathbf{v}}$ with the largest $s_j^{(\ell)}$; break ties by the smallest device index;
  Create $\mathbf{v}^{\mathrm{in}} = (\mathcal{I}_{\mathbf{v}} \cup \{j\}, \mathcal{O}_{\mathbf{v}}, \mathcal{U}_{\mathbf{v}} \setminus \{j\})$ and $\mathbf{v}^{\mathrm{out}} = (\mathcal{I}_{\mathbf{v}}, \mathcal{O}_{\mathbf{v}} \cup \{j\}, \mathcal{U}_{\mathbf{v}} \setminus \{j\})$;
  Compute their upper bounds and insert both children into $\mathcal{A}$;
end
Recover the resource variables using (26);
return the incumbent device selection and its resource configuration;

A.4 Complete Round-Wise PABP Procedure

Algorithm 2: Round-wise federated learning incentive mechanism

Input: $\mathcal{N}$, L, $\mathcal{R}$, K, $\boldsymbol{w}^{(0)}$, training parameters, cost parameters, privacy parameters, and payment-calibration parameters

Output: $\boldsymbol{w}^{(L)}$ and the round-wise allocation, resource, model-quality-contribution, and payment records

Initialize the global model and leave every α_i^0 undefined;
for $\ell = 1, \ldots, L$ **do**
  Announce the task, training configuration, deadline, cost parameters, and current channel parameters;
  Keep the devices that pass the maximum-resource deadline test;
  if $\ell \in \mathcal{R}$ **then**
    Set $x_i^{(\ell)} \leftarrow 1$ for every willing and resource-feasible device and set $x_i^{(\ell)} \leftarrow 0$ otherwise;
    For every participating device, fix $x_i^{(\ell)} = 1$ and compute the minimum-cost resource configuration from (24);
    Compute the CPU frequency and upload power from (25);
    Let participating devices train and upload, evaluate their local models, and update α_i^ℓ using (1);
    Form $\boldsymbol{w}^{(\ell)}$ using the weighted aggregation rule in Section 3.1, set $\pi_i^{(\ell)} \leftarrow 0$, and deliver the aggregate model to participating devices;
  else
    Keep devices with an available $\alpha_i^{\ell-1}$ that can complete the task before the deadline, forming $\mathcal{E}^{(\ell)}$;
    Use pre-generated bid vectors and simulation noise to compute (29) and the sample means used in (30), solve (30), and fix $\boldsymbol{\theta}_i^{*(\ell)}$ before bidding;
    Each device locally samples the announced Gaussian noise for a new bid; if the original bid is unchanged, reuse the stored perturbed bid and submit $\widetilde{b}_i^{(\ell)}$ using the clipping rule in Section 3.2;
    Compute $\phi_i^{(\ell)} = \alpha_i^{\ell-1}\widetilde{b}_i^{(\ell)}$ and call Algorithm 1;
    Selected devices train and upload according to the joint allocation;
    Evaluate local models, update α_i^ℓ using (1), and form $\boldsymbol{w}^{(\ell)}$ using the weighted aggregation rule in Section 3.1;
    Compute $\pi_i^{(\ell)}$ from (22), settle payments, and deliver the aggregate model to selected devices;
  end
end
return the final global model and all round-wise allocation, resource, model-quality-contribution, and payment records;

A.5 Pre-Bid Payment Calibration

The target payment π_i^{ref} involves an expectation over bid noise and an integral. It therefore cannot be computed from a single observed perturbed bid. Before bidding, the server generates S bid vectors $\boldsymbol{b}_r$ and M noise samples for each vector. For device i, the interval $[0, b_{i,r}]$ is divided into G subintervals. At each grid point $z_{i,r,g}$, the server changes only device i's bid and reuses the same noise samples. The clipping rule yields $\widetilde{\boldsymbol{b}}_{r,g,m}$. Algorithm 1 computes the allocation used to estimate the target and actual payments. These estimates determine the MCP parameters before bidding. Appendix C.1 gives the sample sizes.

The trapezoidal rule estimates the target payment π_i^{ref} at $\boldsymbol{b}_r$ as

$$\widehat{\pi}_{i,r}^{\mathrm{ref},(\ell)} = b_{i,r}\widehat{Q}_{i,r}(b_{i,r}) - \sum_{g=1}^{G} \frac{\widehat{Q}_{i,r}(z_{i,r,g-1}) + \widehat{Q}_{i,r}(z_{i,r,g})}{2} \left(z_{i,r,g} - z_{i,r,g-1}\right) - \mathrm{RC}_i^{(\ell)}\widehat{q}_{i,r}(b_{i,r}), \tag{29}$$

Here $\widehat{q}_{i,r}(z_{i,r,g}) = \frac{1}{M}\sum_{m=1}^{M} x_i^{*(\ell)}(\widetilde{\boldsymbol{b}}_{r,g,m})$ and $\widehat{Q}_{i,r}(z_{i,r,g}) = \alpha_i^{\ell-1}\widehat{q}_{i,r}(z_{i,r,g})$. To estimate the conditional expectation of (21) at $\boldsymbol{b}_r$, use the endpoint $g = G$ and write $\widetilde{\boldsymbol{b}}_{r,m} = \widetilde{\boldsymbol{b}}_{r,G,m}$. Set $\phi_i^{(\ell,r,m)} = \alpha_i^{\ell-1}\widetilde{b}_{i,r,m}$. The sample means are

$$\overline{\pi}_{i,r}^{\text{pre}} = \frac{1}{M}\sum_{m=1}^{M} \pi_i^{\text{pre},(\ell)}\big(\widetilde{\boldsymbol{b}}_{r,m}\big),$$

$$\overline{x}_{i,r} = \frac{1}{M}\sum_{m=1}^{M} x_i^{*(\ell)}\big(\widetilde{\boldsymbol{b}}_{r,m}\big),$$

$$\overline{x\phi}_{i,r} = \frac{1}{M}\sum_{m=1}^{M} x_i^{*(\ell)}\big(\widetilde{\boldsymbol{b}}_{r,m}\big)\phi_i^{(\ell,r,m)}.$$

They estimate the pre-calibration payment, winning probability, and winning outcome weighted by the server's value estimate, respectively.

