DECENTRALIZED LOCALLY PRIVATE INFERENCE FOR QUANTILES

Anonymous authors
Paper under double-blind review

ABSTRACT

We study quantile inference under local differential privacy for heterogeneous data streams distributed over a peer-to-peer network without a trusted central server. The proposed procedure targets a weighted mixture quantile using one-bit randomized-response scores and blockwise gossip, with each observation processed once and communication restricted to neighboring nodes. The main challenge is that incomplete synchronization creates persistent network disagreement that interacts with discontinuous quantile scores. We derive graph-dependent error bounds and establish asymptotic normality and a functional central limit theorem for the spatial–temporal average under suitable mixing and growth conditions. These results yield self-normalized confidence intervals without estimating the mixture density or the privacy-inflated score variance. Under stronger balance conditions, node-wise temporal averages inherit the same first-order limits, enabling inference from a single node's trajectory without additional network-wide aggregation. Numerical experiments and a real-data application examine the effects of privacy, communication frequency, and network connectivity on estimation and inference.

1 INTRODUCTION

Quantiles are among the most common distributional targets in statistics and machine learning. They summarize tails and medians rather than averages, are stable under heavy-tailed noise, and underlie quantile regression, distributional risk assessment, robust online monitoring, and prediction interval construction. Classical quantile theory has a long history, including asymptotic representations for sample quantiles and the check-loss formulation of quantile regression (Bahadur, 1966; Serfling, 1980; Koenker & Bassett, 1978; Koenker, 2005). In modern data systems, quantile targets are also central in streaming analytics, where one wants to update distributional summaries without storing the full sample (Greenwald & Khanna, 2001). These settings motivate stochastic-approximation algorithms that use only simple threshold comparisons and remain valid when the raw observations are never pooled (Robbins & Monro, 1951).

Decentralized learning addresses a different bottleneck: data are held by multiple agents that may have no trusted central server. A node exchanges messages only with graph neighbors, and the network learns a target defined by all local distributions. This model appears in sensor networks, mobile and edge devices, collaborative clinical studies, and distributed monitoring systems. Foundational methods combine local subgradient or stochastic-gradient steps with consensus or gossip mixing (Nedić & Ozdaglar, 2009; Duchi et al., 2012; Lian et al., 2017); subsequent work covers changing topologies, local updates, and heterogeneous data (Nedić & Olshevsky, 2015; Koloskova et al., 2020). Statistical inference for decentralized learning is more recent. For smooth M-estimation, Gu & Chen (2024) develop decentralized federated inference, while Bonnerjee et al. (2025) establish Gaussian approximations for decentralized federated learning trajectories. Local-SGD inference and Polyak–Ruppert averaging provide related inferential tools in server-based distributed systems (Ruppert, 1988; Polyak & Juditsky, 1992; Li et al., 2022).

Quantile objectives create additional non-smoothness. Decentralized quantile-regression methods based on ADMM, network averaging, and smoothing have been studied in optimization and signal-processing settings (Wang & Lian, 2023; Wang et al., 2024; Shi et al., 2025). These works are not locally private and typically focus on optimization or estimation rather than nuisance-free online

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

inference for a scalar mixture quantile. Our setting instead combines a discontinuous threshold score, heterogeneous one-bit privacy mechanisms, and fully decentralized communication.

Privacy constraints make the problem more stringent. Differential privacy gives a formal stability guarantee against changes in one record (Dwork, 2006; Dwork & Roth, 2014). In the local model, each record is randomized before transmission, so a trusted curator is not required. Local differential privacy can be traced to randomized response (Warner, 1965) and was formalized for learning and statistical estimation by Kasiviswanathan et al. (2011) and Duchi et al. (2013); see also Kairouz et al. (2016) for extremal local mechanisms and Gaboardi et al. (2019) for locally private confidence intervals. One-bit randomized-response mechanisms are also practically important in large-scale telemetry systems (Erlingsson et al., 2014; Ding et al., 2017). For quantile estimation, recent work gives online and server-based locally private inference procedures (Liu et al., 2023; Cai et al., 2026); central differential privacy for robust statistics and quantiles is related but relies on different trust assumptions (Dwork & Lei, 2009; Gillenwater et al., 2021). The online and server-based LDP quantile methods above do not address inference from a node trajectory when aggregation is implemented through local neighbor communication. The central challenge in peer-to-peer quantile inference is the interaction between network disagreement and discontinuous quantile scores. Server-coordinated quantile procedures synchronize local iterates through aggregation and broadcast (Cai et al., 2026). With blockwise gossip, nodes may begin successive blocks at different thresholds, so disagreement propagates across blocks and affects subsequent private score evaluations.

Our analysis accounts jointly for score discontinuity and data heterogeneity. The quantile indicator score has only square-root L_2 regularity in the threshold, rather than the Lipschitz regularity available for smooth stochastic gradients. Moreover, under heterogeneous node distributions, the local score means need not vanish individually even though they balance in the weighted aggregate at the target quantile. These features are incorporated into a network-mean linearization that explicitly tracks how disagreement propagates across communication rounds and feeds back into subsequent private score evaluations. Controlling these terms on the inferential scale yields distributional limits for the spatial–temporal average and allows the corresponding process limit to be transferred to individual node trajectories. Table S.1 compares the proposed method with related private and decentralized inferential frameworks. Figure 1 summarizes the proposed workflow.

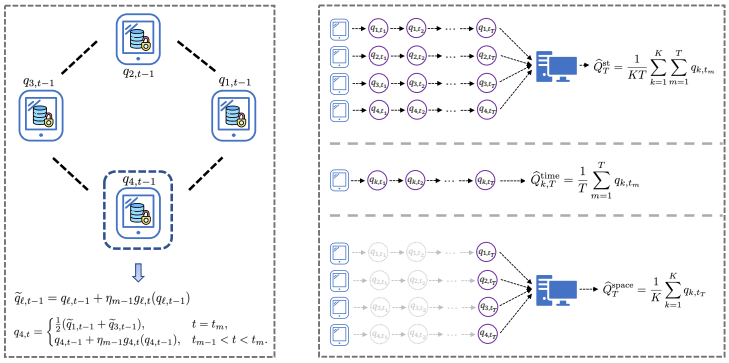

Figure 1: Decentralized locally private quantile inference with $K = 4$ nodes. Left: local updates using privatized one-bit scores and neighbor-to-neighbor gossip. Right: spatial–temporal, node-wise temporal, and terminal spatial averaging; see Section 2.

Our contributions are as follows.

- We develop a one-pass event-level LDP protocol for a weighted mixture quantile over a peer-to-peer network. The protocol combines one-bit randomized response and blockwise gossip, while accommodating heterogeneous node distributions, weights, target levels, and privacy parameters.
- We derive a network-mean linearization with explicit graph-dependent error bounds that account for the interaction between network disagreement, discontinuous private scores, and within-block drift. This yields a central limit theorem, a functional central limit theorem, and a pivotal self-normalized limit without plug-in estimation of the mixture density or the privacy-inflated variance.

- We establish graph–statistics balance conditions under which the spatial–temporal average has the same first-order limit as its server-coordinated counterpart. Under stronger conditions, we transfer the network-mean functional limit to individual node trajectories, enabling self-normalized inference from a single node's trajectory without additional network-wide aggregation. With communication after every local update, we also derive a Chung-type limit for the terminal spatial average.

The rest of the paper is organized as follows. Section 2 defines the target, the local privacy mechanism, and the decentralized recursion. Section 3 gives the assumptions and main inference results. Section 4 evaluates the finite-sample performance of the proposed method through numerical experiments and a real-data application. Section 5 discusses limitations and extensions. Proofs and auxiliary lemmas are collected in Appendix.

2 METHODOLOGY

Necessary notation and definitions. For $v \in \mathbb{R}^K$, write $\|v\|_\infty = \max_{1\le i\le K}|v_i|$. For $A \in \mathbb{R}^{K\times K}$, define $\|A\|_\infty = \max_{1\le i\le K}\sum_{j=1}^K |A_{ij}|$, and $\|A\|_1 = \max_{1\le j\le K}\sum_{i=1}^K |A_{ij}|$. For vector $\boldsymbol{q}_t = (q_{1,t},\dots,q_{K,t})$, write $\bar{q}_t = K^{-1}\sum_{k=1}^K q_{k,t}$. We next recall the definitions of central and local differential privacy.

Definition 1 (Differential privacy (Dwork, 2006; Dwork & Roth, 2014))**.** *A randomized algorithm $\mathcal{A}$ is (ϵ,δ)-differentially private if, for any two datasets S, S' differing in a single individual record and any measurable event E,*

$$\mathbb{P}\{\mathcal{A}(S)\in E\} \le e^{\epsilon}\mathbb{P}\{\mathcal{A}(S')\in E\} + \delta.$$

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

Definition 2 (Local differential privacy (Kasiviswanathan et al., 2011; Duchi et al., 2013))**.** *A randomized map $R : \mathcal{X} \to \mathcal{Y}$ is (ϵ,δ)-locally differentially private if, for every pair of inputs $x, x' \in \mathcal{X}$ and every measurable event $E \subseteq \mathcal{Y}$,*

$$\mathbb{P}\{R(x)\in E\} \le e^{\epsilon}\mathbb{P}\{R(x')\in E\} + \delta.$$

Problem formulation. We consider an online decentralized system with K nodes connected by a peer-to-peer communication network. For the stationary theory, node k observes an i.i.d. stream $\{X_{k,t}\}_{t\ge1}$ from a distribution $\mathcal{P}_k$ with CDF F_k and density f_k. The node distributions may differ. Let $p_k \ge 0$ be prespecified target weights satisfying $\sum_k p_k = 1$, and let $\tau_k \in (0,1)$. These local levels induce the global quantile level $\tau = \sum_{k=1}^K p_k\tau_k$. Our target is the τ-th quantile of the global mixture distribution $F(q) = \sum_{k=1}^K p_k F_k(q)$. We denote the target parameter by $Q^\star$, which is the unique solution of $F(Q^\star) = \tau$. Equivalently, $Q^\star$ can be characterized as the minimizer of a weighted global check-loss objective, that is,

$$Q^\star = \arg\min_{q\in\Theta}\sum_{k=1}^K p_k\mathbb{E}_{X\sim\mathcal{P}_k}\big[(X-q)\{\tau_k - \mathbf{1}(X<q)\}\big]. \tag{1}$$

All population quantities, including K, p_k, F_k, and $Q^\star$, may depend on the horizon T; this triangular-array dependence is suppressed unless needed. The parameter space Θ is assumed to be bounded as in Gu & Chen (2024). When the target is the quantile of a pooled finite population partitioned into node sizes N_k, we use $p_k = N_k/\sum_j N_j$; equal node weights instead define a deliberately region-balanced target.

One-bit LDP score. In practice, we consider an SGD estimator. Then node k needs to calculate the negative stochastic subgradient $\tau_k - \mathbf{1}\{X_{k,t} < q\}$. This quantity involves the threshold bit $\mathbf{1}\{X_{k,t} < q\}$, which may reveal information about the individual observation $X_{k,t}$. To protect this bit, we apply randomized response at each queried threshold, similar to the approach in a single-client setting; see Liu et al. (2023). Specifically, given a truthful-response probability $r_k \in (0,1]$, node k generates $U_{k,t} \sim \text{Bernoulli}(r_k)$ and $V_{k,t} \sim \text{Bernoulli}(1/2)$, and releases $S_{k,t}(q) = \mathbf{1}\{X_{k,t} < q\}U_{k,t} + V_{k,t}(1-U_{k,t})$ instead of $\mathbf{1}\{X_{k,t} < q\}$. We use the debiased privatized score

$$g_{k,t}(q) = \tau_k - \frac{S_{k,t}(q) - (1-r_k)/2}{r_k}.$$

The following proposition gives the conditional channel guarantee and score moments. The sequential transcript guarantee is stated immediately afterward because the threshold is adaptive.

Proposition 1 (Privacy, unbiasedness, and variance)**.** *Fix k, t, and q.*

- *If $0 < r_k < 1$, then $S_{k,t}(q)$ is ε_k-LDP with $\varepsilon_k = \log(1 + r_k) - \log(1 - r_k)$. If $r_k = 1$, the mechanism is the non-private boundary case, which we record by the convention $\varepsilon_k = \infty$.*
- *The privatized score is unbiased: $\mathbb{E}[g_{k,t}(q) \mid X_{k,t}, q] = \tau_k - \mathbf{1}\{X_{k,t} < q\}$.*
- *The variance of the privatized score is $\mathrm{Var}(g_{k,t}(q)) = (r_k^{-2} - \{2F_k(q) - 1\}^2)/4$.*
- *If $r_{\min} := \inf_T \min_{k \le K} r_k > 0$, then $|g_{k,t}(q)| \le r_{\min}^{-1}$ almost surely.*

Decentralized implementation. We subsequently consider solving the optimization problem (1) under LDP over a decentralized federated system. Let $\mathcal{G} = ([K], \mathcal{E})$ be the communication graph, and let $W \in \mathbb{R}^{K \times K}$ be a nonnegative mixing matrix compatible with $\mathcal{G}$, in the sense that $W_{\ell k} > 0$ only when $\ell = k$ or $(\ell, k) \in \mathcal{E}$. Communication occurs at times $0 = t_0 < t_1 < \cdots < t_T$, and $E_m = t_m - t_{m-1}$ denotes the length of the mth communication interval. For $t_{m-1} < t \le t_m$, the iterate at node k is updated as

$$q_{k,t} = \begin{cases} \sum_{\ell=1}^{K} W_{\ell k} \left(q_{\ell,t-1} + \eta_{m-1} K p_\ell \, g_{\ell,t}(q_{\ell,t-1}) \right), & t = t_m, \\ q_{k,t-1} + \eta_{m-1} K p_k \, g_{k,t}(q_{k,t-1}), & t_{m-1} < t < t_m. \end{cases}$$

The complete procedure is summarized in Algorithm 1.

Algorithm 1 Decentralized one-bit LDP quantile recursion

Require: $\{p_k, \tau_k, r_k\}_{k=1}^K$, $\{\eta_m\}$, $\{t_m\}_{m=0}^T$, W.
1: Initialize $q_{k,0} = q_0$ for all $k \in [K]$.
2: **for** $m = 1, \ldots, T$ **do**
3: **for** $t = t_{m-1} + 1, \ldots, t_m$ **do**
4: For all $k \in [K]$ in parallel, draw $U_{k,t} \sim \mathrm{Bernoulli}(r_k)$, $V_{k,t} \sim \mathrm{Bernoulli}(1/2)$.
5: Locally randomize the fresh record to obtain $S_{k,t} = \mathbf{1}\{X_{k,t} < q_{k,t-1}\} U_{k,t} + V_{k,t}(1 - U_{k,t})$.
 $q_{k,t} = q_{k,t-1} + \eta_{m-1} K p_k \, [\tau_k - \{S_{k,t} - (1 - r_k)/2\}/r_k]$.
6: **end for**
7: Communicate with neighbors and set $\boldsymbol{q}_{t_m} \leftarrow \boldsymbol{q}_{t_m} W$.
8: **end for**

Averaging schemes. Following the decentralized-inference literature (Gu & Chen, 2024; Bonnerjee et al., 2025), we consider three averaging schemes. The spatial–temporal average combines the node-wise Polyak–Ruppert trajectories, the node-wise temporal average is computed locally from a single node's trajectory, and the terminal spatial average averages the final node iterates. The two spatial averages require a network-wide reduction, whereas the node-wise temporal average is locally available at each node. Specifically,

$$\widehat{Q}_T^{\mathrm{st}} = \frac{1}{KT} \sum_{k=1}^{K} \sum_{m=1}^{T} q_{k,t_m}, \qquad \widehat{Q}_{k,T}^{\mathrm{time}} = \frac{1}{T} \sum_{m=1}^{T} q_{k,t_m}, \qquad \widehat{Q}_T^{\mathrm{space}} = \frac{1}{K} \sum_{k=1}^{K} q_{k,t_T}.$$

Establishing and comparing their statistical properties is therefore a problem of great importance.

3 Theoretical results

In this section, we establish the asymptotic theory for the proposed estimators. To this end, we impose the following regularity assumptions.

Assumption 1 (Density regularity)**.** *There exists a compact interval $\Theta \subset \mathbb{R}$ and constants $0 < f_- \le f_+ < \infty$ and $L_f < \infty$ such that:*

- *each f_k is bounded and Lipschitz on Θ: $0 \le f_k(x) \le f_+, |f_k(x) - f_k(y)| \le L_f |x - y|, x, y \in \Theta$;*

- *the mixture density* $f(x) = \sum_{k=1}^{K} p_k f_k(x)$ *satisfies* $f(x) \geq f_-, \qquad x \in \Theta$.

Assumption 2 (Privacy, balanced weights, and nondegenerate score variance)**.** *There exist constants* $0 < r_{\min} \leq 1$ *and* $0 < b_1 \leq b_2 < \infty$ *such that* $r_{\min} \leq r_k \leq 1$ *and* $b_1 \leq Kp_k \leq b_2$ *for all* $1 \leq k \leq K$. *Moreover, let* $\kappa^2 = \sum_{k=1}^{K} p_k^2, \overline{S} = \kappa^{-2} \sum_{k=1}^{K} p_k^2 \mathrm{Var}(g_{k,t}(Q^\star))$. *There exist constants* $0 < s_- \leq s_+ < \infty$ *such that* $s_- \leq \overline{S} \leq s_+$.

Assumption 3 (Fixed symmetric spectral network)**.** *The matrix* W *is nonnegative, symmetric, and doubly stochastic,* $W\mathbf{1}_K = \mathbf{1}_K, \mathbf{1}_K^\top W = \mathbf{1}_K^\top, W^\top = W, (W)_{ij} \geq 0$. *Let* $J = K^{-1}\mathbf{1}\mathbf{1}^\top$ *and* $\rho = \|W - J\|_2 = \max_{j \geq 2} |\lambda_j(W)|$, *assume a uniform gap:* $\rho \leq \bar{\rho} < 1$ *for a constant.*

Assumption 4 (Block schedule)**.** *Define the effective block step size as* $\gamma_{m-1} = \eta_{m-1} E_m$. *There exist constants* $c_E, c_\gamma \in (0, \infty)$, $\beta \in [0, 2\alpha - 1)$, *and* $\alpha \in (1/2, 1)$ *such that* $E_m \sim c_E(m + 1)^\beta, \gamma_m \sim c_\gamma(m + 1)^{-\alpha}$. *Define* $t_T = \sum_{m=1}^{T} E_m, H_T = \sum_{m=1}^{T} E_m^{-1}, \nu_T = t_T H_T / T^2$. *Then* $t_T \sim c_E T^{1+\beta}/(1 + \beta), H_T \sim c_E^{-1} T^{1-\beta}/(1 - \beta), \nu_T \to \nu_\beta := 1/(1 - \beta^2)$.

Assumption 1 is a mild regularity condition on the uniform boundedness of the density functions. Assumptions 2 and 3 require relatively balanced sample sizes and network topology; see similar assumptions in Gu & Chen (2024) and Bonnerjee et al. (2025). The finite-graph factors are $\sum_{j=1}^{m} \gamma_{j-1} \rho^{m-j+1}$ and $A_{K,\rho} = \sum_{\ell \geq 1} \min\{2, \sqrt{K}\rho^\ell\}$. A fixed connected sparse graph with a lazy nonnegative doubly stochastic mixing matrix satisfies $\rho < 1$. Assumption 4 requires that the effective step sizes decay slowly and the communication intervals increase slowly. It relaxes the common assumption in the literature that E_m is constant (Gu & Chen, 2024; Bonnerjee et al., 2025).

Write $\mathfrak{n}_T = t_T / \sum_{k=1}^{K} p_k^2$ as the effective sample size. We first study the properties of $\widehat{Q}_T^{\mathrm{st}}$.

Theorem 1 (Linear expansion for $\widehat{Q}_T^{\mathrm{st}}$)**.** *Under Assumptions 1–4,*

$$\sqrt{\mathfrak{n}_T}\big(\widehat{Q}_T^{\mathrm{st}} - Q^\star\big) = \frac{\sqrt{\mathfrak{n}_T}}{f(Q^\star)T} \sum_{m=1}^{T} \xi_m + R_T,$$

where $\xi_m = E_m^{-1} \sum_{t=t_{m-1}+1}^{t_m} \sum_{k=1}^{K} p_k \left[g_{k,t}(Q^\star) - \{\tau_k - F_k(Q^\star)\} \right]$, *and* $R_T = \mathcal{O}_p\Big(T^{(\alpha-1)/2} + \sqrt{K} T^{(1+\beta)/2-\alpha} + \sqrt{K}\, T^{(\beta-1)/2} + T^{(\beta-1)/2} \log T\Big) + o_p(1)$. *Consequently, if* $K = o\big(T^{2\alpha-1-\beta}\big)$, *then* $R_T = o_p(1)$.

Theorem 1 gives a linear expansion for the spatial–temporal average. Its leading term is driven by the blockwise martingale score ξ_m, which has the same first-order form as the score driving the centralized LDP quantile recursion (Li et al., 2022; Cai et al., 2026). Therefore, Theorem 1 shows that, when $K = o(T^{2\alpha-1-\beta})$, the spatial–temporal average is first-order equivalent to the centralized recursion. By Theorem 1, we immediately obtain the following central limit theorem.

Corollary 1 (CLT of $\widehat{Q}_T^{\mathrm{st}}$)**.** *Under Assumptions 1–4, suppose* $K = o\left(T^{2\alpha-1-\beta}\right)$ *and*

$$\frac{\nu_\beta}{4f^2(Q^\star)} \frac{\sum_{k=1}^{K} p_k^2 \left[r_k^{-2} - \{2F_k(Q^\star) - 1\}^2 \right]}{\sum_{k=1}^{K} p_k^2} \to \sigma_\star^2 \in (0, \infty).$$

Then,

$$\sqrt{\mathfrak{n}_T}(\widehat{Q}_T^{\mathrm{st}} - Q^\star) \xrightarrow{d} N(0, \sigma_\star^2). \tag{2}$$

Corollary 1 quantifies two trade-offs. Larger r_k weakens privacy but reduces score variance. Larger β reduces communication frequency but increases the variance factor $\nu_\beta = (1 - \beta^2)^{-1}$ and tightens the node-growth condition $K = o(T^{2\alpha-1-\beta})$.

We next establish a functional CLT for the spatial–temporal average, which enables self-normalized inference based only on the already released iterates, avoiding plug-in estimation of the mixture density, and hence requiring no additional privacy budget. A fully online implementation is summarized in Algorithm 2.

