FEDERATED LEARNING OF QUANTILE INFERENCE UNDER LOCAL DIFFERENTIAL PRIVACY

Anonymous authors
Paper under double-blind review

ABSTRACT

In this paper, we investigate federated learning for quantile inference under local differential privacy (LDP). We propose an estimator based on local stochastic gradient descent (SGD), whose local gradients are perturbed via a randomized mechanism with global parameters, making the procedure tolerant of communication and storage constraints without compromising statistical efficiency. Although the quantile loss and its corresponding gradient do not satisfy standard smoothness conditions typically assumed in existing literature, we establish asymptotic normality for our estimator as well as a functional central limit theorem. The proposed method accommodates data heterogeneity and allows each server to operate with an individual privacy budget. Furthermore, we construct confidence intervals for the target value through a self-normalization approach, thereby circumventing the need to estimate additional nuisance parameters. Extensive numerical experiments and real data application validate the theoretical guarantees of the proposed methodology.

1 INTRODUCTION

Quantile estimation and inference are critical tools in various scientific and applied domains. In healthcare, quantile methods facilitate more informed decisions regarding the optimal distribution of scarce medical resources, thus promoting equitable and effective patient care (Yadlowsky et al., 2025). Similarly, quantile techniques have proven highly valuable in policy evaluation, as they reveal heterogeneous effects across different subgroups, nuances typically obscured by traditional average-based analyses (Kallus et al., 2024; Chernozhukov & Fernández-Val, 2011; Chernozhukov & Hansen, 2005). In reliability engineering, quantile-based approaches have significantly improved the assessment of system robustness, particularly under rare or extreme conditions, demonstrating their broad applicability and precision (He et al., 2023; Hu et al., 2022). Moreover, finance widely employs quantile-based metrics such as value-at-risk, essential for managing financial risks in the face of regulatory pressures and market uncertainties (Barbaglia et al., 2023; Chen, 2008; Wang et al., 2012). In general, quantile methods excel at capturing the characteristics of skewed or extreme-valued data, delivering richer insights into complex distributions prevalent in practical scenarios (Angrist et al., 2006; Chen et al., 2023).

Traditional quantile estimation methods have been extensively studied. However, with the rapid increase in massive datasets (Jordan et al., 2019; Hector & Song, 2021; Fan et al., 2023), traditional approaches that rely on analyzing all data on a single machine may no longer be computationally feasible. This challenge has motivated the emergence of federated learning methods (McMahan et al., 2017; Liu et al., 2020; Tian et al., 2023). Federated learning enables multiple distributed clients to collaboratively train a global model without exchanging raw data, effectively addressing computational efficiency and privacy concerns (Konečnỳ et al., 2016). In standard federated learning, a central server coordinates iterative model updates among clients, and under suitable conditions, this process guarantees convergence (Li et al., 2020; Chen et al., 2022). To further enhance communication efficiency, local stochastic gradient descent (SGD) has been proposed, allowing clients to perform multiple local updates before synchronization. Under i.i.d. scenarios, the theoretical optimality of local SGD has been established (Stich, 2018). However, data heterogeneity, which frequently occurs in federated learning, significantly complicates local SGD. A series of studies have investigated this issue by analyzing convergence in worst-case heterogeneous scenarios (Hu et al., 2024), proposing regularization techniques to ensure local models remain close to the global model

(Li et al., 2020), and introducing momentum-based algorithms to stabilize training under non-i.i.d. conditions (Li et al., 2025). Moreover, inference methods have also been developed and analyzed (Li et al., 2022).

Federated learning aggregates individual information to enable efficient collaborative model training. These personal data power indispensable services, from facial recognition unlocking our phones to recommendation systems curating news feeds, but they also carry latent risks. Leaked genetic markers can jeopardize insurance rates, and cleverly crafted prompts can coax large language models into regurgitating fragments of their private training corpora Nasr et al. (2023). Differential privacy (DP) offers a principled defense: by bounding the statistical influence of any single participant, DP ensures that outputs remain virtually unchanged whether or not an individual opts in Dwork et al. (2006). This safeguard, however, dissolves if the data custodian is breached, coerced, or simply misconfigured access controls—scenarios illustrated by repeated healthcare leaks and high-profile cloud missteps. Local differential privacy (LDP) fortifies the pipeline by introducing randomness at the point of collection: users perturb their data locally, send only noisy summaries, and retain the key to their raw information Duchi et al. (2013). Even a fully compromised server therefore receives nothing decipherable. Industry adoption is accelerating: Google's RAPPOR measures Chrome settings, iOS uses LDP to count emoji preferences, and Windows telemetry applies similar techniques to malware prevalence Erlingsson et al. (2014); Ding et al. (2017). Collectively, these systems prove that granular user analytics and uncompromising privacy need not be mutually exclusive; instead, LDP sets a practical, legally robust baseline for responsible data-driven innovation.

DP federated learning has attracted considerable attention recently (e.g., (Liu et al., 2023a; Agarwal et al., 2018; Shi et al., 2022; Ma et al., 2022)). The additional communication layer between local clients and the global server gives rise to distinct privacy requirements. As delineated in (Lowy & Razaviyayn, 2023), one can categorize DP at the individual record level, inter-silo record level, shuffled-model, and user-level, in order of increasing trust assumptions. In particular, LDP posits that each individual does not trust any other party, including their own silo, and therefore must randomize her report before release. Extensive work has focused on this setting (e.g., (Zhao et al., 2020; Shen et al., 2023; Jiang et al., 2022)).

Whereas prior studies of LDP in federated learning (e.g., (Zhao et al., 2020; Shen et al., 2023; Jiang et al., 2022)) primarily address estimation, statistical inference, such as constructing confidence intervals and conducting hypothesis tests, poses additional challenges. Beyond deriving the limiting distribution, inference requires a consistent estimator of the asymptotic variance. For SGD-based methods, this typically involves the Hessian matrix, which exists only for smooth loss functions (Chen et al., 2020). Moreover, because only privatized gradients are observed, one may need extra privacy budget or data-splitting to estimate variance reliably. Finally, existing single-machine LDP quantile algorithms, such as Huang et al. (2021) or Liu et al. (2023b) cannot derive the inference result or do not readily extend to federated settings due to client-heterogeneity in local loss functions.

In this paper, we propose a novel federated learning algorithm for quantile inference under LDP. Our method accommodates client-level heterogeneity in quantile targets, privacy budgets, and data distributions, thereby enhancing the applicability of quantile inference in realistic federated environments. A key innovation is our theoretical analysis of the local SGD quantile estimator. We first design an LDP mechanism that effectively reduces the federated quantile estimation problem to an equivalent non-private setting. Exploiting this reduction, we establish the estimator's asymptotic normality and derive a functional central limit theorem without average-smoothness condition on the loss function. To the best of our knowledge, this constitutes the first weak-convergence result for local SGD when the loss does not satisfy the usual average-smoothness condition (Li et al., 2022; Xie et al., 2024; Zhu et al., 2024). Building on these non-private asymptotic results, we develop an LDP-compliant inference procedure for federated quantile estimation. By employing a self-normalization technique, we avoid direct estimation of the asymptotic variance, instead constructing confidence intervals that automatically eliminate the unknown variance term. To the best of our knowledge, we provide the first inference framework for federated quantile estimation, even without privacy constraints.

The remainder of the paper is organized as follows. Section 2 reviews background and notation. Section 3 presents the asymptotic analysis of the proposed estimator. Section 4 reports extensive numerical experiments and real data application. All technical proofs and additional simulation results are deferred to the Appendix.

2 METHODOLOGIES

First, we recall the definitions of central and local differential privacy. We then describe our problem setting and algorithmic details.

Definition 1 (Central Differential Privacy, CDP (Dwork et al., 2006)). *A randomized algorithm $\mathcal{A}$ operating on a dataset S is* (ϵ, δ)-differentially private *if, for any pair of datasets S and S' differing in a single record and for any measurable set E,*

$$\Pr\big[\mathcal{A}(S) \in E\big] \le e^{\epsilon} \Pr\big[\mathcal{A}(S') \in E\big] + \delta.$$

When $\delta = 0$, $\mathcal{A}$ is called ϵ-DP.

Definition 2 (Local Differential Privacy, LDP (Joseph et al., 2019)). *A family of randomized mappings $R : X \to Y$ is* (ϵ, δ)-locally differentially private *if, for every pair of inputs $x, x' \in X$ and every measurable subset $E \subseteq Y$,*

$$\Pr\big[R(x) \in E\big] \le e^{\epsilon} \Pr\big[R(x') \in E\big] + \delta.$$

Under CDP, a trusted curator collects the raw data and adds noise before release; this model simplifies algorithm design and typically incurs only an $\mathcal{O}(1/n)$ loss in accuracy (Cai et al., 2021), where n denotes the sample size. In contrast, LDP dispenses with any trust assumption: each user i holds a private value X_i, applies a predetermined randomized mechanism R_i satisfying (ϵ, δ)-DP, and submits only the perturbed output. We adopt the non-adaptive LDP model, in which all randomizers $\{R_i\}$ are fixed in advance (Cheu et al., 2019, Definitions 2.3 and 2.6). Consequently, inference must proceed solely from locally privatized data.

In the CDP setting, the privatized estimator $\widehat{\theta}_{\mathrm{CDP}}$ typically satisfies $\widehat{\theta}_{\mathrm{CDP}} - \theta = \mathcal{O}_p\big(n^{-1}\big)$, thus after $\sqrt{n}$-scaling, it shares the same asymptotic distribution as the non-private estimator, and one can recover its asymptotic variance with modest additional privacy cost. Under LDP, however, the error rate degrades to $\widehat{\theta}_{\mathrm{LDP}} - \theta = \mathcal{O}_p\big(n^{-1/2}\big)$, which both alters the limiting law and inflates the asymptotic variance. Moreover, because only privatized data are available, consistently estimating this variance from data collected solely for point estimation is generally infeasible.

We consider a federated learning framework involving K clients, each independently holding a local dataset i.i.d. drawn from an unknown distribution $\mathcal{P}_k$ with cumulative distribution function (CDF) F_k and density function f_k (Li et al., 2022). The goal is to collaboratively estimate the global quantile via weighted loss, i.e., the objective is to solve the following optimization problem:

$$\arg\min_{Q\in\Theta} \mathcal{L}(Q) \overset{\text{def}}{=} \arg\min_{Q\in\Theta} \sum_{k=1}^{K} p_k \mathcal{L}_{\tau_k}(Q) \overset{\text{def}}{=} \arg\min_{Q\in\Theta} \sum_{k=1}^{K} p_k \mathbb{E}_{x_k\sim\mathcal{P}_k}\{\ell_{\tau_k}(x_k, Q)\}, \tag{2.1}$$

where p_k denotes the weight assigned to client k, $\tau_k \in (0,1)$ is local quantile level, x_k is the sample generated from $\mathcal{P}_k$, and $\ell_{\tau_k}(x, Q)$ represents the check loss function defined as:

$$\ell_{\tau_k}(x, Q) = (x - Q)(\tau_k - \mathbb{I}(x < Q)), \tag{2.2}$$

where $\mathbb{I}(\cdot)$ is the indicator function. Let $\tau := \sum_{k=1}^{K} p_k \tau_k \in (0,1)$. For the global minimizer $Q^\star$ of (2.1), it corresponds to the global quantile at level τ of a weighted CDF, i.e., $\sum_{k=1}^{K} p_k F_k(Q^\star) = \tau$. In the following, we denote $F_k(Q^\star) = Q_k$, and considers the parameter space Θ is bounded; see Gu & Chen (2023).

