MainFL: Intertemporal Data Valuation for Robust Auction-based Federated Learning

Bangqi Pan[1,#] **Yun Xin**[1,#] **Jianfeng Lu**[2,*]
Gang Li[3] **Shuqin Cao**[4,*] **Zhou Chendi**[5]

[1]School of Computer Science and Technology, Wuhan University of Science and Technology, China
[2]Hubei Province Key Laboratory of Intelligent Information Processing and Real-time Industrial System, Wuhan University of Science and Technology, China
[3]College of Computer Science, Inner Mongolia University, China
[4]Key Laboratory of Social Computing and Cognitive Intelligence (Dalian University of Technology), Ministry of Education, China
[5]School of Computing and Data Science, University of Hong Kong, China
[*]Corresponding authors: Jianfeng Lu <lujianfeng@wust.edu.cn>, Shuqin Cao <shuqincao@wust.edu.cn>
[#]These authors contributed equally to this work.

Abstract

Auction-based Federated Learning (AFL) provides a principled framework for incentivizing data owners and allocating decentralized training resources to data consumers. However, practical AFL systems face a coupled valuation and allocation challenge: data quality is uncertain before training, post-training feedback can be strategically distorted, and heterogeneous deadlines render late model updates ineffective for learning. We propose MainFL, an intertemporal data valuation mechanism for robust auction-based federated learning. MainFL integrates three components: (i) a bounded mutual-information peer prediction rule that elicits truthful pre-training quality reports without requiring ground-truth labels; (ii) an intertemporal reliability calibration scheme that updates future valuations using post-session validation signals; and (iii) a deadline-induced layered clinching mechanism that allocates quality-weighted data while enforcing budget and feasibility constraints. We establish Bayes–Nash truthfulness for the DO peer-scoring subgame; extending it to full-mechanism DO reporting requires an escrow-compatible score-dominance condition. We also analyze allocation feasibility, budget safety, and DC residual demand under clinching-auction regularity conditions. Across the reported vision settings, MainFL improves welfare by 8.46% on average, with a maximum improvement of 12.1%, and the reported average-accuracy metric by 2.79 percentage points on average over the strongest baseline in each setting, while reducing timeout events.

1 Introduction

In recent years, Federated Learning (FL) has garnered increasing attention from both academia and industry due to its ability to enable collaborative model training while safeguarding user privacy [1, 2]. Nevertheless, practical FL faces coupled heterogeneity: participants hold data of uncertain quality that biases aggregated updates, while stragglers induce staleness that distorts optimization dynamics [3]. Moreover, participants act as rational agents who weigh resource consumption against benefits, and may be reluctant to contribute or truthfully report their profiles without compensation [4]. Consequently, designing mechanisms robust to quality uncertainty and latency heterogeneity is essential to motivate participation and ensure efficient resource allocation toward convergence [5].

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

Auction-based Federated Learning (AFL) has emerged as a dominant paradigm in FL incentive mechanism design, offering a flexible market framework for coordinating data owners (DOs), data consumers (DCs), and a platform that clears allocation and payment decisions [6–8]. A *session* denotes one marketplace execution followed by an FL training job; a *communication round* denotes one optimization step inside that job. Data valuation and payments are determined per session, while convergence is measured across communication rounds.

Existing AFL mechanisms typically rely on reputation or contribution scores to guide recruitment and matching. However, this paradigm encounters three critical limitations. *(1) Lack of pre-training valuation.* Reputation mechanisms are inherently post-hoc assessments, rendering them ineffective for pricing DOs that participate for the first time. Consequently, there is a distinct lack of historical data to support decision-making during the initial recruitment phase [9]. *(2) Unreliable post-training updates.* Competitive DCs may lack incentives to provide truthful contribution assessments, introducing systemic bias into multi-session recruitment. *(3) Deadline-aware matching.* Heterogeneous time constraints are frequently overlooked, although late updates are equivalent to unusable training contributions in deadline-constrained FL [10]. These limitations jointly undermine the efficiency of AFL equilibria, degrading both social welfare and model performance.

To address these problems, we propose **MainFL**, a valuation and allocation mechanism for AFL that integrates pre-training truthful valuation, intertemporal reliability calibration, and deadline-aware matching. For pre-training valuation, MainFL uses a bounded peer-prediction score based on pointwise mutual information (PMI) to elicit truthful DO quality reports without relying on trusted labels or DC-provided feedback. For intertemporal valuation, MainFL updates DO reliability using post-session training signals and uses it to calibrate future data quality weights. For deadline-aware matching, MainFL exploits the factorization between model size and DO latency to construct deadline-induced layers and runs a budget-aware clinching auction within each layer.

Our contributions are summarized as follows.

- We design a ground-truth-free pre-training valuation rule based on bounded mutual-information peer prediction. The finite reporting interface yields deterministic score bounds, enabling nonnegative and revenue-funded DO rewards.
- We introduce an intertemporal valuation rule that combines posterior quality estimates with post-session reliability, reducing reliance on strategically biased DC evaluations.
- We develop a deadline-induced layered clinching allocation mechanism that satisfies capacity, budget, and deadline constraints by construction and distinguish restricted per-layer DC incentives from strategic bidding across layers.
- We provide theoretical guarantees and empirical evidence showing that MainFL improves social welfare, model accuracy, and timeout mitigation over representative AFL baselines.

2 Related Work

2.1 Auction Mechanisms in FL

Existing FL incentives adopt several procurement-style designs in which DOs bid into platform-initiated recruitment [4, 7, 11–14], but often ignore richer market structure and timing. Alternative formulations emphasize DC-side bidding strategies [5, 15–18], but frequently omit explicit feasibility guarantees for heterogeneous completion times. Double auctions allow two-sided bidding [19], but often treat recruitment as one-shot procurement. Deadline-sensitive FL studies [20–22] analyze time constraints and convergence without auction allocation. MainFL integrates deadline-aware matching into the auction.

2.2 Data Heterogeneity and Valuation in FL

Feature alignment [23], one-shot prompt refinement [24], and geometry-aware continual correction [25] address training heterogeneity, complementing MainFL's valuation and allocation focus.

Shapley values [26–28] and reputation scores [29] estimate data contributions but may need evaluation signals or participation history before recruitment [30]. MainFL combines pre-training peer scoring with optional post-session reliability calibration to reduce reliance on strategic consumer

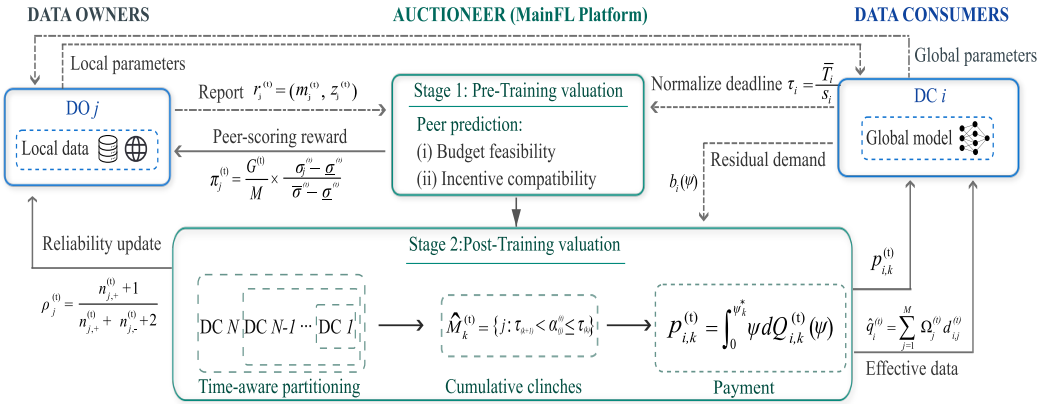

Figure 1: Flowchart of MainFL.

feedback. Reporting guarantees require the stated statistical assumptions and, for full-mechanism DO truthfulness, score dominance.

3 Problem Formalization

We consider a platform-mediated AFL market with M DOs and N DCs over sessions $t = 1, \ldots, H$. At the start of each session, the platform elicits pre-training reports from DOs and bids from DCs, clears allocation and payments, and only then launches the underlying FL training job. In session t, DO j owns local data $\mathcal{D}_j^{(t)}$ and can be assigned to at most one DC. Let $x_{i,j}^{(t)} \in \{0, 1\}$ denote assignment and $d_{i,j}^{(t)} \in [0, |\mathcal{D}_j^{(t)}|]$ denote the raw amount of data used by DC i from DO j:

$$\sum_{i=1}^{N} x_{i,j}^{(t)} \le 1, \qquad 0 \le d_{i,j}^{(t)} \le x_{i,j}^{(t)} |\mathcal{D}_j^{(t)}|. \tag{1}$$

Each DC has a non-decreasing concave value function $V_i(\cdot)$ over weighted data and a budget B_i. If $\hat{q}_i^{(t)}$ is the effective data received and $p_i^{(t)}$ is the auction payment, its quasi-linear utility is

$$U_i^{(t)} = V_i(\hat{q}_i^{(t)}) - p_i^{(t)}. \tag{2}$$

3.1 Pre-training Signals and Beta–Binomial Instantiation

Before training, the platform cannot directly verify data quality. Instead, it collects a *finite* inspection report from each DO, which will later be compared against the reports of randomly paired peers. Formally, each DO submits

$$r_j^{(t)} = (m_j^{(t)}, z_j^{(t)}) \in \mathcal{R}^{(t)}, \tag{3}$$

where $m_j^{(t)}$ is an inspection size and $z_j^{(t)}$ is the reported number of corrupted or noisy inspected samples. We assume reports are correlated only through a latent session-level noise state $\theta^{(t)}$, and are conditionally independent given it:

$$p(r_1^{(t)}, \ldots, r_M^{(t)} \mid \theta^{(t)}) = \prod_{j=1}^{M} p(r_j^{(t)} \mid \theta^{(t)}). \tag{4}$$

For concreteness, we instantiate this by a Beta–Binomial model with $\theta^{(t)} \sim \mathrm{Beta}(a_0^{(t)}, b_0^{(t)})$, and $z_j^{(t)} \mid \theta^{(t)}, m_j^{(t)} \sim \mathrm{Bin}(m_j^{(t)}, \theta^{(t)})$. Standard conjugacy then gives the posterior

$$p^{(t)}(\theta \mid r_j^{(t)}) = \mathrm{Beta}\Big(a_0^{(t)} + z_j^{(t)},\, b_0^{(t)} + m_j^{(t)} - z_j^{(t)}\Big). \tag{5}$$

The pre-training quality weight is the posterior cleanliness estimate $\omega_j^{(t)} = 1 - \mathbb{E}[\theta^{(t)} \mid r_j^{(t)}]$, which evaluates to $\frac{b_0^{(t)}+m_j^{(t)}-z_j^{(t)}}{a_0^{(t)}+b_0^{(t)}+m_j^{(t)}}$. When historical signals are available, we use a reliability summary $\rho_j^{(t-1)} \in [0,1]$ accumulated before session t. MainFL combines the two signals linearly as

$$\Omega_j^{(t)} = (1-\chi)\omega_j^{(t)} + \chi\rho_j^{(t-1)}, \tag{6}$$

where $\chi \in [0,1]$ and $\chi = 0$ recovers a pure pre-training valuation (used in DOs' first session). The effective quality-weighted data received by DC i is

$$\hat{q}_i^{(t)} = \sum_{j=1}^{M} \Omega_j^{(t)} d_{i,j}^{(t)}. \tag{7}$$

3.2 Deadlines and Welfare

Let $s_i > 0$ be the model size of DC i, and let c_j and r_j be DO j's compute capability and uplink rate. Motivated by compute and communication latency models [31], we use the factorized round-level completion time for pair (i,j) as a DC-specific *model-size* factor s_i times a DO-specific *per-unit-size* factor

$$T_{i,j}^{(t)} = s_i \frac{|\mathcal{D}_j^{(t)}|\mu}{\kappa_{\text{cpu}} c_j} + \frac{s_i}{r_j}. \tag{8}$$

For convenience, we denote $\alpha_j^{(t)} = \frac{|\mathcal{D}_j^{(t)}|\mu}{\kappa_{\text{cpu}} c_j} + \frac{1}{r_j}$. If DO j serves DC i, the platform accounts for resource cost

$$C_{i,j}^{(t)}(d) = \xi_j^{(t)} d + \zeta_j T_{i,j}^{(t)}. \tag{9}$$

Payments are transfers and do not enter system welfare. The platform seeks feasible resource allocation solutions with high social welfare:

$$\max_{\{x_{i,j}^{(t)}, d_{i,j}^{(t)}\}} \quad W^{(t)} = \sum_{i=1}^{N} V_i(\hat{q}_i^{(t)}) - \sum_{i=1}^{N}\sum_{j=1}^{M} x_{i,j}^{(t)} C_{i,j}^{(t)}(d_{i,j}^{(t)}) \tag{10a}$$

$$\text{s.t.} \quad \sum_i x_{i,j}^{(t)} \le 1, \qquad 0 \le d_{i,j}^{(t)} \le x_{i,j}^{(t)} |\mathcal{D}_j^{(t)}|, \tag{10b}$$

$$\hat{q}_i^{(t)} \le \gamma_i, \qquad p_i^{(t)} \le B_i, \tag{10c}$$

$$x_{i,j}^{(t)} T_{i,j}^{(t)} \le x_{i,j}^{(t)} \overline{T}_i. \tag{10d}$$

4 MainFL Mechanism

MainFL consists of three tightly coupled components that jointly address valuation, reliability, and allocation in a single session. First, Sec. 4.1 introduces a mutual-information peer score: using only DOs' self-inspection reports, the platform scores each DO against a random peer and issues a revenue-funded reward. Second, Sec. 4.2 converts post-training update signals into a Beta-smoothed reliability estimate, which is then fed back into the quality weight $\Omega_j^{(t)}$ defined in Eq. (6). Third, Sec. 4.3 groups DOs into deadline-induced layers and runs a budget-aware ascending clinching auction within each layer, enforcing the stated feasibility constraints by construction. Figure 1 summarizes the pipeline, and Algorithm 1 gives a unified procedural description.