Theorem 2 (Functional CLT of $\widehat{Q}_T^{\mathrm{st}}$)**.** *Let* $s_0 = 0$, $s_m = \left(\sum_{i=1}^{m} E_i^{-1}\right) / \left(\sum_{i=1}^{T} E_i^{-1}\right)$, *and* $h(s, T) = \max\{m \in \{0, 1, \ldots, T\} : s_m \leq s\}$, $s \in [0, 1]$. *Define the partial-sum process*

$$Q_T(s) = \frac{\sqrt{\mathfrak{n}_T}}{T} \sum_{m=1}^{h(s,T)} (\bar{q}_{t_m} - Q^\star), \qquad s \in [0, 1].$$

Under the assumptions of Corollary 1, $Q_T(\cdot) \xrightarrow{d} \sigma_\star B(\cdot)$ in $\ell^\infty[0,1]$, where B is standard Brownian motion. Define the observable self-normalizer

$$\widehat{V}_T = \sum_{m=1}^{T} (s_m - s_{m-1}) \left[\frac{\sqrt{\mathfrak{n}_T}}{T} \left\{ \sum_{i=1}^{m} \bar{q}_{t_i} - \frac{m}{T} \sum_{i=1}^{T} \bar{q}_{t_i} \right\} \right]^2 . \tag{3}$$

Then

$$\frac{Q_T(1)}{\sqrt{\widehat{V}_T}} \xrightarrow{d} \frac{B(1)}{\left\{ \int_0^1 \left[B(r) - r^{1/(1-\beta)} B(1) \right]^2 dr \right\}^{1/2}}. \tag{4}$$

Consequently, for any $\varphi \in (0,1)$, the interval $\left[\widehat{Q}_T^{\mathrm{st}} \pm C_{\varphi,\beta}^{\mathrm{SN}} \sqrt{\widehat{V}_T / \mathfrak{n}_T} \right]$ is asymptotically valid at level $1-\varphi$, where $C_{\varphi,\beta}^{\mathrm{SN}}$ is the $(1-\varphi)$-quantile of $\left| B(1) / \left\{ \int_0^1 \left[B(r) - r^{1/(1-\beta)} B(1) \right]^2 dr \right\}^{1/2} \right|$.

Algorithm 2 Online self-normalized inference

Input: synchronized averages $\{\bar{q}_{t_m}\}_{m=1}^T$, block lengths $\{E_m\}_{m=1}^T$, and the critical value $C_{\varphi,\beta}^{\mathrm{SN}}$.

Initialize: $\widehat{Q}_0^{\mathrm{st}} = 0$, $A_0 = 0$, $B_0 = 0$, $C_0 = 0$, and $H_0 = 0$.

1: **for** $m = 1, \ldots, T$ **do**

2: Update the running PR average: $\widehat{Q}_m^{\mathrm{st}} = (m-1)\widehat{Q}_{m-1}^{\mathrm{st}}/m + \bar{q}_{t_m}/m$.

3: Update the scalar accumulators $A_m = A_{m-1} + m^2(\widehat{Q}_m^{\mathrm{st}})^2/E_m$, $B_m = B_{m-1} + m^2\widehat{Q}_m^{\mathrm{st}}/E_m$, $C_m = C_{m-1} + m^2/E_m$, $H_m = H_{m-1} + 1/E_m$.

4: Compute the self-normalizer: $\widehat{V}_m = \mathfrak{n}_m \left(A_m - 2B_m\widehat{Q}_m^{\mathrm{st}} + C_m(\widehat{Q}_m^{\mathrm{st}})^2 \right) / (m^2 H_m)$ with $\mathfrak{n}_m = t_m / \sum_{k=1}^K p_k^2$.

5: **end for**

Return: $\left[\widehat{Q}_T^{\mathrm{st}} - C_{\varphi,\beta}^{\mathrm{SN}} \sqrt{\widehat{V}_T/\mathfrak{n}_T},\ \widehat{Q}_T^{\mathrm{st}} + C_{\varphi,\beta}^{\mathrm{SN}} \sqrt{\widehat{V}_T/\mathfrak{n}_T} \right]$.

Next, we establish the theoretical properties of the node-wise temporal average $\widehat{Q}_{k,T}^{\mathrm{time}}$.

Theorem 3 (Theoretical properties of $\widehat{Q}_{k,T}^{\mathrm{time}}$)**.** *Under Assumptions 1–4, for any fixed $1 \le k \le K$,*

$$\sqrt{\mathfrak{n}_T} \left(\widehat{Q}_{k,T}^{\mathrm{time}} - \widehat{Q}_T^{\mathrm{st}} \right) = \mathcal{O}_{a.s.} \left(\log(eK) \frac{\sqrt{\mathfrak{n}_T}}{T} \sum_{m=1}^{T} \gamma_{m-1} \right). \tag{5}$$

Moreover, with

$$Q_{k,T}(s) = \frac{\sqrt{\mathfrak{n}_T}}{T} \sum_{m=1}^{h(s,T)} (q_{k,t_m} - Q^\star),$$

we have the process transfer bound

$$\sup_{s \in [0,1]} |Q_{k,T}(s) - Q_T(s)| = \mathcal{O}_{a.s.} \left(\log(eK) \frac{\sqrt{\mathfrak{n}_T}}{T} \sum_{m=1}^{T} \gamma_{m-1} \right).$$

Consequently, under the assumptions of Theorem 2 and $K\{\log(eK)\}^2 = o\big(T^{2\alpha-1-\beta}\big)$. We have $\sqrt{\mathfrak{n}_T}\big(\widehat{Q}_{k,T}^{\mathrm{time}} - Q^\star\big) \xrightarrow{d} \mathcal{N}(0, \sigma_\star^2)$, $Q_{k,T}(\cdot) \xrightarrow{d} \sigma_\star B(\cdot)$. The node-wise self-normalized statistic obtained from (3) *by replacing every $\bar{q}_{t_m}$ with q_{k,t_m} has the same limiting pivotal law as* (4)*.*

By Theorem 3, each node-wise temporal average $\widehat{Q}_{k,T}^{\mathrm{time}}$ has the same limiting behavior as the spatial–temporal average $\widehat{Q}_T^{\mathrm{st}}$, provided that $K\{\log(eK)\}^2 = o\big(T^{2\alpha-1-\beta}\big)$. In contrast to $\widehat{Q}_T^{\mathrm{st}}$, the node-wise temporal average is available locally at each node and requires no additional network-wide aggregation. The spatial averages can be obtained through an additional decentralized reduction if needed.

Finally, we characterize the limiting behavior of the *terminal spatial average* $\widehat{Q}_T^{\mathrm{space}}$. It follows the classical Chung-type behavior of stochastic approximation (Chung, 1954): it converges more slowly than averaged estimators, but tolerates faster node growth.

Theorem 4 (Theoretical properties of $\widehat{Q}_T^{\mathrm{space}}$)**.** *Under Assumptions 1–4, if $K = o(T^{\alpha-\beta})$, then*

$$\frac{\sqrt{E_T}\big(\widehat{Q}_T^{\mathrm{space}} - Q^\star\big)}{\kappa\sqrt{\gamma_{T-1}}\{\overline{S}/(2f(Q^\star))\}^{1/2}} \xrightarrow{d} N(0,1).$$

In particular, when $E_m \equiv 1$ and $\eta_t = D(t+1)^{-\alpha}$,

$$(\kappa\{D\overline{S}/(2f(Q^\star))\}^{1/2})^{-1} T^{\alpha/2}\big(\widehat{Q}_T^{\mathrm{space}} - Q^\star\big) \xrightarrow{d} N(0,1).$$

Remark (Comparison of the three averaging schemes)**.** *The three estimators are compared as follows. Recall that T is the number of communication blocks, $E_m \sim c_E(m+1)^\beta$ is the number of fresh observations processed per node in block m, $t_T = \sum_{m=1}^{T} E_m$ is the per-node observation count, and $\mathfrak{n}_T = t_T / \sum_k p_k^2$ is the effective inferential scale.*

Estimator	*Requirement on K*	*Convergence rate*	*Self-normalized inference*	*Requires network-wide post-processing aggregation*
$\widehat{Q}_T^{\mathrm{st}}$	$K = o\big(T^{2\alpha-1-\beta}\big)$	$\mathfrak{n}_T^{-1/2}$	✓	✓
$\widehat{Q}_{k,T}^{\mathrm{time}}$	$K\{\log(eK)\}^2 = o\big(T^{2\alpha-1-\beta}\big)$	$\mathfrak{n}_T^{-1/2}$	✓	×
$\widehat{Q}_T^{\mathrm{space}}$	$K = o\big(T^{\alpha-\beta}\big)$	$\kappa\sqrt{\gamma_{T-1}/E_T}$	×	✓

4 Experiments

4.1 Simulation setup

We first evaluate the proposed decentralized locally private quantile procedure on synthetic data streams. Throughout the simulations, we use equal mixture weights, $p_k = 1/K$, and common local target levels, $\tau_k \equiv \tau$. We vary the number of nodes and the quantile level over $K \in \{10, 30, 50\}$ and $\tau \in \{0.3, 0.5, 0.8\}$. Each experiment is replicated $R = 500$ times using 110 Intel® Xeon® Platinum 8352V with 48 cores and 96GB memory.

Data generation. We consider one homogeneous design and two heterogeneous designs. In the homogeneous design, denoted by **Homo**, the observations are generated as $X_{k,t} \sim N(0,1)$ for all nodes k. In the heterogeneous-location design, denoted by **Hete-L**, we generate $X_{k,t} \sim N(\mu_k, 1)$, where the locations μ_k are equally spaced on $[-1, 1]$. In the heterogeneous-distribution design, denoted by **Hete-D**, 30% of the nodes follow $N(0,1)$, 30% follow $U(-1,1)$, and 40% follow $\mathrm{Cauchy}(0,1)$. We consider three privacy regimes. The first two privacy regimes use $r_k \equiv 0.25$ and $r_k \equiv 0.9$, corresponding to $\varepsilon = \log(1.25/0.75) \approx 0.511$ and $\varepsilon = \log(1.9/0.1) \approx 2.944$, respectively. The third regime is heterogeneous, with r_k equally spaced on $[0.25, 0.9]$ across nodes.

Communication network and schedules. Following Gu & Chen (2024), we use a cyclic local-neighbor graph with a lazy Metropolis–Hastings mixing matrix. Specifically, the mixing matrix is defined as

$$W_{k\ell} = \begin{cases} 1/2, & k = \ell, \\ 1/(2K_{\mathrm{neigh}}), & \ell \text{ is one of the } K_{\mathrm{neigh}} \text{ ring neighbors of } k, \\ 0, & \text{otherwise.} \end{cases}$$

We retain $K_{\mathrm{neigh}} = (3/5)K$. We compare two communication schedules: $E_m \equiv 1$ and $E_m \equiv 5$, denoted by C_1 and C_5, respectively. We use schedule-specific values of the effective step size γ_m, with the detailed choices reported in Appendix C.1.

Estimators and benchmarks. We compare three proposed decentralized estimators: the spatial–temporal average $\widehat{Q}_T^{\mathrm{st}}$, the node-wise temporal average $\widehat{Q}_{k,T}^{\mathrm{time}}$, and the terminal spatial average $\widehat{Q}_T^{\mathrm{space}}$. We also include two centralized server benchmarks based on the corresponding server-aggregated recursion: $\widehat{Q}_T^{\mathrm{cent},s}$ and $\widehat{Q}_T^{\mathrm{cent},t(\mathrm{st})}$. The first benchmark, $\widehat{Q}_T^{\mathrm{cent},s}$, is the terminal public iterate. The second benchmark, $\widehat{Q}_T^{\mathrm{cent},t(\mathrm{st})}$, is the temporal average of the common public iterate. Because all nodes share the same public iterate in the centralized recursion, this temporal average also equals the spatial–temporal average. Table 1 reports the core results for $K = 10$. Additional results exhibit qualitatively similar patterns; see Appendix C.1.

4.2 SIMULATION RESULTS

Several interesting patterns emerge from the results. First, the proposed decentralized estimators become more accurate as the effective inferential sample size increases: the MAE decreases with larger K or t_T. The MAE also decreases as the truthful-response probabilities r_k increase, since larger r_k reduces the randomized-response variance inflation. For fixed communication rounds, larger communication intervals yield larger t_T and hence smaller MAE. The ECPs of $\widehat{Q}_T^{\mathrm{st}}$ and $\widehat{Q}_{k,T}^{\mathrm{time}}$ are close to the nominal level in most regimes, supporting the self-normalized inference in Theorem 2. Second, $\widehat{Q}_{k,T}^{\mathrm{time}}$ and $\widehat{Q}_T^{\mathrm{st}}$ have nearly identical MAE and ECP in the main experiments, in agreement with Theorem 3. By contrast, $\widehat{Q}_T^{\mathrm{space}}$ has larger MAE than the temporal averages in most regimes, reflecting the slower stochastic-approximation scale in Theorem 4. Finally, the decentralized temporal averages are comparable to their centralized counterparts. Overall, these results demonstrate the practical usefulness of the proposed decentralized locally private quantile inference method.

Finally, we conduct additional experiments to examine (1) sparse and slowly mixing networks; (2) baseline comparisons under common fresh-record budgets; (3) directed and asynchronous communication; (4) transient edge dropout; and (5) target-preserving and moving-target distribution shifts; see Appendix C for details. In the sparse-network comparisons, slower mixing increases disagreement and paired decentralized–centralized gaps, while estimation errors remain close to the centralized benchmarks. In the heterogeneous-location design, local-only estimation and averaging local quantiles incur substantially larger errors for the mixture quantile.

4.3 REAL DATA APPLICATION

We next evaluate the proposed method on the Government Salary dataset Plečko et al. (2024). The dataset contains more than 200,000 individual records from the 2018 American Community Survey, including annual salary in USD and a range of demographic and socioeconomic covariates. We take annual salary as the response and use economic region to define the node groups. This application is well suited to the decentralized locally private setting: salary is a sensitive financial variable, and regional data holders may be unwilling to transmit raw records to a central server. We treat all 204,309 records as a fixed population and evaluate repeated-sampling performance over 500 replications. The main analysis uses $K = 7$ regional nodes with mixing parameter $\rho = 0.6002$. We use the C_5 schedule with warm-up as specified in Appendix C.7, $T = 2{,}567$ rounds, and $r_k \equiv r \in \{0.6, 0.9\}$. Table 2 reports the results at $\tau = 0.8$, for which the fixed-population quantile is $Q_{\mathrm{pop}} = 80{,}000$ USD.

The spatial–temporal estimator provides accurate salary-quantile estimates together with informative confidence intervals. At $r = 0.9$, its MAE is USD 291 and empirical coverage is 0.972. Under the stronger privacy level $r = 0.6$, MAE is USD 640 and coverage is 0.944. The corresponding CI lengths are USD 3,665 and USD 3,596, respectively. In both privacy regimes, temporal averaging substantially reduces error relative to the terminal spatial estimator. These results illustrate the practical value of the proposed procedure for estimating a common salary threshold from locally privatized regional data. Additional experiments use degree-2 cycle networks with $K = 14, 21,$ and 28, so that each node communicates with only two neighbors. At $r = 0.9$, the spatial–temporal estimator has MAE between USD 409 and USD 425 and empirical coverage between 0.936 and 0.952. Full results and network diagnostics are reported in Appendix C.7.

5 CONCLUDING REMARKS

This paper develops locally private quantile inference for heterogeneous data streams over peer-to-peer networks. Our analysis controls the interaction between network disagreement and discontinuous private scores, establishing conditions under which the spatial–temporal average has the same first-order limit as its server-coordinated counterpart. Under stronger balance conditions, the functional limit transfers to node-wise temporal averages, enabling inference from a single node's trajectory without additional network-wide aggregation. The resulting self-normalized confidence intervals avoid estimating the mixture density and privacy-inflated score variance. Numerical experiments and the salary application illustrate the method's practical utility and the effects of privacy and communication on estimation and inference.

Table 1: Empirical coverage probabilities at the 95% nominal level (mean absolute errors $\downarrow$ ($\times 10^{-3}$)) under the homogeneous and heterogeneous-distribution designs for C_1 and C_5. The number of nodes is fixed at $K = 10$. In Homo, data for every node are generated from $\mathcal{N}(0,1)$. In Hete D, nodes are assigned to $\mathcal{N}(0,1)$, $\mathcal{U}(-1,1)$, and $\mathcal{C}(0,1)$ groups. "hetero" indicates node-specific truthful response rates r_k equally spaced on $[0.25, 0.9]$. For the terminal spatial average, we report only the MAEs.

		C1			C5	Centralized (C1)	
τ	r_k	$\widehat{Q}_T^{\mathrm{space}}$	$\widehat{Q}_T^{\mathrm{st}}$	$\widehat{Q}_{k,T}^{\mathrm{time}}$	$\widehat{Q}_{k,T}^{\mathrm{time}}$	$\widehat{Q}_T^{\mathrm{cent},s}$	$\widehat{Q}_T^{\mathrm{cent},t(\mathrm{st})}$
				Homo | $T = 100{,}000$			
0.3	0.25	9.5	0.94(5.3)	0.94(5.3)	0.98(3.6)	9.5	0.94(5.3)
0.3	hetero	8	0.93(2.9)	0.93(2.9)	0.96(1.4)	8	0.93(2.9)
0.3	0.9	4.9	0.95(1.3)	0.95(1.3)	0.98(0.71)	4.9	0.95(1.3)
0.5	0.25	8.8	0.92(4)	0.92(4)	0.90(1.9)	8.8	0.91(4)
0.5	hetero	7.4	0.93(2.3)	0.93(2.3)	0.91(0.99)	7.4	0.93(2.3)
0.5	0.9	5.1	0.94(1.2)	0.94(1.2)	0.95(0.47)	5.1	0.94(1.2)
0.8	0.25	11	0.95(7.6)	0.95(7.6)	1.00(7.5)	11	0.96(7.6)
0.8	hetero	9.2	0.96(3.3)	0.96(3.3)	0.99(2.1)	9.2	0.96(3.3)
0.8	0.9	5	0.97(1.6)	0.97(1.6)	1.00(1.1)	5	0.97(1.6)
				Homo | $T = 200{,}000$			
0.3	0.25	7.8	0.90(3.4)	0.90(3.4)	0.98(2.1)	7.8	0.90(3.4)
0.3	hetero	6.1	0.94(1.9)	0.94(1.9)	0.94(0.9)	6.1	0.94(1.9)
0.3	0.9	3.8	0.94(0.89)	0.94(0.89)	0.96(0.43)	3.8	0.94(0.89)
0.5	0.25	7.4	0.92(2.8)	0.92(2.8)	0.92(1.3)	7.4	0.92(2.8)
0.5	hetero	6.2	0.93(1.6)	0.93(1.6)	0.91(0.73)	6.2	0.92(1.6)
0.5	0.9	3.9	0.95(0.8)	0.95(0.8)	0.94(0.35)	3.9	0.95(0.79)
0.8	0.25	9.2	0.95(4.7)	0.95(4.7)	1.00(3.8)	9.2	0.95(4.7)
0.8	hetero	7.1	0.94(2.3)	0.94(2.3)	0.98(1.2)	7.1	0.94(2.3)
0.8	0.9	3.8	0.95(1.1)	0.95(1.1)	0.98(0.61)	3.8	0.95(1.1)
				Hete D | $T = 100{,}000$			
0.5	0.25	9	0.91(3.9)	0.91(3.9)	0.90(1.8)	9	0.91(3.9)
0.5	hetero	7.6	0.94(2.1)	0.94(2.1)	0.93(1)	7.6	0.94(2.1)
0.5	0.9	4.9	0.94(1.1)	0.94(1.1)	0.94(0.47)	4.9	0.94(1.1)
				Hete D | $T = 200{,}000$			
0.5	0.25	7.4	0.91(2.9)	0.91(2.9)	0.89(1.3)	7.4	0.91(2.9)
0.5	hetero	5.9	0.96(1.5)	0.96(1.5)	0.93(0.75)	5.9	0.96(1.5)
0.5	0.9	3.9	0.94(0.82)	0.94(0.82)	0.93(0.36)	3.9	0.94(0.82)

Table 2: Government Salary: repeated-sampling performance for the seven-node regional network under C_5 at $\tau = 0.8$.

r	T	Spatial–temporal ECP (MAE)	CI length	Terminal MAE
0.6	2,567	0.944 (640)	3,596	2,308
0.9	2,567	0.972 (291)	3,665	3,908

Notes. Results use 500 replications. ECP is empirical coverage, with MAE in parentheses. MAE and CI length are in USD. Terminal denotes the terminal spatial estimator.

The present theory concerns stationary streams and fixed, synchronous networks with symmetric doubly stochastic mixing and a uniform spectral gap. Extending the analysis to growing slowly mixing graphs, time-varying or directed networks, asynchronous communication, and evolving quantiles is an important direction for future work. Simultaneous and conditional quantile inference additionally require joint privacy accounting and multivariate inferential tools. Adaptive privacy allocation across nodes offers another direction for balancing local privacy requirements and statistical efficiency.

AI USE STATEMENT

Generative AI tools were used to revise wording and improve the clarity and readability of the manuscript; to identify potentially relevant literature, references, and search directions, and to assist with research execution, including drafting, debugging, and refining code. For mathematical proofs, AI was used primarily to review and check the authors arguments for potential errors, inconsistencies, or gaps. AI-assisted manuscript revisions, literature-search suggestions, feedback on research ideas and proofs, and generated or modified code were reviewed and verified by the authors before incorporation into the final work. The authors independently assessed the accuracy and appropriateness of these outputs and retain full responsibility for the methodology, proofs, implementation, analysis, and conclusions presented in the paper.

REPRODUCIBILITY STATEMENT

The simulation code and an accompanying `README` file are submitted with the manuscript. We provide a dedicated reproduction script for each table, allowing the reported simulation results to be reproduced individually. The `README` provides instructions for running the code.

REFERENCES

Jayadev Acharya and Ziteng Sun. Communication complexity in locally private distribution estimation and heavy hitters. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 51–60. PMLR, 2019. URL https://proceedings.mlr.press/v97/acharya19c.html.

R. R. Bahadur. A note on quantiles in large samples. *The Annals of Mathematical Statistics*, 37(3):577–580, 1966.

Soham Bonnerjee, Sayar Karmakar, and Wei Biao Wu. Sharp gaussian approximations for decentralized federated learning. In *Advances in Neural Information Processing Systems*, 2025.

Leheng Cai, Qirui Hu, and Shuyuan Wu. Federated learning of quantile inference under local differential privacy. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=a5bFKVtTyF.

Kai Lai Chung. On a stochastic approximation method. *The Annals of Mathematical Statistics*, 25(3):463–483, 1954.

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

John C. Duchi, Alekh Agarwal, and Martin J. Wainwright. Dual averaging for distributed optimization: Convergence analysis and network scaling. *IEEE Transactions on Automatic Control*, 57(3):592–606, 2012.

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. Differential privacy. In *Automata, Languages and Programming*, pp. 1–12, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.

Cynthia Dwork and Jing Lei. Differential privacy and robust statistics. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pp. 371–380, 2009.

Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4):211–407, 2014.

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*, CCS '14, pp. 1054–1067. ACM, 2014.

Marco Gaboardi, Ryan Rogers, and Or Sheffet. Locally private mean estimation: Z-test and tight confidence intervals. In *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pp. 2545–2554. PMLR, 2019.

Jennifer Gillenwater, Matthew Joseph, and Alex Kulesza. Differentially private quantiles. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 3713–3722. PMLR, 2021.