As noted in introduction, to improve the communication efficiency, we consider a local SGD based estimator, for communication iteration sets $\mathcal{I} = \{t_0, t_1, \ldots, t_T\}$, the global server will receive the local iterations and broadcast the update to K clients, otherwise, the iterations are only conducted in each local clients, i.e., for $k = 1, \ldots K$,

$$q_{t+1}^k = \begin{cases} q_t^k - \eta_t \left\{\mathbb{I}(x_t^k < q_t^k) - \tau_k\right\}, & t \notin \mathcal{I}, \\ \sum_{k=1}^{K} p_k \left[q_t^k - \eta_t \left\{\mathbb{I}(x_t^k < q_t^k) - \tau_k\right\}\right], & t \in \mathcal{I}. \end{cases}$$

Here η_t is the pre-determined learning rate, and x_t^k represents an independent realization of $\mathcal{P}_k$, The final estimator is Polyak-Ruppert type, which averages the historical iterations,

$$\widetilde{Q}_T = \frac{1}{t_T} \sum_{m=1}^{T} \sum_{k=1}^{K} p_k q_{t_m}^k.$$

The communication and statistical efficiency are determined by the interval length $E_m := t_m - t_{m-1}$ for $m \in \mathbb{N}^+$. If $E_m = 1$, the local clients must communicate with the global server at every iteration. In this scenario, the approach reduces to parallel SGD, which, as noted by Li et al. (2022), may achieve the Cramér-Rao lower bound and thus serve as an efficient estimator for certain smooth loss functions. Conversely, if $E_m = n$, implying only one communication at the last iteration, the estimator degenerates to a divide-and-conquer estimator. Such an estimator is consistent only when $\tau_k \equiv \tau$ for all $k = 1, \ldots, K$ and some common $\tau \in (0, 1)$. In this case, minimizing the loss function (2.1) becomes a distributed learning problem. However, as pointed out by Gu & Chen (2023), the divide-and-conquer estimator may still be statistically inefficient for certain weight choices. Therefore, a careful balance must be struck between communication and statistical efficiency. For a general positive interval $E_m > 0$, the local SGD method allows us to find an appropriate choice of E_m that can ensure an optimal trade-off between these efficiencies.

On the other hand, the data collected from each client may be subject to privacy protection policies, particularly in surveys involving sensitive information such as income or health status. For the local quantile loss function equation 2.2, we observe that the structure of its gradient resembles a binary response. This motivates us to incorporate an LDP mechanism based on randomized response and permutation, following the framework of Liu et al. (2023b), with a truthful response rate $r_k \in (0, 1]$. Specifically, the mechanism allows each local client to either return a true gradient with probability r_k or a synthetic Bernoulli random variable with probability $1-r_k$. This iterative mechanism ensures ϵ_k-LDP, where the privacy parameter is given by $\epsilon_k = \log(1 + r_k) - \log(1 - r_k)$, as established in Liu et al. (2023b).

It is worth noting that the method of Liu et al. (2023b) was originally developed in a single-machine setting. Extending it directly to federated learning raises new challenges, since federated systems inherently involve the issue of heterogeneity. We illustrate with a simple example. Consider collaboratively estimating the national median annual income using state-level data from the United States, where each state is treated as a client. First, income distributions typically vary across states (see Figure 1(i)). Second, privacy preferences can differ across states due to cultural norms and development levels (Milberg et al., 2000; Bellman et al., 2004). Figure 1(ii) shows how the released information can vary under different privacy budgets. Due to such heterogeneity, a naive combination of local LDP estimators from Liu et al. (2023b) may result in severely biased results. To address this problem, we propose a novel federated quantile estimation algorithm under LDP, equipped with a carefully designed local SGD updating rule. This method accommodates heterogeneous data distributions, quantile targets, and privacy budgets across clients while maintaining a common global target. The complete procedure is summarized in Algorithm 1, and we denote the resulting estimator as $\widehat{Q}_T$.

Algorithm 1: Federated quantile estimation with local SGD under LDP

Input: step sizes $\{\eta_m\}_{m=1}^T$, target quantile $\tau \in (0, 1)$, truthful response rates $\{r_k\}_{k=1}^K$, communication set $\mathcal{I} = \{t_0, t_1, \ldots, t_T\}$.

Initialization: set $q_0^k = q_0 \sim \mathcal{N}(0, 1)$ for all $1 \le k \le K$, let $\widehat{Q}_0 \leftarrow 0$.

for $m = 1$ to T **do**
  for $k = 1$ to K (distributedly) **do**
    for $t = t_{m-1} + 1$ to t_m **do** ▷ Local updates
      $u_t^k \sim Bernoulli(r), \quad v_t^k \sim Bernoulli(0.5)$
      $s_t^k = \mathbb{I}(x_t^k > q_{t-1}^k)\mathbb{I}(u_t^k = 1) + v_t^k \mathbb{I}(u_t^k = 0)$
      $q_t^k \leftarrow q_{t-1}^k + \dfrac{1 - r_k + 2\tau r_k}{2r_k}\eta_{m-1}\mathbb{I}(s_t^k = 1) - \dfrac{1 + r_k - 2\tau r_k}{2r_k}\eta_{m-1}\mathbb{I}(s_t^k = 0)$
    end for
  end for
  $\bar{q}_{t_m} \leftarrow \sum_{k=1}^K p_k q_{t_m}^k$; $q_{t_m}^k \leftarrow \bar{q}_{t_m}$ for all $1 \le k \le K$. ▷ Aggregation and synchronization.
  $\widehat{Q}_m \leftarrow ((m-1)\widehat{Q}_{m-1} + \bar{q}_{t_m})/m$.
end for

Return: $\widehat{Q}_T$.

In our proposed Algorithm 1, each iteration integrates global information (the global quantile τ) with client-specific privacy budget (r_k), thereby correcting bias arising from the aggregation of heterogeneous local LDP mechanisms and loss functions. The following theorem shows that this procedure effectively reduces the LDP inference problem to its non-private analogue.

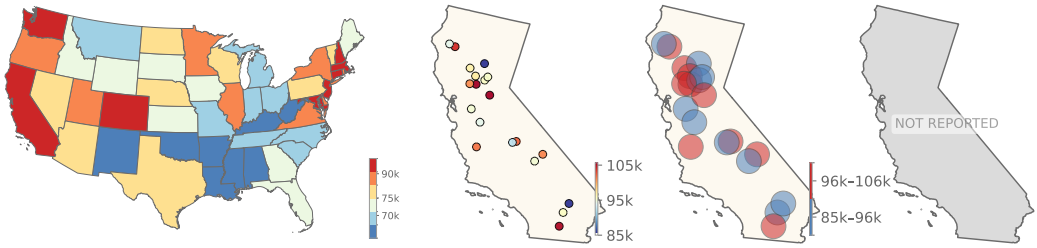

(i) Data heterogeneity (ii) Privacy heterogeneity

Figure 1: Illustration of client heterogeneity. Income data source: U.S. Census Bureau (`https://data.census.gov/table/ACSST5Y2023.S1901?g=010XX00US0400000`). Panel (i) shows median annual income by state. Panel (ii) shows three income disclosure schemes under different privacy budgets: (a) each individual release true income; (b) each individual release an income interval; and (c) withholding release.

Theorem 2.1. *Denote $\widetilde{\tau}_k = r_k\tau + (1-r_k)/2$. For a privacy budget $\epsilon_k = \log(1+r_k) - \log(1-r_k)$, there exists a dataset consisting of i.i.d. samples drawn from some distribution $\widetilde{\mathcal{P}}_k, 1 \le k \le K$, such that solving the federated loss equation 2.1 with ϵ_k-LDP using data drawn from $\mathcal{P}_k$ is equivalent to solving the following non-private problem:*

$$\arg\min_Q \mathcal{L}(Q) = \arg\min_Q \widetilde{\mathcal{L}}(Q) := \arg\min_Q \sum_{k=1}^{K} p_k \mathbb{E}_{x_k \sim \widetilde{\mathcal{P}}_k} \left\{ r_k^{-1} \ell_{\widetilde{\tau}_k}(x_k, Q) \right\}. \tag{2.3}$$

Therefore, by Theorem 2.1, the LDP federated quantile estimation problem can be reformulated as a non-DP federated quantile estimation task under modified distributions and shifted quantile levels. The main challenge then becomes analyzing the statistical properties of the resulting non-DP estimator, particularly in the presence of the non-smooth quantile loss function.

3 Asymptotic Analysis

In this section, we focus on the asymptotic analysis of the proposed LDP estimator and the practical construction of confidence intervals. Before presenting the main results, we first introduce several necessary assumptions.

Assumption 1. *For some constant $C_f > 0$, $f_k(\cdot)$, $1 \le k \le K$, is uniformly bounded by C_f.*

Assumption 2. *Define the effective step $\gamma_m = \eta_m E_m$, which is non-increasing in m and satisfies that $\sum_{m=1}^{\infty} \gamma_m^2 < \infty$, $\sum_{m=1}^{\infty} \gamma_m = \infty$, and $(\gamma_m - \gamma_{m+1})/\gamma_m = \mathcal{O}(\gamma_m)$.*

Assumption 3. *The sequence $\{E_m\}_{m\ge1}$ satisfies that*

(a) $\{E_m\}_{m\ge1}$ is either uniformly bounded or non-decreasing.

(b) There exist some $\delta > 0$ and $\nu \ge 1$ such that

$$\limsup_{T\to\infty} \frac{1}{T^2} \left(\sum_{m=0}^{T-1} E_m^{1+\delta} \right) \left(\sum_{m=0}^{T-1} E_m^{-1-\delta} \right) < \infty, \lim_{T\to\infty} \frac{1}{T^2} \left(\sum_{m=0}^{T-1} E_m \right) \left(\sum_{m=0}^{T-1} E_m^{-1} \right) = \nu.$$

(c) Denote $t_T = \sum_{m=0}^{T-1} E_m$, satisfying

$$\lim_{T\to\infty} \frac{\sqrt{t_T}}{T} \sum_{m=0}^{T} \gamma_m = 0, \quad \lim_{T\to\infty} \frac{\sqrt{t_T}}{T} \frac{1}{\sqrt{\gamma_T}} = 0$$

Assumption 1 is a mild and regular condition concerning the uniform boundedness of density functions. Assumptions 2 and 3 require that the effective step sizes decay slowly and the communication intervals increase slowly; see also Li et al. (2022).