The server uses these estimates to choose the correction parameters in (21) by solving

$$\min_{\boldsymbol{\theta}_i^{(\ell)},\widehat{\delta}_i^{(\ell)}} \quad \widehat{\delta}_i^{(\ell)} \tag{30a}$$

$$\text{s.t.} \quad \left|\overline{\pi}_{i,r}^{\text{pre}} + \theta_{i1}^{(\ell)}\overline{x}_{i,r} + \theta_{i2}^{(\ell)}\overline{x\phi}_{i,r} - \widehat{\pi}_{i,r}^{\text{ref},(\ell)}\right| \le \widehat{\delta}_i^{(\ell)}, \qquad r = 1,\dots,S, \tag{30b}$$

$$\theta_{i1}^{(\ell)} \ge 0, \tag{30c}$$

$$\theta_{i1}^{(\ell)} + \theta_{i2}^{(\ell)}\alpha_i^{\ell-1}B^{\text{bid}} \ge 0, \tag{30d}$$

$$\widehat{\delta}_i^{(\ell)} \ge 0. \tag{30e}$$

The objective minimizes the largest estimated payment deviation over the S calibration bids. The absolute-value constraint is equivalent to two linear inequalities. The correction is affine in ϕ. Its nonnegativity over $0 \le \phi \le \alpha_i^{\ell-1}B^{\text{bid}}$ therefore follows from the two endpoint constraints (30c) and (30d). Thus, (30) is a linear program. The resulting $\boldsymbol{\theta}_i^{*(\ell)}$ is fixed before devices bid.

A.6 Formal Privacy Analysis of Bid Randomization

This section analyzes the privacy guarantee for a single perturbed-bid release under the local randomization in Section 3.2. The guarantee compares output distributions for adjacent original bids.

For original bids $b_i, b_i' \in [0, B^{\text{bid}}]$, define two bids as adjacent when $|b_i - b_i'| \le \Delta_b$. If every pair in the bid domain is adjacent, then $\Delta_b = B^{\text{bid}}$. A mechanism satisfies ρ-zCDP if the Rényi divergence between its output distributions for any adjacent inputs is at most $\lambda\rho$ for every order $\lambda > 1$(Bun & Steinke, 2016).

Single Release. For a new original bid, the device releases $\mathcal{M}_i^{(\ell)}(b_i^{(\ell)}) = \Pi_{[0,B^{\text{bid}}]}(b_i^{(\ell)} + \eta_i^{(\ell)})$, where $\eta_i^{(\ell)} \sim \mathcal{N}(0, (\sigma_i^{(\ell)})^2)$ and $\sigma_i^{(\ell)} > 0$. The Gaussian mechanism, followed by interval projection, satisfies $\rho_i^{(\ell)}$-zCDP with

$$\rho_i^{(\ell)} = \frac{\Delta_b^2}{2(\sigma_i^{(\ell)})^2}. \tag{31}$$

Interval projection cannot increase the Rényi divergence. The server's allocation and payment are also post-processing of the perturbed bid when all other inputs to these computations are fixed or independent of the protected original bid. Under this condition, they do not increase $\rho_i^{(\ell)}$(Bun & Steinke, 2016).

Repeated Observation. In the fixed-bid setting of the server inference experiment, the device reuses the same perturbed bid in later rounds, as specified in Algorithm 2. Repeated copies of this single release do not increase the bound in (31), provided no new bid-dependent information is released. If the device instead draws fresh noise and releases another perturbed bid, the single-release bound does not cover the joint releases, which require composition accounting(Bun & Steinke, 2016).

Equation (31) does not measure reconstruction error under a particular attacker prior and estimator. The main-paper experiment therefore reports original-bid inference error separately and does not present it as a direct validation of the zCDP parameter.

B Theoretical Proofs and Welfare Evaluation

Except for the joint-allocation optimality proof, which retains round indices, the proofs below for Sections 4.3–4.4 fix a single regular round and omit round superscripts on b_i, $\widetilde{b}_i$, x_i^*, α_i, RC_i, and related functions.

B.1 Optimality of Joint Allocation for Fixed Perturbed Bids

Proposition 1. For any fixed vector of perturbed bids, Algorithm 1 finds a globally optimal device selection for (16). Equation (26) gives the corresponding optimal resource allocation. Together, they maximize $SW_{\mathrm{E}}^{(\ell)}$, the social welfare objective evaluated using perturbed bids. This result does not claim optimality for social welfare evaluated using true values.

Proof. Fix the vector of perturbed bids, so each net score $s_i^{(\ell)}$ is fixed. Each branch either selects or excludes one device. The search tree therefore covers all binary allocations. At a node $\mathbf{v}$, $G_{\mathbf{v}}$ is the objective value contributed by devices already selected. If fewer than K devices have been selected, $UB_{\mathbf{v}}$ adds every positive score among the undecided devices. It is therefore an upper bound on any feasible allocation below that node. If K devices have been selected, $UB_{\mathbf{v}} = G_{\mathbf{v}}$. The incumbent value LB comes from a feasible allocation, so a node with $UB_{\mathbf{v}} \le LB$ can be pruned. Because the candidate set is finite, the algorithm terminates with a globally optimal selection. For each selected device, (24) gives the minimum-cost feasible resource allocation. Equation (26) recovers this allocation. Hence the resulting selection and resource allocation maximize $SW_{\mathrm{E}}^{(\ell)}$.