Michael Greenwald and Sanjeev Khanna. Space-efficient online computation of quantile summaries. *ACM SIGMOD Record*, 30(2):58–66, 2001.

Jia Gu and Song Xi Chen. Statistical inference for decentralized federated learning. *The Annals of Statistics*, 52(6):2931–2955, 2024.

Peter Kairouz, Sewoong Oh, and Pramod Viswanath. Extremal mechanisms for local differential privacy. *Journal of Machine Learning Research*, 17(17):1–51, 2016.

Shiva Prasad Kasiviswanathan, Homin K. Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011.

Roger Koenker. *Quantile Regression*. Cambridge University Press, Cambridge, 2005.

Roger Koenker and Gilbert Bassett. Regression quantiles. *Econometrica*, 46(1):33–50, 1978.

Anastasia Koloskova, Nicolas Loizou, Sadra Boreiri, Martin Jaggi, and Sebastian Stich. A unified theory of decentralized sgd with changing topology and local updates. In *International conference on machine learning*, pp. 5381–5393. PMLR, 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.

Xiangru Lian, Ce Zhang, Huan Zhang, Cho-Jui Hsieh, Wei Zhang, and Ji Liu. Can decentralized algorithms outperform centralized algorithms? a case study for decentralized parallel stochastic gradient descent. In *Advances in Neural Information Processing Systems*, volume 30, pp. 5330–5340, 2017.

Yi Liu, Qirui Hu, Lei Ding, and Linglong Kong. Online local differential private quantile inference via self-normalization. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 21698–21714. PMLR, 2023.

Angelia Nedić and Alex Olshevsky. Distributed optimization over time-varying directed graphs. *IEEE Transactions on Automatic Control*, 60(3):601–615, 2015.

Angelia Nedić and Asuman Ozdaglar. Distributed subgradient methods for multi-agent optimization. *IEEE Transactions on Automatic Control*, 54(1):48–61, 2009.

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

Boris T. Polyak and Anatoli B. Juditsky. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30(4):838–855, 1992.

Herbert Robbins and Sutton Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, 22(3):400–407, 1951.

David Ruppert. Efficient estimators from a slowly convergent Robbins–Monro process. Technical Report 781, School of Operations Research and Industrial Engineering, Cornell University, 1988.

Robert J. Serfling. *Approximation Theorems of Mathematical Statistics*. John Wiley & Sons, 1980.

Jianwei Shi, Yue Wang, Zhongyi Zhu, and Heng Lian. Decentralized learning of quantile regression: A smoothing approach. *Journal of Computational and Graphical Statistics*, 34(3):1091–1101, 2025.

Yue Wang and Heng Lian. On linear convergence of ADMM for decentralized quantile regression. *IEEE Transactions on Signal Processing*, 71:3945–3955, 2023.

Yue Wang, Wenqi Lu, and Heng Lian. Distributed statistical estimation in quantile regression over a network. *Signal Processing*, 222:109512, 2024.

Stanley L. Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.

A TECHNICAL PROOFS

Throughout the appendix, all quantities may depend on T, but the dependence is suppressed unless confusion may arise. Define the network mean, the mean error, and the scalar mean field as

$$\theta_m = \bar{q}_{t_m}, \qquad e_m = \theta_m - Q^\star, \qquad h(q) = \sum_{k=1}^{K} p_k\{\tau_k - F_k(q)\} = \tau - F(q).$$

Thus $f(Q^\star) = f(Q^\star)$ and $h'(Q^\star) = -f(Q^\star)$ whenever the derivative is used. Let $\mathcal{F}_0$ denote the initial sigma-field. For within-block arguments, define the filtration

$$\mathcal{G}_{m,t} = \sigma(\mathcal{F}_{m-1}, \{X_{j,s}, U_{j,s}, V_{j,s} : j \le K,\ t_{m-1} < s \le t\}), \qquad t_{m-1} < t \le t_m,$$

and set $\mathcal{F}_m = \mathcal{G}_{m,t_m}$. Let $\sigma_k^2 = \mathrm{Var}(g_{k,t}(Q^\star))$, which does not depend on t. Let

$$\Delta_T = \frac{\sqrt{\mathfrak{n}_T}}{T}\sum_{m=1}^{T}\gamma_{m-1}, \qquad \Lambda_T = \frac{\sqrt{\mathfrak{n}_T}}{T}\sum_{m=1}^{T}\gamma_{m-1}^2, \qquad J = K^{-1}\mathbf{1}_K\mathbf{1}_K^\top.$$

Under Assumption 2, $\kappa^2 \asymp K^{-1}$.

A.1 TECHNICAL LEMMAS

Lemma 1. *Under Assumptions 1–2, for all $q, q' \in \Theta$, all $k \le K$, and all t,*

$$\mathbb{E}\,|g_{k,t}(q) - g_{k,t}(q')| \le \frac{f_+}{r_{\min}}\,|q - q'|\,, \tag{6}$$

$$\mathbb{E}\{g_{k,t}(q) - g_{k,t}(q')\}^2 \le \frac{f_+}{r_{\min}^2}\,|q - q'|\,. \tag{7}$$

Proof of Lemma 1. Fix k, t and abbreviate $g(q) = g_{k,t}(q)$. Couple $g(q)$ and $g(q')$ through the same record $X_{k,t}$ and the same randomized-response coins. Under this coupling, the two outputs can differ only if the truthful indicators differ: $\mathbf{1}\{X_{k,t} < q\} \ne \mathbf{1}\{X_{k,t} < q'\}$, which implies $X_{k,t} \in (q \wedge q', q \vee q']$.

Because $g(q)$ changes only by switching between its two support values,

$$|g(q) - g(q')| \le \frac{1}{r_k} \le \frac{1}{r_{\min}}.$$

Therefore

$$|g(q) - g(q')| \le \frac{1}{r_{\min}}\mathbf{1}\{X_{k,t} \in (q \wedge q', q \vee q']\}.$$

Taking expectations and using Assumption 1,

$$\mathbb{E}\,|g(q) - g(q')| \le \frac{1}{r_{\min}}\mathbb{P}(X_{k,t} \in (q \wedge q', q \vee q']) \le \frac{f_+}{r_{\min}}\,|q - q'|\,.$$

This proves (6). Squaring the pointwise bound gives

$$\{g(q) - g(q')\}^2 \le \frac{1}{r_{\min}^2}\mathbf{1}\{X_{k,t} \in (q \wedge q', q \vee q']\},$$

and taking expectations yields (7). □

Lemma 2 (Spectral row transfer and convolution)**.** *Under Assumption 3, for every integer $\ell \ge 0$, where e_k is the kth canonical basis vector,*

$$\max_{1\le k\le K}\left\|e_k^\top(W^\ell - J)\right\|_1 \le \min\{2, \sqrt{K}\,\rho^\ell\}. \tag{8}$$

Consequently, for $\gamma_m \asymp (m+1)^{-\alpha}$ with $\alpha \in (0,1)$,

$$\sum_{j=1}^{m}\gamma_{j-1}\min\{2, \sqrt{K}\,\rho^{m-j+1}\} \le C\log(eK)\gamma_{m-1}, \qquad 1 \le m \le T, \tag{9}$$

where C depends only on α, c_γ and the uniform spectral-gap bound $\bar{\rho}$.

Proof of Lemma 2. Because W is nonnegative and doubly stochastic, both $e_k^\top W^\ell$ and $K^{-1}\mathbf{1}_K^\top$ are probability vectors. Hence their ℓ_1 distance is at most 2. Symmetry gives the spectral identity $\left\|W^\ell - J\right\|_2 \le \rho^\ell$, and therefore

$$\left\|e_k^\top (W^\ell - J)\right\|_1 \le \sqrt{K}\, \left\|e_k^\top (W^\ell - J)\right\|_2 \le \sqrt{K}\, \rho^\ell.$$

Combining the two bounds proves (8).

Let $a_\ell = \min\{2, \sqrt{K}\rho^\ell\}$. Assumption 3 implies $\sum_{\ell \ge 0} a_\ell \le C_{\bar\rho} \log(eK)$. For the convolution, write $\ell = m - j + 1$. If $\ell \le m/2$, regular variation gives $\gamma_{m-\ell} \le C\gamma_m$ and hence the contribution is at most $C \log(eK)\gamma_m$. If $\ell > m/2$, then either $m = O(\log(eK))$, in which case $\sum_{u=0}^{m-1} \gamma_u \le C \log(eK)\gamma_m$, or m is larger than a constant multiple of $\log(eK)$, in which case the exponential factor in a_ℓ dominates the polynomial ratios in $\gamma_{m-\ell}/\gamma_m$. This proves (9). □

Lemma 3. *Under Assumptions 1–4, there exists $C < \infty$, independent of T, such that for every $m \le T$,*

$$\frac{1}{K} \sum_{k=1}^{K} (q_{k,t_m} - \bar q_{t_m})^2 \le C\gamma_{m-1}^2 \qquad \textit{almost surely}, \tag{10}$$

and the row-wise bound

$$\max_{k \le K} |q_{k,t_m} - \bar q_{t_m}| \le C \log(eK)\gamma_{m-1} \qquad \textit{almost surely} \tag{11}$$

holds. Consequently,

$$\frac{\sqrt{t_T}}{\kappa T} \sum_{m=1}^{T} \max_{k \le K} |q_{k,t_m} - \bar q_{t_m}| \le C \frac{\log(eK)}{\kappa} \frac{\sqrt{t_T}}{T} \sum_{m=1}^{T} \gamma_{m-1} \qquad \textit{almost surely}. \tag{12}$$

Proof of Lemma 3. Let $\boldsymbol{q}_m = (q_{1,t_m}, \dots, q_{K,t_m})$. Define the row-vector disagreement $\boldsymbol{d}_m = \boldsymbol{q}_m(I - J)$. Within block j, write the vector of scaled local scores as $\boldsymbol{u}_t = \big(Kp_1 g_{1,t}(q_{1,t-1}), \dots, Kp_K g_{K,t}(q_{K,t-1})\big)$ and $\boldsymbol{a}_j = \eta_{j-1} \sum_{t=t_{j-1}+1}^{t_j} \boldsymbol{u}_t (I - J)$. Since all nodes are initialized at the same value, $\boldsymbol{d}_0 = 0$. Moreover $JW = J$, so iteration over communication rounds gives

$$\boldsymbol{d}_m = \sum_{j=1}^{m} \boldsymbol{a}_j (W^{m-j+1} - J).$$

By Proposition 1 and Assumption 2, $|Kp_k g_{k,t}(q_{k,t-1})| \le b_2/r_{\min}$, so $\|\boldsymbol{a}_j\|_\infty \le C\eta_{j-1}E_j = C\gamma_{j-1}$ and $K^{-1/2}\|\boldsymbol{a}_j\|_2 \le C\gamma_{j-1}$. For the average squared consensus error, symmetry gives $\left\|W^\ell - J\right\|_2 \le \rho^\ell$. Hence

$$K^{-1/2}\|\boldsymbol{d}_m\|_2 \le \sum_{j=1}^{m} K^{-1/2}\|\boldsymbol{a}_j\|_2 \left\|W^{m-j+1} - J\right\|_2 \le C \sum_{j=1}^{m} \gamma_{j-1}\rho^{m-j+1} \le C\gamma_{m-1},$$

where the last step follows because $\rho \le \bar\rho < 1$ and $\gamma_m \asymp (m+1)^{-\alpha}$ imply $\sum_{j=1}^{m} \rho^{m-j+1}\gamma_{j-1} \le C_{\bar\rho}\gamma_{m-1}$. This is obtained by splitting the sum at $\lfloor m/2 \rfloor$. This proves (10). Moreover, for the row-wise bound, for every coordinate k,

$$\begin{aligned} |d_{m,k}| &\le \sum_{j=1}^{m} \|\boldsymbol{a}_j\|_\infty \left\|e_k^\top (W^{m-j+1} - J)\right\|_1 \\ &\le C \sum_{j=1}^{m} \gamma_{j-1} \min\{2, \sqrt{K}\rho^{m-j+1}\} \\ &\le C \log(eK)\gamma_{m-1}, \end{aligned}$$

where the last step is Lemma 2. This proves (11). Summing (11) over m gives (12). □

Lemma 4. *Under Assumptions 1–2, there exists $C < \infty$ such that for every block m,*

$$\max_{t_{m-1} < t \le t_m} \max_{k \le K} \left|q_{k,t} - q_{k,t_{m-1}}\right| \le C\gamma_{m-1} \qquad \textit{almost surely}.$$

Proof. During block m, each local step has size

$$|q_{k,t} - q_{k,t-1}| = \eta_{m-1} K p_k \, |g_{k,t}(q_{k,t-1})| \le \eta_{m-1} \frac{b_2}{r_{\min}}.$$

Summing over at most E_m steps gives the result. □

Recall that $\theta_m = \bar{q}_{t_m}$, and $e_m = \theta_m - Q^\star$. Because W is doubly stochastic, θ_m is preserved by the communication step. Hence, summing the local updates over nodes and over block m yields

$$\theta_m = \theta_{m-1} + \gamma_{m-1}\{h(\theta_{m-1}) + \widetilde{\xi}_m + \delta_m^{\mathrm{dec}}\}, \tag{13}$$

where the proof-level martingale difference is

$$\widetilde{\xi}_m = E_m^{-1} \sum_{t=t_{m-1}+1}^{t_m} \sum_{k=1}^{K} p_k \left(g_{k,t}(\theta_{m-1}) - \mathbb{E}[g_{k,t}(\theta_{m-1}) \mid \mathcal{F}_{m-1}] \right),$$

$$\delta_m^{\mathrm{dec}} = E_m^{-1} \sum_{t=t_{m-1}+1}^{t_m} \sum_{k=1}^{K} p_k \{ g_{k,t}(q_{k,t-1}) - g_{k,t}(\theta_{m-1}) \}.$$

The leading martingale in Theorem 1 keeps the original notation

$$\xi_m = E_m^{-1} \sum_{t=t_{m-1}+1}^{t_m} \sum_{k=1}^{K} p_k \left[g_{k,t}(Q^\star) - \{\tau_k - F_k(Q^\star)\} \right].$$

Lemma 5. *Under Assumptions 1–4:*

(i) $\{\widetilde{\xi}_m, \mathcal{F}_m\}$ *and* $\{\xi_m, \mathcal{F}_m\}$ *are martingale difference sequences, and* $\left|\widetilde{\xi}_m\right| + |\xi_m| \le 4 r_{\min}^{-1}$ *almost surely.*

(ii)

$$\mathrm{Var}(\widetilde{\xi}_m \mid \mathcal{F}_{m-1}) = \frac{1}{E_m} \sum_{k=1}^{K} p_k^2 \mathrm{Var}(g_{k,1}(\theta_{m-1}) \mid \theta_{m-1}), \tag{14}$$

and

$$\left| \mathrm{Var}(\widetilde{\xi}_m \mid \mathcal{F}_{m-1}) - \frac{\kappa^2 \overline{S}}{E_m} \right| \le \frac{C\kappa^2}{E_m} \, |e_{m-1}| \, . \tag{15}$$

Moreover,

$$\mathrm{Var}(\xi_m \mid \mathcal{F}_{m-1}) = \frac{\kappa^2 \overline{S}}{E_m}. \tag{16}$$

(iii)

$$\mathbb{E}[\left|\delta_m^{\mathrm{dec}}\right| \mid \mathcal{F}_{m-1}] + \mathbb{E}[(\delta_m^{\mathrm{dec}})^2 \mid \mathcal{F}_{m-1}] \le C\gamma_{m-1} \qquad \textit{almost surely}. \tag{17}$$

Proof of Lemma 5. Part (i) follows immediately from the definition of ξ_m and the bound $|g_{k,t}(\cdot)| \le r_{\min}^{-1}$.

For part (ii), conditional on $\mathcal{F}_{m-1}$, the variables

$$\{g_{k,t}(\theta_{m-1}) - \mathbb{E}[g_{k,t}(\theta_{m-1}) \mid \mathcal{F}_{m-1}]\}_{k,t}$$

are independent across k and t. Therefore

$$\begin{aligned} \mathrm{Var}(\widetilde{\xi}_m \mid \mathcal{F}_{m-1}) &= \frac{1}{E_m^2} \sum_{t=t_{m-1}+1}^{t_m} \sum_{k=1}^{K} p_k^2 \mathrm{Var}(g_{k,t}(\theta_{m-1}) \mid \theta_{m-1}) \\ &= \frac{1}{E_m} \sum_{k=1}^{K} p_k^2 \mathrm{Var}(g_{k,1}(\theta_{m-1}) \mid \theta_{m-1}), \end{aligned}$$

which proves (14). Using Proposition 1,

$$\mathrm{Var}(g_{k,1}(q) \mid q) - \sigma_k^2 = \frac{1}{4}\Big(\{2F_k(Q^\star) - 1\}^2 - \{2F_k(q) - 1\}^2\Big).$$

Since F_k is Lipschitz with constant f_+ on Θ,

$$\big|\mathrm{Var}(g_{k,1}(q) \mid q) - \sigma_k^2\big| \le C\,|q - Q^\star|\,.$$

Multiplying by p_k^2 and summing over k yields

$$\left|\mathrm{Var}(\widetilde{\xi}_m \mid \mathcal{F}_{m-1}) \; - \frac{1}{E_m}\sum_{k=1}^{K} p_k^2\sigma_k^2\right| \le \frac{C\kappa^2}{E_m}\,|\theta_{m-1} - Q^\star|\,,$$

which is (15).

For part (iii), write

$$\delta_m^{\mathrm{dec}} = \frac{1}{E_m}\sum_{t=t_{m-1}+1}^{t_m}\sum_{k=1}^{K} p_k\Delta_{k,t}, \qquad \Delta_{k,t} = g_{k,t}(q_{k,t-1}) - g_{k,t}(\theta_{m-1}).$$

Fix $t_{m-1} < t \le t_m$. Conditional on $\mathcal{G}_{m,t-1}$, both $q_{k,t-1}$ and θ_{m-1} are measurable, while the current sample $X_{k,t}$ and the current privacy coins are independent of $\mathcal{G}_{m,t-1}$. Hence the conditional versions of (6)–(7) yield

$$\mathbb{E}\big[|\Delta_{k,t}| \mid \mathcal{G}_{m,t-1}\big] \le \frac{f_+}{r_{\min}}\,|q_{k,t-1} - \theta_{m-1}|\,, \tag{18}$$

$$\mathbb{E}\big[\Delta_{k,t}^2 \mid \mathcal{G}_{m,t-1}\big] \le \frac{f_+}{r_{\min}^2}\,|q_{k,t-1} - \theta_{m-1}|\,. \tag{19}$$

Now decompose, for each $t_{m-1} < t \le t_m$,

$$|q_{k,t-1} - \theta_{m-1}| \le \big|q_{k,t_{m-1}} - \theta_{m-1}\big| + \big|q_{k,t-1} - q_{k,t_{m-1}}\big|\,.$$

The first term is controlled in weighted average by the spectral mean-square consensus bound:

$$\sum_{k=1}^{K} p_k\,\big|q_{k,t_{m-1}} - \theta_{m-1}\big| \le \left(\sum_{k=1}^{K} p_k(q_{k,t_{m-1}} - \theta_{m-1})^2\right)^{1/2} \le C\gamma_{m-1},$$

where we used $p_k \le b_2/K$ and (10). The second term is bounded by $C\gamma_{m-1}$ uniformly by Lemma 4. Hence

$$\sum_{k=1}^{K} p_k\mathbb{E}\big[|q_{k,t-1} - \theta_{m-1}| \mid \mathcal{F}_{m-1}\big] \le C\gamma_{m-1} \qquad \text{almost surely.} \tag{20}$$

Combining (18) with (20) and averaging over t gives

$$\mathbb{E}\big[\big|\delta_m^{\mathrm{dec}}\big| \mid \mathcal{F}_{m-1}\big] \le C\gamma_{m-1} \qquad \text{almost surely.}$$

For the second moment, Jensen's inequality, (19), and (20) imply

$$\begin{aligned}
\mathbb{E}\big[(\delta_m^{\mathrm{dec}})^2 \mid \mathcal{F}_{m-1}\big] &\le \frac{1}{E_m}\sum_{t=t_{m-1}+1}^{t_m}\mathbb{E}\left[\left(\sum_{k=1}^{K} p_k\Delta_{k,t}\right)^2 \;\Big|\; \mathcal{F}_{m-1}\right] \\
&\le \frac{1}{E_m}\sum_{t=t_{m-1}+1}^{t_m}\sum_{k=1}^{K} p_k\mathbb{E}\big[\Delta_{k,t}^2 \mid \mathcal{F}_{m-1}\big] \\
&\le \frac{C}{E_m}\sum_{t=t_{m-1}+1}^{t_m}\sum_{k=1}^{K} p_k\mathbb{E}\big[|q_{k,t-1} - \theta_{m-1}| \mid \mathcal{F}_{m-1}\big] \le C\gamma_{m-1} \qquad \text{a.s.}
\end{aligned}$$

This proves (17). □

Lemma 6. *Let $\{u_m\}$ be nonnegative and suppose that for all large m,*

$$u_m \le (1 - c\gamma_{m-1})u_{m-1} + a\frac{\gamma_{m-1}^2}{E_m} + b\gamma_{m-1}^3,$$

where $a, b, c > 0$, $\gamma_m \sim c_\gamma(m+1)^{-\alpha}$, and $E_m \sim c_E m^\beta$ with $\alpha \in (1/2, 1)$ and $\beta \in [0, 1)$. Then, for all large m,

$$u_m \le C\left(a\frac{\gamma_{m-1}}{E_m} + b\gamma_{m-1}^2\right).$$

Proof. Let $v_m = a\gamma_{m-1}/E_m + b\gamma_{m-1}^2$. Regular variation gives

$$\frac{\gamma_{m-2}/E_{m-1}}{\gamma_{m-1}/E_m} = 1 + O(m^{-1}), \qquad \frac{\gamma_{m-2}^2}{\gamma_{m-1}^2} = 1 + O(m^{-1}).$$

Since $\alpha < 1$, $m^{-1} = o(\gamma_{m-1})$. Hence there is an index M such that, for all $m \ge M$,

$$\frac{\gamma_{m-2}}{E_{m-1}} \le \left(1 + \frac{c}{4}\gamma_{m-1}\right)\frac{\gamma_{m-1}}{E_m}, \qquad \gamma_{m-2}^2 \le \left(1 + \frac{c}{4}\gamma_{m-1}\right)\gamma_{m-1}^2.$$

Therefore

$$v_{m-1} \le \left(1 + \frac{c}{4}\gamma_{m-1}\right)v_m.$$

Choose C so that $u_M \le Cv_M$ and $C \ge 4/c$. Suppose $u_{m-1} \le Cv_{m-1}$. Then

$$\begin{aligned} u_m &\le (1 - c\gamma_{m-1})Cv_{m-1} + a\frac{\gamma_{m-1}^2}{E_m} + b\gamma_{m-1}^3 \\ &\le C(1 - c\gamma_{m-1})\left(1 + \frac{c}{4}\gamma_{m-1}\right)v_m + a\frac{\gamma_{m-1}^2}{E_m} + b\gamma_{m-1}^3. \end{aligned}$$