Theorem 3.1. *Under Assumptions 1-3, as $T \to \infty$, the proposed LDP federated estimator enjoys*

$$\sqrt{t_T}(\widehat{Q}_T - Q^\star) \xrightarrow{d} N\left(0, \nu \frac{\sum_{k=1}^K p_k^2 \left\{r_k^{-2} - (2Q_k - 1)^2\right\}}{4\left(\sum_{k=1}^K p_k f_k(Q^\star)\right)^2}\right).$$

Theorem 3.1 establishes the asymptotic normality of the estimator $\widehat{Q}_T$, which theoretically allows for the theoretical construction of a confidence interval for $Q^\star$. However, the construction involves unknown quantities, such as the individual quantiles Q_k and the density values $f_k(Q^\star)$. Even in cases where $Q_k = \tau_k$ is known, the estimation of $f_k(Q^\star)$ remains challenging. In particular, it is difficult to recover these density values using only the perturbed gradients available from Algorithm 1. Moreover, in SGD-based methods, consistent variance estimation typically relies on the Hessian matrix, which is well-defined only for smooth loss functions, as previously discussed. Therefore, although Theorem 3.1 provides a theoretically valid basis for confidence interval construction, it is not practically implementable due to these limitations.

Inspired by the quantile inference framework for single clients in Liu et al. (2023b), it is necessary to strengthen the pointwise result of Theorem 3.1 to a functional version.

Theorem 3.2. *Under Assumptions 1–3, as $T \to \infty$, we have*

$$\mathcal{Q}_T(s) := \frac{\sqrt{t_T}}{T} \sum_{m=1}^{h(s,T)} (\bar{q}_{t_m} - Q^\star) \xrightarrow{d} \frac{\sqrt{\nu \sum_{k=1}^K p_k^2 \left\{r_k^{-2} - (2Q_k - 1)^2\right\}}}{2\sum_{k=1}^K p_k f_k(Q^\star)} B(s),$$

where $t_T = \sum_{m=0}^{T-1} E_m$, $\bar{q}_{t_m} = \sum_{k=1}^K p_k r_k q_{t_m}^k$, $B(\cdot)$ is a standard Brownian motion on $[0,1]$, and

$$h(s,T) = \max\left\{ n \in \mathbb{Z}_{>0} \,\middle|\, s \sum_{m=1}^{T} \frac{1}{E_m} \geq \sum_{m=1}^{n} \frac{1}{E_m} \right\}, \quad \textit{for } s \in (0,1].$$

Theorem 3.2 establishes a functional central limit theorem (FCLT) for $Q_T(s)$ over $s \in (0,1]$, showing that it converges weakly in the $\ell^\infty[0,1]$ (the space of bounded real-valued functions) to a Brownian motion, which is our another major theoretical contribution. Note that the sample quantile loss doesn't satisfy the common L-average smooth conditions for weakly convergence result, such in (Li et al., 2022; Xie et al., 2024; Zhu et al., 2024), leading to extra challenge in deriving the almost sure and $\mathcal{L}^2$ convergence rates of $\bar{q}_{t_m}$, which are essential for handling the asymptotically negligible terms. Theorem 3.1 arises as a special case of Theorem 3.2 when $s = 1$. Building on Theorem 3.2, we proceed to construct a self-normalized test statistic and derive its asymptotic pivotal distribution via the continuous mapping theorem.

Define $r_0 = 0$ and, for $m \geq 1$, $r_m = \left(\sum_{i=1}^m 1/E_i\right)\left(\sum_{i=1}^T 1/E_i\right)^{-1}$, which ensures that

$$\mathcal{Q}_T(r_m) = \frac{\sqrt{t_T}}{T} \sum_{i=1}^{m} (\bar{q}_{t_i} - Q^\star), \quad \text{and in particular,} \quad \mathcal{Q}_T(1) = \frac{\sqrt{t_T}}{T} \sum_{i=1}^{T} (\bar{q}_{t_i} - Q^\star).$$

Following the arguments in (Shao, 2015), once a functional central limit theorem such as Theorem 3.2 is established, one can construct a self-normalized statistic that asymptotically enjoys a pivotal distribution. Specifically, define

$$\mathcal{V}_T = \sum_{m=1}^{T} (r_m - r_{m-1}) \left(\mathcal{Q}_T(r_m) - \frac{m}{T}\mathcal{Q}_T(1)\right)^2. \tag{3.1}$$

Corollary 3.1. *Suppose Assumptions 1-3 hold and $g(r_m) \asymp m/T$ for some continuous function g on $[0,1]$. Then, as $T \to \infty$,*

$$\frac{\mathcal{Q}_T(1)}{\sqrt{\mathcal{V}_T}} \xrightarrow{d} \frac{B(1)}{\sqrt{\int_0^1 \left(B(r) - g(r)B(1)\right)^2 dr}}.$$

Corollary 3.1 presents the asymptotic distribution of the self-normalized statistic $Q_T(1)/\mathcal{V}_T$, which is distribution-free. As a result, there is no need to allocate additional DP budget to estimate nuisance parameters when constructing confidence intervals.

The selection of the self-normalizer is not unique, and an appropriate norm of the Gaussian process $B(r) - g(r)B(1)$ can yield similar results to those in Corollary 3.1. For example, using the supremum norm and the $\mathcal{L}_1$ norm, one can define alternative self-normalizers as follows:

$$\mathcal{V}_T' = \sup_{1\le m\le T} \left|\mathcal{Q}_T(r_m) - \frac{m}{T}\mathcal{Q}_T(1)\right|, \quad \mathcal{V}_T'' = \sum_{m=1}^{T} (r_m - r_{m-1}) \left|\mathcal{Q}_T(r_m) - \frac{m}{T}\mathcal{Q}_T(1)\right|,$$

which are related to the processes $\sup_{0\le r\le 1} |B(r)-g(r)B(1)|$ and $\int_0^1 |B(r)-g(r)B(1)|\, dr$, respectively. However, the self-normalizer defined in equation equation 3.1 enjoys greater computational efficiency, as the $\mathcal{L}_2$ norm can be computed in an online manner, as described in Algorithm 2. Let $\widehat{\mathcal{V}}_T$ denote the estimator of the self-normalizer in equation 3.1, and let $v_{\alpha/2,g}$ be the $(1-\alpha/2)$ quantile of the random variable $B(1)/\left(\int_0^1 \left(B(r) - g(r)B_1(1)\right)^2 dr\right)^{1/2}$. The following corollary ensures the asymptotic validity of the constructed LDP confidence interval.

The following Corollary 3.2 ensures the asymptotic validity of the constructed LDP confidence interval.

Corollary 3.2. *Suppose the same conditions in Theorem 3.2 hold, as $T \to \infty$, one has that*

$$\mathbb{P}\left(\widehat{Q}_T - v_{\frac{\alpha}{2},g}\sqrt{\widehat{\mathcal{V}}_T} \le Q^\star \le \widehat{Q}_T + v_{\frac{\alpha}{2},g}\sqrt{\widehat{\mathcal{V}}_T}\right) \to 1-\alpha$$

Algorithm 2: Online Inference

Input: step sizes $\{\eta_m\}_{m=1}^T$, target quantile $\tau \in (0,1)$, truthful response rates $\{r_k\}_{k=1}^K$, communication set $\mathcal{I} = \{t_0, t_1, \ldots, t_T\}$.

Initialization: set $q_0^k \sim \mathcal{N}(0,1)$ for all k, let $\mathcal{V}_0^a \leftarrow 0$, $\mathcal{V}_0^b \leftarrow 0$, $\mathcal{V}_0^s \leftarrow 0$, $\mathcal{V}_0^p \leftarrow 0$, and $Q_0 \leftarrow 0$.

for $m = 1$ to T **do**

Obtain $\widehat{Q}_m$ from Algorithm 1.

$\mathcal{V}_m^a \leftarrow \mathcal{V}_{m-1}^a + m^2 Q_m^2/E_m$, ▷ $E_m = t_m - t_{m-1}$

$\mathcal{V}_m^b \leftarrow \mathcal{V}_{m-1}^b + m^2 Q_m/E_m$,

$\mathcal{V}_m^s \leftarrow \mathcal{V}_{m-1}^s + 1/E_m$,

$\mathcal{V}_m^p \leftarrow \mathcal{V}_{m-1}^p + m^2/E_m$.

$\widehat{\mathcal{V}}_m \leftarrow \frac{1}{m^2\mathcal{V}_m^s}\left(\mathcal{V}_m^a - 2\mathcal{V}_m^b Q_m + \mathcal{V}_m^p Q_m^2\right)$. ▷ Online inference.

end for

Return: Confidence interval $\left[\widehat{Q}_T - v_{\frac{\alpha}{2},g}\sqrt{\widehat{\mathcal{V}}_T},\ \ \widehat{Q}_T + v_{\frac{\alpha}{2},g}\sqrt{\widehat{\mathcal{V}}_T}\right]$.

4 Experiments

4.1 Simulation setup

We first evaluate our proposed method through extensive simulation studies using synthetic data. In all experiments, we fixed $p_k = 1/K$ for $1 \le k \le K$, the number of clients is fixed at $K = 10$. The quantile levels examined range from 0.3 to 0.8, and the truthful response rates vary between 0.25 and 0.9. We focus on the following four scenarios of heterogeneity:

- **heterogeneous quantile levels:** We investigate two distinct scenarios: (1) Case τ_{low}: lower quantile levels, where each client is assigned a unique quantile level τ_k ranging uniformly from 0.3 to 0.5; and (2) Case τ_{high}: higher quantile levels, where τ_k ranges uniformly from 0.5 to 0.8.
- **heterogeneous response rates.** Each client has a unique truthful response rate r_k, ranging uniformly from 0.25 to 0.9.

- **heterogeneous locations (Hete L).** Data for each client k are independently generated from $\mathcal{N}(\mu_k, 1)$, where $\mu_k \sim \mathcal{N}(0, 1)$.
- **heterogeneous distribution families (Hete D).** Data are generated independently across ten clients, with three drawing from $\mathcal{N}(0, 1)$, three from the uniform distribution $\mathcal{U}(-1, 1)$, and four from a standard Cauchy distribution $\mathcal{C}(0, 1)$.