B.2 Nonnegativity of the Executable Service Payment

Claim. The executable service payment in (22) is nonnegative for winning devices and zero for non-winning devices.

Proof. With other perturbed bids fixed, the path integral for the binary monotone allocation gives the winner's critical payment. The algorithm keeps only positive net-score improvements, so this payment covers $\mathrm{RC}_i^{(\ell)}$. The endpoint constraints (30c) and (30d) then imply that the service payment for a winner is nonnegative. For a losing device, monotonicity makes the allocation zero along the integration path. The additional term in (21) is also multiplied by $x_i^* = 0$, so the service payment is zero.

B.3 Monotonicity of Allocation

Lemma 1. Fix the other devices' perturbed bids, the model-quality contributions known at the start of the round, the resource parameters, the candidate set, and the deterministic tie-breaking rule. If (16) is solved exactly, device i' s allocation x_i^* is nondecreasing in its own perturbed bid $\widetilde{b}_i$.

Proof. Let device i have perturbed bids $\widetilde{b}_i^- \le \widetilde{b}_i^+$, scores $\phi_i^- \le \phi_i^+$, and optimal allocations $\boldsymbol{x}^-$ and $\boldsymbol{x}^+$. The candidate set, capacity, and minimum resource cost do not depend on the

bid. Only the coefficient of $x_i\phi_i$ changes. Adding the two optimality inequalities gives

$$(\phi_i^+ - \phi_i^-)\left(x_i^*(\widetilde{b}_i^+) - x_i^*(\widetilde{b}_i^-)\right) \geq 0. \tag{32}$$

If the score increases strictly, (32) rules out a change from 1 to 0. If the bids are equal, the optimization problems and deterministic tie-breaking rule are identical. Thus x_i^* is nondecreasing.

Corollary 1. Take two original bids $z^- \leq z^+$ and couple them with the same noise vector $\boldsymbol{\eta}$. The clipping map in Section 3.2 is nondecreasing, so $\widetilde{b}_i(z^-) \leq \widetilde{b}_i(z^+)$. Lemma 1 then gives

$$x_i^*(\widetilde{b}_i(z^-), \widetilde{b}_{-i}) \leq x_i^*(\widetilde{b}_i(z^+), \widetilde{b}_{-i}).$$

Taking expectations over all noise samples shows that the conditional winning probability $q_i(z; b_{-i})$ is nondecreasing in z. Since α_i is fixed before bidding, $Q_i = \alpha_i q_i$ is also nondecreasing.

B.4 Approximate Expected Incentive Compatibility and Individual Rationality

Proof Overview. Suppose the conditional expected payment differs from the target payment by at most $\delta_i^{(\ell)}$ for every bid profile. The corresponding expected utilities then differ by at most $\delta_i^{(\ell)}$ for any reported bid. Comparing truthful bidding with a deviation introduces at most two such errors. This gives $2\delta_i^{(\ell)}$-approximate expected IC in (11). Under truthful bidding, only one error enters the IR bound. This gives $\delta_i^{(\ell)}$-approximate expected IR in (12). These expectations are over bid noise for fixed bids, and the bounds hold for each device separately.

Formal Definition of Payment Error. The worst-case conditional expected deviation of the calibrated payment from the target payment over the full bid domain is defined as

$$\delta_i^{(\ell)} = \sup_{\boldsymbol{b} \in \mathbf{B}^N} \left| \mathbb{E}_{\boldsymbol{\eta}^{(\ell)}}\left[\pi_i^{(\ell)}(\widetilde{\boldsymbol{b}}) \mid \boldsymbol{b}\right] - \pi_i^{\mathrm{ref},(\ell)}(b_i; b_{-i}^{(\ell)}) \right|. \tag{33}$$

The supremum ranges over the full bid domain for each device. The expectation is over the mechanism's bid noise. Thus, $\delta_i^{(\ell)}$ bounds the expected payment error for each device. It is not an average across devices or a bound for every noise realization.

Fix α_i, resource parameters, and b_{-i}, and write $Q_i(z; b_{-i})$ as $Q_i(z)$. By (24) and (26), the winning resource cost when reporting b_i is $x_i^*\mathrm{RC}_i$, so the expected utility is

$$U_i^{\mathrm{exec}}(v_i, b_i; b_{-i}) = v_i Q_i(b_i) - \mathbb{E}\left[\pi_i(\widetilde{\boldsymbol{b}}) + x_i^*\mathrm{RC}_i \mid b_i, b_{-i}\right]. \tag{34}$$

Substituting the target payment π_i^{ref} and the device-borne expected resource cost gives

$$U_i^{\mathrm{ref}}(v_i, b_i; b_{-i}) = (v_i - b_i)Q_i(b_i) + \int_0^{b_i} Q_i(z)\,\mathrm{d}z. \tag{35}$$

Under truthful bidding $b_i = v_i$,

$$U_i^{\mathrm{ref}}(v_i, v_i; b_{-i}) = \int_0^{v_i} Q_i(z)\,\mathrm{d}z \geq 0, \tag{36}$$

Because $Q_i(z) \geq 0$, the target utility is nonnegative, so the target payment rule satisfies expected IR. To verify expected IC, suppose the device reports $b_i' < v_i$. From (35),

$$\begin{aligned} &U_i^{\mathrm{ref}}(v_i, v_i; b_{-i}) - U_i^{\mathrm{ref}}(v_i, b_i'; b_{-i}) \\ &= \int_{b_i'}^{v_i} \left[Q_i(z) - Q_i(b_i')\right]\,\mathrm{d}z \geq 0. \end{aligned} \tag{37}$$