For all large m,

$$(1 - c\gamma_{m-1})\left(1 + \frac{c}{4}\gamma_{m-1}\right) \le 1 - \frac{c}{2}\gamma_{m-1}.$$

Since

$$\gamma_{m-1}v_m = a\frac{\gamma_{m-1}^2}{E_m} + b\gamma_{m-1}^3,$$

we obtain

$$u_m \le Cv_m - \frac{cC}{2}\left(a\frac{\gamma_{m-1}^2}{E_m} + b\gamma_{m-1}^3\right) + \left(a\frac{\gamma_{m-1}^2}{E_m} + b\gamma_{m-1}^3\right) \le Cv_m,$$

where the last step uses $C \ge 4/c$. This closes the induction. □

Lemma 7. *Under Assumptions 1–4, there exists $C < \infty$ such that for all large m,*

$$h(\theta_{m-1}) = -f(Q^\star)e_{m-1} + r_m^{\mathrm{lin}}, \qquad \left|r_m^{\mathrm{lin}}\right| \le Ce_{m-1}^2, \tag{21}$$

$$\mathbb{E}e_m^2 \le C\left(\kappa^2\frac{\gamma_{m-1}}{E_m} + \gamma_{m-1}^2\right). \tag{22}$$

Therefore

$$e_m = \mathcal{O}_p\left(\kappa\sqrt{\frac{\gamma_{m-1}}{E_m}} + \gamma_{m-1}\right). \tag{23}$$

Proof of Lemma 7. By Assumption 1, the mixture density is Lipschitz on Θ. Since $h(q) = \tau - F(q)$, Taylor's formula for the distribution function gives

$$h(\theta_{m-1}) = -f(Q^\star)e_{m-1} + r_m^{\mathrm{lin}}, \qquad |r_m^{\mathrm{lin}}| \le Ce_{m-1}^2.$$

Also, because $f \ge f_-$ on Θ,

$$e_{m-1}h(\theta_{m-1}) = -(\theta_{m-1} - Q^\star)\{F(\theta_{m-1}) - F(Q^\star)\} \le -f_-e_{m-1}^2.$$

Using the exact recursion (13),

$$e_m = e_{m-1} + \gamma_{m-1}\{h(\theta_{m-1}) + \widetilde{\xi}_m + \delta_m^{\mathrm{dec}}\}.$$

After squaring and conditioning on $\mathcal{F}_{m-1}$, the cross term involving $\widetilde{\xi}_m$ and the deterministic factor e_{m-1} vanishes:

$$\mathbb{E}(e_m^2 \mid \mathcal{F}_{m-1}) = e_{m-1}^2 + 2\gamma_{m-1}e_{m-1}h(\theta_{m-1}) + 2\gamma_{m-1}e_{m-1}\mathbb{E}(\delta_m^{\mathrm{dec}} \mid \mathcal{F}_{m-1}) + \gamma_{m-1}^2\mathbb{E}\left[\{h(\theta_{m-1}) + \widetilde{\xi}_m + \delta_m^{\mathrm{dec}}\}^2 \mid \mathcal{F}_{m-1}\right].$$

Young's inequality and Lemma 5 imply

$$2\gamma_{m-1}|e_{m-1}||\mathbb{E}(\delta_m^{\mathrm{dec}} \mid \mathcal{F}_{m-1})| \le \frac{f_-}{2}\gamma_{m-1}e_{m-1}^2 + C\gamma_{m-1}^3.$$

Furthermore, $|h(\theta_{m-1})| \le C|e_{m-1}|$, so for all large m,

$$\gamma_{m-1}^2 h(\theta_{m-1})^2 \le C\gamma_{m-1}^2 e_{m-1}^2 \le \frac{f_-}{4}\gamma_{m-1}e_{m-1}^2.$$

We first record the uniform conditional variance bound

$$\mathrm{Var}(\widetilde{\xi}_m \mid \mathcal{F}_{m-1}) \le \frac{C\kappa^2}{E_m}. \tag{24}$$

Indeed, by (14), Proposition 1, and Assumption 2,

$$\mathrm{Var}(\widetilde{\xi}_m \mid \mathcal{F}_{m-1}) = \frac{1}{E_m}\sum_{k=1}^K p_k^2 \mathrm{Var}(g_{k,1}(\theta_{m-1}) \mid \theta_{m-1}) \le \frac{C}{E_m}\sum_{k=1}^K p_k^2 = \frac{C\kappa^2}{E_m}.$$

Using $(a+b+c)^2 \le 3(a^2+b^2+c^2)$, (24), and (17),

$$\begin{aligned}\mathbb{E}(e_m^2 \mid \mathcal{F}_{m-1}) &\le (1 - c\gamma_{m-1})e_{m-1}^2 + C\gamma_{m-1}^3 + C\gamma_{m-1}^2\mathrm{Var}(\widetilde{\xi}_m \mid \mathcal{F}_{m-1}) \\ &\le (1 - c\gamma_{m-1})e_{m-1}^2 + C\gamma_{m-1}^3 + C\kappa^2\frac{\gamma_{m-1}^2}{E_m}.\end{aligned}$$

Taking expectations and applying Lemma 6 with $a = C\kappa^2$ and $b = C$ gives (22). Finally, (23) follows from Markov's inequality. □

Lemma 8. *Under Assumptions 1–4,*

$$\mathbb{E}\left[\sup_{0\le H\le T}\left|\frac{\sqrt{\mathfrak{n}_T}}{T}\sum_{m=1}^H(\widetilde{\xi}_m - \xi_m)\right|^2\right] \le C\frac{t_T}{T^2}\sum_{m=1}^T E_m^{-1}\mathbb{E}|e_{m-1}| = o(1). \tag{25}$$

Consequently,

$$\sup_{0\le H\le T}\left|\frac{\sqrt{\mathfrak{n}_T}}{T}\sum_{m=1}^H(\widetilde{\xi}_m - \xi_m)\right| = o_p(1).$$

Proof. Conditional on $\mathcal{F}_{m-1}$, both $\widetilde{\xi}_m$ and ξ_m are centered; therefore $\widetilde{\xi}_m - \xi_m$ is also centered. The conditional version of Lemma 1 gives

$$\begin{aligned}\mathrm{Var}(\widetilde{\xi}_m - \xi_m \mid \mathcal{F}_{m-1}) &\le \frac{C}{E_m^2}\sum_{t=t_{m-1}+1}^{t_m}\sum_{k=1}^K p_k^2\mathbb{E}\{g_{k,t}(\theta_{m-1}) - g_{k,t}(Q^\star)\}^2 \\ &\le \frac{C}{E_m}\sum_{k=1}^K p_k^2|\theta_{m-1} - Q^\star| = \frac{C\kappa^2}{E_m}|e_{m-1}|.\end{aligned}$$

By martingale orthogonality and Doob's L^2 inequality,

$$\mathbb{E}\left[\sup_{0\le H\le T}\left|\frac{\sqrt{\mathfrak{n}_T}}{T}\sum_{m=1}^{H}(\widetilde{\xi}_m-\xi_m)\right|^2\right]$$
$$\le C\frac{\mathfrak{n}_T}{T^2}\sum_{m=1}^{T}\mathbb{E}\{\mathrm{Var}(\widetilde{\xi}_m-\xi_m\mid\mathcal{F}_{m-1})\}\le C\frac{t_T}{T^2}\sum_{m=1}^{T}E_m^{-1}\mathbb{E}|e_{m-1}|.$$

Lemma 7 implies

$$\mathbb{E}|e_{m-1}|\le C\left\{\kappa\sqrt{\frac{\gamma_{m-1}}{E_m}}+\gamma_{m-1}\right\}.$$

Since $t_T\asymp T^{1+\beta}$, $E_m\asymp m^{\beta}$, and $\gamma_m\asymp m^{-\alpha}$, the first contribution is bounded by

$$C\kappa T^{\beta-1}\sum_{m=1}^{T}m^{-\beta-(\alpha+\beta)/2}=\begin{cases}O\{\kappa T^{-(\alpha+\beta)/2}\}, & \alpha/2+3\beta/2<1,\\ O\{\kappa T^{\beta-1}\log T\}, & \alpha/2+3\beta/2=1,\\ O\{\kappa T^{\beta-1}\}, & \alpha/2+3\beta/2>1,\end{cases}$$

which is $o(1)$. The second contribution is bounded by

$$CT^{\beta-1}\sum_{m=1}^{T}m^{-\alpha-\beta}=\begin{cases}O(T^{-\alpha}), & \alpha+\beta<1,\\ O(T^{\beta-1}\log T), & \alpha+\beta=1,\\ O(T^{\beta-1}), & \alpha+\beta>1,\end{cases}$$

which is also $o(1)$, because $\alpha>0$ and $\beta<1$. This proves (25).

□

A.2 Proof of Proposition 1

Fix k,t,q and abbreviate $B=\mathbf{1}\{X_{k,t}<q\}\in\{0,1\}$, $S=S_{k,t}(q)$, and $r=r_k$. When $0<r<1$, conditional on $B=b\in\{0,1\}$,

$$\mathbb{P}(S=1\mid B=b)=rb+\frac{1-r}{2},\qquad \mathbb{P}(S=0\mid B=b)=r(1-b)+\frac{1-r}{2}.$$

Hence for $s\in\{0,1\}$ and $b\neq b'$,

$$\frac{\mathbb{P}(S=s\mid B=b)}{\mathbb{P}(S=s\mid B=b')}\le\frac{(1+r)/2}{(1-r)/2}=\frac{1+r}{1-r},$$

which proves ε_k-LDP with $\varepsilon_k=\log\frac{1+r_k}{1-r_k}$. When $r_k=1$, we have $S=B$, so the mechanism is the truthful-release boundary case and is non-private.

For unbiasedness, note that

$$\mathbb{E}(S\mid X_{k,t},q)=r_k\mathbf{1}\{X_{k,t}<q\}+\frac{1-r_k}{2}.$$

Substituting this into $g_{k,t}(q)$ gives $\mathbb{E}[g_{k,t}(q)\mid X_{k,t},q]=\tau_k-\mathbf{1}\{X_{k,t}<q\}$. Taking expectations over $X_{k,t}$ yields $\mathbb{E}[g_{k,t}(q)\mid q]=\tau_k-F_k(q)$. Next, conditional on q, S is Bernoulli with success probability

$$p_k(q)=r_kF_k(q)+\frac{1-r_k}{2}=\frac{1}{2}+r_k\Big(F_k(q)-\frac{1}{2}\Big).$$

Therefore

$$\mathrm{Var}(g_{k,t}(q)\mid q)=\frac{\mathrm{Var}(S\mid q)}{r_k^2}=\frac{p_k(q)\{1-p_k(q)\}}{r_k^2}=\frac{r_k^{-2}-\{2F_k(q)-1\}^2}{4}.$$

Finally, $g_{k,t}(q)$ takes only two values,

$$\tau_k+\frac{1-r_k}{2r_k}\qquad\text{and}\qquad\tau_k-\frac{1+r_k}{2r_k}.$$

Since $\tau_k\in(0,1)$ and $r_k\ge r_{\min}$,

$$|g_{k,t}(q)|\le\frac{1+r_k}{2r_k}\le\frac{1}{r_k}\le\frac{1}{r_{\min}}.$$

This completes the proof.

A.3 Proof of Theorem 1 and Corollary 1

Proof of Theorem 1. From (13) and (21),

$$e_m - e_{m-1} = \gamma_{m-1}\{-f(Q^\star)e_{m-1} + \widetilde{\xi}_m + \delta_m^{\mathrm{dec}} + r_m^{\mathrm{lin}}\}.$$

Rearranging gives the exact identity

$$e_{m-1} = f(Q^\star)^{-1}\widetilde{\xi}_m + f(Q^\star)^{-1}\delta_m^{\mathrm{dec}} + f(Q^\star)^{-1}r_m^{\mathrm{lin}} - f(Q^\star)^{-1}\frac{e_m - e_{m-1}}{\gamma_{m-1}}. \tag{26}$$

Summing (26) over $m = 1, \dots, T$ and using $\widehat{Q}_T^{\mathrm{st}} - Q^\star = T^{-1}\sum_{m=1}^T e_m$,

$$\sqrt{\mathfrak{n}_T}(\widehat{Q}_T^{\mathrm{st}} - Q^\star) = \frac{\sqrt{\mathfrak{n}_T}}{f(Q^\star)T}\sum_{m=1}^T \widetilde{\xi}_m + R_{\delta,T} + R_{\mathrm{lin},T} + R_{\mathrm{Abel},T} + R_{\mathrm{end},T}, \tag{27}$$

where

$$R_{\delta,T} = \frac{\sqrt{\mathfrak{n}_T}}{f(Q^\star)T}\sum_{m=1}^T \delta_m^{\mathrm{dec}}, \quad R_{\mathrm{lin},T} = \frac{\sqrt{\mathfrak{n}_T}}{f(Q^\star)T}\sum_{m=1}^T r_m^{\mathrm{lin}},$$
$$R_{\mathrm{Abel},T} = -\frac{\sqrt{\mathfrak{n}_T}}{f(Q^\star)T}\sum_{m=1}^T \frac{e_m - e_{m-1}}{\gamma_{m-1}}, \quad R_{\mathrm{end},T} = \frac{\sqrt{\mathfrak{n}_T}}{T}(e_T - e_0).$$

The endpoint term satisfies

$$R_{\mathrm{end},T} = O_p\left(\frac{\sqrt{\mathfrak{n}_T}}{T}\right) = O_p\left(\frac{\sqrt{t_T}}{\kappa T}\right),$$

because the iterates remain in the compact set Θ.

Step 1: decentralized remainder. By Lemma 5,

$$\mathbb{E}|R_{\delta,T}| \le C\frac{\sqrt{\mathfrak{n}_T}}{T}\sum_{m=1}^T \gamma_{m-1} = C\frac{\sqrt{t_T}}{\kappa T}\sum_{m=1}^T \gamma_{m-1} = C\Delta_T.$$

Thus

$$R_{\delta,T} = O_p(\Delta_T). \tag{28}$$

Step 2: nonlinear linearization remainder. By (21) and Lemma 7,

$$\mathbb{E}|r_m^{\mathrm{lin}}| \le C\mathbb{E}e_{m-1}^2 \le C\left(\kappa^2\frac{\gamma_{m-1}}{E_m} + \gamma_{m-1}^2\right).$$

Hence

$$\begin{aligned}\mathbb{E}|R_{\mathrm{lin},T}| &\le C\frac{\sqrt{t_T}}{\kappa T}\sum_{m=1}^T\left(\kappa^2\frac{\gamma_{m-1}}{E_m} + \gamma_{m-1}^2\right) = C\kappa\frac{\sqrt{t_T}}{T}\sum_{m=1}^T\frac{\gamma_{m-1}}{E_m} + C\frac{\sqrt{t_T}}{\kappa T}\sum_{m=1}^T\gamma_{m-1}^2 \\ &\le C\Delta_T + C\Lambda_T. \end{aligned} \tag{29}$$

The last inequality uses $E_m \ge 1$, $\kappa \le 1$, and the definitions of Δ_T and Λ_T.

Step 3: Abel term. Let

$$S_T = \sum_{m=1}^T \frac{e_m - e_{m-1}}{\gamma_{m-1}}.$$

The Abel identity gives

$$S_T = \frac{e_T}{\gamma_{T-1}} - \frac{e_0}{\gamma_0} + \sum_{m=1}^{T-1}\left(\frac{1}{\gamma_{m-1}} - \frac{1}{\gamma_m}\right)e_m.$$

The initial contribution is

$$O_p\left(\frac{\sqrt{\mathfrak{n}_T}}{T}\right) = O_p\left(\frac{\sqrt{t_T}}{\kappa T}\right).$$

For the terminal contribution, Cauchy's inequality and Lemma 7 give

$$\mathbb{E}\left[\frac{\sqrt{\mathfrak{n}_T}}{T}\frac{|e_T|}{\gamma_{T-1}}\right] \le C\frac{\sqrt{t_T}}{\kappa T\gamma_T}\left(\kappa\sqrt{\frac{\gamma_T}{E_T}} + \gamma_T\right) \le C\left(\frac{\sqrt{t_T}}{T\sqrt{\gamma_T E_T}} + \frac{\sqrt{t_T}}{\kappa T}\right).$$

For the variation term, regular variation implies $\left|\gamma_{m-1}^{-1} - \gamma_m^{-1}\right| \le Cm^{\alpha-1}$. Using Lemma 7 again,

$$\begin{aligned}
&\mathbb{E}\left[\frac{\sqrt{\mathfrak{n}_T}}{T}\sum_{m=1}^{T-1}\left|\gamma_{m-1}^{-1} - \gamma_m^{-1}\right||e_m|\right] \\
&\quad\le C\frac{\sqrt{t_T}}{T}\sum_{m=1}^{T-1} m^{\alpha-1}\sqrt{\frac{\gamma_m}{E_m}} + C\frac{\sqrt{t_T}}{\kappa T}\sum_{m=1}^{T-1} m^{\alpha-1}\gamma_m \\
&\quad\le C\{T^{(\alpha-1)/2} + T^{(\beta-1)/2}\log T\} + C\frac{\sqrt{t_T}}{\kappa T}\log T.
\end{aligned}$$

Because $\sum_{m=1}^{T}\gamma_{m-1} \asymp T^{1-\alpha}$, the last term is dominated by Δ_T, since $T^{\alpha-1}\log T \to 0$. Therefore

$$\frac{\sqrt{\mathfrak{n}_T}}{T}S_T = O_p\left(\Delta_T + \Lambda_T + \frac{\sqrt{t_T}}{\kappa T} + \frac{\sqrt{t_T}}{T\sqrt{\gamma_T E_T}} + T^{(\alpha-1)/2} + T^{(\beta-1)/2}\log T\right). \tag{30}$$

Step 4: collecting rates. Combining (27)–(30) gives the expansion with $\widetilde{\xi}_m$. Lemma 8 then replaces $\widetilde{\xi}_m$ by ξ_m at $o_p(1)$ cost. By Assumption 2 and $E_m \asymp m^\beta$,

$$\Delta_T \asymp \sqrt{K}\,T^{(1+\beta)/2-\alpha}, \qquad \Lambda_T \asymp \sqrt{K}\,T^{(\beta-1)/2},$$

and

$$\frac{\sqrt{t_T}}{T\sqrt{\gamma_T E_T}} \asymp T^{(\alpha-1)/2}, \qquad \frac{\sqrt{t_T}}{\kappa T} \asymp \sqrt{K}\,T^{(\beta-1)/2}.$$

Thus the final remainder is

$$R_T = O_p\left(\sqrt{K}\,T^{(1+\beta)/2-\alpha} + \sqrt{K}\,T^{(\beta-1)/2} + T^{(\alpha-1)/2} + T^{(\beta-1)/2}\log T\right),$$

which proves the final statement. If $K = o(T^{2\alpha-1-\beta})$, then all terms above are $o_p(1)$, because $2\alpha - 1 - \beta < 1 - \beta$ when $\alpha < 1$. □

We now turn to the asymptotic normality of the leading martingale term.

Proof of Corollary 1. By Theorem 1 and the balance condition $K = o(T^{2\alpha-1-\beta})$,

$$\sqrt{\mathfrak{n}_T}(\widehat{Q}_T^{\mathrm{st}} - Q^\star) = M_T + o_p(1), \qquad M_T = \frac{\sqrt{\mathfrak{n}_T}}{f(Q^\star)T}\sum_{m=1}^{T}\xi_m.$$

Define

$$X_{m,T} = \frac{\sqrt{\mathfrak{n}_T}}{f(Q^\star)T}\xi_m, \qquad M_T = \sum_{m=1}^{T} X_{m,T}.$$

By (16), the predictable quadratic variation is deterministic:

$$\sum_{m=1}^{T}\mathbb{E}(X_{m,T}^2 \mid \mathcal{F}_{m-1}) = \frac{t_T}{f(Q^\star)^2\kappa^2T^2}\sum_{m=1}^{T}\frac{\kappa^2\overline{S}}{E_m} = \frac{\nu_T\overline{S}}{f(Q^\star)^2}.$$

For Lindeberg's condition, $|\xi_m| \le C$ almost surely, and therefore

$$\max_{m\le T}|X_{m,T}| \le C\frac{\sqrt{t_T}}{\kappa T} \asymp C\sqrt{K}\,T^{(\beta-1)/2}.$$

The condition $K = o(T^{2\alpha-1-\beta})$ implies $K = o(T^{1-\beta})$, because $\alpha < 1$. Hence $\max_m |X_{m,T}| \to 0$, and the conditional Lindeberg condition is automatic. The martingale CLT gives the self-normalized form

$$\frac{M_T}{(\nu_T \overline{S}/f(Q^\star)^2)^{1/2}} \xrightarrow{d} N(0,1).$$

Slutsky's lemma yields

$$\frac{\sqrt{\mathfrak{n}_T}(\widehat{Q}_T^{\mathrm{st}} - Q^\star)}{(\nu_T \overline{S}/f(Q^\star)^2)^{1/2}} \xrightarrow{d} N(0,1).$$

If $\nu_T \overline{S}/f(Q^\star)^2 \to \sigma_\star^2 \in (0,\infty)$, then the unstandardized CLT (2) follows. □

A.4 Proof of Theorem 2

Lemma 9 (Uniform Abel transfer)**.** *Under the assumptions of Corollary 1,*

$$\sup_{0\le H\le T} \left| \frac{\sqrt{\mathfrak{n}_T}}{T} \sum_{m=1}^{H} \frac{e_m - e_{m-1}}{\gamma_{m-1}} \right| = o_p(1).$$

Proof. For every $H \le T$, the stopped Abel identity is

$$\sum_{m=1}^{H} \frac{e_m - e_{m-1}}{\gamma_{m-1}} = \frac{e_H}{\gamma_{H-1}} - \frac{e_0}{\gamma_0} + \sum_{m=1}^{H-1} \left(\frac{1}{\gamma_{m-1}} - \frac{1}{\gamma_m} \right) e_m. \tag{31}$$

We control the two terms on the right uniformly in H.

Step 1: uniform boundary term. The scalar recursion can be written as

$$e_m = (1 - f(Q^\star)\gamma_{m-1}) e_{m-1} + \gamma_{m-1}\{\widetilde{\xi}_m + \delta_m^{\mathrm{dec}} + r_m^{\mathrm{lin}}\}.$$

Expanding this stable recursion gives, for any H,

$$e_H = \Pi_{H,0} e_0 + \sum_{i=1}^{H} \gamma_{i-1} \Pi_{H,i} \{\widetilde{\xi}_i + \delta_i^{\mathrm{dec}} + r_i^{\mathrm{lin}}\}, \qquad \Pi_{H,i} = \prod_{j=i+1}^{H} (1 - f(Q^\star)\gamma_{j-1}).$$

Since $f(Q^\star) \ge f_-$ and $\gamma_j \to 0$, $|\Pi_{H,i}| \le \exp\{-c \sum_{j=i+1}^{H} \gamma_{j-1}\}$ for all large indices. The standard stable-recursion maximal inequality, obtained by combining this kernel bound with Burkholder's inequality for the martingale part and the absolute-sum bounds for the forcing terms, yields

$$\sup_{1\le H\le T} \frac{\sqrt{\mathfrak{n}_T}}{T} \frac{|e_H|}{\gamma_H} = O_p\left(\Delta_T + \Lambda_T + \frac{\sqrt{t_T}}{\kappa T} + \frac{\sqrt{t_T}}{T\sqrt{\gamma_T E_T}} + T^{(\beta-1)/2} \log T \right). \tag{32}$$

Here the martingale part uses $\mathbb{E}(\widetilde{\xi}_m^2 \mid \mathcal{F}_{m-1}) \le C\kappa^2/E_m$, the forcing term δ_m^{dec} uses (28), and the forcing term r_m^{lin} uses (29). The initial condition is exponentially damped and is absorbed into the same bound.