We set the step size γ_m as: $\gamma_m = 20\bar{r}/(m^{0.51} + 100)$, with $\gamma_m = E_m\eta_m$ and $\bar{r} = K^{-1}\sum_{k=1}^{K} r_k$. Following Li et al. (2022), we implement a warm-up phase, setting the communication interval $E_m = 1$ for the first 5% of iterations. After the warm-up period, we redefine the interval sequence $\{E_m\}$ based on a new sequence $\{E'_m\}$, specifically: $E_m = E'_{m-0.05\cdot T}$. We examine three different interval strategies for E'_m: (1) C1: $E'_m \equiv 1$ (equivalent to parallel SGD), (2) C5: $E'_m \equiv 5$, and (3) Log: $E'_m = \lceil \log_2(m + 1) \rceil$. The initial parameter estimates are set to $q_0^k = q_0 \sim \mathcal{N}(0, 1)$ for all clients k. All experimental settings are replicated $R = 1,000$ times. The simulations are conducted on computational resources comprising 36 Intel Xeon Gold 6271 CPUs, with a total of 128GB RAM and 500GB storage.

4.2 Simulation results

We first illustrate the performance of our proposed method by presenting sample iteration trajectories for estimation and inference. Specifically, we randomly select one simulation run and plot the resulting estimates and corresponding confidence intervals against t_T (Figure S.1). The results demonstrate that our approach accurately captures the true quantile value and provides reliable inference. Subsequently, we fix the total sample size t_T at $10,000$ and $50,000$ and evaluate the finite sample performance under different settings. Let $\widehat{Q}_T^{(r)}$ denote the quantile estimator and $\mathrm{CI}^{(r)}$ represent the corresponding 95% confidence interval obtained from Algorithm 2 in the r-th simulation. We consider two metrics: the mean absolute error (MAE), defined as $R^{-1}\sum_{r=1}^{R} |\widehat{Q}_T^{(r)} - Q^*|$, and the empirical coverage probability (ECP), defined as $R^{-1}\sum_{r=1}^{R} \mathbb{I}(Q^* \in \mathrm{CI}^{(r)})$. For comparison, we also consider two alternative methods: (1) the SGD with DP updates (Song et al., 2013) (DP-SGD), which adds noise directly to the gradients instead of introducing DP through randomized response. To align with the original paper's setup, we focus on the case with $C = 1$. In this regime, the gradient-descent update in Algorithm 1 becomes

$$q_t^k \leftarrow q_{t-1}^k + \eta_{t-1}\Big\{\tau_k \mathbb{I}(x_t^k > q_{t-1}^k) - (1 - \tau_k)\mathbb{I}(x_t^k < q_{t-1}^k) + Z_t^k\Big\},$$

where Z_t^k is drawn from a Laplace distribution. A simple calculation shows that Z_t^k has mean zero and scale parameter $1/\log\{(1 + r_k)/(1 - r_k)\}$. (2) the divide-and-conquer (DC) method, which corresponds to the special case $E_m = n$. Here we use step size $\eta_t = 2\bar{r}/(t^{0.51} + 100)$ (Goyal et al., 2017). The numerical results for all of the methods are reported in Tables S.1 and S.2.

From Tables S.1 and S.2, we observe that our method consistently achieves ECP close to or exceeding the nominal 95% level across all scenarios. As either the total sample size t_T or the truthful response rate increases, the MAE decreases, which aligns with our theoretical results. Comparing the three interval strategies, we find that the C1 strategy (parallel SGD) yields the smallest MAE, as it has the highest communication frequency. Comparing with the two competing methods, we find that the DC approach results in the largest errors. Notably, in certain heterogeneous cases, such as Hete L with $\tau = 0.8$, the DC estimator exhibits significant bias and an ECP far below the nominal 95% level. In contrast, our proposed estimators successfully achieve approximately 95% empirical coverage in these cases. Moreover, while DP-SGD attains empirical coverage probabilities close to or even exceeding 95% in most settings, its MAE remain uniformly larger than those of our method. To further illustrate the communication efficiency of our method, we also consider scenarios with a fixed number of communication rounds T. The results are summarized in Tables S.3 and S.4. We observe that our proposed method continues to provide valid inference. Additionally, under fixed communication rounds, the Log strategy generally achieves the best performance, yielding the smallest MAE.

4.3 Real data

In this subsection, we empirically evaluate the effectiveness of our proposed method using a representative real-world dataset widely employed in privacy research: Government Salary Dataset

(Plečko et al., 2024). This dataset is sourced from the 2018 American Community Survey conducted by the U.S. Census Bureau and contains over 200,000 records, with annual salary (in USD) as the response variable. Since annual salary represents sensitive financial information (Gillenwater et al., 2021), we treat it as requiring privacy protection. To incorporate the dataset's inherent geographic structure, we partition the sample according to the feature "economic region." The three smallest regions are merged into a single "Others" category, yielding seven regions in total, each regarded as one client. Because region-level sample sizes vary, we apply oversampling to balance the data, resulting in $t_T = 53,960$ observations per client. All other hyperparameters follow the settings in Section 4.1. For analysis, we apply a log transformation to the response variable and subsequently back-transform it.

We target quantile levels $\tau_k \equiv \tau \in \{0.3, 0.5, 0.8\}$ and consider response rate ranges from 0.6 to 0.9. For reference, we also compute the full-sample quantiles without LDP. The resulting estimators and confidence-interval lengths are summarized in Table 1. As shown, higher response rates r and more communication rounds generally produce shorter confidence intervals, consistent with our simulation findings. In most cases, the empirical quantiles fall within our reported intervals, highlighting the practical utility of our method for real data.

Quantile (τ)	Rate (r)	C1	C5	Log	Empirical
0.3	0.6	33367 (1742)	33184 (6697)	33030 (12093)	
0.3	hetero	33418 (1424)	33229 (5291)	33140 (9788)	34000
0.3	0.9	33547 (1548)	33403 (4443)	33239 (7828)	
0.5	0.6	48454 (2255)	48212 (6315)	47951 (11361)	
0.5	hetero	48462 (1435)	48290 (4973)	48091 (9025)	50000
0.5	0.9	48610 (1454)	48494 (3851)	48311 (6863)	
0.8	0.6	78586 (2066)	78168 (6646)	77995 (13144)	
0.8	hetero	78390 (1291)	78054 (5862)	77722 (11101)	80000
0.8	0.9	78657 (1138)	78300 (4677)	78084 (8928)	

Table 1: Estimation results (interval lengths) on the real dataset across varying quantile levels and response rates. "Empirical" denotes the full-sample quantile estimator without LDP. "hetero" indicates client-specific truthful response rates r_k range from 0.6 to 0.9.

5 Concluding remark

We propose a federated-learning algorithm for quantile inference under LDP that flexibly accommodates client-level heterogeneity in quantile targets, privacy budgets, and data distributions. In addition, one innovation that should be emphasized is that our developed theoretical results of local SGD quantile estimator. We first design an LDP mechanism that can transform the LDP federated quantile estimation into the non-DP case, and then derive the asymptotic normality and functional central limit theorem of the proposed estimator under non-DP cases. It is first weak-convergence result for local SGD without the usual average-smoothness assumption in existing literature. Building on these non-private asymptotic results, we develop a self-normalized inference procedure that constructs valid confidence intervals under LDP without requiring direct estimation of the asymptotic variance.

Despite these advances, our method has several limitations. First, it relies on additional regularity assumptions to handle arbitrary client-level data heterogeneity. Second, as noted in (Shao, 2015), self-normalization yields heavier-tailed limit distributions than the Gaussian, which can produce conservative confidence intervals or reduced power in hypothesis testing. Finally, our framework depends on a central server for aggregation and synchronization, which may not be available in fully decentralized environments. Addressing these challenges and extending the algorithm to decentralized settings remain important directions for future research.

REPRODUCIBILITY STATEMENT

All numerical experiments and real-data analyses are fully reproducible via the code included in the submitted anonymized supplementary materials.

REFERENCES

Naman Agarwal, Ananda Theertha Suresh, Felix Xinnan X Yu, Sanjiv Kumar, and Brendan McMahan. cpsgd: Communication-efficient and differentially-private distributed sgd. *Advances in Neural Information Processing Systems*, 31, 2018.

Joshua Angrist, Victor Chernozhukov, and Iván Fernández-Val. Quantile regression under misspecification, with an application to the us wage structure. *Econometrica*, 74(2):539–563, 2006.

Luca Barbaglia, Sergio Consoli, and Sebastiano Manzan. Forecasting with economic news. *Journal of Business & Economic Statistics*, 41(3):708–719, 2023.

Steven Bellman, Eric J Johnson, Stephen J Kobrin, and Gerald L Lohse. International differences in information privacy concerns: A global survey of consumers. *The Information Society*, 20(5): 313–324, 2004.

T Tony Cai, Yichen Wang, and Linjun Zhang. The cost of privacy: Optimal rates of convergence for parameter estimation with differential privacy. *The Annals of Statistics*, 49(5):2825–2850, 2021.

Likai Chen, Georg Keilbar, and Wei Biao Wu. Recursive quantile estimation: Non-asymptotic confidence bounds. *Journal of Machine Learning Research*, 24(91):1–25, 2023.

Song Xi Chen. Nonparametric estimation of expected shortfall. *Journal of financial econometrics*, 6(1):87–107, 2008.

Xi Chen, Jason D Lee, Xin T Tong, and Yichen Zhang. Statistical inference for model parameters in stochastic gradient descent. *The Annals of Statistics*, 48(1):251–273, 2020.

Xi Chen, Weidong Liu, and Yichen Zhang. First-order newton-type estimator for distributed estimation and inference. *Journal of the American Statistical Association*, 117(540):1858–1874, 2022.

Victor Chernozhukov and Iván Fernández-Val. Inference for extremal conditional quantile models, with an application to market and birthweight risks. *The Review of Economic Studies*, 78(2): 559–589, 2011.

Victor Chernozhukov and Christian Hansen. An iv model of quantile treatment effects. *Econometrica*, 73(1):245–261, 2005.

Albert Cheu, Adam Smith, Jonathan Ullman, David Zeber, and Maxim Zhilyaev. Distributed differential privacy via shuffling. In *Advances in Cryptology–EUROCRYPT 2019: 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19–23, 2019, Proceedings, Part I 38*, pp. 375–403. Springer, 2019.

Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. Collecting telemetry data privately. *Advances in Neural Information Processing Systems*, 30, 2017.

John C Duchi, Michael I Jordan, and Martin J Wainwright. Local privacy and statistical minimax rates. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pp. 429–438. IEEE, 2013.

Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pp. 486–503. Springer, 2006.

Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*, pp. 1054–1067, 2014.

Jianqing Fan, Yongyi Guo, and Kaizheng Wang. Communication-efficient accurate statistical estimation. *Journal of the American Statistical Association*, 118(542):1000–1010, 2023.