4.1 Bounded Mutual-Information Peer Scoring

Since the platform has no oracle access to data quality, we score each DO's report by how well it *predicts* the report of a randomly matched peer drawn from the same session. For two reports r_a and r_b, define the posterior-agreement kernel

$$\mathrm{K}(r_a, r_b) = \int_{\Theta} \frac{p(\theta \mid r_a) p(\theta \mid r_b)}{p(\theta)} d\theta = \frac{p(r_a, r_b)}{p(r_a)p(r_b)}, \tag{11}$$

where the second equality follows from Eq. (4) and Bayes' rule (Appendix F). Thus $\log \mathrm{K}(r_a, r_b)$ is a *pointwise mutual information* between two reports that share the latent θ; it is large when the two reports reinforce each other's posterior and small (or negative) when they do not. The Beta–Binomial closed form is given in Appendix F. MainFL randomly pairs DO j with a peer $k \neq j$ and assigns the log-PMI score

$$\sigma_j^{(t)} = \log \mathrm{K}(r_j^{(t)}, r_k^{(t)}). \tag{12}$$

Because $\mathcal{R}^{(t)}$ is finite and the Beta prior has strictly positive parameters, every pair of reports has strictly positive joint and marginal probabilities. Consequently, the platform can *precompute* deterministic bounds $\underline{\sigma}^{(t)} \leq \sigma_j^{(t)} \leq \overline{\sigma}^{(t)}$ by enumerating the (finite) score table off-line. These bounds let us affinely rescale the raw log-PMI into a nonnegative reward that can be financed entirely from DC-side revenue. Let $R^{(t)} = \sum_i p_i^{(t)}$ be DC-side revenue. Before collecting reports, the platform commits a report-independent pool $0 \leq G^{(t)} \leq \lambda R^{(t-1)}$, with $R^{(0)} := R_0 \geq 0$ a platform-funded seed and $\lambda \in (0, 1)$. The peer-scoring reward to DO j is

$$\pi_j^{(t)} = \frac{G^{(t)}}{M} \cdot \frac{\sigma_j^{(t)} - \underline{\sigma}^{(t)}}{\overline{\sigma}^{(t)} - \underline{\sigma}^{(t)}}, \tag{13}$$

and $\pi_j^{(t)} = 0$ otherwise. By construction $\pi_j^{(t)} \in [0, G^{(t)}/M]$, so the total peer reward is bounded by the previously escrowed pool $G^{(t)}$. The procedural form of peer scoring is given in Appendix E.

4.2 Post-session Reliability

After training, MainFL converts each recruited DO's update into a binary reliability signal and Bayes-smooths it into a persistent score for the next session. Let $w^{(t)}$ be the pre-aggregation model and let $\Delta w_j^{(t)} = w_{j,\mathrm{local}}^{(t)} - w^{(t)}$ be the unweighted local update. The validation probe uses a fixed coefficient $\eta > 0$:

$$\Delta_j^{(t)} = \ell(w^{(t)}; \mathcal{V}^{(t)}) - \ell(w^{(t)} + \eta \Delta w_j^{(t)}; \mathcal{V}^{(t)}), \tag{14}$$

where $\mathcal{V}^{(t)}$ is a small public, platform-hosted, or securely evaluated validation interface. The binary reliability observation is

$$y_j^{(t)} = \mathbf{1}\{\Delta_j^{(t)} \geq \varepsilon\}. \tag{15}$$

Intuitively, $y_j^{(t)} = 1$ says that DO j's update improves the validation loss by at least ε. For recruited DOs, MainFL maintains Beta counters $n_{j,+}^{(t)} = n_{j,+}^{(t-1)} + y_j^{(t)}$, $n_{j,-}^{(t)} = n_{j,-}^{(t-1)} + 1 - y_j^{(t)}$, and uses the Laplace-smoothed posterior mean

$$\rho_j^{(t)} = \frac{n_{j,+}^{(t)} + 1}{n_{j,+}^{(t)} + n_{j,-}^{(t)} + 2}. \tag{16}$$

If no validation interface is available, the platform can set $\chi = 0$ in Eq. (6), so MainFL degenerates to pure pre-training peer-scored valuation while retaining the peer-scoring subgame guarantee; full-mechanism DO incentives still require score dominance.

4.3 Deadline-Induced Layered Clinching

For each DC we define the normalized deadline $\tau_i = \overline{T}_i / s_i$, which, by the factorization $T_{i,j}^{(t)} = s_i \alpha_j^{(t)}$ in Eq. (8), reduces the deadline condition to the single-index inequality $\alpha_j^{(t)} \leq \tau_i$. We sort the DCs in decreasing order of τ, so that $\tau_{(1)} \geq \cdots \geq \tau_{(N)}$, and partition the DOs into *deadline layers*

$$\hat{\mathcal{M}}_k^{(t)} = \{j : \tau_{(k+1)} < \alpha_j^{(t)} \leq \tau_{(k)}\}, \qquad k = 1, \ldots, N, \tag{17}$$

with the convention $\tau_{(N+1)} = 0$. Layer k is feasible for exactly the DC prefix $\mathcal{F}_k = \{(1), \ldots, (k)\}$: every DO in $\hat{\mathcal{M}}_k^{(t)}$ meets the deadline of these k DCs, and assigning it to any other DC would violate that DC's deadline. The layers therefore form a laminar family, on which an independent clinching auction can be run per layer without breaking deadline feasibility (Appendix J).

Inside layer k, MainFL sells divisible quality-weighted supply

$$S_k^{(t)} = \sum_{j \in \hat{\mathcal{M}}_k^{(t)}} \Omega_j^{(t)} |\mathcal{D}_j^{(t)}|, \tag{18}$$

and lets each eligible DC $i \in \mathcal{F}_k$ submit residual demand $b_i(\psi)$ over effective units as the price ψ ascends. Write $\Gamma_i(\psi)$ for the remaining capacity and $B_i^{\mathrm{rem}}(\psi)$ for the remaining budget at price ψ. Because demand must respect both capacity and budget, we use the budget-truncated demand

$$\tilde{b}_i(\psi) = \min\left\{ b_i(\psi),\ \Gamma_i(\psi),\ \frac{B_i^{\mathrm{rem}}(\psi)}{\psi} \right\}. \tag{19}$$

If $S^{\mathrm{rem}}(\psi)$ is the remaining layer supply, the incremental clinch awarded to DC i at price ψ is the part of its demand that *cannot* be absorbed by the remaining demand of competitors:

$$\Delta Q_i(\psi) = \min\left\{ \tilde{b}_i(\psi),\ \left[S^{\mathrm{rem}}(\psi) - \sum_{n \in \mathcal{F}_k,\ n \neq i} \tilde{b}_n(\psi) \right]_+ \right\}. \tag{20}$$

This is the standard Ausubel-style clinching rule [32]; the intuition is that DC i wins any unit that is not under contention. The payment is the usual path integral along the clinched-quantity trajectory,

$$p_{i,k}^{(t)} = \int_0^{\psi_k^\star} \psi \, dQ_{i,k}^{(t)}(\psi), \tag{21}$$

where $Q_{i,k}^{(t)}(\psi)$ is the nondecreasing cumulative allocation path, with increments given by Eq. (20), and $\psi_k^\star$ is the terminal price. On a price grid this payment is $\sum_\ell \psi_\ell \Delta Q_{i,k,\ell}$, with supply, capacity, and budget updated after each step. The payment identity, together with demand truncation by $B_i^{\mathrm{rem}}(\psi)/\psi$, is what lets us guarantee budget feasibility (Appendix K).

After clinching, let $Q_{i,k}^{(t)}$ denote the target quota and $\widehat{Q}_{i,k}^{(t)} = \sum_{j \in \hat{\mathcal{M}}_k^{(t)}} \Omega_j^{(t)} d_{i,j}^{(t)}$ its DO-level realization. Owners are processed in decreasing $\Omega_j^{(t)}$ with fixed tie-breaking, assigned once, and capped at the consumer's remaining quota. Thus $0 \le \widehat{Q}_{i,k}^{(t)} \le Q_{i,k}^{(t)}$; equality requires jointly realizable quotas. Aggregate supply alone does not guarantee exact realization under the one-DO-one-DC constraint. A sufficient supply-reserve extension is given in Appendix N; it is not used in the reported experiments.

The procedure of Algorithm 1 runs in five phases. Lines 3–5 normalize deadlines into $\tau_i = \overline{T}_i / s_i$ and build the laminar layers $\hat{\mathcal{M}}_k^{(t)}$ of Sec. 4.3. Lines 7–8 extract the eligible DCs and the quality-weighted supply of each layer. Lines 9–10 run a budget-aware ascending clinching auction on *effective units*, absorbing the heterogeneity of $\Omega_j^{(t)}$ into a common numeraire. Line 11 discretizes the resulting quotas into one-to-one DO–DC assignments using a fixed assignment order; realization and quota shortfall are discussed in Appendix D. Line 12 redistributes at most a $(1-\lambda)$ fraction of each DC's layer payment to the DOs assigned in that layer according to their realized effective contributions $\Omega_j^{(t)} d_{i,j}^{(t)}$. Theorem K.1 establishes assignment, capacity, budget, and modeled-deadline feasibility; exact delivery and DC individual rationality require the additional conditions below.

5 Theoretical Guarantees

We now state the main theoretical guarantees that MainFL enjoys. Since the log-PMI score only depends on a DO's report through its predictive distribution, reports inducing the same predictive distribution over peers are indistinguishable to the scoring rule; truthfulness is thus stated up to this equivalence.

Theorem 5.1. *Truthful reporting forms a Bayes–Nash equilibrium of the peer-scoring subgame. For every session,* $\pi_j^{(t)} \ge 0$, $\sum_{j=1}^M \pi_j^{(t)} \le G^{(t)} \le \lambda R^{(t-1)}$.

For a correctly specified peer-report distribution and a report-independent positive pool, the result follows from the KL-divergence identity in Appendix G. A zero pool gives only weak indifference.

Algorithm 1 MainFL allocation procedure

Input: DCs $\{(s_i, \overline{T}_i, \gamma_i, B_i)\}_{i=1}^{N}$, DOs $\{(|\mathcal{D}_j^{(t)}|, c_j, r_j)\}_{j=1}^{M}$, calibrated qualities $\{\Omega_j^{(t)}\}$, revenue share $\lambda \in (0,1)$
Output: assignment $\{x_{i,j}^{(t)}\}$, traded amounts $\{d_{i,j}^{(t)}\}$, DC payments $\{p_i^{(t)}\}$, rewards $\{\pi_j^{(t,\mathrm{alloc})}\}$
Compute $\alpha_j^{(t)}$ by Eq. (8); sort DCs by $\tau_i = \overline{T}_i / s_i$ descending
Construct deadline layers $\{\hat{\mathcal{M}}_k^{(t)}\}_{k=1}^{N}$ by Eq. (17)
Initialize $x_{i,j}^{(t)} \leftarrow 0$, $d_{i,j}^{(t)} \leftarrow 0$, $p_i^{(t)} \leftarrow 0$, $\pi_j^{(t,\mathrm{alloc})} \leftarrow 0$, $\bar{\gamma}_i \leftarrow \gamma_i$, $\bar{B}_i \leftarrow B_i$
for $k = 1$ to N **do**
  Set eligible DCs $\mathcal{F}_k \leftarrow \{(1), \ldots, (k)\}$
  Compute effective supply $S_k^{(t)} \leftarrow \sum_{j \in \hat{\mathcal{M}}_k^{(t)}} \Omega_j^{(t)} |\mathcal{D}_j^{(t)}|$ // cf. Eq. (18)
  Run budget-aware ascending clinching over $S_k^{(t)}$ among $\mathcal{F}_k$ using remaining $(\bar{\gamma}_i, \bar{B}_i)$
  Obtain effective quotas $\{Q_{i,k}^{(t)}\}_{i \in \mathcal{F}_k}$ and layer payments $\{p_{i,k}^{(t)}\}_{i \in \mathcal{F}_k}$
  Quota discretization. In fixed DC order, assign unused DOs in decreasing quality order, cap the final contribution at the remaining target, and retire every assigned DO. Set $x_{i,j}^{(t)} = 1$ for positive trades; realized quantities satisfy $\widehat{Q}_{i,k}^{(t)} \le Q_{i,k}^{(t)}$.
  DO allocation rewards. For each $i \in \mathcal{F}_k$ with $Q_{i,k}^{(t)} > 0$ and each DO j assigned to i in this layer, accrue $\pi_j^{(t,\mathrm{alloc})} \mathrel{+}= (1-\lambda)\, p_{i,k}^{(t)} \cdot \frac{\Omega_j^{(t)} d_{i,j}^{(t)}}{Q_{i,k}^{(t)}}$
  Update $p_i^{(t)} \leftarrow p_i^{(t)} + p_{i,k}^{(t)}$, $\bar{B}_i \leftarrow \bar{B}_i - p_{i,k}^{(t)}$, $\bar{\gamma}_i \leftarrow \bar{\gamma}_i - Q_{i,k}^{(t)}$
end for

Full-mechanism DO reporting additionally requires score dominance. Let $D_\sigma = \overline{\sigma} - \underline{\sigma} > 0$, let $\delta_{\min} > 0$ be the minimum nonzero predictive KL gap, and let $\overline{A}$ bound a deviation's expected non-score gain. A sufficient condition for non-equivalent deviations is

$$\frac{G^{(t)} \delta_{\min}}{M D_\sigma} \ge \overline{A}, \qquad G^{(t)} \le \lambda R^{(t-1)}. \tag{22}$$

Escrow compatibility is not automatic. Predictive-equivalent deviations separately require no positive non-score gain; Corollary H.1 gives the conditional full-mechanism statement and proof.