Step 2: uniform variation term. Because $\gamma_m \asymp m^{-\alpha}$, $\left|\gamma_{m-1}^{-1} - \gamma_m^{-1}\right| \le C m^{\alpha-1}$. For the stopped sum, the supremum is bounded by the full absolute sum:

$$\sup_{H\le T} \frac{\sqrt{\mathfrak{n}_T}}{T} \sum_{m=1}^{H-1} \left|\gamma_{m-1}^{-1} - \gamma_m^{-1}\right| |e_m| \le \frac{\sqrt{\mathfrak{n}_T}}{T} \sum_{m=1}^{T-1} \left|\gamma_{m-1}^{-1} - \gamma_m^{-1}\right| |e_m|.$$

The same calculation as in the proof of Theorem 1, using Lemma 7, gives

$$\mathbb{E}\left[\frac{\sqrt{\mathfrak{n}_T}}{T} \sum_{m=1}^{T-1} \left|\gamma_{m-1}^{-1} - \gamma_m^{-1}\right| |e_m| \right] \le C\{T^{(\alpha-1)/2} + T^{(\beta-1)/2} \log T\} + C \frac{\sqrt{t_T}}{\kappa T} \log T.$$

The last term is $O(\Delta_T)$, because $T^{\alpha-1} \log T \to 0$.

Step 3: conclusion. Combining (31)– (32),

$$\sup_{0\le H\le T}\left|\frac{\sqrt{\mathfrak{n}_T}}{T}\sum_{m=1}^{H}\frac{e_m-e_{m-1}}{\gamma_{m-1}}\right| = O_p\left(\Delta_T+\Lambda_T+T^{(\alpha-1)/2}+T^{(\beta-1)/2}\log T\right).$$

Under $K=o(T^{2\alpha-1-\beta})$, the right-hand side is $o_p(1)$. □

Proof of Theorem 2. Define the fixed-target martingale process

$$M_T(s)=\frac{\sqrt{\mathfrak{n}_T}}{f(Q^\star)T}\sum_{m=1}^{h(s,T)}\xi_m,\qquad s\in[0,1].$$

Its predictable quadratic variation is

$$\langle M_T\rangle(s)=\frac{t_T}{f(Q^\star)^2\kappa^2T^2}\sum_{m=1}^{h(s,T)}\frac{\kappa^2\overline{S}}{E_m}=\frac{\nu_T\overline{S}}{f(Q^\star)^2}s_{h(s,T)}.$$

Because $E_m\asymp m^\beta$ with $\beta<1$,

$$\max_{m\le T}(s_m-s_{m-1})=\frac{\max_{m\le T}E_m^{-1}}{\sum_{i=1}^{T}E_i^{-1}}\to 0.$$

Thus $\sup_{s\in[0,1]}|s_{h(s,T)}-s|\to 0$. Also,

$$\sup_s|\Delta M_T(s)|\le C\frac{\sqrt{t_T}}{\kappa T}\asymp C\sqrt{K}\,T^{(\beta-1)/2}\to 0,$$

where the last step follows from $K=o(T^{2\alpha-1-\beta})$ and $\alpha<1$. The martingale functional CLT gives

$$\frac{M_T(\cdot)}{(\nu_T\overline{S}/f(Q^\star)^2)^{1/2}}\xrightarrow{d}B(\cdot)\qquad\text{in }\ell^\infty[0,1]. \tag{33}$$

It remains to transfer this limit from M_T to Q_T. Let

$$\widetilde{M}_T(s)=\frac{\sqrt{\mathfrak{n}_T}}{f(Q^\star)T}\sum_{m=1}^{h(s,T)}\widetilde{\xi}_m.$$

Lemma 8 gives

$$\sup_{s\in[0,1]}|\widetilde{M}_T(s)-M_T(s)|=o_p(1).$$

For a fixed s, put $H=h(s,T)$. Summing (26) from $m=1$ to H yields

$$\begin{aligned}Q_T(s)-\widetilde{M}_T(s)=&\frac{\sqrt{\mathfrak{n}_T}}{f(Q^\star)T}\sum_{m=1}^{H}\delta_m^{\mathrm{dec}}+\frac{\sqrt{\mathfrak{n}_T}}{f(Q^\star)T}\sum_{m=1}^{H}r_m^{\mathrm{lin}}\\&-\frac{\sqrt{\mathfrak{n}_T}}{f(Q^\star)T}\sum_{m=1}^{H}\frac{e_m-e_{m-1}}{\gamma_{m-1}}+\frac{\sqrt{\mathfrak{n}_T}}{T}(e_H-e_0).\end{aligned}$$

The stopped absolute sums of δ_m^{dec} and r_m^{lin} are bounded by their full absolute sums, so (28) and (29) give

$$\sup_s\left|\frac{\sqrt{\mathfrak{n}_T}}{T}\sum_{m=1}^{h(s,T)}\delta_m^{\mathrm{dec}}\right|+\sup_s\left|\frac{\sqrt{\mathfrak{n}_T}}{T}\sum_{m=1}^{h(s,T)}r_m^{\mathrm{lin}}\right|=o_p(1).$$

Lemma 9 controls the Abel term uniformly. For the endpoint term, compactness of Θ implies

$$\sup_{H\le T}\frac{\sqrt{\mathfrak{n}_T}}{T}|e_H-e_0|=O_p\left(\frac{\sqrt{t_T}}{\kappa T}\right)=o_p(1).$$

Therefore

$$\sup_{s\in[0,1]} |Q_T(s) - M_T(s)| = o_p(1),$$

and (33) implies the standardized FCLT for Q_T.

Under Assumption 4, $E_m \sim c_E(m+1)^\beta$ with $\beta < 1$. Let

$$R_m = \sum_{i=1}^{m} E_i^{-1}, \qquad R_0 = 0.$$

By the uniform convergence theorem for regularly varying sequences,

$$\sup_{0\le m\le T} \left| \frac{R_m}{R_T} - \left(\frac{m}{T}\right)^{1-\beta} \right| \to 0.$$

Since $s_m = R_m/R_T$, the uniform continuity of $x \mapsto x^{1/(1-\beta)}$ on $[0,1]$ gives

$$\sup_{0\le m\le T} \left| s_m^{1/(1-\beta)} - \frac{m}{T} \right| \to 0.$$

For self-normalization, define

$$\Psi_\beta(x) = \frac{x(1)}{\left\{\int_0^1 \left[x(r) - r^{1/(1-\beta)}x(1)\right]^2 dr\right\}^{1/2}}.$$

The denominator is positive almost surely for Brownian motion, so Ψ_β is continuous at almost every Brownian path. At the grid points,

$$\frac{\sqrt{\mathfrak{n}_T}}{T}\left\{\sum_{i=1}^{m} \bar{q}_{t_i} - \frac{m}{T}\sum_{i=1}^{T} \bar{q}_{t_i}\right\} = Q_T(s_m) - \frac{m}{T}Q_T(1),$$

so the unknown $Q^\star$ cancels exactly. By the clock calculation above,

$$\sup_{0\le m\le T} \left| s_m^{1/(1-\beta)} - \frac{m}{T} \right| \to 0.$$

Since $Q_T(\cdot) = O_p(1)$ in $\ell^\infty[0,1]$, replacing m/T by $s_m^{1/(1-\beta)}$ changes the Riemann sum by at most

$$O_p(1) \sup_{m\le T} \left| s_m^{1/(1-\beta)} - \frac{m}{T} \right| = o_p(1).$$

Consequently, using $\max_{m\le T}(s_m - s_{m-1}) \to 0$, $\widehat{V}_T$ is asymptotically equivalent to

$$\int_0^1 \left\{Q_T(r) - r^{1/(1-\beta)}Q_T(1)\right\}^2 dr.$$

Applying the continuous mapping theorem to the standardized FCLT gives

$$\frac{Q_T(1)}{\sqrt{\widehat{V}_T}} \xrightarrow{d} \frac{B(1)}{\left\{\int_0^1 \left[B(r) - r^{1/(1-\beta)}B(1)\right]^2 dr\right\}^{1/2}}.$$

Since $Q_T(1) = \sqrt{\mathfrak{n}_T}(\widehat{Q}_T^{\mathrm{st}} - Q^\star)$, the confidence interval follows. □

A.5 Proof of Theorem 3

By definition,

$$\widehat{Q}_{k,T}^{\mathrm{time}} - \widehat{Q}_T^{\mathrm{st}} = \frac{1}{T}\sum_{m=1}^{T} (q_{k,t_m} - \bar{q}_{t_m}).$$

Lemma 3 gives the row-wise bound

$$|q_{k,t_m} - \bar{q}_{t_m}| \le C\log(eK)\gamma_{m-1} \qquad \text{almost surely.}$$

Hence

$$\left|\widehat{Q}_{k,T}^{\mathrm{time}} - \widehat{Q}_T^{\mathrm{st}}\right| \le \frac{C\log(eK)}{T}\sum_{m=1}^{T}\gamma_{m-1}.$$

Multiplying by $\sqrt{\mathfrak{n}_T}$ proves (5). The same stopped-sum argument gives

$$\sup_{s\in[0,1]}|Q_{k,T}(s) - Q_T(s)| \le C\log(eK)\frac{\sqrt{\mathfrak{n}_T}}{T}\sum_{m=1}^{T}\gamma_{m-1}.$$

If $\log(eK)\sqrt{\mathfrak{n}_T}\sum_{m=1}^{T}\gamma_{m-1}/T \to 0$ holds, the right-hand side is $o(1)$, so the spatial–temporal CLT and FCLT transfer to the node-wise temporal average by Slutsky's lemma. The node-wise self-normalized statistic is then handled by the same continuous mapping argument as in Theorem 2, with q_{k,t_m} replacing $\bar{q}_{t_m}$.

By Assumption 2 and $E_m \asymp m^\beta$,

$$\log(eK)\frac{\sqrt{\mathfrak{n}_T}}{T}\sum_{m=1}^{T}\gamma_{m-1} \asymp \log(eK)\sqrt{K}\,T^{(1+\beta)/2-\alpha}.$$

Therefore $\log(eK)\sqrt{\mathfrak{n}_T}\sum_{m=1}^{T}\gamma_{m-1}/T \to 0$ is implied by $K\{\log(eK)\}^2 = o(T^{2\alpha-1-\beta})$.

A.6 Proof of Theorem 4

Recall that $e_m = \theta_m - Q^\star$. Since $\widehat{Q}_T^{\mathrm{space}} = \theta_T$, it suffices to study the last mean error e_T. The mean recursion is

$$e_m = (1 - f(Q^\star)\gamma_{m-1})e_{m-1} + \gamma_{m-1}\widetilde{\xi}_m + \gamma_{m-1}\delta_m^{\mathrm{dec}} + \gamma_{m-1}r_m^{\mathrm{lin}}. \tag{34}$$

Lemma (Last-iterate coefficient bounds with diverging blocks)**.** *Let*

$$a_T = \sqrt{\gamma_{T-1}/E_T}, \qquad c_{T,m} = \frac{\gamma_{m-1}}{a_T}\prod_{j=m+1}^{T}(1 - f(Q^\star)\gamma_{j-1}), \qquad 1 \le m \le T.$$

Under $\gamma_m \sim c_\gamma(m+1)^{-\alpha}$ *and* $E_m \sim c_E(m+1)^\beta$*,*

$$\sum_{m=1}^{T}\frac{c_{T,m}^2}{E_m} = \frac{1+o(1)}{2f(Q^\star)}, \tag{35}$$

$$\max_{m\le T}|c_{T,m}| \le C\sqrt{\gamma_{T-1}E_T}, \tag{36}$$

$$\sum_{m=1}^{T}|c_{T,m}|\gamma_{m-1} \le C\sqrt{\gamma_{T-1}E_T}, \tag{37}$$

$$\sum_{m=1}^{T}|c_{T,m}|\frac{\gamma_{m-1}}{E_m} \le C\sqrt{\frac{\gamma_{T-1}}{E_T}}, \tag{38}$$

$$\sum_{m=1}^{T}|c_{T,m}|\gamma_{m-1}^2 \le C\gamma_{T-1}\sqrt{\gamma_{T-1}E_T}, \tag{39}$$

$$\sum_{m=1}^{T}\frac{c_{T,m}^2}{E_m}\sqrt{\frac{\gamma_{m-1}}{E_m}} \le C\sqrt{\frac{\gamma_{T-1}}{E_T}}, \qquad \sum_{m=1}^{T}\frac{c_{T,m}^2}{E_m}\gamma_{m-1} \le C\gamma_{T-1}. \tag{40}$$

Proof. Let $L_T = \lfloor T^\alpha \log T\rfloor$. Since $f(Q^\star) \ge f_- > 0$ and $\gamma_j \to 0$, for all large T,

$$\prod_{j=m+1}^{T}(1 - f(Q^\star)\gamma_{j-1}) \le \exp\left\{-c\sum_{j=m+1}^{T}\gamma_{j-1}\right\}.$$

If $m \le T - L_T$, then

$$\sum_{j=m+1}^{T} \gamma_{j-1} \ge \sum_{j=T-L_T+1}^{T} \gamma_{j-1} \asymp L_T \gamma_T \asymp \log T,$$

so all contributions outside the terminal window are polynomially negligible after enlarging the constant in L_T.

Inside the terminal window write $m = T - \ell$, $0 \le \ell \le L_T$. Because $L_T/T \to 0$, uniformly over this window,

$$\gamma_{T-\ell-1} = \gamma_{T-1}\{1 + o(1)\}, \qquad E_{T-\ell} = E_T\{1 + o(1)\}.$$

Also, $\sum_{j=T-\ell+1}^{T} \gamma_{j-1}^2 = O(L_T T^{-2\alpha}) = o(1)$, so

$$\prod_{j=T-\ell+1}^{T} (1 - f(Q^\star)\gamma_{j-1}) = \exp\left\{-f(Q^\star) \sum_{j=T-\ell+1}^{T} \gamma_{j-1}\right\}\{1 + o(1)\}$$

uniformly in ℓ. Therefore

$$\begin{aligned}\frac{c_{T,T-\ell}^2}{E_{T-\ell}} &= \frac{\gamma_{T-\ell-1}^2}{a_T^2 E_{T-\ell}} \exp\left\{-2f(Q^\star) \sum_{j=T-\ell+1}^{T} \gamma_{j-1}\right\}\{1 + o(1)\} \\ &= \gamma_{T-1} \exp\left\{-2f(Q^\star) \sum_{j=T-\ell+1}^{T} \gamma_{j-1}\right\}\{1 + o(1)\},\end{aligned}$$

because $a_T^2 = \gamma_{T-1}/E_T$. With $u_{T,0} = 0$ and $u_{T,\ell} = \sum_{r=0}^{\ell-1} \gamma_{T-r-1}$, the mesh size is $\max_\ell |u_{T,\ell+1} - u_{T,\ell}| = O(\gamma_T) \to 0$ and $u_{T,L_T} \to \infty$. Hence

$$\sum_{\ell=0}^{L_T} \frac{c_{T,T-\ell}^2}{E_{T-\ell}} = \frac{1 + o(1)}{2f(Q^\star)},$$

which proves (35).

The same display gives the envelope

$$|c_{T,T-\ell}| \le C\sqrt{\gamma_{T-1} E_T}\exp\{-c\ell\gamma_{T-1}\}, \qquad 0 \le \ell \le L_T,$$

while outside the window the product is polynomially smaller. The maximum bound (36) follows. The remaining bounds follow by summing this envelope. For instance,

$$\sum_{m=1}^{T} |c_{T,m}|\gamma_{m-1} \le C\sqrt{\gamma_{T-1}E_T}\,\gamma_{T-1} \sum_{\ell=0}^{\infty} e^{-c\ell\gamma_{T-1}} \le C\sqrt{\gamma_{T-1}E_T},$$

and

$$\sum_{m=1}^{T} |c_{T,m}|\frac{\gamma_{m-1}}{E_m} \le C\sqrt{\gamma_{T-1}E_T}\,\frac{\gamma_{T-1}}{E_T} \sum_{\ell=0}^{\infty} e^{-c\ell\gamma_{T-1}} \le C\sqrt{\frac{\gamma_{T-1}}{E_T}}.$$

The other inequalities are obtained identically, using $c_{T,T-\ell}^2/E_{T-\ell} \le C\gamma_{T-1}e^{-c\ell\gamma_{T-1}}$. □

Proof of Theorem 4. Iterating (34) and dividing by κa_T gives

$$\frac{e_T}{\kappa a_T} = \sum_{m=1}^{T} c_{T,m}\frac{\widetilde{\xi}_m}{\kappa} + R_T^{(1)} + R_T^{(2)} + R_T^{(3)},$$

where the three remainders correspond to the initial condition, δ_m^{dec}, and r_m^{lin}, respectively:

$$R_T^{(1)} = \frac{e_0}{\kappa a_T}\prod_{j=1}^{T}(1 - f(Q^\star)\gamma_{j-1}), \quad R_T^{(2)} = \sum_{m=1}^{T} c_{T,m}\frac{\delta_m^{\mathrm{dec}}}{\kappa}, \quad R_T^{(3)} = \sum_{m=1}^{T} c_{T,m}\frac{r_m^{\mathrm{lin}}}{\kappa}.$$

The initial product is exponentially small because $\sum_{j=1}^{T}\gamma_{j-1} \asymp T^{1-\alpha}$, so $R_T^{(1)} = o(1)$. For the decentralized remainder, Lemma 5 and (37) imply

$$\mathbb{E}|R_T^{(2)}| \le \frac{C}{\kappa}\sum_{m=1}^{T}|c_{T,m}|\gamma_{m-1} \le C\frac{\sqrt{\gamma_{T-1}E_T}}{\kappa}.$$

Thus $R_T^{(2)} = o_p(1)$ if $\sqrt{\gamma_{T-1}E_T}/\kappa \to 0$. For the nonlinear remainder, Lemma 7 gives

$$\mathbb{E}|r_m^{\mathrm{lin}}| \le C\left(\kappa^2\frac{\gamma_{m-1}}{E_m} + \gamma_{m-1}^2\right).$$

Using (38) and (39),

$$\begin{aligned}\mathbb{E}|R_T^{(3)}| &\le C\kappa\sum_{m=1}^{T}|c_{T,m}|\frac{\gamma_{m-1}}{E_m} + \frac{C}{\kappa}\sum_{m=1}^{T}|c_{T,m}|\gamma_{m-1}^2\\ &\le C\kappa\sqrt{\frac{\gamma_{T-1}}{E_T}} + C\gamma_{T-1}\frac{\sqrt{\gamma_{T-1}E_T}}{\kappa} = o(1).\end{aligned}$$

For the leading martingale term, set $Y_{m,T} = c_{T,m}\widetilde{\xi}_m/\kappa$. Then $\{Y_{m,T}, \mathcal{F}_m\}$ is a martingale array. Its predictable variance is

$$\sum_{m=1}^{T}\mathbb{E}(Y_{m,T}^2 \mid \mathcal{F}_{m-1}) = \sum_{m=1}^{T}c_{T,m}^2\frac{\mathrm{Var}(\widetilde{\xi}_m \mid \mathcal{F}_{m-1})}{\kappa^2} = \overline{S}\sum_{m=1}^{T}\frac{c_{T,m}^2}{E_m} + R_{V,T}.$$

By (15),

$$|R_{V,T}| \le C\sum_{m=1}^{T}\frac{c_{T,m}^2}{E_m}|e_{m-1}|.$$

Using Lemma 7 and (40),

$$\mathbb{E}|R_{V,T}| \le C\sum_{m=1}^{T}\frac{c_{T,m}^2}{E_m}\left(\kappa\sqrt{\frac{\gamma_{m-1}}{E_m}} + \gamma_{m-1}\right) \le C\kappa\sqrt{\frac{\gamma_{T-1}}{E_T}} + C\gamma_{T-1} \to 0.$$

Together with (35), this gives

$$\sum_{m=1}^{T}\mathbb{E}(Y_{m,T}^2 \mid \mathcal{F}_{m-1}) \xrightarrow{p} \frac{\overline{S}}{2f(Q^\star)}.$$

The Lindeberg condition follows from

$$\max_{m\le T}|Y_{m,T}| \le C\frac{\max_m|c_{T,m}|}{\kappa} \le C\frac{\sqrt{\gamma_{T-1}E_T}}{\kappa} \to 0.$$

Therefore the martingale CLT yields

$$\frac{e_T}{\kappa\sqrt{\gamma_{T-1}/E_T}}\Big/\left(\frac{\overline{S}}{2f(Q^\star)}\right)^{1/2} \xrightarrow{d} N(0,1).$$

Since $\widehat{Q}_T^{\mathrm{space}} = \theta_T$, this proves the standardized statement for the terminal spatial average. By Assumption 2, $\kappa^2 \asymp K^{-1}$. Together with $\gamma_{T-1} \asymp T^{-\alpha}$ and $E_T \asymp T^{\beta}$, the condition $\sqrt{\gamma_{T-1}E_T}/\kappa \to 0$ is equivalent to $K = o(T^{\alpha-\beta})$. When $E_m \equiv 1$ and $\eta_t = D(t+1)^{-\alpha}$, we have $E_T = 1$ and $\gamma_{T-1} = \eta_{T-1} \sim DT^{-\alpha}$, which gives the stated special case. □

B RELATION TO EXISTING INFERENTIAL FRAMEWORKS.

Table S.1 summarizes the relationship between the proposed method and related frameworks for private and decentralized inference. Our setting combines locally private quantile scores with neighbor-only communication and develops graph-dependent inference under incomplete synchronization.

Table S.1: Comparison of related inferential frameworks.

Framework	Target and communication architecture	Inferential contribution
Online LDP quantile inference (Liu et al., 2023)	Scalar quantile from a single private data stream	One-bit stochastic approximation and self-normalized online inference.
Server-coordinated federated LDP quantile inference (Cai et al., 2026)	Weighted mixture quantile with heterogeneous clients and exact aggregation and broadcast at communication rounds	Asymptotic theory for nonsmooth local SGD, including a functional central limit theorem and self-normalized confidence intervals.
Peer-to-peer inference for smooth M-estimation (Gu & Chen, 2024)	General smooth M-estimation with communication over a peer-to-peer graph	Graph-dependent asymptotic inference for decentralized stochastic approximation under smooth-score conditions.
Proposed method	Weighted mixture quantile from heterogeneous private streams with neighbor-only gossip	Graph-dependent control of residual disagreement and transfer of the network-mean process limit to node-wise trajectories.