Jennifer Gillenwater, Matthew Joseph, and Alex Kulesza. Differentially private quantiles. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 3713–3722. PMLR, 18–24 Jul 2021.

Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large minibatch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017.

Jia Gu and Song Xi Chen. Distributed statistical inference under heterogeneity. *Journal of Machine Learning Research*, 24(387):1–57, 2023.

Xuming He, Xiaoou Pan, Kean Ming Tan, and Wen-Xin Zhou. Smoothed quantile regression with large-scale inference. *Journal of Econometrics*, 232(2):367–388, 2023.

Emily C Hector and Peter X-K Song. A distributed and integrated method of moments for high-dimensional correlated data analysis. *Journal of the American Statistical Association*, 116(534): 805–818, 2021.

Jiaqiao Hu, Yijie Peng, Gongbo Zhang, and Qi Zhang. A stochastic approximation method for simulation-based quantile optimization. *INFORMS Journal on Computing*, 34(6):2889–2907, 2022.

Jie Hu, Yi-Ting Ma, and Do-Young Eun. Does worst-performing agent lead the pack? analyzing agent dynamics in unified distributed sgd. *Advances in Neural Information Processing Systems*, 37:72754–72789, 2024.

Ziyue Huang, Yuting Liang, and Ke Yi. Instance-optimal mean estimation under differential privacy. *Advances in Neural Information Processing Systems*, 34:25993–26004, 2021.

Xue Jiang, Xuebing Zhou, and Jens Grossklags. Signds-fl: Local differentially private federated learning with sign-based dimension selection. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 13(5):1–22, 2022.

Michael I Jordan, Jason D Lee, and Yun Yang. Communication-efficient distributed statistical inference. *Journal of the American Statistical Association*, 114(526):668–681, 2019.

Matthew Joseph, Jieming Mao, Seth Neel, and Aaron Roth. The role of interactivity in local differential privacy. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 94–105, 2019. doi: 10.1109/FOCS.2019.00015.

Nathan Kallus, Xiaojie Mao, and Masatoshi Uehara. Localized debiased machine learning: Efficient inference on quantile treatment effects and beyond. *Journal of Machine Learning Research*, 25 (16):1–59, 2024.

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.

Boyuan Li, Shaohui Zhang, and Qiuying Han. Federated learning with joint server-client momentum. *Scientific Reports*, 15(1):15626, 2025.

Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *Proceedings of Machine learning and systems*, 2:429–450, 2020.

Xiang Li, Jiadong Liang, Xiangyu Chang, and Zhihua Zhang. Statistical estimation and online inference via local sgd. In *Conference on Learning Theory*, pp. 1613–1661. PMLR, 2022.

Chun Liu, Youliang Tian, Jinchuan Tang, Shuping Dang, and Gaojie Chen. A novel local differential privacy federated learning under multi-privacy regimes. *Expert systems with applications*, 227: 120266, 2023a.

Wei Liu, Li Chen, Yunfei Chen, and Wenyi Zhang. Accelerating federated learning via momentum gradient descent. *IEEE Transactions on Parallel and Distributed Systems*, 31(8):1754–1766, 2020.

Yi Liu, Qirui Hu, Lei Ding, and Linglong Kong. Online local differential private quantile inference via self-normalization. In *International Conference on Machine Learning*, pp. 21698–21714. PMLR, 2023b.

Andrew Lowy and Meisam Razaviyayn. Private federated learning without a trusted server: Optimal algorithms for convex losses. In *The Eleventh International Conference on Learning Representations*, 2023.

Xu Ma, Xiaoqian Sun, Yuduo Wu, Zheli Liu, Xiaofeng Chen, and Changyu Dong. Differentially private byzantine-robust federated learning. *IEEE Transactions on Parallel and Distributed Systems*, 33(12):3690–3701, 2022.

Horia Mania, Xinghao Pan, Dimitris Papailiopoulos, Benjamin Recht, Kannan Ramchandran, and Michael I Jordan. Perturbed iterate analysis for asynchronous stochastic optimization. *SIAM Journal on Optimization*, 27(4):2202–2229, 2017.

Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pp. 1273–1282. PMLR, 2017.

Sandra J Milberg, H Jeff Smith, and Sandra J Burke. Information privacy: Corporate management and national regulation. *Organization science*, 11(1):35–57, 2000.

Milad Nasr, Nicholas Carlini, Jonathan Hayase, Matthew Jagielski, A Feder Cooper, Daphne Ippolito, Christopher A Choquette-Choo, Eric Wallace, Florian Tramèr, and Katherine Lee. Scalable extraction of training data from (production) language models. *arXiv preprint arXiv:2311.17035*, 2023.

Drago Plečko, Nicolas Bennett, and Nicolai Meinshausen. fairadapt: Causal reasoning for fair data preprocessing. *Journal of Statistical Software*, 110:1–35, 2024.

Herbert Robbins and David Siegmund. A convergence theorem for non negative almost supermartingales and some applications. In *Optimizing methods in statistics*, pp. 233–257. Elsevier, 1971.

Xiaofeng Shao. Self-normalization for time series: a review of recent developments. *Journal of the American Statistical Association*, 110(512):1797–1817, 2015.

Xiaoying Shen, Hang Jiang, Yange Chen, Baocang Wang, and Le Gao. Pldp-fl: Federated learning with personalized local differential privacy. *Entropy*, 25(3):485, 2023.

Siping Shi, Chuang Hu, Dan Wang, Yifei Zhu, and Zhu Han. Distributionally robust federated learning for differentially private data. In *2022 IEEE 42nd International Conference on Distributed Computing Systems (ICDCS)*, pp. 842–852. IEEE, 2022.

Shuang Song, Kamalika Chaudhuri, and Anand D Sarwate. Stochastic gradient descent with differentially private updates. In *2013 IEEE global conference on signal and information processing*, pp. 245–248. IEEE, 2013.

Sebastian U Stich. Local sgd converges fast and communicates little. *arXiv preprint arXiv:1805.09767*, 2018.

Weijie J Su and Yuancheng Zhu. Uncertainty quantification for online learning and stochastic approximation via hierarchical incremental gradient descent. *arXiv preprint arXiv:1802.04876*, 2018.

Ye Tian, Haolei Weng, and Yang Feng. Unsupervised federated learning: A federated gradient em algorithm for heterogeneous mixture models with robustness against adversarial attacks. *arXiv preprint arXiv:2310.15330*, 2023.

Huixia Judy Wang, Deyuan Li, and Xuming He. Estimation of high conditional quantiles for heavy-tailed distributions. *Journal of the American Statistical Association*, 107(500):1453–1464, 2012.

Chuhan Xie, Kaicheng Jin, Jiadong Liang, and Zhihua Zhang. Asymptotic time-uniform inference for parameters in averaged stochastic approximation. *arXiv preprint arXiv:2410.15057*, 2024.

Steve Yadlowsky, Scott Fleming, Nigam Shah, Emma Brunskill, and Stefan Wager. Evaluating treatment prioritization rules via rank-weighted average treatment effects. *Journal of the American Statistical Association*, 120(549):38–51, 2025.

Yang Zhao, Jun Zhao, Mengmeng Yang, Teng Wang, Ning Wang, Lingjuan Lyu, Dusit Niyato, and Kwok-Yan Lam. Local differential privacy-based federated learning for internet of things. *IEEE Internet of Things Journal*, 8(11):8836–8853, 2020.

Wanrong Zhu, Zhipeng Lou, Ziyang Wei, and Wei Biao Wu. High confidence level inference is almost free using parallel stochastic optimization. *arXiv preprint arXiv:2401.09346*, 2024.

A Tables and Figures

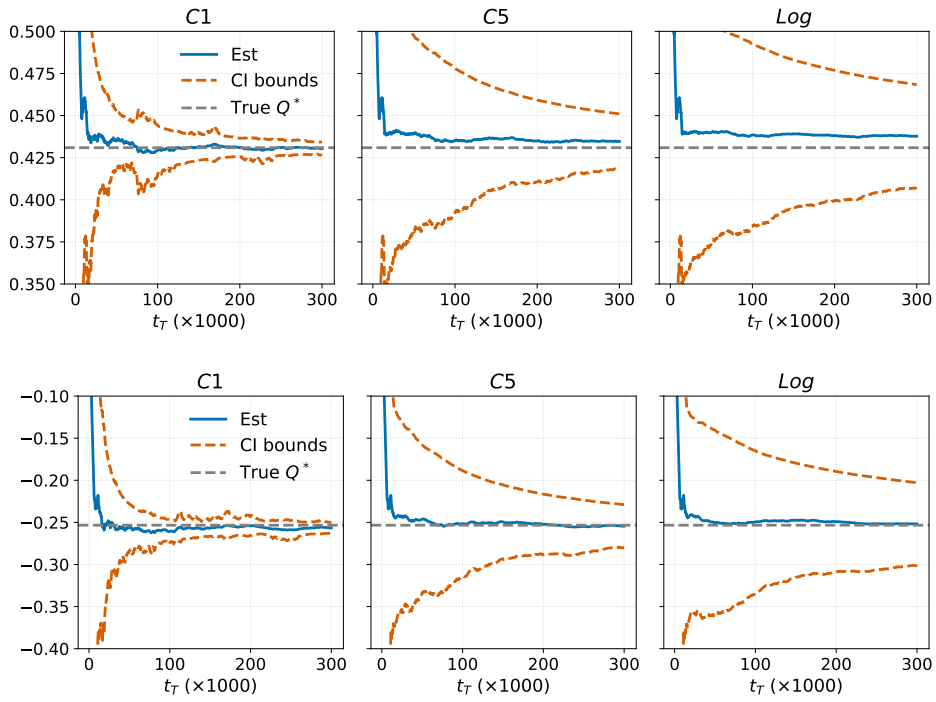

Figure S.1: Sample trajectories of the iterative estimator and corresponding confidence intervals under heterogeneous distributions (Hete L, with $r_k = 0.9$ and $\tau = 0.5$, left panel) and heterogeneous quantile levels (τ_{low}, with heterogeneous response rates, right panel). The horizontal dotted line indicates the true quantile value Q^*.