On the integration interval $z \geq b_i'$, Corollary 1 makes the integrand nonnegative. If the device reports $b_i' > v_i$,

$$U_i^{\mathrm{ref}}(v_i, v_i; b_{-i}) - U_i^{\mathrm{ref}}(v_i, b_i'; b_{-i}) = \int_{v_i}^{b_i'} [Q_i(b_i') - Q_i(z)] \, \mathrm{d}z \geq 0. \tag{38}$$

The integrand is also nonnegative for overreporting. Therefore, truthful bidding maximizes expected utility under the target payment rule. Together with (36), this proves exact expected IC and IR. Truthful bidding need not be the unique optimal report.

We now prove approximate expected IC and IR for the executable service payment. Executable and target payments use the same expected allocation value and resource cost, so their utility difference comes only from the payment. Hence, for every report b_i,

$$\left| U_i^{\mathrm{exec}}(v_i, b_i; b_{-i}) - U_i^{\mathrm{ref}}(v_i, b_i; b_{-i}) \right| \leq \delta_i. \tag{39}$$

Applying (39) to the truthful bid and any alternative bid, and then using expected IC of the target payment rule, gives

$$\begin{aligned} U_i^{\mathrm{exec}}(v_i, v_i; b_{-i}) &\geq U_i^{\mathrm{ref}}(v_i, v_i; b_{-i}) - \delta_i \\ &\geq U_i^{\mathrm{ref}}(v_i, b_i'; b_{-i}) - \delta_i \\ &\geq U_i^{\mathrm{exec}}(v_i, b_i'; b_{-i}) - 2\delta_i. \end{aligned} \tag{40}$$

This proves $2\delta_i$-approximate expected IC. From (36) and (39),

$$U_i^{\mathrm{exec}}(v_i, v_i; b_{-i}) \geq U_i^{\mathrm{ref}}(v_i, v_i; b_{-i}) - \delta_i \geq -\delta_i, \tag{41}$$

Thus, the executable service payment satisfies δ_i-approximate expected IR.

B.5 Numerical Estimation and Cross-Round Aggregation of the True Welfare Ratio

Fix the candidate devices, true values, initial quality coefficients, and resource conditions in round ℓ. Only the bid noise varies. Define the feasible set of selections and resource allocations by

$$\begin{aligned} \mathcal{F}^{(\ell)} = \Bigg\{ (\boldsymbol{x}^{(\ell)}, \boldsymbol{f}^{(\ell)}, \boldsymbol{p}^{(\ell)}) \;\Bigg|\; & x_i^{(\ell)} \in \{0,1\}, \; \sum_{i=1}^{N} x_i^{(\ell)} \leq K, \\ & 0 \leq f_i^{(\ell)} \leq F_i, \quad 0 \leq p_i^{(\ell)} \leq P_i, \\ & T_i^{\mathrm{tr},(\ell)}(f_i^{(\ell)}) + T_i^{\mathrm{up},(\ell)}(p_i^{(\ell)}) \leq T^d, \; \forall i \Bigg\}. \end{aligned} \tag{42}$$

The full-information benchmark in the denominator is the computable optimum

$$SW_{\mathrm{FI}}^{*,(\ell)} = \max_{(\boldsymbol{x}^{(\ell)}, \boldsymbol{f}^{(\ell)}, \boldsymbol{p}^{(\ell)}) \in \mathcal{F}^{(\ell)}} \sum_{i=1}^{N} x_i^{(\ell)} \left[v_i^{(\ell)} \alpha_i^{\ell-1} - \mathcal{C}_i^{\mathrm{tr},(\ell)}(f_i^{(\ell)}) - \mathcal{C}_i^{\mathrm{up},(\ell)}(p_i^{(\ell)}) \right]. \tag{43}$$

For truthful pre-perturbation bids, the expected achieved social welfare ratio is

$$\gamma_{\mathrm{SW}}^{(\ell)} = \frac{\mathbb{E}_{\boldsymbol{\eta}^{(\ell)}} \left[\sum_{i=1}^{N} x_i^{*,(\ell)}\left(\widetilde{\boldsymbol{b}}^{(\ell)}\right) \left(v_i^{(\ell)} \alpha_i^{\ell-1} - \mathcal{C}_i^{\mathrm{tr},(\ell)}\left(f_i^{*,(\ell)}\left(\widetilde{\boldsymbol{b}}^{(\ell)}\right)\right) - \mathcal{C}_i^{\mathrm{up},(\ell)}\left(p_i^{*,(\ell)}\left(\widetilde{\boldsymbol{b}}^{(\ell)}\right)\right)\right)\right]}{SW_{\mathrm{FI}}^{*,(\ell)}}. \tag{44}$$

Here, $(\boldsymbol{x}^{*,(\ell)}, \boldsymbol{f}^{*,(\ell)}, \boldsymbol{p}^{*,(\ell)})$ is the allocation selected by PABP from $\widetilde{\boldsymbol{b}}^{(\ell)}$. The denominator is the optimum in (43), computed with the true values and the same feasible set. The numerator evaluates PABP's allocation using true values and its realized resource costs.

For M_w independent noise vectors, we estimate (44) by replacing its expectation with the sample mean.

For multiple regular rounds, we use the welfare-weighted aggregate ratio

$$\Gamma_{\mathrm{SW}} = \frac{\sum_{\ell \notin \mathcal{R}} \mathbb{E}_{\boldsymbol{\eta}^{(\ell)}} \left[SW^{(\ell)} \right]}{\sum_{\ell \notin \mathcal{R}} SW_{\mathrm{FI}}^{*,(\ell)}}. \tag{45}$$

The ratio aggregates expected welfare across regular rounds, weighted by each round's full-information optimum.