C ADDITIONAL EXPERIMENTS AND DISCUSSIONS

In this section, we report additional experiments that complement Section 4. We first present the full simulation results, followed by comparisons across network structures and baseline methods, sensitivity analyses for communication and distribution shifts, and additional real-data results. The simulation designs and estimators are defined in Section 4.1; the settings for each experiment are given below.

C.1 OTHER RESULTS IN SECTION 4

In this subsection, we report the full simulation results that complement the main experiments in Section 4. Tables S.2–S.3, Tables S.4–S.5, and Tables S.6–S.7 summarize the results for $K = 10$, $K = 30$, and $K = 50$, respectively. For each K, we report both the homogeneous/heterogeneous-distribution designs and the heterogeneous-location design. For these simulations, the effective step size at communication round m is $\gamma_m = ((m + m_0)^{-\alpha_\gamma})(c_\gamma \bar{r})$ with $\bar{r} = K^{-1} \sum_{k=1}^{K} r_k$. The values of $(c_\gamma, \alpha_\gamma, m_0)$ are $(1.8, 0.65, 500)$ for C1, $(1.2, 0.65, 200)$ for C5, and $(1.6, 0.55, 100)$ for $P(1/5)$. Here $P(1/5)$ denotes $E_m = \lceil (m + 1)^{1/5} \rceil$. Its choice $\alpha_\gamma = 0.55$ falls outside Assumption 4, since $\beta = 1/5 > 2\alpha_\gamma - 1 = 0.1$; we report only point-estimation errors for this setting.

All synthetic-data experiments initialize the node iterates and centralized comparator at zero. The step-size choices above also apply to the additional simulations: C_5 for network connectivity, directed/asynchronous communication, and temporal shifts; C_1 for edge dropout; and the corresponding C_1 or C_5 rule for the iterative baselines. In asynchronous runs, the index is the active node's local block count. NP uses the same rule with $r_k = 1$, whereas HIST has no stochastic-approximation step size. Real-data tuning is given in Section C.7.

Table S.2: Empirical coverage probabilities at the 95% nominal level (mean absolute errors ↓ ($\times 10^{-3}$)) under the homogeneous and heterogeneous-distribution designs for C_1 and C_5, with MAEs only for $P(1/5)$. The number of nodes is fixed at $K = 10$. In Homo, data for every node are generated from $\mathcal{N}(0,1)$. In Hete D, nodes are assigned to $\mathcal{N}(0,1)$, $\mathcal{U}(-1,1)$, and $\mathcal{C}(0,1)$ groups. "hetero" indicates node-specific truthful response rates r_k equally spaced on $[0.25, 0.9]$. For the terminal spatial estimator, we report only the MAEs.

		C1					C5					$P(1/5)$ (MAE)				
τ	r_k	$\widehat{Q}^{s}$	$\widehat{Q}^{st}$	$\widehat{Q}^{t}$	$\widehat{Q}^{\text{cent},s}$	$\widehat{Q}^{\text{cent},t(st)}$	$\widehat{Q}^{s}$	$\widehat{Q}^{st}$	$\widehat{Q}^{t}$	$\widehat{Q}^{\text{cent},s}$	$\widehat{Q}^{\text{cent},t(st)}$	$\widehat{Q}^{s}$	$\widehat{Q}^{st}$	$\widehat{Q}^{t}$	$\widehat{Q}^{\text{cent},s}$	$\widehat{Q}^{\text{cent},t(st)}$
Homo | $T = 100{,}000$																
0.3	0.25	9.5	0.94(5.3)	0.94(5.3)	9.5	0.94(5.3)	3.6	0.98(3.6)	0.98(3.6)	3.6	0.98(3.6)	5.3	1.7	1.7	5.3	1.7
0.3	hetero	8	0.93(2.9)	0.93(2.9)	8	0.93(2.9)	3	0.96(1.4)	0.96(1.4)	3	0.96(1.4)	4.4	0.9	0.9	4.3	0.9
0.3	0.9	4.9	0.95(1.3)	0.95(1.3)	4.9	0.95(1.3)	1.7	0.98(0.71)	0.98(0.71)	1.7	0.98(0.71)	2.6	0.44	0.44	2.6	0.44
0.5	0.25	8.8	0.92(4)	0.92(4)	8.8	0.91(4)	3.3	0.90(1.9)	0.90(1.9)	3.3	0.90(1.9)	4.4	1.4	1.4	4.4	1.4
0.5	hetero	7.4	0.93(2.3)	0.93(2.3)	7.4	0.93(2.3)	2.7	0.91(0.99)	0.91(0.99)	2.7	0.91(0.99)	4	0.74	0.74	4	0.74
0.5	0.9	5.1	0.94(1.2)	0.94(1.2)	5.1	0.94(1.2)	1.7	0.95(0.47)	0.95(0.47)	1.7	0.95(0.47)	2.5	0.37	0.37	2.5	0.37
0.8	0.25	11	0.95(7.6)	0.95(7.6)	11	0.96(7.6)	4.1	1.00(7.5)	1.00(7.5)	4.1	1.00(7.5)	5.7	2.1	2.1	5.7	2.1
0.8	hetero	9.2	0.96(3.3)	0.96(3.3)	9.2	0.96(3.3)	3.1	0.99(2.1)	0.99(2.1)	3.1	0.99(2.1)	4.6	1.1	1.1	4.6	1.1
0.8	0.9	5	0.97(1.6)	0.97(1.6)	5	0.97(1.6)	1.7	1.00(1.1)	1.00(1.1)	1.7	1.00(1.1)	2.5	0.5	0.5	2.5	0.5
Homo | $T = 200{,}000$																
0.3	0.25	7.8	0.90(3.4)	0.90(3.4)	7.8	0.90(3.4)	2.7	0.98(2.1)	0.98(2.1)	2.7	0.98(2.1)	3.7	1	1	3.7	1
0.3	hetero	6.1	0.94(1.9)	0.94(1.9)	6.1	0.94(1.9)	2.4	0.94(0.9)	0.94(0.9)	2.4	0.94(0.9)	3.4	0.55	0.55	3.4	0.55
0.3	0.9	3.8	0.94(0.89)	0.94(0.89)	3.8	0.94(0.89)	1.3	0.96(0.43)	0.96(0.43)	1.3	0.96(0.43)	1.9	0.27	0.27	1.9	0.27
0.5	0.25	7.4	0.92(2.8)	0.92(2.8)	7.4	0.92(2.8)	2.6	0.92(1.3)	0.92(1.3)	2.6	0.91(1.3)	3.5	0.94	0.94	3.5	0.94
0.5	hetero	6.2	0.93(1.6)	0.93(1.6)	6.2	0.92(1.6)	2.2	0.91(0.73)	0.91(0.73)	2.2	0.91(0.73)	3	0.52	0.52	3	0.52
0.5	0.9	3.9	0.95(0.8)	0.95(0.8)	3.9	0.95(0.79)	1.3	0.94(0.35)	0.94(0.35)	1.3	0.94(0.35)	1.9	0.26	0.26	1.9	0.26
0.8	0.25	9.2	0.95(4.7)	0.95(4.7)	9.2	0.95(4.7)	3.2	1.00(3.8)	1.00(3.8)	3.2	1.00(3.8)	4.4	1.4	1.4	4.4	1.4
0.8	hetero	7.1	0.94(2.3)	0.94(2.3)	7.1	0.94(2.3)	2.5	0.98(1.2)	0.98(1.2)	2.5	0.98(1.2)	3.4	0.76	0.76	3.4	0.76
0.8	0.9	3.8	0.95(1.1)	0.95(1.1)	3.8	0.95(1.1)	1.3	0.98(0.61)	0.98(0.61)	1.3	0.98(0.6)	1.9	0.31	0.31	1.9	0.31
Hete D | $T = 100{,}000$																
0.5	0.25	9	0.91(3.9)	0.91(3.9)	9	0.91(3.9)	3.4	0.89(1.8)	0.90(1.8)	3.4	0.90(1.8)	5	1.4	1.4	5	1.4
0.5	hetero	7.6	0.94(2.1)	0.94(2.1)	7.6	0.94(2.1)	2.7	0.92(1)	0.93(1)	2.7	0.93(1)	3.9	0.75	0.75	3.9	0.75
0.5	0.9	4.9	0.94(1.1)	0.94(1.1)	4.9	0.94(1.1)	1.8	0.94(0.47)	0.94(0.47)	1.8	0.94(0.47)	2.6	0.36	0.36	2.6	0.36
Hete D | $T = 200{,}000$																
0.5	0.25	7.4	0.91(2.9)	0.91(2.9)	7.4	0.91(2.9)	2.7	0.89(1.3)	0.89(1.3)	2.7	0.89(1.3)	3.4	0.95	0.95	3.4	0.95
0.5	hetero	5.9	0.96(1.5)	0.96(1.5)	5.9	0.96(1.5)	2.3	0.93(0.75)	0.93(0.75)	2.3	0.93(0.75)	3.1	0.5	0.5	3.1	0.5
0.5	0.9	3.9	0.94(0.82)	0.94(0.82)	3.9	0.94(0.82)	1.3	0.93(0.36)	0.93(0.36)	1.3	0.93(0.36)	1.8	0.26	0.26	1.8	0.26

Table S.3: Empirical coverage probabilities at the 95% nominal level (mean absolute errors ↓ ($\times 10^{-3}$)) under the heterogeneous-location design for C_1 and C_5, with MAEs only for $P(1/5)$. The number of nodes is fixed at $K = 10$. Data for node k are generated from $\mathcal{N}(\mu_k, 1)$ with deterministic locations μ_k equally spaced on $[-1, 1]$. "hetero" indicates node-specific truthful response rates r_k equally spaced on $[0.25, 0.9]$. For the terminal spatial estimator, we report only the MAEs.

		C1					C5					$P(1/5)$ (MAE)				
τ	r_k	$\widehat{Q}^{s}$	$\widehat{Q}^{st}$	$\widehat{Q}^{t}$	$\widehat{Q}^{\mathrm{cent},s}$	$\widehat{Q}^{\mathrm{cent},t(st)}$	$\widehat{Q}^{s}$	$\widehat{Q}^{st}$	$\widehat{Q}^{t}$	$\widehat{Q}^{\mathrm{cent},s}$	$\widehat{Q}^{\mathrm{cent},t(st)}$	$\widehat{Q}^{s}$	$\widehat{Q}^{st}$	$\widehat{Q}^{t}$	$\widehat{Q}^{\mathrm{cent},s}$	$\widehat{Q}^{\mathrm{cent},t(st)}$
Hete L | $T = 100{,}000$																
0.3	0.25	11	0.93(6.8)	0.93(6.8)	11	0.94(6.8)	3.8	1.00(5.9)	1.00(5.9)	3.8	1.00(5.8)	5.5	2.2	2.2	5.5	2.1
0.3	hetero	8.6	0.93(3.3)	0.93(3.3)	8.6	0.93(3.3)	3.2	0.99(1.9)	0.99(1.9)	3.2	0.99(1.8)	4.7	1.2	1.5	4.7	1.1
0.3	0.9	5	0.95(1.6)	0.95(1.7)	5	0.96(1.5)	1.7	0.99(1.1)	0.99(1.1)	1.7	0.99(0.93)	2.5	0.84	1.6	2.5	0.49
0.5	0.25	10	0.89(5.1)	0.89(5.1)	10	0.89(5.1)	3.6	0.88(2.2)	0.88(2.2)	3.6	0.88(2.2)	5.4	1.7	1.7	5.5	1.7
0.5	hetero	8.2	0.93(2.6)	0.93(2.6)	8.2	0.92(2.6)	3	0.90(1.2)	0.89(1.3)	3	0.91(1.2)	4.2	0.9	1.3	4.2	0.89
0.5	0.9	5	0.94(1.2)	0.91(1.4)	5	0.92(1.2)	1.9	0.94(0.55)	0.91(0.7)	1.9	0.93(0.55)	2.5	0.42	1.7	2.5	0.41
0.8	0.25	12	0.98(9.8)	0.98(9.8)	12	0.98(9.8)	4.6	1.00(12)	1.00(12)	4.6	1.00(12)	6.2	2.9	2.9	6.2	2.8
0.8	hetero	9.2	0.96(4.1)	0.96(4.1)	9.2	0.97(4)	3.4	1.00(3.1)	1.00(3.1)	3.4	1.00(3)	4.7	1.6	1.7	4.7	1.4
0.8	0.9	4.9	0.97(2.1)	0.97(2.2)	4.9	0.97(2)	1.9	1.00(1.8)	1.00(1.8)	1.9	1.00(1.6)	2.7	1.2	1.5	2.7	0.6
Hete L | $T = 200{,}000$																
0.3	0.25	8.8	0.95(4.2)	0.95(4.2)	8.8	0.95(4.2)	3.3	0.99(3.1)	0.99(3.1)	3.3	0.99(3)	4	1.4	1.4	4	1.3
0.3	hetero	6.7	0.94(2.2)	0.94(2.2)	6.7	0.94(2.2)	2.6	0.96(1.1)	0.96(1.2)	2.6	0.96(1.1)	3.6	0.77	0.96	3.5	0.7
0.3	0.9	3.8	0.96(1)	0.96(1)	3.8	0.96(0.95)	1.4	0.97(0.6)	0.97(0.65)	1.4	0.98(0.54)	1.9	0.55	1.1	1.9	0.31
0.5	0.25	8.4	0.92(3.5)	0.91(3.5)	8.4	0.92(3.5)	2.8	0.88(1.6)	0.88(1.6)	2.8	0.88(1.6)	4	1.1	1.1	4	1.1
0.5	hetero	6.9	0.93(2)	0.93(2)	6.9	0.92(2)	2.5	0.94(0.8)	0.94(0.82)	2.5	0.94(0.8)	3.3	0.6	0.92	3.4	0.6
0.5	0.9	3.7	0.95(0.89)	0.93(0.98)	3.8	0.95(0.89)	1.3	0.95(0.4)	0.90(0.48)	1.3	0.96(0.4)	2.1	0.28	1.2	2.1	0.28
0.8	0.25	9	0.97(5.9)	0.97(5.9)	9	0.97(5.9)	3.5	1.00(6.1)	1.00(6.1)	3.5	1.00(6)	4.8	1.9	1.9	4.8	1.8
0.8	hetero	7.6	0.95(2.8)	0.95(2.8)	7.6	0.94(2.7)	3	0.99(1.9)	0.98(1.9)	3	0.98(1.8)	3.8	0.95	1.1	3.7	0.84
0.8	0.9	3.7	0.97(1.3)	0.97(1.3)	3.7	0.97(1.2)	1.4	0.99(0.93)	0.99(0.94)	1.4	0.99(0.81)	2	0.71	0.98	1.9	0.34

Table S.4: Empirical coverage probabilities at the 95% nominal level (mean absolute errors ↓ ($\times 10^{-3}$)) under the homogeneous and heterogeneous-distribution designs for C_1 and C_5, with MAEs only for $P(1/5)$. The number of nodes is fixed at $K = 30$. In Homo, data for every node are generated from $\mathcal{N}(0,1)$. In Hete D, nodes are assigned to $\mathcal{N}(0,1)$, $\mathcal{U}(-1,1)$, and $\mathcal{C}(0,1)$ groups. "hetero" indicates node-specific truthful response rates r_k equally spaced on $[0.25, 0.9]$. For the terminal spatial estimator, we report only the MAEs.

		C1					C5					$P(1/5)$ (MAE)				
τ	r_k	$\widehat{Q}^{s}$	$\widehat{Q}^{st}$	$\widehat{Q}^{t}$	$\widehat{Q}^{\text{cent},s}$	$\widehat{Q}^{\text{cent},t(st)}$	$\widehat{Q}^{s}$	$\widehat{Q}^{st}$	$\widehat{Q}^{t}$	$\widehat{Q}^{\text{cent},s}$	$\widehat{Q}^{\text{cent},t(st)}$	$\widehat{Q}^{s}$	$\widehat{Q}^{st}$	$\widehat{Q}^{t}$	$\widehat{Q}^{\text{cent},s}$	$\widehat{Q}^{\text{cent},t(st)}$
Homo | $T = 100{,}000$																
0.3	0.25	5.6	0.98(3.6)	0.98(3.6)	5.6	0.98(3.6)	2	1.00(3.6)	1.00(3.6)	2	1.00(3.6)	2.8	1.1	1.1	2.8	1.1
0.3	hetero	4.5	0.97(1.5)	0.97(1.5)	4.5	0.97(1.5)	1.6	0.99(1.1)	0.99(1.1)	1.6	0.99(1.1)	2.3	0.49	0.49	2.3	0.49
0.3	0.9	2.7	0.97(0.88)	0.97(0.88)	2.7	0.97(0.88)	0.95	1.00(0.56)	1.00(0.56)	0.95	1.00(0.56)	1.5	0.26	0.26	1.5	0.26
0.5	0.25	4.9	0.91(2.3)	0.91(2.3)	4.9	0.91(2.3)	1.9	0.87(1.2)	0.87(1.2)	1.9	0.87(1.2)	2.7	0.77	0.77	2.7	0.77
0.5	hetero	4.2	0.95(1.2)	0.95(1.2)	4.2	0.95(1.2)	1.5	0.91(0.57)	0.91(0.57)	1.5	0.91(0.57)	2.3	0.41	0.41	2.3	0.41
0.5	0.9	2.6	0.96(0.67)	0.96(0.67)	2.6	0.95(0.67)	0.97	0.93(0.29)	0.93(0.29)	0.97	0.93(0.29)	1.4	0.22	0.22	1.4	0.22
0.8	0.25	5.9	0.99(6.1)	0.99(6.1)	5.9	0.99(6.1)	2.4	1.00(7.2)	1.00(7.2)	2.4	1.00(7.2)	3.3	1.6	1.6	3.3	1.6
0.8	hetero	5.1	0.98(2.3)	0.98(2.3)	5.1	0.98(2.3)	1.8	1.00(1.9)	1.00(1.9)	1.8	1.00(1.9)	2.5	0.69	0.69	2.5	0.69
0.8	0.9	2.6	1.00(1.2)	1.00(1.2)	2.6	1.00(1.2)	0.97	1.00(1)	1.00(1)	0.97	1.00(1)	1.4	0.35	0.35	1.5	0.35
Homo | $T = 200{,}000$																
0.3	0.25	4.4	0.97(2.1)	0.97(2.1)	4.4	0.97(2.1)	1.5	1.00(1.8)	1.00(1.8)	1.5	1.00(1.8)	2.1	0.67	0.67	2.1	0.67
0.3	hetero	3.4	0.94(1)	0.94(1)	3.4	0.94(1)	1.3	0.99(0.58)	0.99(0.58)	1.3	0.99(0.58)	1.7	0.33	0.33	1.7	0.33
0.3	0.9	2.2	0.97(0.53)	0.97(0.53)	2.2	0.97(0.54)	0.8	0.99(0.31)	0.99(0.31)	0.8	0.99(0.31)	1	0.16	0.16	1	0.16
0.5	0.25	4.4	0.92(1.6)	0.92(1.6)	4.4	0.92(1.6)	1.6	0.91(0.72)	0.91(0.72)	1.6	0.91(0.72)	2.1	0.56	0.56	2.1	0.56
0.5	hetero	3.3	0.96(0.89)	0.96(0.89)	3.3	0.96(0.89)	1.3	0.93(0.4)	0.93(0.4)	1.3	0.93(0.4)	1.8	0.3	0.3	1.8	0.3
0.5	0.9	2.2	0.95(0.45)	0.95(0.45)	2.2	0.95(0.45)	0.81	0.95(0.21)	0.95(0.21)	0.81	0.95(0.21)	1.1	0.14	0.14	1.1	0.14
0.8	0.25	5	0.98(3.3)	0.98(3.3)	5	0.98(3.3)	1.9	1.00(3.7)	1.00(3.7)	1.9	1.00(3.7)	2.3	0.9	0.9	2.3	0.9
0.8	hetero	3.8	0.97(1.4)	0.97(1.4)	3.8	0.97(1.4)	1.6	1.00(0.94)	1.00(0.94)	1.6	1.00(0.94)	1.9	0.44	0.44	1.9	0.44
0.8	0.9	2.3	0.97(0.7)	0.97(0.7)	2.3	0.97(0.7)	0.76	1.00(0.55)	1.00(0.55)	0.76	1.00(0.55)	1.1	0.21	0.21	1.1	0.21
Hete D | $T = 100{,}000$																
0.5	0.25	5.3	0.91(2.5)	0.91(2.5)	5.3	0.90(2.5)	1.9	0.89(1.1)	0.89(1.1)	1.9	0.89(1.1)	2.6	0.85	0.85	2.6	0.85
0.5	hetero	4.4	0.92(1.3)	0.92(1.3)	4.4	0.92(1.3)	1.6	0.90(0.55)	0.90(0.55)	1.6	0.90(0.55)	2.3	0.42	0.42	2.3	0.42
0.5	0.9	2.8	0.94(0.64)	0.94(0.64)	2.8	0.94(0.64)	1	0.95(0.27)	0.95(0.27)	1	0.95(0.27)	1.4	0.23	0.23	1.4	0.23
Hete D | $T = 200{,}000$																
0.5	0.25	4.2	0.91(1.7)	0.91(1.7)	4.2	0.91(1.7)	1.5	0.91(0.77)	0.91(0.77)	1.5	0.91(0.77)	2.1	0.49	0.49	2.1	0.49
0.5	hetero	3.4	0.92(0.9)	0.92(0.9)	3.4	0.92(0.9)	1.2	0.94(0.39)	0.94(0.39)	1.2	0.94(0.39)	1.6	0.28	0.28	1.6	0.28
0.5	0.9	2.5	0.94(0.47)	0.94(0.47)	2.5	0.94(0.47)	0.76	0.93(0.21)	0.93(0.21)	0.76	0.93(0.21)	1.1	0.15	0.15	1.1	0.15

Table S.5: Empirical coverage probabilities at the 95% nominal level (mean absolute errors ↓ ($\times 10^{-3}$)) under the heterogeneous-location design for C_1 and C_5, with MAEs only for $P(1/5)$. The number of nodes is fixed at $K = 30$. Data for node k are generated from $\mathcal{N}(\mu_k, 1)$ with deterministic locations μ_k equally spaced on $[-1, 1]$. "hetero" indicates node-specific truthful response rates r_k equally spaced on $[0.25, 0.9]$. For the terminal spatial estimator, we report only the MAEs.