Quantile (τ)	Rate (r)	C1	C5	Log	DP-SGD (C1)	DC
			$t_T = 10000$			
0.5	0.25	0.949(0.0133)	0.967(0.0244)	0.992(0.0360)	0.949(0.0191)	0.939(0.2503)
0.5	hetero	0.963(0.0071)	0.989(0.0112)	0.997(0.0161)	0.955(0.0100)	1.000(0.0497)
0.5	0.9	0.995(0.0023)	1.000(0.0054)	1.000(0.0082)	0.980(0.0036)	1.000(0.0158)
τ_{low}	0.25	0.947(0.0136)	0.982(0.0253)	0.990(0.0369)	0.940(0.0200)	0.969(0.2616)
τ_{low}	hetero	0.962(0.0072)	0.993(0.0113)	0.997(0.0162)	0.949(0.0105)	0.999(0.0530)
τ_{low}	0.9	0.999(0.0020)	1.000(0.0055)	1.000(0.0083)	0.985(0.0036)	1.000(0.0162)
τ_{high}	0.25	0.939(0.0145)	0.987(0.0268)	0.986(0.0399)	0.952(0.0210)	0.984(0.2771)
τ_{high}	hetero	0.968(0.0076)	0.987(0.0126)	0.999(0.0182)	0.956(0.0111)	1.000(0.0516)
τ_{high}	0.9	0.996(0.0023)	1.000(0.0067)	1.000(0.0102)	0.980(0.0038)	1.000(0.0172)
			$t_T = 50000$			
0.5	0.25	0.956(0.0056)	0.982(0.0081)	0.996(0.0122)	0.949(0.0081)	0.988(0.0571)
0.5	hetero	0.960(0.0032)	0.979(0.0044)	0.992(0.0064)	0.950(0.0046)	1.000(0.0115)
0.5	0.9	1.000(0.0018)	0.988(0.0027)	0.990(0.0036)	0.983(0.0021)	1.000(0.0038)
τ_{low}	0.25	0.957(0.0061)	0.981(0.0083)	0.994(0.0125)	0.944(0.0091)	0.993(0.0594)
τ_{low}	hetero	0.953(0.0036)	0.981(0.0046)	0.990(0.0066)	0.934(0.0054)	0.999(0.0121)
τ_{low}	0.9	1.000(0.0019)	1.000(0.0026)	0.989(0.0038)	0.988(0.0024)	1.000(0.0057)
τ_{high}	0.25	0.968(0.0059)	0.986(0.0086)	0.997(0.0133)	0.946(0.0089)	0.999(0.0620)
τ_{high}	hetero	0.953(0.0032)	0.990(0.0045)	0.998(0.0065)	0.952(0.0047)	0.993(0.0154)
τ_{high}	0.9	0.998(0.0010)	0.999(0.0023)	1.000(0.0034)	0.977(0.0016)	0.938(0.0132)

Table S.1: Empirical coverage probabilities (mean absolute errors) under varying quantile levels and response rates, with different t_T and fixed $K = 10$ clients and data generated from $\mathcal{N}(0, 1)$. In Case τ_{low}, each client uses a unique quantile level τ_k ranging uniformly from $[0.3, 0.5]$; in Case τ_{high}, τ_k is ranging from $[0.5, 0.8]$. "hetero" indicates client-specific truthful response rates r_k range from $[0.25, 0.9]$.

Quantile (τ)	Rate (r)	C1	C5	Log	DP-SGD (C1)	DC
			Hete L — $t_T = 10000$			
0.3	0.25	0.958(0.0184)	0.981(0.0311)	0.990(0.0452)	0.942(0.0260)	0.985(0.3066)
0.3	hetero	0.949(0.0096)	0.982(0.0150)	0.993(0.0205)	0.947(0.0142)	0.898(0.1302)
0.3	0.9	1.000(0.0029)	1.000(0.0066)	1.000(0.0100)	0.981(0.0049)	0.215(0.1273)
0.5	0.25	0.950(0.0165)	0.984(0.0315)	0.988(0.0465)	0.953(0.0224)	1.000(0.2822)
0.5	hetero	0.952(0.0085)	0.991(0.0155)	0.998(0.0221)	0.955(0.0119)	1.000(0.0525)
0.5	0.9	0.996(0.0025)	0.999(0.0078)	1.000(0.0120)	0.984(0.0041)	1.000(0.0186)
0.8	0.25	0.966(0.0237)	0.995(0.0512)	0.992(0.0791)	0.957(0.0328)	0.892(0.6152)
0.8	hetero	0.962(0.0122)	0.995(0.0227)	0.996(0.0347)	0.943(0.0186)	0.709(0.2684)
0.8	0.9	0.990(0.0042)	1.000(0.0116)	1.000(0.0185)	0.968(0.0065)	0.049(0.2098)
			Hete L — $t_T = 50000$			
0.3	0.25	0.937(0.0089)	0.981(0.0111)	0.990(0.0165)	0.916(0.0135)	0.949(0.1328)
0.3	hetero	0.911(0.0056)	0.981(0.0056)	0.997(0.0080)	0.885(0.0083)	0.093(0.1282)
0.3	0.9	0.977(0.0034)	1.000(0.0019)	1.000(0.0030)	0.908(0.0041)	0.000(0.1290)
0.5	0.25	0.958(0.0069)	0.988(0.0098)	0.995(0.0147)	0.949(0.0099)	1.000(0.0609)
0.5	hetero	0.964(0.0035)	0.994(0.0048)	0.996(0.0069)	0.957(0.0052)	0.997(0.0145)
0.5	0.9	1.000(0.0010)	1.000(0.0016)	1.000(0.0026)	0.993(0.0018)	0.979(0.0143)
0.8	0.25	0.956(0.0102)	0.991(0.0144)	0.998(0.0226)	0.931(0.0160)	0.799(0.2829)
0.8	hetero	0.950(0.0055)	0.992(0.0072)	0.997(0.0112)	0.923(0.0092)	0.014(0.2034)
0.8	0.9	1.000(0.0013)	1.000(0.0053)	0.999(0.0082)	0.985(0.0026)	0.000(0.1929)
			Hete D — $t_T = 10000$			
0.5	0.25	0.949(0.0132)	0.985(0.0243)	0.986(0.0354)	0.953(0.0183)	0.904(0.2496)
0.5	hetero	0.966(0.0069)	0.990(0.0117)	0.989(0.0172)	0.955(0.0098)	0.999(0.0488)
0.5	0.9	1.000(0.0023)	1.000(0.0074)	1.000(0.0117)	0.991(0.0035)	1.000(0.0163)
			Hete D — $t_T = 50000$			
0.5	0.25	0.958(0.0057)	0.980(0.0082)	0.993(0.0127)	0.943(0.0081)	0.981(0.0589)
0.5	hetero	0.966(0.0030)	0.988(0.0046)	0.998(0.0073)	0.950(0.0041)	1.000(0.0111)
0.5	0.9	0.999(0.0008)	1.000(0.0029)	1.000(0.0052)	0.990(0.0014)	1.000(0.0037)

Table S.2: Empirical coverage probabilities (mean absolute errors) under heterogeneous distributions for different t_T. The number of clients K is fixed at 10. In Hete L, data for each client k are independently generated from $\mathcal{N}(\mu_k, 1)$, where $\mu_k \sim \mathcal{N}(0, 1)$. In Hete D, data are generated from $\mathcal{N}(0, 1)$, $\mathcal{U}(-1, 1)$, and $\mathcal{C}(0, 1)$ across different clients. "hetero" indicates client-specific truthful response rates r_k range from $[0.25, 0.9]$.

Quantile (τ)	Rate (r)	C1	C5	Log
		$T = 5000$		
0.5	0.25	0.954(0.0189)	0.974(0.0129)	0.986(0.0112)
0.5	hetero	0.959(0.0103)	0.976(0.0065)	0.995(0.0052)
0.5	0.9	0.999(0.0033)	1.000(0.0040)	1.000(0.0026)
τ_{low}	0.25	0.957(0.0200)	0.974(0.0137)	0.991(0.0116)
τ_{low}	hetero	0.957(0.0108)	0.977(0.0067)	0.993(0.0053)
τ_{low}	0.9	1.000(0.0033)	1.000(0.0040)	1.000(0.0029)
τ_{high}	0.25	0.956(0.0212)	0.975(0.0128)	0.993(0.0123)
τ_{high}	hetero	0.961(0.0112)	0.984(0.0062)	0.996(0.0056)
τ_{high}	0.9	0.998(0.0037)	0.997(0.0028)	1.000(0.0031)
		$T = 10000$		
0.5	0.25	0.949(0.0133)	0.968(0.0078)	0.987(0.0061)
0.5	hetero	0.963(0.0071)	0.978(0.0037)	0.991(0.0030)
0.5	0.9	0.995(0.0023)	0.999(0.0020)	0.999(0.0014)
τ_{low}	0.25	0.947(0.0136)	0.972(0.0078)	0.984(0.0064)
τ_{low}	hetero	0.962(0.0072)	0.985(0.0038)	0.983(0.0033)
τ_{low}	0.9	0.999(0.0020)	1.000(0.0016)	0.967(0.0018)
τ_{high}	0.25	0.939(0.0145)	0.974(0.0086)	0.985(0.0066)
τ_{high}	hetero	0.968(0.0076)	0.988(0.0043)	0.985(0.0032)
τ_{high}	0.9	0.996(0.0023)	0.999(0.0031)	0.996(0.0014)

Table S.3: ECP (MAE) under varying quantile levels and response rates, with different T and fixed $K = 10$ clients and data generated from $\mathcal{N}(0, 1)$. In Case τ_{low}, each client uses a unique quantile level τ_k ranging uniformly from $[0.3, 0.5]$; in Case τ_{high}, τ_k is ranging from $[0.5, 0.8]$. "hetero" indicates client-specific truthful response rates r_k range from $[0.25, 0.9]$.

Quantile (τ)	Rate (r)	C1	C5	Log
		Hete L — $T = 5000$		
0.3	0.25	0.942(0.0271)	0.960(0.0168)	0.975(0.0151)
0.3	hetero	0.962(0.0131)	0.966(0.0086)	0.987(0.0067)
0.3	0.9	0.998(0.0043)	0.959(0.0063)	1.000(0.0033)
0.5	0.25	0.954(0.0254)	0.973(0.0154)	0.990(0.0153)
0.5	hetero	0.963(0.0120)	0.981(0.0072)	0.991(0.0071)
0.5	0.9	0.992(0.0042)	0.998(0.0032)	1.000(0.0034)
0.8	0.25	0.954(0.0375)	0.982(0.0242)	0.998(0.0248)
0.8	hetero	0.968(0.0181)	0.988(0.0109)	0.998(0.0116)
0.8	0.9	0.985(0.0108)	0.999(0.0070)	0.982(0.0094)
		Hete L — $T = 10000$		
0.3	0.25	0.958(0.0184)	0.966(0.0102)	0.981(0.0083)
0.3	hetero	0.949(0.0096)	0.965(0.0050)	0.979(0.0040)
0.3	0.9	1.000(0.0029)	0.979(0.0022)	0.867(0.0036)
0.5	0.25	0.950(0.0165)	0.974(0.0094)	0.985(0.0085)
0.5	hetero	0.952(0.0085)	0.976(0.0045)	0.991(0.0039)
0.5	0.9	0.996(0.0025)	0.985(0.0018)	1.000(0.0016)
0.8	0.25	0.966(0.0237)	0.983(0.0163)	0.990(0.0149)
0.8	hetero	0.962(0.0122)	0.988(0.0088)	0.974(0.0090)
0.8	0.9	0.990(0.0042)	0.997(0.0087)	0.645(0.0095)
		Hete D — $T = 5000$		
0.5	0.25	0.954(0.0195)	0.974(0.0129)	0.987(0.0109)
0.5	hetero	0.965(0.0098)	0.974(0.0075)	0.993(0.0049)
0.5	0.9	1.000(0.0037)	0.989(0.0060)	1.000(0.0026)
		Hete D — $T = 10000$		
0.5	0.25	0.949(0.0132)	0.968(0.0078)	0.982(0.0064)
0.5	hetero	0.966(0.0069)	0.973(0.0039)	0.972(0.0034)
0.5	0.9	1.000(0.0023)	0.999(0.0014)	0.966(0.0023)

Table S.4: ECP (MAE) under heterogeneous distributions for different T. The number of clients K is fixed at 10. In Hete L, data for each client k are independently generated from $\mathcal{N}(\mu_k, 1)$, where $\mu_k \sim \mathcal{N}(0, 1)$. In Hete D, data are generated from $\mathcal{N}(0, 1)$, $\mathcal{U}(-1, 1)$, and $\mathcal{C}(0, 1)$ across different clients. "hetero" indicates client-specific truthful response rates r_k range from $[0.25, 0.9]$.