Assumption 5.2 (Restricted per-layer incentive setting)**.** Consider an isolated layer with its entering state fixed, private nondecreasing concave values over effective data, nonbinding budgets, and exact quota realization. The auction uses the standard exact ascending clinching process; continuation choices across layers are excluded.

Theorem 5.3 (Per-layer DC incentives)**.** *Under Assumption 5.2, truthful residual demand is a best response to truthful rivals in the isolated layer, and the resulting trade is individually rational. Budget feasibility under the truncated updates holds separately (Appendix K).*

This is a restricted per-layer result [32], not a full-mechanism DC incentive guarantee. Binding budgets, unrealized quotas, interdependent values, finite price grids, and strategic deferral across layers require separate analysis (Appendices L and M).

6 Experiments

6.1 Experimental Setup

Datasets and Models. We evaluate MainFL on standard vision benchmarks FMNIST [33], CIFAR-10 [34], and CINIC-10 [35], plus the LEAF Shakespeare next-character task [36]. Vision experiments further include MNIST [37] alongside FMNIST, CIFAR-10, and CINIC-10 in Tables 1 and 2. For vision jobs we use ResNet-18 [38] with batch normalization replaced by group normalization. For Shakespeare [36] we use a two-layer LSTM encoder with a softmax next-token head.

Baselines. We evaluate MainFL against several SOTA FL auction mechanisms, each instantiated with the closest available procurement-market designs in their original papers. Specifically, **FedCA**

Table 1: Social Welfare in AFL system

Budget	Methods	MNIST		FMNIST		CIFAR-10		CINIC-10	
		Data	Welfare	Data	Welfare	Data	Welfare	Data	Welfare
100	FedCA	9187	85.73 ± 8.62	9756	88.91 ± 7.34	8584	62.89 ± 8.47	18976	129.85 ± 9.58
	AFLM	8821	94.65 ± 7.34	9489	96.21 ± 6.58	8205	68.74 ± 7.63	18742	143.36 ± 8.14
	FedDRL	12518	141.73 ± 5.38	13645	147.26 ± 6.42	11829	124.52 ± 6.75	22364	222.54 ± 8.31
	MBOS-AFL	12933	147.38 ± 6.27	14092	151.52 ± 6.84	12156	128.65 ± 7.41	22687	234.13 ± 9.56
	ARAS-AFL	12165	138.46 ± 7.52	13214	143.64 ± 7.29	11478	120.38 ± 8.17	22019	211.95 ± 9.62
	MainFL	14209	**162.45 ± 6.63**	15473	**167.84 ± 6.27**	13581	**144.26 ± 7.58**	23825	**245.47 ± 9.14**
200	FedCA	16912	107.68 ± 9.27	17548	112.43 ± 8.51	15224	79.86 ± 9.42	30965	153.92 ± 10.73
	AFLM	16029	116.47 ± 8.54	17088	120.32 ± 8.76	14615	85.26 ± 9.38	30561	171.35 ± 9.43
	FedDRL	24082	190.46 ± 7.54	26145	193.62 ± 8.48	22823	165.75 ± 8.23	38459	272.76 ± 9.51
	MBOS-AFL	24813	196.73 ± 8.15	26891	200.46 ± 8.32	23466	169.24 ± 9.57	38927	286.53 ± 10.24
	ARAS-AFL	23256	185.58 ± 9.61	25439	189.73 ± 9.25	22174	159.32 ± 10.46	37912	260.84 ± 11.37
	MainFL	26834	**214.65 ± 8.82**	29177	**221.57 ± 8.14**	25708	**186.43 ± 9.76**	41183	**294.82 ± 10.65**
400	FedCA	31183	124.72 ± 10.68	32815	129.87 ± 9.46	27584	92.95 ± 11.72	46057	149.42 ± 12.58
	AFLM	30018	133.75 ± 9.54	31763	138.27 ± 9.48	26872	99.43 ± 10.62	45275	159.58 ± 11.45
	FedDRL	47022	228.45 ± 8.63	50189	231.26 ± 9.37	42655	199.54 ± 9.15	61284	273.37 ± 11.72
	MBOS-AFL	48293	234.56 ± 9.17	51015	239.63 ± 9.74	43561	205.35 ± 10.28	62042	284.82 ± 12.56
	ARAS-AFL	45684	220.84 ± 10.35	48627	226.47 ± 10.52	41439	191.76 ± 11.64	60166	261.25 ± 13.43
	MainFL	51475	**252.63 ± 9.46**	54836	**262.35 ± 9.84**	46814	**226.87 ± 10.53**	65193	**301.74 ± 12.65**
800	FedCA	43752	146.58 ± 11.87	45216	150.73 ± 11.46	39587	109.86 ± 13.78	68174	182.54 ± 14.58
	AFLM	42269	154.63 ± 11.42	43823	158.26 ± 10.85	38291	115.74 ± 12.43	67018	193.45 ± 13.67
	FedDRL	49566	267.53 ± 10.68	51749	272.36 ± 11.54	46884	235.85 ± 11.27	76231	365.82 ± 13.86
	MBOS-AFL	50345	275.28 ± 11.76	52382	281.54 ± 11.92	47413	241.63 ± 12.58	76905	379.76 ± 14.45
	ARAS-AFL	48674	259.37 ± 12.84	50831	266.75 ± 12.46	46026	227.42 ± 13.35	75692	350.04 ± 15.73
	MainFL	52891	**297.45 ± 11.36**	55217	**306.84 ± 11.62**	49133	**263.57 ± 12.74**	79468	**398.53 ± 14.28**

[16] adopts a return-on-investment (ROI)-based method to determine the optimal price for each data resource. **MBOS-AFL** [17] strategically allocates budgets across multiple training sessions for long-horizon cost control. **ARAS-AFL** [29] dynamically adjusts the allocation strategy via Lyapunov optimization, jointly considering client reputation and queue stability. **FedDRL** [39] proposes a long-term Vickrey–Clarke–Groves (VCG) auction mechanism adaptive to the FL update strategy. **AFLM** [14] uses a greedy algorithm to select winners over multiple auction rounds and computes payments via VCG rules.

Configurations. By default each session uses $M = 100$ DOs and $N = 10$ DCs, a Dirichlet non-IID split with $\alpha = 0.5$, $|\mathcal{D}_j| \in [300, 1500]$, and inspection size $m_j = 200$. We report traded data, welfare in Eq. (10a), the accuracy summaries in Table 2, and DO-side timeout events. When device costs are unavailable we set $\xi_j^{(t)} = \zeta_j = 0$, so welfare tracks quality-weighted value.

6.2 Recruitment results

Table 1 presents the recruitment outcomes and social welfare achieved by different mechanisms under varying average budgets. It is evident that across all four datasets and budget configurations, MainFL consistently outperforms all baseline methods in terms of social welfare. Specifically, compared to the second-best approach, MainFL achieves a substantial improvement in social welfare, ranging from 2.9% to 12.1% across different settings. MBOS-AFL and FedDRL also achieve higher welfare than FedCA and AFLM in the reported settings.

Comparing with the strongest baseline separately in each table cell, Table 1 yields mean relative welfare gain 8.46% across its 16 dataset–budget cells (maximum 12.13% from the printed values). Table 2 yields mean gains in the reported highest/average accuracy metrics of $2.16/2.79$ percentage points across 32 dataset–budget–noise cells, with maxima $4.41/5.57$ points. Its 7.4% figure is the mean relative gain in average accuracy within the 20% noise regime, not a maximum absolute accuracy gain. Table 3 yields mean timeout reduction 77.39% in the 12 cells with nonzero best-baseline counts; the other four are zero-count ties.

6.3 Performance Evaluation

Table 2 illustrates the comparison of model performance under noisy scenarios. In settings with different levels of noise injection ($p = 10\%$ and $p = 20\%$), MainFL consistently excels in improving model accuracy compared to existing methods. Specifically, under the 10% noise setting, the maximum and average accuracy of MainFL surpass the best baseline performance by approximately 3.0%

Table 2: Model performance in AFL system (Highest/Average accuracy in %).

Budget	Methods	MNIST		FMNIST		CIFAR-10		CINIC-10	
		$p = 10\%$	$p = 20\%$	$p = 10\%$	$p = 20\%$	$p = 10\%$	$p = 20\%$	$p = 10\%$	$p = 20\%$
100	FedCA	84.56 / 81.94	82.18 / 79.07	71.28 / 67.85	67.13 / 63.24	45.47 / 40.58	35.64 / 30.31	28.64 / 24.87	23.57 / 20.02
	AFLM	86.15 / 83.42	83.89 / 80.96	73.04 / 69.61	68.87 / 64.73	47.68 / 42.35	37.12 / 31.49	29.74 / 25.91	24.67 / 20.83
	FedDRL	87.81 / 85.76	85.93 / 83.38	75.59 / 72.32	71.14 / 67.76	52.17 / 47.03	44.28 / 39.15	32.16 / 28.49	26.92 / 23.14
	MBOS-AFL	87.39 / 85.15	85.42 / 82.71	74.78 / 71.65	70.43 / 67.19	50.84 / 46.12	42.66 / 37.81	31.58 / 27.97	26.35 / 22.69
	ARAS-AFL	87.54 / 85.38	85.76 / 83.13	75.01 / 71.97	70.82 / 67.44	51.49 / 46.56	43.51 / 38.48	31.83 / 28.15	26.62 / 22.87
	MainFL	**89.17 / 87.42**	**87.75 / 85.69**	**76.83 / 74.61**	**74.48 / 72.16**	**54.85 / 51.24**	**48.69 / 44.72**	**34.01 / 31.08**	**29.86 / 26.43**
200	FedCA	85.97 / 83.85	83.76 / 81.02	72.54 / 68.82	68.15 / 64.31	47.89 / 43.31	37.96 / 33.08	29.96 / 26.51	24.63 / 21.45
	AFLM	87.14 / 85.09	85.02 / 82.37	73.95 / 70.48	69.93 / 65.82	49.91 / 45.16	39.84 / 34.68	30.95 / 27.39	25.74 / 22.36
	FedDRL	88.69 / 86.93	86.81 / 84.64	76.47 / 73.52	72.38 / 69.15	55.36 / 50.62	46.68 / 41.97	33.93 / 30.45	28.19 / 24.61
	MBOS-AFL	88.25 / 86.41	86.48 / 84.12	75.73 / 73.09	71.56 / 68.43	53.84 / 49.67	45.19 / 40.75	33.12 / 29.88	27.56 / 24.19
	ARAS-AFL	88.37 / 86.64	86.62 / 84.39	76.04 / 73.25	71.81 / 68.66	54.48 / 49.92	45.73 / 41.14	33.47 / 30.01	27.85 / 24.38
	MainFL	**89.83 / 88.26**	**88.54 / 86.82**	**77.85 / 75.93**	**75.79 / 73.34**	**57.21 / 53.75**	**50.36 / 46.69**	**35.82 / 33.17**	**31.43 / 28.35**
400	FedCA	86.91 / 84.98	84.95 / 82.71	73.25 / 69.81	68.96 / 65.42	49.82 / 45.13	40.39 / 35.56	31.12 / 27.78	25.83 / 22.42
	AFLM	87.72 / 86.08	85.83 / 83.56	74.45 / 71.29	70.41 / 66.64	51.37 / 46.82	41.93 / 36.95	32.09 / 28.52	26.68 / 23.16
	FedDRL	89.26 / 87.43	87.65 / 85.71	76.92 / 74.48	73.06 / 69.97	57.48 / 53.14	48.65 / 44.13	35.62 / 32.69	30.45 / 27.21
	MBOS-AFL	89.01 / 87.19	87.24 / 85.28	76.38 / 73.81	72.43 / 69.25	56.16 / 52.09	47.34 / 42.82	34.93 / 31.95	29.76 / 26.58
	ARAS-AFL	89.15 / 87.22	87.38 / 85.33	76.51 / 74.07	72.69 / 69.46	56.82 / 52.48	47.91 / 43.35	35.14 / 32.17	30.08 / 26.79
	MainFL	**90.29 / 88.93**	**89.11 / 87.65**	**78.47 / 76.62**	**76.55 / 74.28**	**59.13 / 55.86**	**51.84 / 47.97**	**37.69 / 35.12**	**33.15 / 30.23**
800	FedCA	87.48 / 85.82	85.96 / 83.81	73.68 / 70.85	69.82 / 66.51	50.94 / 46.52	42.07 / 37.29	32.03 / 29.17	27.11 / 23.84
	AFLM	88.16 / 86.64	86.62 / 84.77	74.79 / 71.95	70.93 / 67.68	52.41 / 48.13	43.65 / 38.56	33.02 / 30.19	28.17 / 24.74
	FedDRL	89.78 / 88.23	88.21 / 86.49	77.74 / 75.66	74.12 / 71.25	59.63 / 56.49	51.97 / 48.24	37.38 / 34.82	32.96 / 30.11
	MBOS-AFL	89.43 / 87.89	87.96 / 86.04	77.15 / 75.02	73.47 / 70.63	58.74 / 55.65	51.12 / 47.19	36.69 / 34.18	32.14 / 29.45
	ARAS-AFL	89.51 / 87.97	88.05 / 86.12	77.38 / 75.14	73.79 / 70.86	59.06 / 55.92	51.48 / 47.63	36.93 / 34.37	32.41 / 29.69
	MainFL	**90.57 / 89.34**	**89.68 / 88.11**	**79.16 / 77.63**	**77.42 / 75.59**	**61.35 / 58.28**	**53.64 / 50.17**	**39.22 / 36.85**	**34.86 / 31.93**