C Additional Experiments and Results

C.1 Experimental Setup

Table 3 reports the training budgets for the three main datasets. Data partitions, model-initialization seeds, mechanism noise, and diagnostic seeds are fixed in advance for each dataset. All methods on a given dataset share the same partition and initialization seed. The server validation set is not used for device training. Checkpoints are selected only by validation accuracy, and the test set is not used for model selection or hyperparameter tuning. All classification tasks use cross-entropy loss.

Table 3: Primary experimental settings.

Dataset	**Training budget**				**Monte Carlo samples**	
	Rounds	**Local epochs**	**Batch size**	**Learning rate**	M_w	M_u
MNIST	50	1	128	0.05	2000	1000
Fashion-MNIST	50	1	128	0.05	1000	500
CIFAR-10	100	2	64	0.02	1000	500

C.2 Model Performance under Bid Perturbation

Bid Noise and Test Accuracy. Figure 6 shows test accuracies of 98.56%, 98.65%, and 98.61% for $\sigma = 0$, 0.12, and 0.30, respectively. The differences are smaller than the variation across repeated runs and are consistent with the comparison between PABP and Non-private auction in Table 5. Thus, the tested levels of bid noise have little effect on test accuracy.

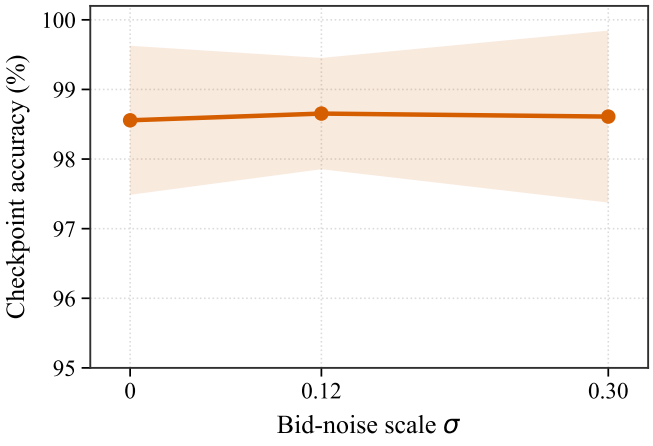

Figure 6: Test accuracy under different bid-noise scales.

The bid-noise sweep, label-heterogeneity sweep, and device-selection and resource-allocation ablation use 30 global rounds and 3 independent paired runs. We set $N = 20$, $K = 5$, the refresh interval to 5, and use VGG9 with the same resource model. The true-bid inference attack covers all three datasets. For each dataset, we generate 5 independent sets of true bids for 20 devices and resample the noise 100 times for each set. Table 4 specifies the conditions and purpose of these experiments.

Table 4: Sensitivity and ablation settings.

Purpose	**Dataset**	**Partition**	**Varied factor**	**Fixed operating point**
Learning utility	MNIST	Dir(0.4)	$\sigma \in \{0, 0.12, 0.30\}$	30 rounds, 3 runs
Inferring true bids	All three	5 independent 20-device bid sets	observed rounds 1–80	100 noise draws per set
Heterogeneity	Fashion-MNIST	Dir(0.1), Dir(0.4), IID	label partition	$\sigma = 0.12$
Selection/resource ablation	MNIST	Dir(0.4)	selection/resource component	$\sigma = 0.12$

Section 5.1 of the main paper gives the device limits, communication parameters, and resource-allocation method used in each round. Payment calibration uses 12 bid vectors generated before bidding, 25 integration grid points, and 128 independent noise samples per vector. For a given bid vector, the same noise samples are used at every grid point and in the endpoint payment estimate. Calibration, payment-execution, and validation noise are generated independently. The sample sizes M_w and M_u refer to welfare estimation and utility evaluation, respectively. Confidence intervals across training runs describe a separate source of variation.

Finite-Sample Validation of Payment Error. We draw N_{val} independent validation bid vectors and a separate set of noise samples. For each device and round, we compute

$$\widehat{\delta}_i^{\text{val},(\ell)} = \max_{1 \le k \le N_{\text{val}}} \Big| \widehat{\mathbb{E}}_{\text{val}} \Big[\pi_i^{(\ell)}(\widetilde{\boldsymbol{b}}^{(\ell)}) \mid \boldsymbol{b}^{(\ell)} = \boldsymbol{b}_{\text{val}}^{[k]} \Big] - \widehat{\pi}_{i,\text{val}}^{\text{ref},(\ell)} \left(b_{i,\text{val}}^{[k]}; b_{-i,\text{val}}^{[k]} \right) \Big|. \tag{46}$$

Here, $\widehat{\mathbb{E}}_{\text{val}}$ is the average over validation-noise samples. We estimate $\widehat{\pi}_{i,\text{val}}^{\text{ref},(\ell)}$ using the winning-probability estimate and the construction in (29). Equation (46) uses $N_{\text{val}} = 20$ bid vectors for each device and round. It measures payment error on this finite validation set. Without a probability bound extending the result to all bids, it cannot replace the theoretical $\delta_i^{(\ell)}$ in (33).

C.3 Main Comparisons and Stability

Table 5: Test accuracy and server-visible bids.