		C1					C5					$P(1/5)$ (MAE)				
τ	r_k	$\widehat{Q}^{s}$	$\widehat{Q}^{st}$	$\widehat{Q}^{t}$	$\widehat{Q}^{\mathrm{cent},s}$	$\widehat{Q}^{\mathrm{cent},t(st)}$	$\widehat{Q}^{s}$	$\widehat{Q}^{st}$	$\widehat{Q}^{t}$	$\widehat{Q}^{\mathrm{cent},s}$	$\widehat{Q}^{\mathrm{cent},t(st)}$	$\widehat{Q}^{s}$	$\widehat{Q}^{st}$	$\widehat{Q}^{t}$	$\widehat{Q}^{\mathrm{cent},s}$	$\widehat{Q}^{\mathrm{cent},t(st)}$
Hete L | $T = 100{,}000$																
0.3	0.25	6.3	0.98(4.9)	0.98(4.9)	6.3	0.98(4.8)	2.4	1.00(5.6)	1.00(5.6)	2.4	1.00(5.5)	3.2	1.4	1.5	3.2	1.3
0.3	hetero	4.9	0.97(2)	0.96(2.1)	4.9	0.97(2)	1.8	1.00(1.6)	1.00(1.6)	1.8	1.00(1.4)	2.4	0.88	1.4	2.4	0.63
0.3	0.9	2.9	0.96(1.2)	0.97(1.3)	2.9	0.97(1)	0.96	1.00(1)	1.00(1)	0.96	1.00(0.8)	1.5	0.91	2.1	1.4	0.3
0.5	0.25	5.4	0.88(2.9)	0.88(2.9)	5.4	0.88(2.9)	2	0.90(1.2)	0.89(1.2)	2	0.90(1.2)	2.9	0.99	1.1	2.9	0.99
0.5	hetero	4.4	0.92(1.4)	0.92(1.5)	4.4	0.92(1.4)	1.7	0.93(0.67)	0.89(0.76)	1.7	0.92(0.67)	2.3	0.5	1.5	2.3	0.5
0.5	0.9	2.8	0.94(0.69)	0.89(1.1)	2.8	0.94(0.69)	1	0.94(0.31)	0.85(0.68)	1	0.95(0.31)	1.5	0.24	2.3	1.5	0.24
0.8	0.25	6.9	1.00(8.3)	1.00(8.3)	6.9	1.00(8.3)	2.6	1.00(12)	1.00(12)	2.6	1.00(12)	3.5	2.4	2.4	3.5	2.2
0.8	hetero	5.4	0.98(3.2)	0.98(3.2)	5.4	0.99(3)	2	1.00(3)	1.00(3)	2	1.00(2.8)	2.8	1.3	1.5	2.8	0.83
0.8	0.9	3	0.99(1.8)	0.99(1.9)	3	0.99(1.6)	1.1	1.00(1.8)	1.00(1.8)	1.1	1.00(1.5)	1.6	1.2	1.8	1.6	0.42
Heterogeneous-location normal | $T = 200{,}000$																
0.3	0.25	4.7	0.97(2.9)	0.96(2.9)	4.7	0.97(2.8)	1.9	1.00(2.8)	1.00(2.8)	1.9	1.00(2.7)	2.5	0.86	0.91	2.5	0.78
0.3	hetero	3.7	0.96(1.3)	0.95(1.3)	3.7	0.96(1.2)	1.5	0.99(0.86)	1.00(0.88)	1.5	0.99(0.79)	1.8	0.54	0.97	1.8	0.39
0.3	0.9	2.3	0.97(0.71)	0.96(0.84)	2.3	0.97(0.63)	0.8	0.99(0.52)	0.99(0.58)	0.79	1.00(0.41)	1.1	0.59	1.4	1.1	0.19
0.5	0.25	4.6	0.88(2)	0.88(2.1)	4.6	0.88(2)	1.5	0.89(0.88)	0.89(0.89)	1.5	0.89(0.88)	2.2	0.62	0.73	2.2	0.62
0.5	hetero	3.7	0.94(1.1)	0.92(1.1)	3.7	0.94(1.1)	1.4	0.94(0.45)	0.92(0.51)	1.4	0.94(0.45)	1.8	0.33	1	1.8	0.33
0.5	0.9	2.1	0.95(0.48)	0.90(0.73)	2.1	0.95(0.48)	0.84	0.95(0.22)	0.87(0.45)	0.84	0.95(0.22)	1.1	0.16	1.6	1.1	0.16
0.8	0.25	5.3	0.99(4.4)	0.99(4.4)	5.3	0.99(4.3)	2	1.00(5.9)	1.00(5.9)	2	1.00(5.8)	2.7	1.4	1.4	2.7	1.2
0.8	hetero	4.1	0.98(1.9)	0.97(1.9)	4.1	0.98(1.8)	1.5	1.00(1.5)	1.00(1.5)	1.5	1.00(1.4)	2	0.78	0.97	2	0.51
0.8	0.9	2.3	0.99(1)	0.99(1.1)	2.3	0.99(0.86)	0.84	1.00(0.92)	1.00(0.92)	0.83	1.00(0.75)	1.1	0.82	1.2	1.1	0.25

Table S.6: Empirical coverage probabilities at the 95% nominal level (mean absolute errors ↓ ($\times 10^{-3}$)) under the homogeneous and heterogeneous-distribution designs for C_1 and C_5, with MAEs only for $P(1/5)$. The number of nodes is fixed at $K = 50$. In Homo, data for every node are generated from $\mathcal{N}(0,1)$. In Hete D, nodes are assigned to $\mathcal{N}(0,1)$, $\mathcal{U}(-1,1)$, and $\mathcal{C}(0,1)$ groups. "hetero" indicates node-specific truthful response rates r_k equally spaced on $[0.25, 0.9]$. For the terminal spatial estimator, we report only the MAEs.

		C1					C5					$P(1/5)$ (MAE)				
τ	r_k	$\widehat{Q}^{s}$	$\widehat{Q}^{st}$	$\widehat{Q}^{t}$	$\widehat{Q}^{\text{cent},s}$	$\widehat{Q}^{\text{cent},t(st)}$	$\widehat{Q}^{s}$	$\widehat{Q}^{st}$	$\widehat{Q}^{t}$	$\widehat{Q}^{\text{cent},s}$	$\widehat{Q}^{\text{cent},t(st)}$	$\widehat{Q}^{s}$	$\widehat{Q}^{st}$	$\widehat{Q}^{t}$	$\widehat{Q}^{\text{cent},s}$	$\widehat{Q}^{\text{cent},t(st)}$
Homo | $T = 100{,}000$																
0.3	0.25	4.1	1.00(3.1)	1.00(3.1)	4.1	1.00(3.1)	1.7	1.00(3.4)	1.00(3.4)	1.7	1.00(3.4)	2.3	0.88	0.88	2.3	0.88
0.3	hetero	3.6	0.98(1.3)	0.98(1.3)	3.6	0.98(1.3)	1.2	1.00(1)	1.00(1)	1.2	1.00(1)	1.8	0.41	0.41	1.8	0.42
0.3	0.9	2.1	0.99(0.7)	0.99(0.7)	2.1	0.99(0.7)	0.75	1.00(0.55)	1.00(0.55)	0.75	1.00(0.55)	1.2	0.22	0.22	1.2	0.22
0.5	0.25	3.9	0.91(1.8)	0.91(1.8)	3.9	0.91(1.8)	1.5	0.88(0.82)	0.88(0.82)	1.5	0.89(0.82)	2.2	0.63	0.63	2.2	0.63
0.5	hetero	3.2	0.94(0.95)	0.94(0.95)	3.2	0.94(0.95)	1.2	0.94(0.42)	0.94(0.42)	1.2	0.94(0.42)	1.8	0.33	0.33	1.8	0.33
0.5	0.9	2.1	0.95(0.51)	0.95(0.51)	2.1	0.95(0.51)	0.77	0.95(0.22)	0.95(0.22)	0.77	0.95(0.22)	1.1	0.17	0.17	1.1	0.17
0.8	0.25	5.1	1.00(5.8)	1.00(5.8)	5.1	1.00(5.8)	1.7	1.00(7.4)	1.00(7.4)	1.7	1.00(7.4)	2.6	1.4	1.4	2.6	1.4
0.8	hetero	3.9	0.99(2.1)	0.99(2.1)	3.9	0.99(2.1)	1.3	1.00(1.9)	1.00(1.9)	1.3	1.00(1.9)	2	0.57	0.57	2	0.57
0.8	0.9	2.1	1.00(1.1)	1.00(1.1)	2.1	1.00(1.1)	0.75	1.00(1)	1.00(1)	0.75	1.00(1)	1.2	0.29	0.29	1.2	0.29
Homo | $T = 200{,}000$																
0.3	0.25	3.3	0.98(1.9)	0.98(1.9)	3.3	0.98(1.8)	1.3	1.00(1.8)	1.00(1.8)	1.3	1.00(1.8)	1.7	0.54	0.54	1.7	0.54
0.3	hetero	2.9	0.96(0.87)	0.96(0.87)	2.9	0.96(0.87)	1	0.99(0.52)	0.99(0.52)	1	0.99(0.52)	1.4	0.26	0.26	1.4	0.26
0.3	0.9	1.7	0.96(0.46)	0.96(0.46)	1.7	0.96(0.46)	0.63	1.00(0.3)	1.00(0.3)	0.63	1.00(0.3)	0.85	0.13	0.13	0.85	0.13
0.5	0.25	3.1	0.92(1.2)	0.92(1.2)	3.1	0.92(1.2)	1.2	0.89(0.59)	0.89(0.59)	1.2	0.89(0.59)	1.6	0.42	0.42	1.6	0.42
0.5	hetero	2.6	0.94(0.69)	0.94(0.69)	2.6	0.94(0.69)	1.1	0.93(0.3)	0.93(0.3)	1.1	0.93(0.3)	1.4	0.22	0.22	1.4	0.22
0.5	0.9	1.6	0.93(0.37)	0.93(0.37)	1.6	0.93(0.37)	0.62	0.94(0.15)	0.94(0.15)	0.62	0.94(0.15)	0.87	0.11	0.11	0.86	0.11
0.8	0.25	3.8	0.99(3.1)	0.99(3.1)	3.8	0.99(3.1)	1.5	1.00(3.7)	1.00(3.7)	1.5	1.00(3.7)	1.8	0.77	0.77	1.8	0.78
0.8	hetero	2.9	0.98(1.2)	0.98(1.2)	2.9	0.98(1.2)	1.1	1.00(0.95)	1.00(0.95)	1.1	1.00(0.95)	1.6	0.36	0.36	1.6	0.36
0.8	0.9	1.7	0.98(0.61)	0.98(0.61)	1.7	0.98(0.61)	0.6	1.00(0.49)	1.00(0.49)	0.6	1.00(0.49)	0.89	0.17	0.17	0.89	0.17
Hete D | $T = 100{,}000$																
0.5	0.25	4.2	0.91(1.8)	0.91(1.8)	4.2	0.92(1.8)	1.5	0.87(0.89)	0.87(0.89)	1.5	0.87(0.89)	2.1	0.61	0.61	2.1	0.61
0.5	hetero	3.3	0.93(0.93)	0.93(0.93)	3.3	0.93(0.92)	1.2	0.91(0.46)	0.91(0.46)	1.2	0.91(0.46)	1.8	0.31	0.31	1.8	0.31
0.5	0.9	2.1	0.95(0.49)	0.95(0.49)	2.1	0.95(0.49)	0.82	0.94(0.21)	0.94(0.21)	0.82	0.94(0.21)	1.1	0.17	0.17	1.1	0.17
Hete D | $T = 200{,}000$																
0.5	0.25	3.3	0.93(1.3)	0.93(1.3)	3.3	0.93(1.3)	1.2	0.87(0.6)	0.87(0.6)	1.2	0.87(0.6)	1.6	0.43	0.43	1.6	0.43
0.5	hetero	2.5	0.94(0.69)	0.94(0.69)	2.5	0.94(0.69)	0.96	0.95(0.29)	0.95(0.29)	0.96	0.95(0.29)	1.3	0.21	0.21	1.3	0.21
0.5	0.9	1.7	0.93(0.38)	0.93(0.38)	1.7	0.93(0.38)	0.63	0.94(0.16)	0.93(0.16)	0.63	0.93(0.16)	0.84	0.11	0.11	0.84	0.11

Table S.7: Empirical coverage probabilities at the 95% nominal level (mean absolute errors ↓ ($\times 10^{-3}$)) under the heterogeneous-location design for C_1 and C_5, with MAEs only for $P(1/5)$. The number of nodes is fixed at $K = 50$. Data for node k are generated from $\mathcal{N}(\mu_k, 1)$ with deterministic locations μ_k equally spaced on $[-1, 1]$. "hetero" indicates node-specific truthful response rates r_k equally spaced on $[0.25, 0.9]$. For the terminal spatial estimator, we report only the MAEs.

		C1					C5					$P(1/5)$ (MAE)				
τ	r_k	$\widehat{Q}^{s}$	$\widehat{Q}^{st}$	$\widehat{Q}^{t}$	$\widehat{Q}^{\mathrm{cent},s}$	$\widehat{Q}^{\mathrm{cent},t(st)}$	$\widehat{Q}^{s}$	$\widehat{Q}^{st}$	$\widehat{Q}^{t}$	$\widehat{Q}^{\mathrm{cent},s}$	$\widehat{Q}^{\mathrm{cent},t(st)}$	$\widehat{Q}^{s}$	$\widehat{Q}^{st}$	$\widehat{Q}^{t}$	$\widehat{Q}^{\mathrm{cent},s}$	$\widehat{Q}^{\mathrm{cent},t(st)}$
Hete L | $T = 100{,}000$																
0.3	0.25	4.8	0.99(4.5)	0.99(4.5)	4.8	0.99(4.4)	1.7	1.00(5.5)	1.00(5.5)	1.7	1.00(5.4)	2.5	1.3	1.4	2.5	1.2
0.3	hetero	3.5	0.97(1.8)	0.97(1.9)	3.5	0.98(1.7)	1.4	1.00(1.6)	1.00(1.6)	1.4	1.00(1.4)	1.9	0.81	1.5	1.9	0.51
0.3	0.9	2.2	0.98(1.1)	0.98(1.3)	2.2	0.98(0.92)	0.8	1.00(0.98)	1.00(1)	0.81	1.00(0.77)	1.1	0.93	2.2	1.1	0.24
0.5	0.25	4.2	0.91(2.1)	0.91(2.1)	4.2	0.91(2.1)	1.6	0.88(0.97)	0.87(0.98)	1.6	0.88(0.97)	2.4	0.75	0.96	2.4	0.75
0.5	hetero	3.5	0.93(1.1)	0.91(1.3)	3.5	0.93(1.1)	1.3	0.91(0.52)	0.88(0.64)	1.3	0.91(0.52)	1.8	0.37	1.6	1.8	0.38
0.5	0.9	2	0.93(0.55)	0.84(1.1)	2	0.93(0.55)	0.77	0.94(0.23)	0.85(0.7)	0.76	0.94(0.23)	1.1	0.21	2.4	1.1	0.21
0.8	0.25	5.4	1.00(8.4)	1.00(8.4)	5.4	1.00(8.3)	1.9	1.00(12)	1.00(12)	1.9	1.00(11)	2.6	2.4	2.4	2.6	2.1
0.8	hetero	4.2	1.00(2.9)	1.00(2.9)	4.2	1.00(2.7)	1.4	1.00(2.9)	1.00(2.9)	1.4	1.00(2.7)	2.1	1.2	1.5	2.1	0.73
0.8	0.9	2.1	1.00(1.7)	1.00(1.8)	2.1	1.00(1.4)	0.89	1.00(1.8)	1.00(1.8)	0.88	1.00(1.5)	1.2	1.3	1.9	1.2	0.38
Hete L | $T = 200{,}000$																
0.3	0.25	3.7	0.97(2.5)	0.97(2.5)	3.7	0.97(2.4)	1.4	1.00(2.8)	1.00(2.8)	1.4	1.00(2.8)	1.8	0.77	0.84	1.8	0.68
0.3	hetero	2.9	0.96(1.1)	0.96(1.2)	2.9	0.96(1.1)	1.1	1.00(0.79)	1.00(0.8)	1.1	1.00(0.71)	1.6	0.5	1	1.6	0.3
0.3	0.9	1.7	0.98(0.61)	0.98(0.77)	1.7	0.98(0.5)	0.63	1.00(0.54)	1.00(0.58)	0.63	1.00(0.4)	0.87	0.61	1.5	0.86	0.15
0.5	0.25	3.5	0.93(1.5)	0.93(1.5)	3.5	0.92(1.5)	1.3	0.90(0.65)	0.91(0.66)	1.3	0.90(0.65)	1.8	0.48	0.63	1.8	0.48
0.5	hetero	2.6	0.92(0.79)	0.92(0.87)	2.6	0.92(0.79)	1	0.94(0.35)	0.90(0.43)	1	0.94(0.35)	1.5	0.25	1.1	1.5	0.25
0.5	0.9	1.7	0.93(0.41)	0.86(0.72)	1.7	0.93(0.41)	0.62	0.93(0.17)	0.87(0.46)	0.62	0.94(0.17)	0.83	0.12	1.7	0.83	0.12
0.8	0.25	4.1	1.00(4.3)	1.00(4.3)	4.1	1.00(4.2)	1.5	1.00(5.8)	1.00(5.8)	1.5	1.00(5.7)	2.1	1.3	1.3	2.1	1.1
0.8	hetero	3.1	0.99(1.6)	0.98(1.7)	3.1	0.98(1.6)	1.2	1.00(1.5)	1.00(1.5)	1.2	1.00(1.4)	1.4	0.74	0.95	1.4	0.44
0.8	0.9	1.8	1.00(0.96)	1.00(1)	1.8	1.00(0.78)	0.64	1.00(0.91)	1.00(0.91)	0.64	1.00(0.73)	0.88	0.82	1.3	0.87	0.22

The full results show lower estimation errors at larger communication horizons in most configurations and illustrate how privacy levels and communication schedules affect estimation and, for C_1 and C_5, interval coverage.

C.2 Effect of network connectivity

We examine the effect of network connectivity using the Hete-L design with $K \in \{30, 50\}$, $\tau = 0.5$, $r = 0.9$, the C_5 schedule, and $T \in \{10^5, 2 \times 10^5\}$ communication rounds. We compare the dense cyclic graph in Section 4.1, a connected degree-3 graph, and a degree-2 cycle. Each configuration uses 500 paired replications, with the decentralized and centralized procedures sharing the same observations and randomized-response draws.

For each K, the degree-3 graph is generated by randomly pairing three stubs per node, rejecting self-loops, repeated edges, and disconnected graphs. The accepted graph is held fixed across replications. With adjacency matrix A, we use $W = I/2 + A/6$, assigning self-weight $1/2$ and weight $1/6$ to each neighbor.

In addition to ECP and MAE, we report the round-averaged RMS disagreement and the paired gaps between the decentralized and centralized estimators:

$$\bar{D}_T = \frac{1}{T}\sum_{m=1}^{T}\left\{\frac{1}{K}\sum_{k=1}^{K}(q_{k,t_m} - \bar{q}_{t_m})^2\right\}^{1/2},$$

$$\Delta_{\mathrm{cent},T} = \sqrt{\mathfrak{n}_T}\left|\widehat{Q}_T^{\mathrm{st}} - \widehat{Q}_T^{\mathrm{cent},t}\right|, \qquad \Delta_{\mathrm{term},T} = \left|\widehat{Q}_T^{\mathrm{space}} - \widehat{Q}_T^{\mathrm{cent},s}\right|.$$

Here $\mathfrak{n}_T = 5KT$. Table S.8 reports the results, with medians over replications for the disagreement and paired gaps.

Table S.8 shows that sparser communication increases disagreement and the paired gaps, while estimation errors remain close to the centralized benchmarks. For $K = 50$ and $T = 10^5$, changing from the dense graph to the degree-2 cycle increases $\bar{D}_T$ from 0.757×10^{-3} to 53.354×10^{-3}, the scaled temporal gap from 0.034 to 0.285, and the terminal gap from 0.015×10^{-3} to 0.127×10^{-3}. At $T = 2 \times 10^5$, the cycle's temporal gap decreases to 0.230 and its spatial–temporal ECP increases from 0.906 to 0.930. These results illustrate how additional observations improve estimation and inference when communication is limited to two neighbors per node.

Table S.8: Network-mixing stress test under Hete-L with $\tau = 0.5$, $r = 0.9$, and C_5.

(a) Estimation and inference

K	Graph	ρ	T ($\times 10^5$)	Spatial–temporal ECP (MAE)	Centralized temporal ECP (MAE)	Terminal MAE	Centralized terminal MAE
30	Dense	0.7150	1	0.924 (0.317)	0.924 (0.317)	1.059	1.056
			2	0.940 (0.207)	0.940 (0.207)	0.807	0.807
30	Degree-3	0.9536	1	0.918 (0.320)	0.924 (0.317)	1.055	1.056
			2	0.940 (0.208)	0.940 (0.207)	0.807	0.807
30	Degree-2 cycle	0.9891	1	0.924 (0.316)	0.924 (0.317)	1.052	1.056
			2	0.948 (0.207)	0.940 (0.207)	0.801	0.807
50	Dense	0.7301	1	0.922 (0.239)	0.922 (0.239)	0.760	0.759
			2	0.922 (0.176)	0.932 (0.177)	0.587	0.586
50	Degree-3	0.9583	1	0.922 (0.240)	0.922 (0.239)	0.760	0.759
			2	0.924 (0.176)	0.932 (0.177)	0.588	0.586
50	Degree-2 cycle	0.9961	1	0.906 (0.243)	0.922 (0.239)	0.760	0.759
			2	0.930 (0.170)	0.932 (0.177)	0.588	0.586

(b) Disagreement and paired estimator gaps

K	Graph	ρ	T ($\times 10^5$)	$10^3 \bar{D}_T$	$\Delta_{\mathrm{cent},T}$	$10^3 \Delta_{\mathrm{term},T}$
30	Dense	0.7150	1	0.713	0.032	0.018
			2	0.467	0.027	0.011
30	Degree-3	0.9536	1	1.186	0.121	0.022
			2	0.779	0.108	0.014
30	Degree-2 cycle	0.9891	1	22.003	0.164	0.099
			2	14.727	0.148	0.058
50	Dense	0.7301	1	0.757	0.034	0.015
			2	0.496	0.030	0.009
50	Degree-3	0.9583	1	1.856	0.107	0.023
			2	1.220	0.110	0.013
50	Degree-2 cycle	0.9961	1	53.354	0.285	0.127
			2	36.678	0.230	0.080

Notes. Results use 500 paired replications. The horizons $T = 10^5$ and $T = 2 \times 10^5$ are communication rounds. ECP is reported with MAE in parentheses; all MAEs are multiplied by 10^3. Terminal refers to the terminal spatial estimator. Panel (b) reports medians over replications: $\bar{D}_T$ and $\Delta_{\mathrm{term},T}$ are multiplied by 10^3, while $\Delta_{\mathrm{cent},T}$ retains its inferential scaling as defined in the text.

C.3 Comparison with baseline methods

We compare the proposed node-wise temporal estimator (PROP) with five baseline methods. The local-only LDP method (LO) estimates a quantile separately at each node, and the local-then-consensus method (LTC) averages the local LDP quantile estimates through consensus. The centralized LDP method (CENT) uses exact weighted aggregation and broadcast at communication times. The non-private decentralized method (NP) sets $r_k = 1$ while retaining the graph and communication schedule. The histogram method (HIST) estimates local bin frequencies using the one-bit randomizer and decoder of Acharya & Sun (2019), averages them through consensus, and inverts the resulting global CDF.