Table 3: Total DO timeout events per session ($T = 200$ rounds, aggregated over $N = 10$ DC jobs; see Appendix B).

Method	Dataset:	MNIST				FMNIST				CIFAR-10				CINIC-10			
	Budget =	100	400	600	800	100	400	600	800	100	400	600	800	100	400	600	800
MBOS-AFL		18	58	72	88	12	48	61	74	210	285	302	348	315	528	596	692
ARAS-AFL		5	14	27	31	4	9	18	22	85	124	190	238	155	265	412	481
AFLM		0	**0**	17	43	0	**0**	11	28	62	156	214	275	188	410	531	657
FedDRL		34	92	104	115	29	78	86	96	175	340	478	610	290	742	1104	1936
MainFL		**0**	**0**	**0**	**0**	**0**	**0**	**0**	**9**	**0**	**0**	**67**	145	42	95	138	182

and 4.6% on average, respectively. The advantage becomes even more pronounced in the harsher 20% noise setting, where the maximum and average accuracy are improved by approximately 5.7% and 7.4%, respectively, compared to the best baseline. This is consistent with the use of quality-aware valuation. MainFL uses peer reports and optional platform validation to reduce reliance on strategic DC evaluations.

6.4 Task Time Analysis

Table 3 shows that MainFL consistently mitigates timeout risks across all tasks and recruitment budgets. By layering deadlines to shift mass away from chronically violated matches, it markedly reduces straggler-induced timeouts, especially on harder vision tasks. For example, on MNIST with $B = 800$, MainFL incurs 0 timeouts, compared to 115 for FedDRL and 43 for AFLM; on CINIC-10 with $B = 800$, it reduces timeouts to 182, versus 1936 for FedDRL and 657 for AFLM. These results confirm MainFL's ability to cluster DOs and manage temporal heterogeneity under varying budgets, while the reported timeout counts remain nonzero in some settings.

7 Conclusion

In this paper, we presented MainFL, an intertemporal data valuation and deadline-aware matching mechanism for AFL. MainFL elicits truthful pre-training quality reports from DOs through bounded mutual-information peer prediction, calibrates future valuation through post-session reliability, and allocates effective data supply through deadline-induced layered clinching. Theoretical analysis establishes DO peer-scoring Bayes–Nash truthfulness. The allocation preserves the stated upper-bound constraints, while DC truthfulness and individual rationality are restricted to the isolated-layer setting. Empirical results across multiple datasets validate that MainFL improves auction efficiency and model accuracy while mitigating straggler-induced timeouts.

Impact Statement

MainFL studies compensation and deadline-aware allocation in federated-learning markets. Its potential benefits include more effective use of training resources and reduced dependence on strategic consumer feedback. Its guarantees rely on the stated statistical and auction assumptions. Biased validation data can disadvantage data owners whose distributions differ from the validation population; coalition manipulation and identity changes also remain deployment concerns.

Acknowledgments and Disclosures

Funding. This work was supported in part by the National Natural Science Foundation of China (Nos. 62372343, 62402352, 62072411), in part by the Hubei Provincial Natural Science Foundation of China (No.2026AFB510, 2026AFC0108), in part by the Open Fund of Key Laboratory of Social Computing and Cognitive Intelligence (Dalian University of Technology), Ministry of Education (No. SCCI2024TB02).

Competing interests. The authors declare no known competing financial interests.

References

[1] Jakub Konečnỳ, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*, 2016.

[2] Jinliang Yuan, Shangguang Wang, Hongyu Li, Daliang Xu, Yuanchun Li, Mengwei Xu, and Xuanzhe Liu. Towards energy-efficient federated learning via int8-based training on mobile dsps. In *Proceedings of the ACM Web Conference 2024*, WWW '24, pages 2786–2794, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400701719.

[3] Jian Tang, Xiuhua Li, Hui Li, Penghua Li, Xiaofei Wang, and Victor C. M. Leung. Joint class-balanced client selection and bandwidth allocation for cost-efficient federated learning in mobile edge computing networks. *IEEE Transactions on Mobile Computing*, 24(7):5681–5698, 2025. doi: 10.1109/TMC.2025.3539284.

[4] Lianbo Ma, Ying Qian, Guo Yu, Zhetao Li, Liang Wang, Qing Li, Xingwei Wang, and Guangjie Han. Tbcim: Two-level blockchain-aided edge resource allocation mechanism for federated learning service market. *IEEE Transactions on Networking*, 33(6):3408–3423, 2025. doi: 10.1109/TON.2025.3589017.

[5] Xiaoli Tang, Han Yu, Zengxiang Li, and Xiaoxiao Li. A bias-free revenue-maximizing bidding strategy for data consumers in auction-based federated learning. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, IJCAI '24, 2024.

[6] Xiaoli Tang, Ying-Peng Tang, Zhuang Qi, Chuan Sun, Qianyu Li, Han Yu, Jie Zhang, and Qiang Yang. Robust bidding strategies under censored feedback for auction-based federated learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–12, 2026. doi: 10.1109/TPAMI.2026.3736862.

[7] Jingwen Zhang, Yuezhou Wu, and Rong Pan. Incentive mechanism for horizontal federated learning based on reputation and reverse auction. In *Proceedings of the Web Conference 2021*, WWW '21, pages 947–956, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450383127. doi: 10.1145/3442381.3449888. URL `https://doi.org/10.1145/3442381.3449888`.

[8] Renhao Lu, Weizhe Zhang, Yan Wang, Qiong Li, Xiaoxiong Zhong, Hongwei Yang, and Desheng Wang. Auction-based cluster federated learning in mobile edge computing systems. *IEEE Transactions on Parallel and Distributed Systems*, 34(4):1145–1158, 2023. doi: 10.1109/TPDS.2023.3240767.

[9] Qingmin Liu. Information acquisition and reputation dynamics. *The Review of Economic Studies*, 78(4):1400–1425, 02 2011. ISSN 0034-6527. doi: 10.1093/restud/rdq039. URL `https://doi.org/10.1093/restud/rdq039`.

[10] Maojun Zhang, Yang Li, Dongzhu Liu, Richeng Jin, Guangxu Zhu, Caijun Zhong, and Tony Q.S. Quek. Joint compression and deadline optimization for wireless federated learning. *IEEE Transactions on Mobile Computing*, 23(7):7939–7951, 2024. doi: 10.1109/TMC.2023.3344108.

[11] Tra Huong Thi Le, Nguyen H. Tran, Yan Kyaw Tun, Minh N. H. Nguyen, Shashi Raj Pandey, Zhu Han, and Choong Seon Hong. An incentive mechanism for federated learning in wireless cellular networks: An auction approach. *IEEE Transactions on Wireless Communications*, 20 (8):4874–4887, 2021. doi: 10.1109/TWC.2021.3062708.

[12] Yutao Jiao, Ping Wang, Dusit Niyato, Bin Lin, and Dong In Kim. Toward an automated auction framework for wireless federated learning services market. *IEEE Transactions on Mobile Computing*, 20(10):3034–3048, 2021. doi: 10.1109/TMC.2020.2994639.

[13] Yongheng Deng, Feng Lyu, Ju Ren, Huaqing Wu, Yuezhi Zhou, Yaoxue Zhang, and Xuemin Shen. Auction: Automated and quality-aware client selection framework for efficient federated learning. *IEEE Transactions on Parallel and Distributed Systems*, 33(8):1996–2009, 2022. doi: 10.1109/TPDS.2021.3134647.

[14] 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.

[15] Xiaoli Tang and Han Yu. Utility-maximizing bidding strategy for data consumers in auction-based federated learning. In *2023 IEEE International Conference on Multimedia and Expo (ICME)*, pages 330–335. IEEE, 2023.

[16] Xiaoli Tang and Han Yu. A cost-aware utility-maximizing bidding strategy for auction-based federated learning. *IEEE Transactions on Neural Networks and Learning Systems*, 36(7): 12866–12879, 2025.

[17] Xiaoli Tang, Han Yu, Zengxiang Li, and Xiaoxiao Li. Multi-session budget optimization for forward auction-based federated learning. In *Forty-second International Conference on Machine Learning*, 2025. URL `https://openreview.net/forum?id=bFB0N8ABIr`.

[18] Xiaoli Tang, Han Yu, and Xiaoxiao Li. A reinforcement learning-based bidding strategy for data consumers in auction-based federated learning. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL `https://openreview.net/forum?id=zIbNGkaYij`.

[19] Peng Sun, Guocheng Liao, Xu Chen, and Jianwei Huang. A socially optimal data marketplace with differentially private federated learning. *IEEE/ACM Transactions on Networking*, 32(3): 2221–2236, 2024. doi: 10.1109/TNET.2024.3351864.

[20] Mingzhe Chen, H. Vincent Poor, Walid Saad, and Shuguang Cui. Convergence time optimization for federated learning over wireless networks. *IEEE Transactions on Wireless Communications*, 20(4):2457–2471, 2021. doi: 10.1109/TWC.2020.3042530.

[21] Ruobei Wang, Ruiting Zhou, Jieling Yu, Bo Li, and Yuqing Li. Heterogeneous federated learning frameworks for balancing job completion time and model accuracy. *IEEE Transactions on Networking*, pages 1–14, 2025. doi: 10.1109/TON.2025.3603254.

[22] Dong Yang, Weiting Zhang, Qiang Ye, Chuan Zhang, Ning Zhang, Chuan Huang, Hongke Zhang, and Xuemin Shen. Detfed: Dynamic resource scheduling for deterministic federated learning over time-sensitive networks. *IEEE Transactions on Mobile Computing*, 23(5):5162–5178, 2024. doi: 10.1109/TMC.2023.3303017.

[23] Zhuang Qi, Lei Meng, Zhaochuan Li, Han Hu, and Xiangxu Meng. Cross-silo feature space alignment for federated learning on clients with imbalanced data. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(19):19986–19994, 2025. doi: 10.1609/aaai.v39i19.34201.

[24] Zhuang Qi, Yu Pan, Lei Meng, Sijin Zhou, Han Yu, Xiaoxiao Li, and Xiangxu Meng. Global prompt refinement with non-interfering attention masking for one-shot federated learning. *Advances in Neural Information Processing Systems*, 38:13487–13511, 2025. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/11e7d0a3341de7f451141248dfbff4d3-Abstract-Conference.html.

[25] Zhuang Qi, Ying-Peng Tang, Lei Meng, Guoqing Chao, Lei Wu, Han Yu, and Xiangxu Meng. From selection to scheduling: Federated geometry-aware correction makes exemplar replay work better under continual dynamic heterogeneity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17494–17504, June 2026.