Method	**Bid visible to server**	**Test accuracy (%)**		
		MNIST	**Fashion-MNIST**	**CIFAR-10**
PABP (Ours)	**Noise-perturbed bid**	**99.12**	**82.11**	**79.75**
Non-private auction	True bid	99.09	82.04	79.34
IAFL-adapted	No monetary bid	97.53	69.76	59.81
QI-DPFL-adapted	True bid and perturbed model update	93.26	71.34	41.37

Table 6 reports the test-accuracy differences over five paired runs on each dataset. The mean difference between PABP and Non-private auction is at most 0.416 percentage points. PABP achieves higher accuracy than IAFL-adapted and QI-DPFL-adapted in all five runs on every dataset. Five runs are insufficient to establish statistical significance.

Table 6: Paired test-accuracy differences.

Dataset	Baseline	PABP − baseline		Exact p
		Δ accuracy (p.p.)	Wins (of 5)	
MNIST	Non-private	$+0.036_{\pm 0.096}$	3	0.4375
	IAFL-adapted	$+\mathbf{1.596}_{\pm 1.609}$	**5**	0.0625
	QI-DPFL-adapted	$+\mathbf{5.868}_{\pm 4.003}$	**5**	0.0625
Fashion-MNIST	Non-private	$+0.072_{\pm 1.165}$	2	0.9375
	IAFL-adapted	$+\mathbf{12.350}_{\pm 6.948}$	**5**	0.0625
	QI-DPFL-adapted	$+\mathbf{10.772}_{\pm 3.488}$	**5**	0.0625
CIFAR-10	Non-private	$+0.416_{\pm 0.558}$	4	0.1250
	IAFL-adapted	$+\mathbf{19.946}_{\pm 6.912}$	**5**	0.0625
	QI-DPFL-adapted	$+\mathbf{38.384}_{\pm 1.144}$	**5**	0.0625

C.4 Welfare Retention

Figure 7 compares the true social welfare of PABP with the full-information optimum. Aggregated over all regular rounds using (45), PABP retains 97.28%–97.82% of the optimal welfare on the three datasets.

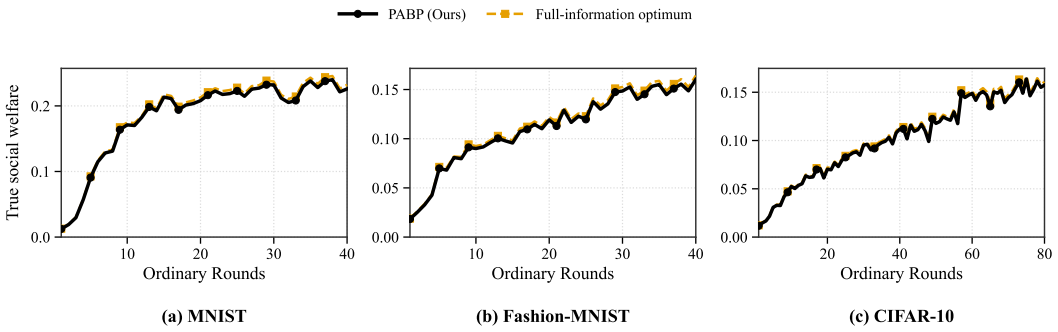

Figure 7: Round-wise true social welfare of PABP and the full-information optimum.

C.5 Payment Calibration and Approximate IC/IR Audit

Figure 8 compares payment errors on validation bids before and after calibration. The mean error decreases from 1.625×10^{-3} to 1.497×10^{-3} on MNIST (7.88%), from 1.174×10^{-3} to 1.046×10^{-3} on Fashion-MNIST (10.90%), and from 1.0025×10^{-3} to 8.9746×10^{-4} on CIFAR-10 (10.48%). Calibration reduces the error in 63.14%, 60.37%, and 62.06% of the device-round pairs, respectively. Thus, the calibrated payment in (22) is closer on average to the target payment than π_i^{pre}.

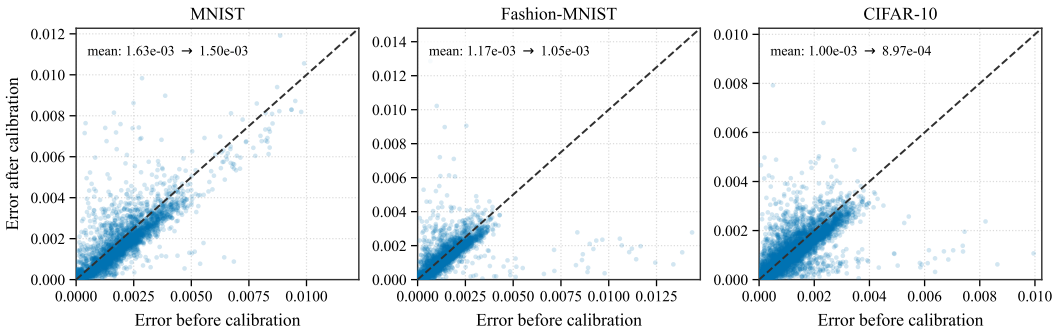

Figure 8: Held-out payment errors before and after payment calibration. The dashed line marks equal error.

Table 7: Approximate IC/IR audit.

Dataset	Cases	Approximate IC			Approximate IR
		Point coverage	95% compatible	Detections	Point coverage
MNIST	3722	99.758%	**99.973%**	1	**100%**
Fashion-MNIST	15	100%	**100%**	0	**100%**
CIFAR-10	15	93.333%	**100%**	0	**100%**

Figure 9 orders the cases in each dataset by $M_{\mathrm{IR}} = U_i^{\mathrm{exec}}(v_i, v_i) + \widehat{\delta}_i^{\mathrm{val}}$. A nonnegative margin means that expected utility under truthful bidding meets the empirical approximate IR bound. All 3,722 MNIST cases and all 15 cases on each of Fashion-MNIST and CIFAR-10 meet this bound. No violation is detected by the 95% intervals.