B TECHNIQUE PROOFS

Proof of Theorem 2.1: Note that the following recursive equation

$$q_t^k = q_{t-1}^k + \frac{1 - r_k + 2\tau r_k}{2r_k}\eta_{m-1}\mathbb{I}(s_t^k = 1) - \frac{1 + r_k - 2\tau r_k}{2r_k}\eta_{m-1}\mathbb{I}(s_t^k = 0),$$

is asymptotically equivalent to

$$q_t^k = q_{t-1}^k + \eta_{m-1}\frac{1}{r_k}\left\{\frac{1 - r_k + 2\tau r_k}{2} - \mathbb{I}(\widetilde{x}_t^k \le q_{t-1}^k)\right\},$$

where

$$\mathbb{P}(\widetilde{x}_t^k = x_t^k) = r_k, \quad \mathbb{P}(\widetilde{x}_t^k = \infty) = \mathbb{P}(\widetilde{x}_t^k = -\infty) = (1 - r_k)/2.$$

Therefore, the minimizer of $\widetilde{\mathcal{L}}$ satisfies

$$\sum_{k=1}^K \frac{p_k}{r_k}\mathbb{P}(\widetilde{x}_t^k \le Q^\star) = \sum_{k=1}^K \{p_k Q_k + p_k(1 - r_k)/(2r_k)\} = \sum_{k=1}^K p_k \frac{1 - r_k + 2\tau r_k}{2r_k}.$$

This yields $\sum_{k=1}^K p_k Q_k = \sum_{k=1}^K p_k \tau = \tau$. The proof is complete.

Proof of Theorem 3.1:

Theorem 3.1 is a direct consequence of Theorem 3.2.

Proof of Theorem 3.2:

We follows the perturbed iterate framework that is derived by Mania et al. (2017) and also used in Li et al. (2022). Define the sequence $\bar{q}_t$ in the following way:

$$\bar{q}_t = \sum_{k=1}^K p_k q_t^k.$$

Define $\zeta^k = (x^k, U^k, V^k)^\top$, with

$$\mathbb{P}(U^k = 1) = r_k, \quad \mathbb{P}(U^k = 0) = 1 - r_k, \qquad \mathbb{P}(V^k = 1) = \mathbb{P}(V^k = 0) = 1/2.$$

For $k = 1, \ldots, K$, let U_t^k and V_t^k be i.i.d. copies of U^k and V^k, respectively. Denote $\zeta_t^k = (x_t^k, U_t^k, V_t^k)^\top$. Define

$$\begin{aligned} G_k(q_{t-1}^k, \zeta_t^k) = &\frac{1 + r_k - 2r_k\tau}{2r_k}\Big[\mathbf{1}\{\xi_t^k \le q_{t-1}^k\}U_t^k + (1 - U_t^k)(1 - V_t^k)\Big] \\ &- \frac{1 - r_k + 2r_k\tau}{2r_k}\Big[\mathbf{1}\{\xi_t^k > q_{t-1}^k\}U_t^k + (1 - U_t^k)V_t^k\Big]. \end{aligned}$$

Elementary calculations show that

$$g_k(q) := \mathbb{E}G(q, \zeta_t^k) = F_k(q) - \tau.$$

Define

$$g(q) = \sum_{k=1}^K p_k g_k(q).$$

Denote

$$\varepsilon_k(q) = G_k(q, \zeta_t^k) - g_k(q).$$

Besides,

$$\mathbb{E}\left(\varepsilon_k^2(q_{t-1}^k)|\mathcal{F}_{t-1}\right) = \frac{1 - r_k^2\{2F_k(q_{t-1}^k) - 1\}^2}{4r_k^2}.$$

By definition, for $t_m \le t < t_{m+1} - 1$, we have

$$q_{t+1}^k = \bar{q}_{t_m} - \eta_m \sum_{i=t_m}^{t} G_k(q_i^k, \zeta_i^k).$$

Define $s_m = \bar{q}_{t_m} - Q^\star$, and recall that $E_m = t_{m+1} - t_m$ and $\gamma_m = \eta_m E_m$. Elementary Iteration from $t = t_m$ to $t_{m+1} - 1$ yields

$$s_{m+1} = s_m - \eta_m \sum_{t=t_m}^{t_{m+1}-1} \sum_{k=1}^{K} p_k G_k(q_{t-1}^k, \zeta_t^k) = s_m - \gamma_m \nu_m,$$

in which

$$\nu_m = \frac{1}{E_m} \sum_{t=t_m}^{t_{m+1}-1} \sum_{k=1}^{K} p_k G_k(q_{t-1}^k, \zeta_t^k).$$

We define

$$h_m := \frac{1}{E_m} \sum_{t=t_m}^{t_{m+1}-1} \sum_{k=1}^{K} p_k G_k\left(\bar{q}_{t_m}; \zeta_t^k\right),$$

and further decompose that

$$\begin{aligned} \nu_m &= G s_m + \left(g\left(\bar{q}_{t_m}\right) - G s_m\right) + \left(h_m - g\left(\bar{q}_{t_m}\right)\right) + \left(\nu_m - h_m\right) \\ &:= G s_m + r_m + \varepsilon_m + \delta_m, \end{aligned}$$

where $G = \sum_{k=1}^{K} p_k r_k f_k(Q^\star)$ is the Hessian at $Q^\star$. It then follows that

$$s_{m+1} = \left(1 - \gamma_m G\right) s_m - \gamma_m \left(r_m + \varepsilon_m + \delta_m\right) := B_m s_m - \gamma_m U_m, \tag{B.1}$$

where $B_m := 1 - \gamma_m G$ and $U_m := r_m + \varepsilon_m + \delta_m$ for short. Recurring (B.1) gives

$$s_{m+1} = \left(\prod_{j=0}^{m} B_j\right) s_0 - \sum_{j=0}^{m} \left(\prod_{i=j+1}^{m} B_i\right) \gamma_j U_j.$$

Here, we use the convention that $\prod_{i=m+1}^{m} B_i = 1$ for any $m \ge 0$. Recall the definition that

$$h(r, T) = \max\left\{ n \in \mathbb{Z}_+ \,\middle|\, r \sum_{m=1}^{T} \frac{1}{E_m} \ge \sum_{m=1}^{n} \frac{1}{E_m} \right\}.$$

Hence,

$$\begin{aligned} \frac{\sqrt{t_T}}{T} \sum_{m=0}^{h(r,T)} s_{m+1} &= \frac{\sqrt{t_T}}{T} \sum_{m=0}^{h(r,T)} \left[\left(\prod_{j=0}^{m} B_j\right) s_0 - \sum_{j=0}^{m} \left(\prod_{i=j+1}^{m} B_i\right) \gamma_j U_j \right] \\ &= \frac{\sqrt{t_T}}{T} \sum_{m=0}^{h(r,T)} \left(\prod_{j=0}^{m} B_j\right) s_0 - \frac{\sqrt{t_T}}{T} \sum_{j=0}^{h(r,T)} \sum_{m=j}^{h(r,T)} \left(\prod_{i=j+1}^{m} B_i\right) \gamma_j U_j. \end{aligned}$$

For any $n \ge j$, define A_j^n as

$$A_j^n = \sum_{l=j}^{n} \left(\prod_{i=j+1}^{l} B_i\right) \gamma_j.$$

With the notation of A_j^n, we can rewrite that

$$\frac{\sqrt{t_T}}{T} \sum_{m=0}^{h(r,T)} s_{m+1} = \frac{\sqrt{t_T}}{T \gamma_0} A_0^{h(r,T)} B_0 s_0 - \frac{\sqrt{t_T}}{T} \sum_{m=0}^{h(r,T)} A_m^{h(r,T)} U_m.$$

Since $U_m = r_m + \varepsilon_m + \delta_m$, then

$$\begin{aligned}\frac{\sqrt{t_T}}{T}\sum_{m=0}^{h(r,T)} s_{m+1} + \frac{\sqrt{t_T}}{T}\sum_{m=0}^{h(r,T)} G^{-1}\varepsilon_m &= \frac{\sqrt{t_T}}{T\gamma_0}A_0^{h(r,T)}B_0 s_0 - \frac{\sqrt{t_T}}{T}\sum_{m=0}^{h(r,T)} A_m^{h(r,T)}\left(r_m + \delta_m\right) \\ &\quad - \frac{\sqrt{t_T}}{T}\sum_{m=0}^{h(r,T)}\left(A_m^T - G^{-1}\right)\varepsilon_m \\ &\quad - \frac{\sqrt{t_T}}{T}\sum_{m=0}^{h(r,T)}\left(A_m^{h(r,T)} - A_m^T\right)\varepsilon_m \\ &=: \mathcal{T}_0 - \mathcal{T}_1 - \mathcal{T}_2 - \mathcal{T}_3.\end{aligned}$$

To complete the proof, we first investigate the partial-sum asymptotic behavior of

$$\frac{\sqrt{t_T}}{T}\sum_{m=0}^{h(r,T)} G^{-1}\varepsilon_m,$$

and then show that the four separate terms: $\sup_{r\in[0,1]}\|\mathcal{T}_0\|, \sup_{r\in[0,1]}\|\mathcal{T}_1\|$, $\sup_{r\in[0,1]}\|\mathcal{T}_2\|$, and $\sup_{r\in[0,1]}\|\mathcal{T}_4\|$ are $o_{\mathbb{P}}(1)$, respectively.