[26] Amirata Ghorbani and James Zou. Data shapley: Equitable valuation of data for machine learning. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pages 2242–2251. PMLR, 09–15 Jun 2019.

[27] Xinyi Xu, Zhaoxuan Wu, Chuan Sheng Foo, and Bryan Kian Hsiang Low. Validation free and replication robust volume-based data valuation. In *Proceedings of the 35th International Conference on Neural Information Processing Systems*, NIPS '21, Red Hook, NY, USA, 2021.

[28] Jinkun Lin, Anqi Zhang, Mathias Lécuyer, Jinyang Li, Aurojit Panda, and Siddhartha Sen. Measuring the effect of training data on deep learning predictions via randomized experiments. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 13468–13504, 17–23 Jul 2022.

[29] Xiaoli Tang and Han Yu. Reputation-aware revenue allocation for auction-based federated learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(19):20832–20840, Apr. 2025. doi: 10.1609/aaai.v39i19.34296. URL https://ojs.aaai.org/index.php/AAAI/article/view/34296.

[30] Zelei Liu, Yuanyuan Chen, Han Yu, Yang Liu, and Lizhen Cui. GTG-Shapley: Efficient and accurate participant contribution evaluation in federated learning, 2021. URL https://arxiv.org/abs/2109.02053.

[31] Jieling Yu, Ruiting Zhou, Bo Li, and Libing Wu. Intelligent frameworks for minimizing job completion time in clustered federated learning. *IEEE Transactions on Networking*, pages 1–16, 2025. doi: 10.1109/TON.2025.3600674.

[32] Lawrence M. Ausubel. An efficient ascending-bid auction for multiple objects. *American Economic Review*, 94(5):1452–1475, 2004.

[33] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.

[34] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009. URL https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf.

[35] Luke N Darlow, Elliot J Crowley, Antreas Antoniou, and Amos J Storkey. Cinic-10 is not imagenet or cifar-10. *arXiv preprint arXiv:1810.03505*, 2018.

[36] Sebastian Caldas, Sai Meher Karthik Duddu, Peter Wu, Tian Li, Jakub Konečný, H. Brendan McMahan, Virginia Smith, and Ameet Talwalkar. LEAF: A benchmark for federated settings, 2018. URL https://arxiv.org/abs/1812.01097.

[37] Yann LeCun, Corinna Cortes, and Christopher J. C. Burges. The mnist dataset. (1998). The MNIST Dataset. [Online]. Available: http://yann.lecun.com/exdb/mnist/, 1998. URL http://yann.lecun.com/exdb/mnist/.

[38] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

[39] Leijie Wu, Song Guo, Zicong Hong, Yi Liu, Wenchao Xu, and Yufeng Zhan. Long-term adaptive vcg auction mechanism for sustainable federated learning with periodical client shifting. *IEEE Transactions on Mobile Computing*, 23(5):6060–6073, 2023.

A Experimental Details

Architectures. Vision jobs use ResNet-18 with batch normalization replaced by group normalization. Shakespeare uses a two-layer LSTM with hidden size 256 and a next-character softmax head, following the LEAF benchmark [36]. **Optimization.** Unless noted otherwise we use SGD with momentum 0.9, weight decay 10^{-5}, and $T = 200$ communication rounds. Vision local training uses MNIST/FMNIST: 5 local epochs and batch size 50; CIFAR-10: 4 epochs and batch 30; CINIC-10: 2 epochs and batch 10. Shakespeare uses 1 local epoch, batch size 20, and backpropagation through time length 50. **Non-IID split.** We partition training shards with a Dirichlet law ($\alpha = 0.5$) and clip each $|\mathcal{D}_j|$ to $[300, 1500]$. **Budgets and deadlines.** DC budgets follow $B \in \{100, 200, 400, 800\}$ in the vision tables. Deadlines $\overline{T}_i$ and model sizes s_i are sampled once per DC and fixed across methods; layers use $\tau_i = \overline{T}_i / s_i$. **Accuracy definition.** For each DC, evaluate its global model on a fixed held-out test set and report the arithmetic mean over DCs. **Fair comparison.** Each method receives the same data partition, model, budgets, deadlines, number of communication rounds, and random seeds; we instantiate each algorithm as an independent marketplace run.

Noise injections. On vision tasks we use the session-level label noise model in Sec. 3.1 (equivalently summarized as mean corruption levels $p \in \{10\%, 20\%\}$ in the large tables). On Shakespeare we simulate two structural perturbations of client text: token corruption (TC) applies random token replacement on each client's local corpus before training reports; sequence perturbation (SP) applies random shuffles or segment-wise reversals of token subsequences. DOs still run the same inspection interface and report corrupted counts z_j, which enter the Beta–Binomial peer kernel as in Sec. 4.

Deployments. All experiments were conducted on a platform equipped with an Intel Core i7-12700K CPU and an NVIDIA GeForce RTX 3090 GPU. The simulations were performed using MATLAB R2022a, while the model training was implemented in Python 3.10 with the PyTorch framework and CUDA 12.1.

B Timeout normalization

For a reported total C in Table 3, the normalized values are $C/200$ events per communication round across all jobs and $C/(200 \cdot 10)$ events per DC-job-round. The timeout table uses $B \in \{100, 400, 600, 800\}$, distinct from $\{100, 200, 400, 800\}$ in the welfare and accuracy tables. For comparison, FedDRL's CINIC-10 total of 1936 at $B = 800$ gives 9.68 events per round across

Table 4: Normalized MainFL timeout counts, directly derived from Table 3. Each cell shows total / per round / per DC-job-round.

Dataset	$B = 100$	$B = 400$	$B = 600$	$B = 800$
MNIST	0 / 0 / 0	0 / 0 / 0	0 / 0 / 0	0 / 0 / 0
FMNIST	0 / 0 / 0	0 / 0 / 0	0 / 0 / 0	9 / .045 / .0045
CIFAR-10	0 / 0 / 0	0 / 0 / 0	67 / .335 / .0335	145 / .725 / .0725
CINIC-10	42 / .210 / .0210	95 / .475 / .0475	138 / .690 / .0690	182 / .910 / .0910

all jobs, or 0.968 per DC-job-round. These are event counts, not probabilities or fractions of recruited owners.

C Scalability Analysis

C.1 Measured allocation scalability

We separately measure the allocation planner on synthetic markets. These measurements exclude peer scoring and offline score-table enumeration. The comparator is the same quota auction with all owners pooled.

Protocol. Let $h = M/(10N)$. Owner raw-data units are independent uniform draws on $[0.3, 1.5]$; cleanliness $q_j \sim \mathrm{Beta}(8, 2)$; latency factors $\alpha_j \sim U[0.15, 1.35]$; and normalized DC deadlines $\tau_i \sim U[0.45, 1.25]$. Consumers have $V_i(q) = a_i\sqrt{q}$, $a_i \sim U[12, 24]\sqrt{h}$, capacity $10h$, and budget $40h$. Inspection counts have $m = 200$, $z_j \sim \mathrm{Bin}(200, 1 - q_j)$, and pre-training weights $(208 - z_j)/210$

Table 5: Extended Shakespeare results. The full text setting contains 4,226,158 items and 1,129 DOs. Accuracy is reported as Highest/Average (%).

B	Method	Token corruption (TC)				Sequence perturbation (SP)			
		Data	Welfare	High.	Avg.	Data	Welfare	High.	Avg.
100	FedCA	526314	141.8	34.82	32.08	504927	132.5	32.07	29.18
	AFLM	558906	153.6	37.54	34.68	536184	143.9	34.28	31.46
	FedDRL	781442	219.7	42.91	40.18	752318	201.6	39.16	36.48
	MBOS-AFL	803175	226.4	42.28	39.55	775909	208.9	38.72	35.92
	ARAS-AFL	759830	211.5	42.56	39.82	731604	193.7	38.88	36.11
	MainFL	**872694**	**256.9**	**45.86**	**43.31**	**845218**	**235.8**	**42.35**	**39.74**
200	FedCA	982441	187.9	38.11	35.64	945203	174.6	35.04	32.28
	AFLM	1065148	202.3	42.36	39.71	1022667	188.8	38.79	36.04
	FedDRL	1428842	295.4	46.71	43.87	1379021	271.2	42.93	40.15
	MBOS-AFL	1475630	304.8	45.93	43.18	1420935	280.6	42.21	39.52
	ARAS-AFL	1389457	286.2	46.84	44.10	1338874	263.7	43.16	40.39
	MainFL	**1621749**	**340.7**	**50.16**	**47.26**	**1564908**	**314.9**	**46.12**	**43.57**
400	FedCA	1694837	260.6	41.20	38.51	1628533	241.9	37.82	35.07
	AFLM	1817962	275.1	44.38	41.63	1750184	256.6	40.91	38.22
	FedDRL	2395418	421.8	49.14	46.22	2310287	390.7	45.28	42.54
	MBOS-AFL	2482305	438.5	48.45	45.54	2391662	404.6	44.53	41.85
	ARAS-AFL	2318564	407.9	48.92	46.07	2234930	376.5	45.06	42.31
	MainFL	**2713468**	**486.1**	**52.39**	**49.61**	**2620191**	**450.8**	**48.23**	**45.51**
800	FedCA	2473184	317.4	42.96	40.19	2387461	292.8	39.47	36.89
	AFLM	2589032	336.8	45.72	43.06	2496147	310.6	42.16	39.64
	FedDRL	3135846	548.2	52.01	49.13	3030265	507.5	48.09	45.47
	MBOS-AFL	3217905	567.9	51.38	48.55	3108818	524.1	47.48	44.82
	ARAS-AFL	3024651	528.6	51.64	48.79	2924720	488.9	47.75	45.10
	MainFL	**3486124**	**622.5**	**54.96**	**52.06**	**3372438**	**579.4**	**50.88**	**48.17**

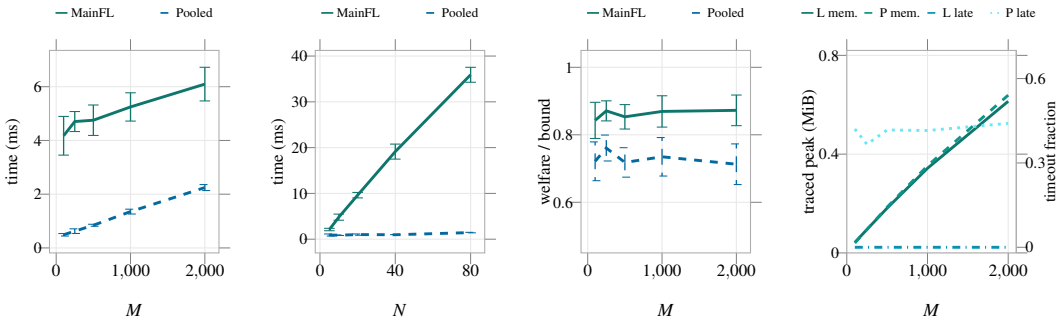

(a) Runtime vs. M. (b) Runtime vs. N. (c) Welfare / bound. (d) Memory / timeouts.

Figure 2: Measured scalability of the synthetic-market allocation planner. Lines in (a)–(c) show means and one standard deviation across ten markets; (d) shows mean traced allocation and the fraction of assigned owners violating their modeled deadline. In (d), L/P denote MainFL/pooled allocation. Timeout fractions concern one session, distinct from the cumulative events in Table 3. No neural training is included.

from a $\mathrm{Beta}(2, 8)$ corruption prior. The allocation uses 100 uniformly spaced prices on $[0, 32]$, fixed tie-breaking, the original total-supply rule, and owner retirement after a positive assignment.

Welfare is evaluated with the same true cleanliness vector across all variants, counting only on-time realized data. Define $S_{\mathrm{true}} = \sum_j q_j |\mathcal{D}_j|$. The common reference $W_{\mathrm{ub}} = \sqrt{S_{\mathrm{true}} \sum_i a_i^2}$ follows from Cauchy–Schwarz after relaxing deadlines, budgets, capacities, and owner exclusivity.

We sweep $M \in \{100, 250, 500, 1000, 2000\}$ at $N = 10$, and $N \in \{5, 10, 20, 40, 80\}$ at $M = 500$. Each point uses ten independently seeded markets (seeds 500–509). Each allocation timing is the median of three repetitions following an untimed warm-up. Memory is peak traced Python/NumPy allocation during a separate call. Measurements use an AMD Ryzen 9 6900HS CPU, Windows 11, Python 3.12.14, and NumPy 2.3.5.