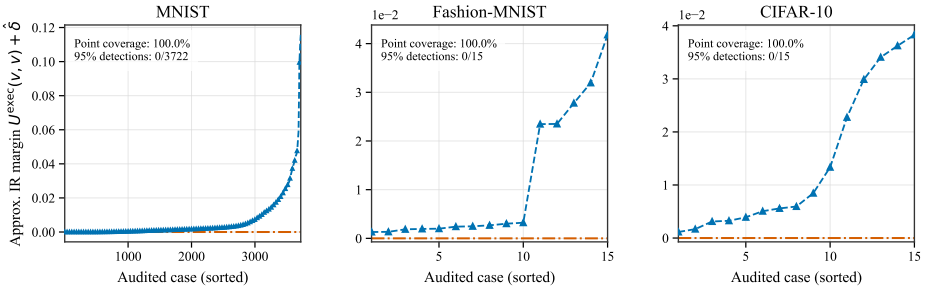

Figure 9: Approximate IR margin under truthful bidding across audited cases.

Table 7 summarizes the expected-utility audit. The 95% interval compatibility rate for approximate IC is 99.973% on MNIST and 100% on the other two datasets. Approximate IR has 100% point coverage on all three datasets.

C.6 Label Heterogeneity and IAFL Adaptation

Figure 10 compares test accuracy under three label partitions. The difference between PABP and Non-private auction is at most 1.03 percentage points, while PABP improves over IAFL-adapted by up to 22.09 percentage points. Bid protection does not cause a noticeable accuracy loss under these partitions.

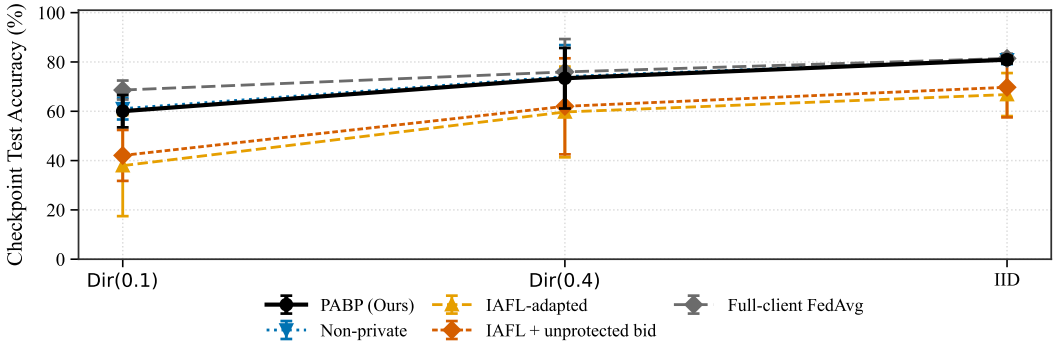

Figure 10: Accuracy under different label partitions.

Table 8 examines the contribution–reward relationship of the IAFL adaptation. The correlations for IAFL-adapted range from 0.697 to 0.910 across the three partitions, and those for the version that uploads original bids range from 0.694 to 0.937. Both preserve a positive correlation(Wu et al., 2024).

Table 8: IAFL adaptation audit.

Partition	Method	Reward-model audit (mean ± 95% CI)	
		Contribution alignment	Test accuracy
Dir(0.1)	IAFL-adapted	$0.7597_{\pm 0.0621}$	$0.2112_{\pm 0.0409}$
	IAFL + unprotected bid	$0.7502_{\pm 0.2231}$	$0.2115_{\pm 0.0186}$
Dir(0.4)	IAFL-adapted	$0.6971_{\pm 0.0509}$	$0.3756_{\pm 0.1261}$
	IAFL + unprotected bid	$0.6941_{\pm 0.4273}$	$0.3612_{\pm 0.1849}$
IID	IAFL-adapted	$0.9101_{\pm 0.0728}$	$0.5229_{\pm 0.0866}$
	IAFL + unprotected bid	$0.9369_{\pm 0.0812}$	$0.5107_{\pm 0.0572}$

Figure 11 reports device-level results on CIFAR-10. The correlation between final contribution and personalized reward-model accuracy is 0.734 ± 0.122. The corresponding correlation for participation rate is 0.628 ± 0.068. Both correlations are positive.

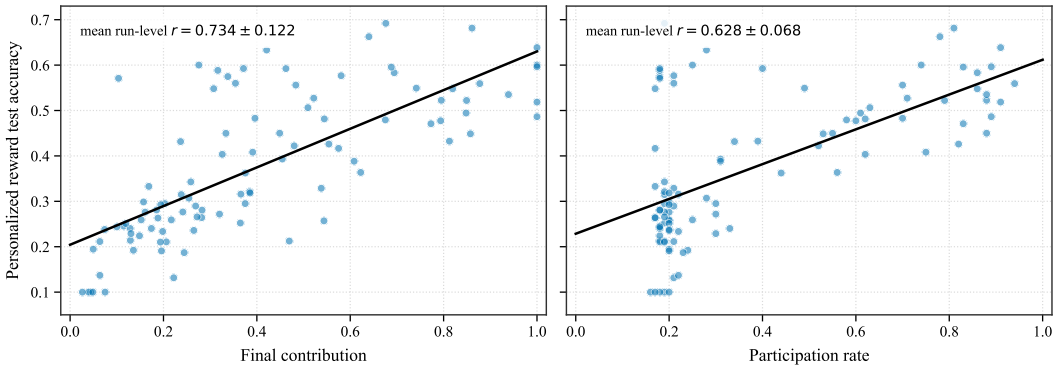

Figure 11: Device-level IAFL reward alignment.