We aim to follow the proof of Theorem 4.2 in Li et al. (2022). However, we find that the average smoothness condition in their Assumption 3.1 is not satisfied, because here we only have

$$\sqrt{\mathbb{E}\left\{G_k(x,\zeta_t^k) - G_k(y,\zeta_t^k)\right\}^2} \lesssim |x-y|^{1/2}. \tag{B.2}$$

Upon close examination of their proof, we find that this condition is crucial in the proof of their key Lemma B.2.

In the following, we re-establish the proof of

$$\mathbb{E}|\bar{q}_{t_m} - Q^\star|^2 \lesssim \gamma_m, \quad \bar{q}_{t_m} \xrightarrow{a.s.} Q^*.$$

under the condition given in (B.2). Consider that

$$\begin{aligned}\mathbb{E}\left(\left|q_{t+1}^k - \bar{q}_{t_m}\right| |\mathcal{F}_{t_m}\right) &= \eta_m \mathbb{E}\left(\left|\sum_{i=t_m}^{t} G_k(q_i^k,\zeta_i^k)\right| |\mathcal{F}_{t_m}\right) \\ &\le \eta_m \sum_{i=t_m}^{t}\mathbb{E}\left(|G_k(q_i^k,\zeta_i^k)||\mathcal{F}_{t_m}\right) \\ &\lesssim \eta_m \sum_{i=t_m}^{t}\left(1 + |q_i^k - \bar{q}_{t_m}| + |\bar{q}_{t_m} - Q^\star|\right),\end{aligned}$$

where the last inequality holds by the following fact

$$\begin{aligned}\mathbb{E}\left(G_k^2(q_i^k,\zeta_i^k)|\mathcal{F}_i\right) &= \mathbb{E}\left(\left|G_k(q_i^k,\zeta_i^k) - g_k(q_i^k)\right|^2 |\mathcal{F}_i\right) + g_k^2(q_i^k) \\ &\le \mathbb{E}\left(\left|G_k(q_i^k,\zeta_i^k) - g_k(q_i^k)\right|^2 |\mathcal{F}_i\right) + 2|g_k(q_i^k) - g_k(Q^\star)|^2 + 2g_k^2(Q^\star) \\ &\lesssim \{1 + 2g_k^2(Q^\star)\} + |q_i^k - Q^\star|^2 \\ &\lesssim 1 + |q_i^k - \bar{q}_{t_m}|^2 + |\bar{q}_{t_m} - Q^\star|^2.\end{aligned}$$

Define

$$V_t = \sum_{k=1}^{K} p_k \mathbb{E}\left(|q_t^k - \bar{q}_{t_m}||\mathcal{F}_{t_m}\right).$$

Hence,

$$V_{t+1} \lesssim \eta_m \sum_{i=t_m}^{t}\left(1 + |\bar{q}_{t_m} - Q^\star| + V_i\right),$$

which further implies that (since $V_{t_m} = 0$)

$$\begin{aligned}\frac{1}{E_m}\sum_{t=t_m}^{t_{m+1}-1} V_t = \frac{1}{E_m}\sum_{t=t_m}^{t_{m+1}-2} V_{t+1} &\lesssim \frac{\eta_m}{E_m}\sum_{t=t_m}^{t_{m+1}-2}\sum_{i=t_m}^{t}\left(1+|\bar{q}_{t_m}-Q^\star|+V_i\right)\\ &= \frac{\eta_m}{E_m}\sum_{t=t_m}^{t_{m+1}-2}\left(t_{m+1}-t-1\right)\left(1+|\bar{q}_{t_m}-Q^\star|+V_t\right)\\ &\lesssim \eta_m\sum_{t=t_m}^{t_{m+1}-2}\left(1+|\bar{q}_{t_m}-Q^\star|+V_t\right)\\ &\lesssim \eta_m E_m\left(1+|\bar{q}_{t_m}-Q^\star|+\frac{1}{E_m}\sum_{t=t_m}^{t_{m+1}-1} V_t\right).\end{aligned}$$

Denote by $\gamma_m = \eta_m E_m$. It follows that

$$\frac{1}{E_m}\sum_{t=t_m}^{t_{m+1}-1} V_t \lesssim \gamma_m\left(1+|\bar{q}_{t_m}-Q^\star|\right). \tag{B.3}$$

Let $\mathcal{G}_k(\cdot)$ denote an antiderivative of $g_k(\cdot)$, and $\mathcal{G}(\cdot)=\sum_{k=1}^K p_k\mathcal{G}_k(\cdot)$. Let $\Delta_m = \mathcal{G}(\bar{q}_{t_m}) - \mathcal{G}(Q^\star)$. The equation (17) in Li et al. (2022) shows that for some constant $L>0$,

$$\begin{aligned}\mathbb{E}\left\{\mathcal{G}(\bar{q}_{t_m+1})|\mathcal{F}_{t_m}\right\} \le \mathcal{G}(\bar{q}_{t_m}) - \gamma_m/2\left|\nabla\mathcal{G}(\bar{q}_{t_m})\right|^2 + \gamma_m^2 L\mathbb{E}(h_m^2|\mathcal{F}_{t_m})\\ + \left(\gamma_m/2+\gamma_m^2 L\right)\mathbb{E}(\delta_m^2|\mathcal{F}_{t_m}),\end{aligned}$$

where

$$h_m = \frac{1}{E_m}\sum_{t=t_m}^{t_{m+1}-1}\sum_{k=1}^K p_k\nabla G_k(\bar{q}_{t_m}^k,\zeta_t^k), \quad \delta_m = \frac{1}{E_m}\sum_{t=t_m}^{t_{m+1}-1}\sum_{k=1}^K p_k\nabla G_k(\bar{q}_t^k,\zeta_t^k).$$

Lemma B.9 of Li et al. (2022) obtains that

$$\mathbb{E}(h_m^2|\mathcal{F}_{t_m}) \le |\nabla\mathcal{G}(\bar{q}_{t_m})|^2 + \frac{C_1}{E_m} + \frac{C_2}{E_m}|\bar{q}_{t_m}-Q^\star|^2.$$

Notice that

$$\mathbb{E}\left\{\left(G_k(q_{t-1}^k,\zeta_t^k) - G_k(\bar{q}_{t_m},\zeta_t^k)\right)^2|\mathcal{F}_{t_m}\right\} \lesssim \mathbb{E}\left(\left|q_t^k-\bar{q}_{t_m}\right||\mathcal{F}_{t_m}\right).$$

Thus,

$$\begin{aligned}\mathbb{E}(\delta_m^2|\mathcal{F}_{t_m}) &\lesssim \frac{1}{E_m}\sum_{t=t_m}^{t_{m+1}-1}\sum_{k=1}^K p_k\mathbb{E}\left\{\left(G_k(q_{t-1}^k,\zeta_t^k) - G_k(\bar{q}_{t_m},\zeta_t^k)\right)^2|\mathcal{F}_{t_m}\right\}\\ &\lesssim \frac{1}{E_m}\sum_{t=t_m}^{t_{m+1}-1}\sum_{k=1}^K p_k\mathbb{E}\left(\left|q_t^k-\bar{q}_{t_m}\right||\mathcal{F}_{t_m}\right)\\ &\lesssim \gamma_m\left(1+|\bar{q}_{t_m}-Q^\star|\right),\end{aligned}$$

where the last inequality holds by (B.3). Therefore, we obtain that

$$\begin{aligned}\mathbb{E}(\Delta_{m+1}|\mathcal{F}_{t_m}) &\le \Delta_m - \gamma_m/2|\nabla\mathcal{G}(\bar{q}_{t_m})|^2 + \gamma_m^2 L\left\{|\nabla\mathcal{G}(\bar{q}_{t_m})|^2 + \frac{C_1}{E_m} + \frac{C_2}{E_m}|\bar{q}_{t_m}-Q^\star|^2\right\}\\ &\quad + \left(\gamma_m/2+\gamma_m^2 L\right)\gamma_m\left(C_1+C_2|\bar{q}_{t_m}-Q^*|\right)\\ &\le (1-c_1\gamma_m+c_2\gamma_m^2)\Delta_m + \left(c_3+c_4\Delta_m^{1/2}\right)\gamma_m^2.\end{aligned}$$

Since we assume that the parameter space is uniformly bounded, it entails that Δ_m is also uniformly bounded. Thus, we have

$$\mathbb{E}(\Delta_{m+1}|\mathcal{F}_{t_m}) \le (1-c_1\gamma_m+c_2\gamma_m^2)\Delta_m + (c_3+c_5)\gamma_m^2.$$

Apply Robbins-Siegmund theorem in Robbins & Siegmund (1971) to obtain $\bar{q}_{t_m} \to Q^\star$ almost surely. Lemma A.10 in Su & Zhu (2018) states that for any positive constants c_1, c_2, if $\gamma_m = o(1)$, $\gamma_{m-1}/\gamma_m = 1 + o(\gamma_m)$, and B_m is a positive sequence, satisfying

$$B_m \leq \frac{\gamma_{m-1}(1 - c_1\gamma_m)}{\gamma_m} B_{m-1} + c_2\gamma_m,$$

then $\sup_m B_m < \infty$. Using this lemma, we immediately obtain that for some positive constant $C > 0$,

$$\sup_{m \geq 1} \frac{\mathbb{E}\Delta_m}{\gamma_{m-1}} < C,$$

which entails that

$$\mathbb{E}|\bar{q}_{t_m} - Q^\star|^2 \lesssim \mathbb{E}\Delta_m \lesssim \gamma_{m-1} \lesssim \gamma_m \{1 + o(\gamma_m)\} \lesssim \gamma_m.$$

To demonstrate that our setting satisfies Assumption 3.2 of Li et al. (2022), we define

$$S_k := \mathbb{E}\varepsilon_k^2(Q^\star) = \frac{1 - r_k^2\{2F_k(Q^\star) - 1\}^2}{4} = \frac{1 - r_k^2\{2Q_k - 1\}^2}{4}.$$

Hence,

$$\left|\mathbb{E}\left(\varepsilon_k^2(q_{t-1}^k)|\mathcal{F}_{t-1}\right) - \mathbb{E}\varepsilon_k^2(Q^\star)\right| \lesssim |q_{t-1}^k - Q^\star|,$$

satisfying Assumption 3.2 in Li et al. (2022). Assumptions 3.3 and 3.4 of Li et al. (2022) are the same as our Assumptions 3-4. By carefully examining their proof, we find that Assumption 1 in [29] is only used to establish the key result that $\mathbb{E}|\bar{q}_{t_m} - Q^\star|^2 \lesssim \gamma_m$. While in our setting, we have already rigorously established this result under equation (B.2). Therefore, their Assumption 1 is not required for our theoretical development beyond this step. Therefore, we could follow the arguments in the proof of Theorem 4.2 in Li et al. (2022) to derive the functional CLT established in our Theorem 3.2.