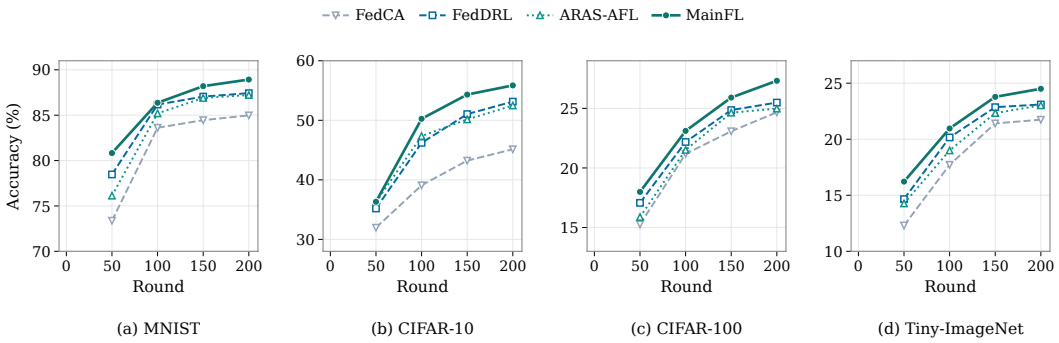

(a) MNIST (b) CIFAR-10 (c) CIFAR-100 (d) Tiny-ImageNet

Figure 3: Test accuracy at the supplied communication-round checkpoints.

Across this sweep, allocation takes milliseconds to tens of milliseconds. Growth in N is more pronounced than in M, consistent with repeated clearing of eligible prefixes; these finite measurements do not establish linear asymptotic scaling. The pooled comparator is faster but can assign owners to consumers whose deadlines they cannot meet. MainFL's zero modeled violations in this study follow from its eligibility construction, and do not explain the nonzero experimental events in Table 3.

C.2 Accuracy at communication rounds

Figure 3 shows increasing accuracy over the recorded checkpoints, with MainFL having the highest displayed accuracy at each checkpoint in all four datasets.

C.3 Theoretical Analysis

Quota realization. Positivity of Ω_j allows a contribution up to $u_j = \Omega_j|\mathcal{D}_j|$ from an owner, but does not make an arbitrary vector of aggregate quotas jointly realizable. Retiring a partially used owner can leave less supply for later consumers, even when $\sum_i Q_{i,k} \le S_k$. The main mechanism therefore guarantees realized quantities no greater than their targets, rather than universal exact realization. Appendix N gives a sufficient optional reserve condition.

Computational complexity. With one peer per DO, score computation costs $O(M)$ after the offline enumeration of the finite score table. Sorting DOs and DCs costs $O(M \log M + N \log N)$. For an ascending price grid of size P, aggregate-demand caching computes all clinches in layer k in $O(Pk)$ time, and summing layers gives $\sum_{k=1}^{N} O(Pk) = O(PN^2)$. Hence the per-session complexity is

$$O(M \log M + N \log N + PN^2). \tag{23}$$

With L independent peers per DO the scoring term becomes $O(LM)$ and the total complexity becomes

$$O(LM + M \log M + N \log N + PN^2). \tag{24}$$

C.4 Large-scale text results

Table 5 reports the submitted Shakespeare results for 1,129 DOs and 4,226,158 text items. These are next-character prediction results; they do not establish runtime or memory scaling. MainFL has the highest reported welfare and accuracy in these settings.

D Implementation details for Algorithm 1

Before stating the peer-scoring routine, we collect three implementation points that are shared by the allocation and scoring algorithms in the main text but were elided there for brevity.

Algorithm 2 Peer scoring and DO rewards (session t)

1: **Input:** reports $\{r_j^{(t)}\}_{j=1}^{M}$, public prior $p^{(t)}(\theta)$, revenue share $\lambda \in (0,1)$, escrowed prior-session revenue $R^{(t-1)}$ (with $R^{(0)} := R_0 \geq 0$ a platform-sponsored seed)
2: **Output:** scores $\{\sigma_j^{(t)}\}$ and peer rewards $\{\pi_j^{(t,\mathrm{peer})}\}$
3: Compute posteriors $p^{(t)}(\theta \mid r_j^{(t)})$ using Eq. (5)
4: Precompute bounds $(\underline{\sigma}^{(t)}, \overline{\sigma}^{(t)})$ by enumerating the finite report set $\mathcal{R}^{(t)}$
5: Commit score pool $G^{(t)} \in [0, \lambda R^{(t-1)}]$ // report-independent; drawn from prior-session escrow
6: **for** $j = 1$ to M **do**
7: Sample peer $k \neq j$ uniformly
8: Set $\sigma_j^{(t)} \leftarrow \log \mathrm{K}(r_j^{(t)}, r_k^{(t)})$ via Eq. (28)
9: **end for**
10: Pay each DO $\pi_j^{(t,\mathrm{peer})}$ by Eq. (13)
11: After session-t clearing, escrow $\lambda R^{(t)}$ to fund $G^{(t+1)}$ in the next session

Deterministic quota discretization. Within each layer, consumers are processed in ascending index order and owners in decreasing Ω_j, with owner-index tie-breaking. Once an owner contributes to a consumer, it leaves the pool, including when only part of its data is used. Contributions stop at the target or at exhaustion of unused owners, so $\widehat{Q}_{i,k} \leq Q_{i,k}$. Remaining capacity in Algorithm 1 is reduced by target quotas, giving a conservative upper-bound invariant. The original target-quota payment and reward rule is retained; if delivery is short, the exact-delivery individual-rationality claim in Theorem 5.3 does not apply. Refund or alternative settlement rules are outside the present analysis.

DO-side revenue envelope. The DO allocation rewards $\{\pi_j^{(t,\mathrm{alloc})}\}$ introduced in line 12 of Algorithm 1 satisfy, by construction and the layer-wise budget-feasibility of clinching,

$$\sum_{j=1}^{M} \pi_j^{(t,\mathrm{alloc})} \;\leq\; (1-\lambda) \sum_{i,k} p_{i,k}^{(t)} \;\leq\; (1-\lambda)\, R^{(t)},$$

where $R^{(t)} = \sum_i p_i^{(t)}$ is the realized session-t revenue. Combined with the peer envelope $\sum_j \pi_j^{(t,\mathrm{peer})} \leq \lambda R^{(t-1)}$ established below, the joint DO-side outlay is revenue-financed across sessions with $\sum_{\tau \leq t} \sum_j (\pi_j^{(\tau,\mathrm{alloc})} + \pi_j^{(\tau,\mathrm{peer})}) \leq \lambda R_0 + \sum_{\tau \leq t} R^{(\tau)}$.

Intertemporal incentive channel. Reliability enters the next session's quality weight with coefficient χ. For a fixed posterior cleanliness estimate, a higher reliability estimate raises Ω_j when $\chi > 0$. Allocation, prices, and recruitment also depend on other participants, so this algebraic monotonicity does not imply monotone or convergent total income. Full-mechanism reporting incentives require the expected-payoff condition in Appendix H.

E Peer scoring algorithms

Algorithm 1 in the main text clears the allocation/payment side of each session; it relies on the calibrated quality weights $\{\Omega_j^{(t)}\}$ computed from the submitted reports and prior reliability. The random peer draw affects the score payment, not these quality weights. The peer-scoring procedure is stated separately because it is conceptually independent of matching: it only consumes the reports $\{r_j^{(t)}\}$ and the committed score pool $G^{(t)}$, and it returns nonnegative rewards $\{\pi_j^{(t,\mathrm{peer})}\}$ together with the log-PMI scores $\{\sigma_j^{(t)}\}$.

Algorithm 2 has four stages. Lines 3–4 precompute the posteriors in Eq. (5) and the deterministic bounds $(\underline{\sigma}^{(t)}, \overline{\sigma}^{(t)})$ by enumerating the finite report interface $\mathcal{R}^{(t)}$. Line 5 commits the peer pool from the previous session's escrow, $G^{(t)} \in [0, \lambda R^{(t-1)}]$, with a platform-sponsored seed R_0 for

$t = 1$; this timing eliminates the circular dependence between $G^{(t)}$ and the yet-unrealized session-t revenue while preserving the report-independence required by Theorem 5.1. Lines 6–9 sample one peer $k \neq j$ uniformly for each DO and evaluate the log-PMI score via the closed-form Beta–Binomial kernel Eq. (28). Line 10 applies the affine rescaling Eq. (13). Line 11 escrows $\lambda R^{(t)}$ to fund $G^{(t+1)}$, closing the intertemporal ledger; the long-run self-financing bound and the compatibility with the DO allocation rewards of Algorithm 1 are recorded in Appendix D. The procedure is linear in M after the one-off score-table precomputation.

F Posterior Agreement Algebra

Fix a session and omit the superscript t. Let R and S be two reports drawn conditionally independently given θ. We first derive the representation used in Eq. (11):

$$\begin{aligned}\mathrm{K}(r,s) &= \int_\Theta \frac{p(\theta \mid r)p(\theta \mid s)}{p(\theta)}\,d\theta = \int_\Theta \frac{p(r \mid \theta)p(\theta)}{p(r)} \cdot \frac{p(s \mid \theta)p(\theta)}{p(s)} \cdot \frac{1}{p(\theta)}\,d\theta \\ &= \frac{1}{p(r)p(s)} \int_\Theta p(r \mid \theta)p(s \mid \theta)p(\theta)\,d\theta = \frac{p(r,s)}{p(r)p(s)}. \end{aligned} \tag{25}$$

By conditional independence in Eq. (4), the integral equals the joint marginal $p(r, s)$. Thus $\log \mathrm{K}(r, s) = \log p(s \mid r) - \log p(s)$ is exactly a pointwise mutual-information score between the two reports.

Under the Beta–Binomial model,

$$p(z \mid m) = \binom{m}{z} \frac{\mathrm{B}(a_0 + z,\ b_0 + m - z)}{\mathrm{B}(a_0, b_0)}, \tag{26}$$

$$p(z_a, z_b \mid m_a, m_b) = \binom{m_a}{z_a}\binom{m_b}{z_b} \frac{\mathrm{B}(a_0 + z_a + z_b,\ b_0 + m_a + m_b - z_a - z_b)}{\mathrm{B}(a_0, b_0)}. \tag{27}$$

Dividing the joint probability by the marginals gives the closed form (with B denoting the beta function):

$$\mathrm{K}(r_a, r_b) = \frac{\mathrm{B}(a_0, b_0)\,\mathrm{B}(a_0 + z_a + z_b,\ b_0 + m_a + m_b - z_a - z_b)}{\mathrm{B}(a_0 + z_a,\ b_0 + m_a - z_a)\,\mathrm{B}(a_0 + z_b,\ b_0 + m_b - z_b)}, \tag{28}$$

Strict positivity of a_0, b_0 implies strictly positive predictive probabilities everywhere on the finite report set, so $\mathrm{K} > 0$ throughout, and both $\underline{\sigma}$ and $\overline{\sigma}$ are finite.

G Proof of the DO Peer-Scoring Theorem

We first prove the reporting-subgame result, then conclude score-pool safety.

Reporting subgame. Fix a DO with true type r and let S be the peer's report. Assume peers report truthfully, so $S \sim p(\cdot \mid r)$ in distribution. If our DO reports $a \in \mathcal{R}$, its conditional expected log-PMI is

$$\mathbb{E}[\log \mathrm{K}(a, S) \mid r] = \sum_s p(s \mid r) \log \frac{p(s \mid a)}{p(s)} = \sum_s p(s \mid r) \log p(s \mid a) - \sum_s p(s \mid r) \log p(s). \tag{29}$$

The second term is independent of a. Adding and subtracting $\sum_s p(s \mid r) \log p(s \mid r)$ rewrites the first term as

$$\sum_s p(s \mid r) \log p(s \mid a) = \sum_s p(s \mid r) \log p(s \mid r) - \mathrm{KL}(p(S \mid r) \,\|\, p(S \mid a))\,. \tag{30}$$

Substituting into Eq. (29) yields

$$\mathbb{E}[\log \mathrm{K}(a, S) \mid r] = \underbrace{\sum_s p(s \mid r) \log \frac{p(s \mid r)}{p(s)}}_{\text{constant in } a} - \mathrm{KL}(p(S \mid r) \,\|\, p(S \mid a))\,. \tag{31}$$

Because KL divergence is nonnegative and vanishes only when the two distributions coincide, every maximizer $a^\star$ of $\mathbb{E}[\log \mathrm{K}(a, S) \mid r]$ satisfies $p(S \mid a^\star) = p(S \mid r)$, i.e. $a^\star$ is posterior-equivalent to r. In particular, $a = r$ is a maximizer. Because the same argument applies symmetrically to each DO conditional on truthful play of the others, truthful (posterior-equivalent) reporting is a Bayes–Nash equilibrium of the peer-scoring subgame.

When $G > 0$ and $\overline{\sigma} > \underline{\sigma}$, the reward is an affine transformation with strictly positive slope and preserves the score maximizers. If $G = 0$, all reports earn the same zero peer reward and truthful reporting is only a weak best response. If the score span is zero, the prescribed reward is zero as well.

Score-pool safety. For any realized score $\sigma_j \in [\underline{\sigma}, \overline{\sigma}]$ we have $(\sigma_j - \underline{\sigma})/(\overline{\sigma} - \underline{\sigma}) \in [0, 1]$, hence $0 \le \pi_j \le G/M$. Summing over j gives $\sum_j \pi_j \le G \le \lambda R^{(t-1)}$, which establishes the budget-safety statement.