The IAFL adaptation preserves the positive contribution–reward relationship, but the version that uploads original bids exposes those bids to the server. PABP does not upload original bids.

C.7 Ablation Study

This section ablates PABP device selection and joint resource allocation on MNIST. Quality-bid only removes resource cost from the selection score. Fixed power optimizes only the CPU frequency, whereas Fixed power + CPU fixes both resource variables. Random selects feasible devices uniformly. All configurations use the same 30 rounds and 3 paired seeds(Li et al., 2025).

Table 9: Ablation study of selection and resource-allocation components on MNIST.

Method	Learning outcome		Service	Resource use	
	Accuracy ↑	True welfare ↑	Selected/round	Cost ↓	Energy (J) ↓
PABP	$\mathbf{0.966}_{\pm 0.030}$	$\mathbf{2.952}_{\pm 1.179}$	$\mathbf{4.96}_{\pm 0.04}$	$\mathbf{0.535}_{\pm 0.180}$	$\mathbf{47.30}_{\pm 15.42}$
Quality-bid only	$0.957_{\pm 0.050}$	$2.844_{\pm 1.080}$	$5.00_{\pm 0.00}$	$0.739_{\pm 0.085}$	$64.30_{\pm 6.12}$
Fixed power	$0.957_{\pm 0.048}$	$2.614_{\pm 0.960}$	$4.14_{\pm 0.27}$	$0.816_{\pm 0.168}$	$72.56_{\pm 13.00}$
Fixed power + CPU	$0.502_{\pm 0.410}$	$0.251_{\pm 0.434}$	$0.58_{\pm 0.97}$	$0.448_{\pm 0.766}$	$42.21_{\pm 71.94}$
Random	$0.900_{\pm 0.138}$	$0.863_{\pm 0.635}$	$5.00_{\pm 0.00}$	$0.569_{\pm 0.112}$	$48.78_{\pm 8.21}$

Table 9 and Figures 12–13 compare the device-selection and resource-allocation settings. At a similar service count, PABP reduces resource cost and energy use relative to Quality-bid

only by 27.6% and 26.4%, respectively, while achieving higher accuracy and true welfare. Fixed resources or random selection reduce the service count or true welfare.

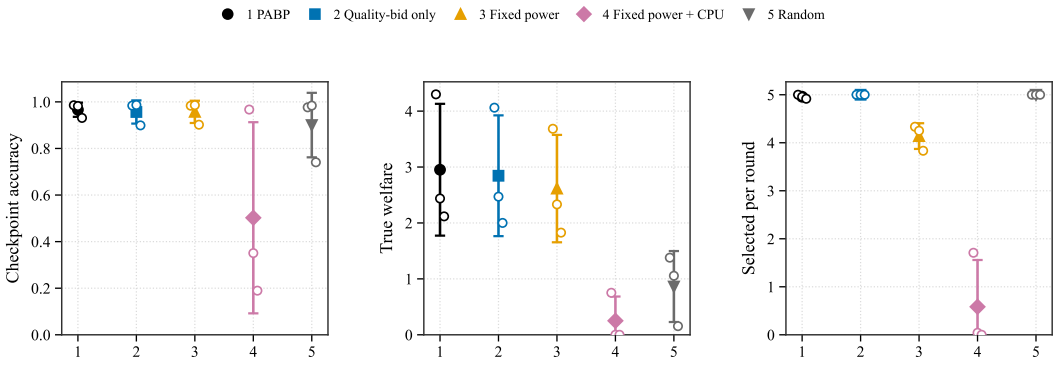

Figure 12: Service outcomes in the ablation study.

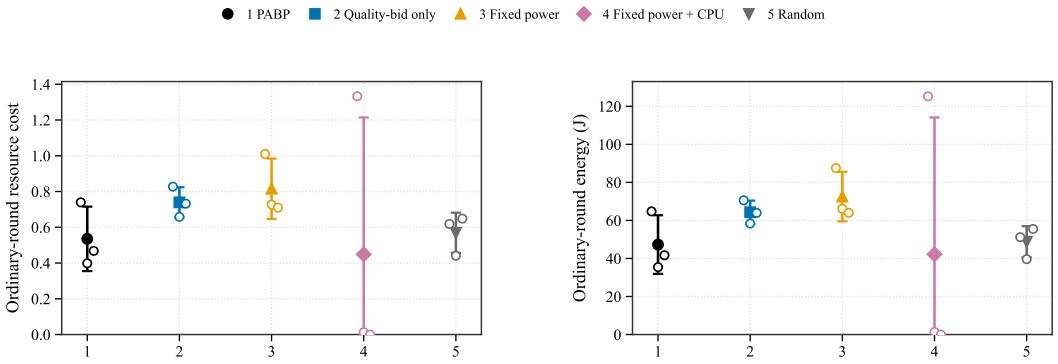

Figure 13: Resource costs in the ablation study.

C.8 Solver Correctness Checks

Table 10 reports the solver checks. The allocation match rate between branch-and-bound and the direct Top-K check is 100%. The maximum cost difference between the continuous resource solution and a dense grid is at most 1.73×10^{-17}, consistent with the optimality result in Proposition 1.

Table 10: Independent numerical checks of the solver.

Dataset	**Max. N**	**Allocation check**		**Resource check**	
		Exact match	**Max. obj. error**	**Feasible cases**	**Max. cost excess**
MNIST	100	100%	2.22×10^{-16}	189/200	1.73×10^{-17}
Fashion-MNIST	40	100%	5.55×10^{-17}	46/50	1.04×10^{-17}
CIFAR-10	40	100%	5.55×10^{-17}	47/50	1.04×10^{-17}