H Score Dominance with Allocation-Side Effects

Fix a session, condition on the public history before reports, and suppose the peer-report distribution used by the score is correctly specified. Write $D_\sigma = \overline{\sigma} - \underline{\sigma} > 0$ and $\eta_{\mathrm{pp}} = G/(MD_\sigma)$. For true report r and deviation a, define

$$\Delta_{\mathrm{pp}}(r, a) = \mathrm{KL}\big(p(S \mid r) \,\|\, p(S \mid a)\big). \tag{32}$$

Let $A_j(a; r_{-j}, h)$ include all non-score utility, including allocation rewards, resource costs, and bounded discounted continuation utility, conditional on history h. Define its conditional expected deviation gain as

$$g_j(r, a \mid h) = \mathbb{E}[A_j(a; r_{-j}, h) - A_j(r; r_{-j}, h) \mid r, h].$$

Corollary H.1 (Conditional full-mechanism DO reporting)**.** *Given truthful reports by the other DOs, truthful reporting maximizes DO j's total conditional expected utility if, for every admissible deviation a,*

$$g_j(r, a \mid h) \le \frac{G}{MD_\sigma} \Delta_{\mathrm{pp}}(r, a). \tag{33}$$

This includes $g_j(r, a \mid h) \le 0$ for predictive-equivalent deviations, whose score gap is zero.

Proof. By the KL identity in Appendix G, deviating loses expected peer reward $G\Delta_{\mathrm{pp}}(r, a)/(MD_\sigma)$. Adding the non-score gain gives total deviation gain $g_j - G\Delta_{\mathrm{pp}}/(MD_\sigma) \le 0$. □

On a finite interface with at least one non-equivalent pair,

$$\delta_{\min} = \min_{r, a:\, p(S|r) \ne p(S|a)} \Delta_{\mathrm{pp}}(r, a) > 0. \tag{34}$$

If $g_j \le \overline{A}$ uniformly, a sufficient threshold for non-equivalent deviations is $G \ge MD_\sigma \overline{A}/\delta_{\min}$. It must coexist with $G \le \lambda R^{(t-1)}$ and the separate condition for zero-gap deviations. A finite threshold need not fit within available escrow. Averaging more peers removes sampling variability but does not increase the expected KL gap or cure insufficient funding.

H.1 Peer averaging and model scope

For L independent uniform peer draws, average their log-PMI scores before applying the same affine normalization. Linearity of expectation preserves the preceding incentive calculation. Averaging all $M-1$ other reports removes pairing randomness conditional on the report profile; signal uncertainty remains. It requires $O(M^2)$ score lookups rather than $O(M)$ for one peer.

The KL proof requires a correctly specified, strictly positive predictive distribution on the finite report interface; Beta–Binomial conjugacy is one instantiation. In a hierarchical model with owner-specific rates drawn independently conditional on a shared hyper-state, the same argument applies to the induced hierarchical marginal ratio $p(r_a, r_b)/(p(r_a)p(r_b))$. It does not justify using a misspecified Beta–Binomial kernel. If reports are independent across owners, $K = 1$ and all score gaps vanish, so peer scoring supplies no strict reporting incentive.

H.2 Limited coalition exposure

For a fixed coalition of size c, uniform matching pairs a member internally with probability $q = (c-1)/(M-1)$. If truthful outside peers provide a conditional expected loss of at least $\delta > 0$ for the deviation being considered, while the score advantage on an internal match is bounded by D_σ, then its expected score gain is at most $qD_\sigma - (1-q)\delta$. Thus $qD_\sigma < (1-q)\delta$ is a sufficient condition against that deviation in the peer-scoring channel. This conditional bound is not coalition-proofness of the allocation mechanism, and provides no protection for zero-gap deviations.

I Reliability Concentration

Suppose the sequence of observed recruited-session signals for owner j is i.i.d. Bernoulli with fixed parameter ϱ_j. At a fixed observation count $n \geq 1$, let $S_{j,n} = \sum_{\ell=1}^{n} Y_{j,\ell}$ and $\rho_{j,n} = (S_{j,n}+1)/(n+2)$. Then

$$\mathbb{E}[\rho_{j,n}] = \varrho_j + \frac{1-2\varrho_j}{n+2}, \tag{35}$$

and Hoeffding's inequality gives

$$\Pr\left(\left|\rho_{j,n} - \frac{n\varrho_j+1}{n+2}\right| \geq \epsilon\right) \leq 2\exp\left(-\frac{2(n+2)^2\epsilon^2}{n}\right). \tag{36}$$

These are fixed-observation-count statements. They must not be read as bounds conditional on an adaptively selected recruitment count. If $N_j(t) \to \infty$, the strong law still gives $\rho_{j,N_j(t)} \to \varrho_j$ almost surely under the stated i.i.d. sequence assumption. Nonstationary validation signals, finite recruitment, and biased validation distributions can invalidate an interpretation of ρ_j as intrinsic data quality; neither income convergence nor fairness follows from this calculation.

J Laminar Deadline Structure

By the factorization in Eq. (8),

$$T_{i,j} \leq \overline{T}_i \quad \Longleftrightarrow \quad s_i\alpha_j \leq \overline{T}_i \quad \Longleftrightarrow \quad \alpha_j \leq \tau_i.$$

Therefore the feasible DO set of DC i is $\tilde{\mathcal{M}}_i = \{j : \alpha_j \leq \tau_i\}$. Ordering $\tau_{(1)} \geq \cdots \geq \tau_{(N)}$ gives a nested chain

$$\tilde{\mathcal{M}}_{(N)} \subseteq \tilde{\mathcal{M}}_{(N-1)} \subseteq \cdots \subseteq \tilde{\mathcal{M}}_{(1)}.$$

Define $\hat{\mathcal{M}}_k = \tilde{\mathcal{M}}_{(k)} \setminus \tilde{\mathcal{M}}_{(k+1)}$ (with $\tilde{\mathcal{M}}_{(N+1)} = \varnothing$); this is precisely Eq. (17). The layers form a disjoint partition of $\bigcup_i \tilde{\mathcal{M}}_i$, and any DO $j \in \hat{\mathcal{M}}_k$ is feasible for every DC in $\mathcal{F}_k = \{(1), \ldots, (k)\}$ and infeasible for every DC outside $\mathcal{F}_k$.

K Feasibility Guarantees

Theorem K.1. *For every session t, Algorithm 1 satisfies: (i) each DO is assigned to at most one DC; (ii) each DC effective-capacity constraint $\hat{q}_i^{(t)} \leq \gamma_i$ holds; (iii) each DC payment satisfies $p_i^{(t)} \leq B_i$; and (iv) every assigned pair satisfies $x_{i,j}^{(t)} T_{i,j}^{(t)} \leq x_{i,j}^{(t)} \overline{T}_i$.*

(i) One-to-one DO assignment. Each DO belongs to at most one layer by Appendix J. Inside a layer, the discretization step processes every DO at most once and assigns it to a single DC. Therefore no DO is assigned to more than one DC across all layers.

(ii) Effective-capacity feasibility. Clinching is truncated by the remaining target capacity, and each layer subtracts $Q_{i,k}$ from that capacity. Hence $\sum_k Q_{i,k} \leq \gamma_i$. Discretization enforces

$$0 \leq \widehat{Q}_{i,k} = \sum_{j \in \hat{\mathcal{M}}_k} \Omega_j d_{i,j} \leq Q_{i,k}. \tag{37}$$

Thus $\hat{q}_i = \sum_k \widehat{Q}_{i,k} \leq \sum_k Q_{i,k} \leq \gamma_i$, without requiring exact quota realization.

(iii) Budget feasibility. At a price-grid step with $\psi > 0$, an admissible increment satisfies $0 \le \Delta Q_i \le B_i^{\mathrm{rem}}/\psi$. Charging $\psi \Delta Q_i$ and immediately updating the remaining budget therefore preserves $B_i^{\mathrm{rem}} \ge 0$. At price zero the charge is zero. Applying this invariant over the price steps and then over layers yields $p_i \le B_i$. The continuous-clock statement follows along an admissible limiting payment path.

(iv) Deadline feasibility. Any assigned pair (i, j) belongs to some layer k with $i \in \mathcal{F}_k$ and $j \in \hat{\mathcal{M}}_k$. By Appendix J, j is feasible for every DC in $\mathcal{F}_k$, hence $T_{i,j} \le \overline{T}_i$.

L Restricted Per-Layer DC Incentives

Fix the entering holdings q_i^0, remaining capacities, and supply of an isolated layer. Normalize $V_i(0) = 0$ and use the incremental value $W_i(q) = V_i(q_i^0 + q) - V_i(q_i^0)$, which is nondecreasing and concave. Write $v_i = V_i'$ (or a monotone marginal-value selection). The residual-demand policy is

$$b_i^\star(\psi) = \sup\{q \in [0, \Gamma_i(\psi)] : \ \psi q \le B_i^{\mathrm{rem}}(\psi), \ v_i(q_i^0 + Q_i(\psi) + q) \ge \psi\}. \tag{38}$$

At price zero the budget restriction is vacuous. If budgets do not bind, this is the sincere demand schedule for a standard ascending clinching auction with private diminishing marginal values. The isolated-layer incentive conclusion in Theorem 5.3 is the standard private-value clinching result [32], applied to W_i. It is not inferred solely from budget truncation, and is not extended here to binding budgets or interdependent values. A finite price-grid implementation is an approximation to this exact process, not an exact incentive theorem.

For individually rational trade, every unit acquired at price ψ has marginal value at least ψ. Under exact delivery,

$$V_i(q_i^0 + Q_i^\star) - V_i(q_i^0) - p_{i,k} = \int_0^{Q_i^\star} v_i(q_i^0 + q)\, dq - \int \psi \, dQ_i(\psi) \ge 0. \tag{39}$$

This inequality concerns the incremental trade in the isolated layer; no statement about strategic continuation follows. If the delivered amount is smaller than $Q_i^\star$ while the target-quota payment is retained, the inequality is not asserted for the realized utility.

M Cross-Layer Incentive Scope

Per-layer incentive compatibility alone does not establish a truthful equilibrium of the sequential composition. Concavity of values and monotonicity of continuation surplus in remaining budget do not rule out profitable deferral to a later, cheaper layer. Accordingly, this paper makes no cross-layer DC incentive-compatibility claim. A continuation-aware analysis or additional restrictions would be needed for such a result.

N Optional Extension: Supply Reserve for Exact Quota Realization

This extension is separate from the total-supply rule in Eq. (18) and was not used to produce the reported experimental tables. Its reserve can be large in sparse layers, so it is not assumed to be a suitable replacement in every market.

In a nonempty layer, write $u_j = \Omega_j |\mathcal{D}_j|$, $S = \sum_j u_j$, $u_{\max} = \max_j u_j$, and $m = |\mathcal{F}|$. Before demand submission, post

$$\widetilde{S} = [S - (m-1)u_{\max}]_+. \tag{40}$$

For an empty layer take $\widetilde{S} = 0$. Run the same quota auction on $\widetilde{S}$, and process quotas and owners in fixed order, retiring each assigned owner even when partially used.

Proposition N.1 (A sufficient reserve)**.** *Every nonnegative quota vector with $\sum_{i=1}^m Q_i \le \widetilde{S}$ is realized exactly by the fixed-order capped assignment, while each owner serves at most one consumer.*

Proof. The case $\widetilde{S} = 0$ has only zero quotas. Otherwise $\widetilde{S} = S - (m-1)u_{\max}$. Each earlier quota boundary discards at most $u_{\max}$ of its last owner. Before filling quota i, the unused supply is

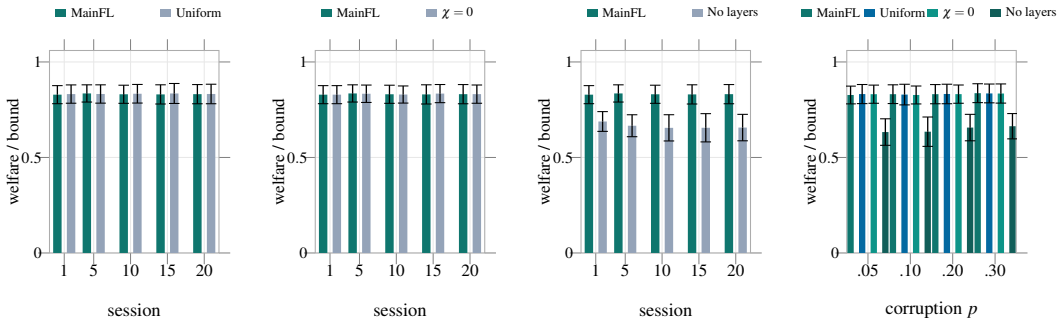

(a) Uniform weights. (b) No calibration. (c) No layers. (d) Noise sweep.

Figure 4: Paired synthetic-market ablations. Bars show mean normalized realized welfare and one standard deviation across twenty seeds; (a)–(c) use $p = 0.20$, and (d) shows session 20.

therefore at least

$$S - \sum_{h<i} Q_h - (i-1)u_{\max} \geq \sum_{h \geq i} Q_h + (m-i)u_{\max} \geq Q_i.$$

Thus the next quota can be filled, and induction fills all quotas. Retiring each assigned owner enforces the one-consumer constraint. □

The reserve is sufficient, not necessary. Positive posted supply requires $S > (m-1)u_{\max}$; a layer with few large owners may otherwise post zero. The construction establishes exact realization only. It does not restore cross-layer truthfulness or general incentive compatibility with binding budgets.

O Paired Synthetic-Market Ablation

This study isolates inputs to the allocation planner. We use the preceding synthetic-market model at $M = 100$, $N = 10$, twenty sessions, and seeds 0–19. For nominal corruption parameter $p \in \{0.05, 0.10, 0.20, 0.30\}$, draw $q_j = \mathrm{clip}(Q_j + 0.20 - p, 0.03, 0.99)$ with $Q_j \sim \mathrm{Beta}(8, 2)$. Market profiles, inspection counts, and potential validation outcomes are shared across variants within each seed. Owner-specific qualities are independent here, so this experiment does not test the shared-latent peer-scoring truthfulness theorem.

Inspection weights are refreshed each session. Recruited owners receive a binary validation observation with success probability q_j^2 if on time, and zero otherwise; Beta-smoothed counters update only after recruitment, which is an explicit market-level proxy. The first session uses $\chi = 0$; later sessions use $\chi = 0.4$. We compare: (i) full valuation and layering; (ii) replacing the pre-training weight by 0.8; (iii) $\chi = 0$ throughout; and (iv) pooled allocation without deadline eligibility. Variant (ii) removes the inspection-based quality input, not merely the peer-payment calculation, which would leave a fixed-report allocation unchanged.

All welfare comparisons use the same true-quality, on-time evaluator and upper bound within each market. Figure 4 reports all twenty seeds without selecting favorable runs. At session 20 with $p = 0.20$, paired unnormalized welfare differences (full minus ablated; mean $\pm$ standard deviation) are -0.56 ± 7.32 for uniform pre-training weights, -0.08 ± 8.56 for no calibration, and 88.32 ± 42.70 for no deadline layers. Thus this synthetic setting shows a clear role for deadline eligibility, while the quality-input and calibration differences are small.

P Supplementary Audits and Sensitivity

P.1 Scope of the additional evidence

The coefficient η in Eq. (14) is a validation-probe coefficient. For the unweighted local update $\Delta w_j = w_{j,\mathrm{local}} - w$, the reference implementation evaluates the local model directly, corresponding to $\eta = 1$. The formula permits other fixed probe coefficients.

Table 6: Maximum conditional expected two-session deviation gain in the finite family.

Funding case	G	Required R_{prev} or R_0	Max. gain	Positive
No current peer pool	0.0000	0.0000	0.4364	114/288
Recurring-revenue upper cap	0.8000	4.0000	0.4072	40/288
Certified sponsor seed	54.0345	270.1724	0.0000	0/288

The integration couples a finite inspection interface, deterministic score bounds, report-independent prior-revenue funding, calibrated effective supply, and deadline-feasible allocation. Random peer selection affects the peer transfer, while the owner's inspection report determines its posterior cleanliness weight. The L-peer and all-other-peer variants and their cost are described in Appendices C.3 and H.

An optional computational extension can group sorted owner latency factors into K groups and test every group using its largest latency factor. Only consumers meeting this conservative group threshold remain eligible. Every retained pair is then feasible, although feasible opportunities may be discarded. The repeated-clearing term is at most $O(PKN)$, in addition to sorting and the other preprocessing costs. The analysis fixes the roster and reporting interface before report collection and funding commitment. Setting $\chi = 0$ is the fallback when calibration history is unavailable.

P.2 A finite two-session unilateral-report audit

There are $M = 3$ owners, $N = 2$ consumers, and inspection size $m = 2$. The admissible report component is $z \in \{0, 1, 2\}$. Each session has a fresh shared $\theta \sim \mathrm{Beta}(2, 8)$ and conditionally independent $z_j \sim \mathrm{Bin}(2, \theta)$. Other owners report truthfully. We enumerate each current unilateral z-deviation and fix truthful reporting in the next session. The utility evaluated is $U_j = \pi_{j,1}^{\mathrm{peer}} + \pi_{j,1}^{\mathrm{alloc}} + 0.9\big(\pi_{j,2}^{\mathrm{peer}} + \pi_{j,2}^{\mathrm{alloc}}\big)$, with resource costs set to zero. Both recruitment-dependent reliability and the report-dependent current revenue feeding the next escrow are included. The next pool is $0.2R_1$; the current pool is fixed before reports.

Two public market layouts, each with seeds 0–7, give sixteen instances. Data and capacity quantities below use normalized units; the inspection count is a separate report-interface size. Consumer values are $a_i\sqrt{q}$ with $a_i \sim U[2.5, 4]$, budgets are 2 each, and normalized deadlines are $(1.2, 0.7)$. The bottleneck layout uses $|\mathcal{D}_j| \sim U[0.8, 1.4]$, latency factors $(0.3, 0.65, 1.05)$, and capacity 1.5 per consumer. The common-layer layout uses $(1.5, 1.4, 0.8)$ multiplied componentwise by independent $U[0.97, 1.03]$ draws, latency factors $(0.3, 0.45, 0.6)$, and capacity 1.1 per consumer. With $\mu = \kappa_{\mathrm{cpu}} = 1$, uplink rate 10 and compute rate $|\mathcal{D}_j|/(\alpha_j - 0.1)$ implement these times. Allocation retains total supply and owner retirement, with seventeen prices on $[0, 4]$. Calibration is zero initially and 0.4 next session.

Potential validation observations are conditionally independent Bernoulli$((1-\theta)^2)$ signals; only recruited owners update their Laplace-smoothed counters. Expectations sum all 27 actual report profiles, eight potential-signal profiles, 27 truthful future report profiles, and uniform peer choices. The signal probabilities are integrated under the conditional Beta posterior. No sampled peer noise is used. The 288 nontrivial owner/type/action comparisons include 100 cases with nonzero next-allocation-income changes.

The minimum predictive KL gap is 0.030631, the score span is 1.203973, and the largest expected non-score gain in this family is 0.436423. Applying the uniform bound in Appendix H gives a sufficient pool of 51.4614 for these current deviations under the fixed truthful continuation; the final row uses a 5% margin. This pool is chosen offline over the entire finite family, not after observing a realized report. Direct enumeration gives the smaller family-specific threshold 27.1533.

P.3 Paired sensitivity to calibration and inspection size

We extend the synthetic-market protocol of Appendix O, with $M = 100$, $N = 10$, twenty sessions, seeds 0–19, nominal corruption $p \in \{0.10, 0.20\}$, the original total-supply rule, and 100 prices on $[0, 32]$. This is an allocation-level study with the same Bernoulli validation proxy, not a CIFAR-10 training experiment. We vary $\chi \in \{0, 0.2, 0.4, 0.6, 0.8\}$ at $m = 200$, and $m \in \{50, 100, 200, 400\}$

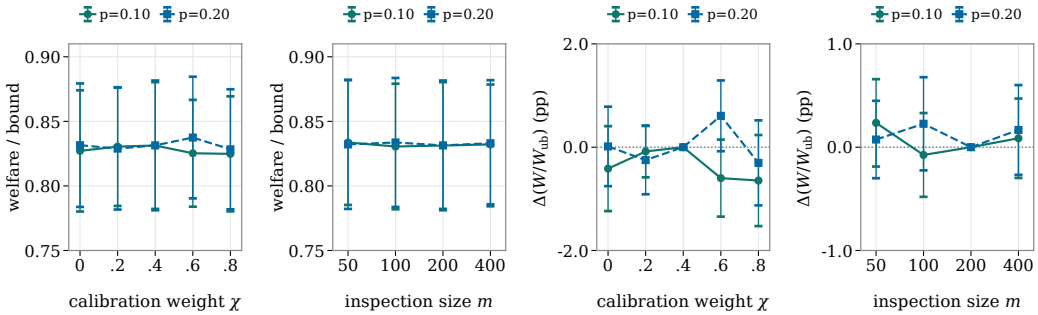

(a) Welfare vs. χ (b) Welfare vs. m (c) Paired effect: χ (d) Paired effect: m

Figure 5: Paired sensitivity at session 20. Panels (a)–(b) show mean W/W_{ub} and one sample standard deviation over twenty seeds. Panels (c)–(d) show paired differences from $\chi = 0.4, m = 200$, expressed in percentage points of W/W_{ub}, with nominal 95% Student-t confidence intervals.

at $\chi = 0.4$; the shared default is counted once. This yields $6{,}400$ observations across eight configurations.

All configurations share market profiles and potential validation outcomes. Inspection samples are nested: the archived $m = 200$ count is the anchor, smaller samples are conditional hypergeometric subsamples, and $m = 400$ adds 200 independent inspections. Each size therefore retains its Binomial marginal law while the default random-number path is preserved. The posterior weight is $(8 + m - z)/(10 + m)$. The first session uses $\chi = 0$ for every configuration. All $1{,}600$ default and no-calibration anchor observations agree with the archived experiment, with zero numeric difference.

Across nondefault settings, mean paired relative welfare changes range from -0.72% to 0.77%, and every nominal paired interval includes zero. The effects are modest within this model and sweep; the experiment does not establish a uniquely optimal χ, monotone gains from larger inspection samples, or trained-accuracy robustness. All modeled deadline violations are zero because deadline eligibility is enforced.

Q Additional Limitations

Full-mechanism DO truthfulness requires escrow-compatible score dominance and a correctly specified peer distribution. Independent reports and predictive-equivalent deviations can eliminate the score margin; random matching is not coalition-proofness. DC incentives have the restricted scope of Assumption 5.2; original quotas may be underfilled. The conservative full-data test can exclude owners feasible at smaller allocations. Nonconcave or compositional preferences, identity changes, arbitrary churn, and inaccurate telemetry are outside the guarantees. Validation can be distributionally biased; setting $\chi = 0$ removes calibration while retaining the peer-scoring subgame result.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and Section 1 distinguish peer-scoring truthfulness, conditional full-mechanism DO reporting, restricted per-layer DC incentives, and descriptive empirical gains. The optional supply reserve is not claimed as the experimental configuration.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: Appendix Q discusses score dominance, model specification, quota shortfall, the scope of DC incentives, validation bias, coalition manipulation, and modeled latency.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: Assumptions are stated in Section 5. The DO score and dominance proofs appear in Appendices G and H; feasibility is proved in Appendix K. Appendix L states the standard result used for the restricted layer setting, and Appendix N proves the optional reserve condition. Cross-layer DC truthfulness is not claimed.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [No]

Justification: The mechanism is specified in equations and pseudocode, and Appendix A describes the training and comparison protocol. Exact training seeds and all tuning details are not exhaustively documented.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility.

In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [No]

   Justification: The benchmark datasets are public and cited. This paper does not provide an independently verified public code release with complete reproduction instructions.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [No]

   Justification: Datasets, architectures, optimizers, round and local-epoch counts, non-IID splitting, and the comparison protocol are reported. The full seed list and learning-rate schedules are not specified.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [No]

   Justification: Table 1 reports welfare as mean plus or minus one standard deviation across independently seeded marketplace runs. The new market-level figures also report mean plus or minus one standard deviation across their stated seeds. Accuracy and timeout entries in the main tables remain point summaries without uncertainty intervals or formal significance tests.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [No]

   Justification: Appendix A lists the reported hardware and software environment, and Appendix C.3 gives analytical mechanism complexity. The new CPU allocation profiles specify timing and memory scope separately. Total training time and compute for the original benchmark runs are not reported.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The experiments use public benchmark datasets or fully synthetic markets, with simulated participants and no recruited human subjects. Peer reporting does not itself provide a formal privacy guarantee.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: The impact statement and Appendix Q discuss potential resource benefits, validation bias, coalition manipulation, and identity and telemetry concerns.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
    - The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
    - If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: The paper proposes an auction mechanism for federated learning and does not release any pre-trained generative model, scraped dataset, or asset with high risk for misuse.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [No]

Justification: The datasets, model architectures, and baseline mechanisms are cited. An exhaustive asset-version and license inventory is not provided.

Guidelines:

- The answer [N/A] means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [N/A]

Justification: No new dataset or pretrained model is released. The contribution is an auction and valuation mechanism described in equations and pseudocode.

Guidelines:

- The answer [N/A] means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [N/A]

Justification: The paper does not involve crowdsourcing or research with human subjects. The benchmark experiments use public datasets with synthetic market participants; the supplementary allocation studies use fully synthetic markets.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [N/A]

Justification: The paper does not involve human subjects or crowdsourced participants, so IRB approval is not applicable.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: LLMs are not components of the proposed valuation or allocation mechanism. This declaration concerns the research methodology, rather than manuscript editing.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.