The comparison uses the Homo and Hete-L designs in Section 4.1, with $K = 10$, $\tau = 0.3$, $r = 0.9$ ($\epsilon = 2.9444$), the C_1 and C_5 schedules, and $T_{\mathrm{comm}} \in \{10^5, 2 \times 10^5\}$. Table S.9 evaluates every method at the common mixture quantile $Q^\star$. Under heterogeneity, the local quantile targeted by LO and the weighted average $\sum_k p_k F_k^{-1}(\tau)$ targeted by LTC generally differ from this target; PROP and HIST estimate the mixture quantile.

For each schedule and horizon, all methods use the same fresh-record budget and private methods use $r = 0.9$. PROP, LTC, and NP also share the degree-6 graph and communication rounds, with 6.0M and 12.0M directed peer-scalar transmissions at the two horizons. LO uses no communication, CENT uses server aggregation, and HIST exchanges multi-bin summaries.

HIST uses a public grid of 31 bins: 29 equal-width bins on $[-4, 4]$ and one overflow bin in each tail. Its one-bit randomizer uses a 32-dimensional Hadamard representation. On the same $K = 10$, degree-6 graph, average consensus runs for 3,125 and 6,250 vector rounds at $T_{\mathrm{comm}} = 10^5$ and 2×10^5, respectively, under both C_1 and C_5.

In the Homo design, LTC performs comparably to PROP, as the local and global quantile targets coincide. In the Hete-L design, PROP has MAE between 0.65×10^{-3} and 1.68×10^{-3} and ECP between 0.952 and 0.994, with performance close to CENT and NP. The errors of LTC and LO are substantially larger, while HIST's MAE is about three to five times that of PROP. These comparisons demonstrate the value of estimating the mixture quantile directly when node distributions are heterogeneous.

Table S.9: Baseline comparison under Homo and Hete-L with $K = 10$ and $\tau = 0.3$.

(a) Communication schedule C_1

T_{comm} ($\times 10^5$)	Method	Homo ECP (MAE)	Homo CI length	Hete-L ECP (MAE)	Hete-L CI length
1	PROP	0.952 (1.27)	8.50	0.952 (1.68)	10.51
	LTC	0.968 (1.20)	8.43	0.000 (107.58)	8.14
	HIST	– (7.88)	–	– (7.37)	–
	LO	0.942 (3.77)	24.01	0.092 (555.04)	24.40
	CENT	0.952 (1.27)	8.51	0.958 (1.52)	10.10
	NP	0.962 (1.10)	7.50	0.950 (1.51)	9.20
2	PROP	0.937 (0.89)	5.56	0.957 (1.05)	6.63
	LTC	0.930 (0.90)	5.58	0.000 (107.29)	5.56
	HIST	– (6.10)	–	– (5.43)	–
	LO	0.940 (2.73)	16.89	0.084 (555.33)	17.03
	CENT	0.938 (0.89)	5.56	0.958 (0.95)	6.38
	NP	0.962 (0.79)	5.07	0.945 (0.97)	6.21

(b) Communication schedule C_5

T_{comm} ($\times 10^5$)	Method	Homo ECP (MAE)	Homo CI length	Hete-L ECP (MAE)	Hete-L CI length
1	PROP	0.980 (0.71)	5.08	0.994 (1.12)	7.79
	LTC	0.970 (0.66)	4.97	0.000 (107.72)	5.02
	HIST	– (5.60)	–	– (3.74)	–
	LO	0.941 (1.72)	10.98	0.075 (555.04)	11.84
	CENT	0.980 (0.71)	5.08	0.994 (0.93)	7.28
	NP	0.982 (0.61)	4.64	0.985 (0.98)	6.70
2	PROP	0.963 (0.43)	3.05	0.973 (0.65)	4.28
	LTC	0.956 (0.44)	3.02	0.000 (107.42)	3.04
	HIST	– (5.02)	–	– (3.18)	–
	LO	0.940 (1.21)	7.60	0.051 (555.29)	8.01
	CENT	0.962 (0.43)	3.05	0.976 (0.54)	3.97
	NP	0.966 (0.39)	2.64	0.970 (0.59)	3.84

Notes. Private methods use $r = 0.9$; NP is the non-private oracle. ECP is evaluated at the common mixture-quantile target $Q^\star$, with MAE in parentheses. MAE and CI length are multiplied by 10^3. HIST reports MAE only; "– denotes an unreported quantity. Methods, targets, and communication budgets are specified in the text.

C.4 DIRECTED AND ASYNCHRONOUS COMMUNICATION

We investigate sensitivity to directional communication and unequal node update rates using the three data-generating designs in Section 4.1. Table S.10 reports the Homo and Hete-D results with $K = 10$, $\tau = 0.5$, $r = 0.9$, the C_5 schedule, and $T \in \{10^5, 2 \times 10^5\}$.

For directed communication, we use the degree-6 cyclic graph with self-weight $1/2$, weights $(1 + \xi)/12$ for the three clockwise neighbors, and $(1 - \xi)/12$ for the three counterclockwise neighbors. We vary $\xi \in \{0, 0.1, 0.2, 0.3\}$, with $\xi = 0$ giving symmetric mixing.

For asynchronous communication, nodes have independent exponential clocks with maximum-to-minimum rate ratio $R \in \{1, 1.25, 1.5, 2\}$. Clock rates are equally spaced on the logarithmic scale from $R^{-1/2}$ to $R^{1/2}$ and randomly permuted over nodes in each replication. At each activation, the active node processes five fresh records, selects one of its six cyclic neighbors uniformly, and performs an atomic equal-weight pairwise average with no communication delay.

Each node initiates exactly T blocks and then stops its clock, but remains available as a communication partner. The run therefore uses KT activation events and $5KT$ records. We record the node iterates after every K events, yielding T common observation times for temporal averaging and interval construction. Node-wise ECP and MAE are averaged equally across nodes and then across replications.

We report estimation accuracy, interval coverage, and the scaled gaps

$$G_{\mathrm{ST}} = \sqrt{\mathfrak{n}_T}|\widehat{Q}_T^{\mathrm{st}} - \widehat{Q}_T^{\mathrm{cent},t}|,$$
$$G_{\mathrm{node}} = \max_k \sqrt{\mathfrak{n}_T}|\widehat{Q}_{k,T}^{\mathrm{time}} - \widehat{Q}_T^{\mathrm{st}}|, \qquad \mathfrak{n}_T = 5KT.$$

For directed communication, MAE, coverage, and both gaps remain stable across the values of ξ in both designs. Under asynchronous communication, point-estimation error also changes little, whereas the node-to-network gap increases with the clock-rate ratio. At $T = 2 \times 10^5$, node-wise MAE ranges from 0.355×10^{-3} to 0.359×10^{-3} under Hete-D and from 0.336×10^{-3} to 0.339×10^{-3} under Homo. In the Homo design, increasing R from 1 to 2 changes node-wise ECP from 0.930 to 0.909 at $T = 10^5$ and from 0.932 to 0.937 at $T = 2 \times 10^5$. Thus, unequal update rates are more visible in the node-to-network comparison than in point-estimation error.

Table S.10: Directed and asynchronous communication under Hete-D and Homo with $K = 10$, $\tau = 0.5$, $r = 0.9$, and C_5.

Parameter	T ($\times 10^5$)	Node-wise temporal ECP (MAE)	Spatial–temporal ECP	G_{ST}	G_{node}
(a) Directed communication: Hete-D; parameter ξ					
0	1	0.941 (0.495)	0.942	0.017	0.007
	2	0.935 (0.359)	0.936	0.016	0.005
0.1	1	0.941 (0.495)	0.940	0.017	0.007
	2	0.936 (0.359)	0.936	0.014	0.005
0.2	1	0.940 (0.495)	0.940	0.018	0.007
	2	0.936 (0.360)	0.936	0.015	0.005
0.3	1	0.940 (0.495)	0.940	0.016	0.007
	2	0.936 (0.360)	0.936	0.014	0.005
(b) Directed communication: Homo; parameter ξ					
0	1	0.928 (0.503)	0.928	0.019	0.007
	2	0.935 (0.336)	0.936	0.015	0.005
0.1	1	0.924 (0.503)	0.924	0.019	0.007
	2	0.934 (0.336)	0.934	0.016	0.005
0.2	1	0.924 (0.503)	0.924	0.019	0.007
	2	0.934 (0.336)	0.934	0.015	0.005
0.3	1	0.925 (0.503)	0.926	0.019	0.007
	2	0.936 (0.336)	0.936	0.015	0.005
(c) Asynchronous communication: Hete-D; parameter R					
1	1	0.940 (0.496)	0.940	0.020	0.007
	2	0.938 (0.359)	0.938	0.016	0.005
1.25	1	0.937 (0.497)	0.936	0.019	0.008
	2	0.935 (0.356)	0.936	0.016	0.007
1.5	1	0.928 (0.500)	0.928	0.019	0.012
	2	0.940 (0.355)	0.940	0.016	0.012
2	1	0.934 (0.506)	0.934	0.019	0.020
	2	0.934 (0.356)	0.934	0.014	0.022
(d) Asynchronous communication: Homo; parameter R					
1	1	0.930 (0.503)	0.930	0.021	0.007
	2	0.932 (0.336)	0.932	0.017	0.005
1.25	1	0.913 (0.509)	0.914	0.020	0.009
	2	0.926 (0.336)	0.926	0.016	0.007
1.5	1	0.908 (0.512)	0.908	0.019	0.012
	2	0.929 (0.338)	0.928	0.017	0.012
2	1	0.909 (0.516)	0.910	0.019	0.019
	2	0.937 (0.339)	0.934	0.016	0.024

Notes. For directed communication, ξ controls directional asymmetry; for asynchronous communication, R is the maximum/minimum clock-rate ratio. Node-wise entries are ECP (MAE), with MAE multiplied by 10^3. The transfer gaps G_{ST} and G_{node} are medians over replications. Each horizon is displayed separately; the experimental protocol is specified in the text.

C.5 TRANSIENT EDGE DROPOUT

We examine the effect of intermittent communication using the Hete-L design with $K = 10$, $\tau = 0.8$, $r = 0.9$, and the C_1 schedule. At each round, every undirected edge is independently dropped with probability $p_{\text{drop}} \in \{0, 0.1, 0.2, 0.3\}$, and dropped edges make no exchange. We use 500 paired replications with common observations, privacy randomization, and nominal record budgets across dropout levels. The horizons are $T \in \{10^5, 2 \times 10^5, 4 \times 10^5\}$ records per node, with $\mathfrak{n}_T = KT$.

The base graph connects each of the $K = 10$ nodes to cyclic offsets $\pm 1, \pm 2, \pm 3$. Each retained undirected edge has weight $1/12$ in both directions, and node k has self-weight $1 - d_{k,m}/12$, where $d_{k,m}$ is its number of retained neighbors in round m. Dropped-edge weights are thus assigned to the diagonal.

Table S.11 reports ECP, MAE, the disagreement $\bar{D}_T$ defined in Section C.2, and the scaled gaps

$$G_{\text{ST},T} = \sqrt{\mathfrak{n}_T}|\widehat{Q}_T^{\text{st}} - \widehat{Q}_T^{\text{cent},t}|,$$
$$G_{\text{node},T}^{\max} = \max_k \sqrt{\mathfrak{n}_T}|\widehat{Q}_{k,T}^{\text{time}} - \widehat{Q}_T^{\text{st}}|.$$

Table S.11: Transient edge dropout under Hete-L with $K = 10$, $\tau = 0.8$, $r = 0.9$, and C_1.

p_{drop}	Retention	T ($\times 10^5$)	$10^3 \bar{D}_T$	$G_{\text{ST},T}$	$G_{\text{node},T}^{\max}$	Spatial–temporal ECP (MAE)	Centralized MAE
0	1.000	1	0.6871	0.245	0.672	0.988 (2.047)	1.911
		2	0.4541	0.230	0.628	0.970 (1.276)	1.213
		4	0.2976	0.212	0.583	0.952 (0.844)	0.814
0.1	0.900	1	0.8106	0.290	0.819	0.988 (2.073)	1.911
		2	0.5358	0.273	0.766	0.966 (1.289)	1.213
		4	0.3511	0.251	0.710	0.948 (0.850)	0.814
0.2	0.800	1	0.9567	0.343	1.002	0.986 (2.111)	1.911
		2	0.6324	0.325	0.938	0.968 (1.306)	1.213
		4	0.4145	0.300	0.870	0.952 (0.858)	0.814
0.3	0.700	1	1.1360	0.420	1.238	0.986 (2.160)	1.911
		2	0.7508	0.393	1.158	0.968 (1.330)	1.213
		4	0.4921	0.364	1.074	0.950 (0.871)	0.814

Notes. Results use 500 paired replications. Under C_1, T is both the number of communication rounds and the number of records per node. Retention is the reported realized edge-retention summary for each dropout level. The disagreement and transfer gaps are medians; $\bar{D}_T$ and both MAEs are multiplied by 10^3. Spatial–temporal entries are ECP (MAE).

Increasing the dropout probability raises disagreement and both scaled gaps. At $T = 10^5$, increasing p_{drop} from 0 to 0.3 changes $10^3 \bar{D}_T$ from 0.6871 to 1.1360, $G_{\text{ST},T}$ from 0.245 to 0.420, and $G_{\text{node},T}^{\max}$ from 0.672 to 1.238. Both gaps decrease with the horizon at every dropout level. At $T = 4 \times 10^5$, spatial–temporal ECP ranges from 0.948 to 0.952 and MAE from 0.844×10^{-3} to 0.871×10^{-3}, compared with 0.814×10^{-3} for the centralized estimator. The results show accurate estimation and coverage close to the nominal level at the largest horizon despite intermittent edge availability.

C.6 SENSITIVITY TO TEMPORAL DISTRIBUTION SHIFTS

We examine sensitivity to changes in the node distributions using the Hete-L design with $K = 10$, $\tau = 0.5$, $r = 0.9$, the C_5 schedule, and $T \in \{10^5, 2 \times 10^5\}$. We compare two linear shifts with the stationary design (S). Let μ_k denote the Hete-L location and let $u_t \in [0, 1]$ increase linearly over time. The target-preserving (TP) shift is

$$X_{k,t} \sim \mathcal{N}((1 + hu_t)\mu_k, 1), \qquad h \in \{.05, .10, .20\},$$

which preserves the network-wide median at zero. The common moving-target (CM) shift is

$$X_{k,t} \sim \mathcal{N}(\mu_k + au_t, 1), \qquad a \in \{.0025, .005, .01, .02\},$$

with instantaneous, time-averaged, and terminal targets au_t, $Q_{\mathrm{avg}} = a/2$, and $Q_T = a$.

The spatial–temporal and node-wise temporal estimators are evaluated at Q_{avg}, and the terminal spatial estimator at Q_T. Table S.12 reports their errors and interval inclusion frequencies, together with

$$G_{\mathrm{ST},T} = \sqrt{\mathfrak{n}_T}|\widehat{Q}_T^{\mathrm{st}} - \widehat{Q}_T^{\mathrm{cent},t}|,$$
$$G_{\mathrm{node},T}^{\mathrm{rms}} = \sqrt{\mathfrak{n}_T}\left\{K^{-1}\sum_k(\widehat{Q}_{k,T}^{\mathrm{time}} - \widehat{Q}_T^{\mathrm{st}})^2\right\}^{1/2},$$
$$G_{\mathrm{term},T} = |\widehat{Q}_T^{\mathrm{space}} - \widehat{Q}_T^{\mathrm{cent},s}|, \qquad \mathfrak{n}_T = 5KT.$$

Under TP shifts, the results remain close to the stationary case. At $T = 2 \times 10^5$, spatial–temporal MAE ranges from 0.393×10^{-3} to 0.397×10^{-3} and inclusion from 0.934 to 0.942, compared with 0.395×10^{-3} and 0.940 under stationarity. Node-wise and terminal errors, as well as the transfer gaps, are also similar. These findings indicate stable empirical performance under the target-preserving distribution changes.

For CM shifts, performance at $a = 0.0025$ remains close to the stationary case, while stronger drift substantially widens the temporal intervals. At $a = 0.02$ and $T = 2 \times 10^5$, the spatial–temporal and node-wise interval lengths reach 23.762×10^{-3} and 23.766×10^{-3} and no longer decrease with the horizon. The paired decentralized–centralized gaps remain close to their stationary values, while the interval widths respond to the moving target.

For time-varying local CDFs $F_{k,t}$, the blockwise score drift at a fixed reference $Q^\star$ is

$$b_{m,T}(Q^\star) = E_m^{-1}\sum_{t=t_{m-1}+1}^{t_m}\sum_k p_k\{\tau_k - F_{k,t}(Q^\star)\}.$$

Extending inference to evolving quantiles requires analyzing this drift on the inferential scale together with uniform regularity. Tracking a terminal quantile also calls for suitable step sizes or rolling windows.

Table S.12: Temporal-shift sensitivity under Hete-L with $K = 10$, $\tau = 0.5$, $r = 0.9$, and C_5.

(a) Temporal estimation and interval inclusion

Shift	Level	T ($\times 10^5$)	Spatial–temporal Inclusion (MAE)	Spatial–temporal CI length	Node-wise temporal Inclusion (MAE)	Node-wise temporal CI length
S	–	1	0.924 (0.564)	3.370	0.900 (0.705)	3.658
		2	0.940 (0.395)	2.429	0.910 (0.482)	2.607
TP	0.05	1	0.924 (0.561)	3.395	0.902 (0.709)	3.679
		2	0.942 (0.393)	2.442	0.914 (0.484)	2.619
TP	0.10	1	0.934 (0.561)	3.419	0.903 (0.712)	3.698
		2	0.934 (0.396)	2.437	0.910 (0.489)	2.614
TP	0.20	1	0.940 (0.574)	3.440	0.897 (0.726)	3.713
		2	0.936 (0.397)	2.460	0.906 (0.494)	2.630
CM	0.0025	1	0.958 (0.566)	4.235	0.917 (0.707)	4.443
		2	0.972 (0.400)	3.537	0.948 (0.485)	3.645
CM	0.005	1	0.980 (0.573)	6.385	0.973 (0.713)	6.483
		2	0.998 (0.411)	6.102	0.996 (0.492)	6.133
CM	0.010	1	1.000 (0.606)	11.931	1.000 (0.740)	11.953
		2	1.000 (0.446)	11.915	1.000 (0.518)	11.924
CM	0.020	1	1.000 (0.743)	23.604	1.000 (0.846)	23.613
		2	1.000 (0.562)	23.762	1.000 (0.619)	23.766

(b) Transfer gaps and terminal estimation

Shift	Level	T ($\times 10^5$)	$G_{\mathrm{ST},T}$	$G^{\mathrm{rms}}_{\mathrm{node},T}$	Terminal MAE	$10^3 G_{\mathrm{term},T}$
S	–	1	0.034	1.154	1.683	0.027
		2	0.026	1.069	1.423	0.016
TP	0.05	1	0.038	1.168	1.705	0.030
		2	0.025	1.082	1.426	0.017
TP	0.10	1	0.034	1.182	1.704	0.030
		2	0.029	1.094	1.433	0.019
TP	0.20	1	0.032	1.209	1.712	0.030
		2	0.027	1.119	1.439	0.019
CM	0.0025	1	0.033	1.154	1.686	0.029
		2	0.026	1.069	1.425	0.019
CM	0.005	1	0.034	1.154	1.697	0.028
		2	0.027	1.069	1.435	0.018
CM	0.010	1	0.032	1.154	1.730	0.028
		2	0.028	1.069	1.467	0.016
CM	0.020	1	0.031	1.154	1.860	0.026
		2	0.024	1.069	1.580	0.018

Notes. S, TP, and CM denote stationary, target-preserving, and common moving-target designs, respectively. Level is h for TP and a for CM; "– marks the stationary control. Inclusion is ECP for S and empirical inclusion otherwise. Node-wise summaries are node averages. MAE and CI length are multiplied by 10^3; transfer gaps are medians, with $G_{\mathrm{term},T}$ multiplied by 10^3. Temporal estimators are evaluated at Q_{avg} and the terminal spatial estimator at Q_T.

C.7 ADDITIONAL REAL-DATA RESULTS

We complement the seven-node analysis in Section 4.3 with degree-2 cycle networks of $K = 14, 21,$ and 28 nodes, where each node communicates with its two neighbors. The $K = 7$ control uses a degree-4 graph. The corresponding mixing parameters are $\rho = 0.9505, 0.9778,$ and 0.9875, compared with 0.6002 for $K = 7$. We use 500 repeated-sampling replications, $r \in \{0.6, 0.9\}$, and the fixed-population target $Q_{\mathrm{pop}} = 80{,}000$ USD at $\tau = 0.8$.

The seven regional groups are fixed before repeated sampling. Within each region, a fixed random permutation is split into two or three nearly equal parts for $K = 14$ or 21, respectively. For $K = 28$, each of the two regional halves is further split in two. Each replication samples training records without replacement within each node population, leaving the remaining records as a disjoint reference sample. The per-node training sizes are $n = 10{,}696, 5{,}348, 3{,}565,$ and $2{,}674$ for $K = 7, 14, 21,$ and 28, respectively; for $K = 7$, equal-sized samples from the two regional halves are pooled. The target weights are $p_k = N_k/204{,}309$, where N_k is the full node-population size, and coverage is evaluated against the fixed full-population quantile.

Estimation uses $z = \log(1 + \mathrm{salary}/1000) - 4$, with all iterates initialized at zero. Estimates and interval endpoints are transformed back to USD. For the reported C_5 results, the first $w = \lfloor 0.05n \rfloor$ records per node use one-record blocks, followed by five-record blocks and a final shorter block when needed. Warm-up iterates are included in the temporal averages. The resulting communication horizons are $T = 2{,}567, 1{,}284, 856,$ and 642 for the four network sizes, respectively. With t_m denoting the cumulative per-node record count, the block step size is

$$\gamma_m = \frac{a}{\max\{1, t_m - w + 1\}^{0.51} + 100},$$

where $a = 12$ for $r = 0.6$ and $a = 72$ for $r = 0.9$. Each within-block update uses γ_m / E_m.

Table S.13 reports the full results. In addition to estimation error and coverage, we report the median absolute paired gap

$$\Delta_{\mathrm{c}}^{\mathrm{abs}} = \mathrm{median}_{b=1,\ldots,500} \left| \widehat{Q}_{T,b}^{\mathrm{st}} - \widehat{Q}_{T,b}^{\mathrm{cent},t} \right|,$$

and D_{med}, the median over replications of $\bar{D}_T$ defined in Section C.2.

Table S.13: Government Salary: repeated-sampling performance under C_5 at $\tau = 0.8$.

K	ρ	T	Spatial–temporal ECP (MAE)	CI length	Terminal MAE	Bias	Δ_c^{abs}	D_{med}
(a) Truthful-response probability $r = 0.6$								
7	0.6002	2,567	0.944 (640)	3,596	2,308	-484	120	0.0196
14	0.9505	1,284	0.910 (928)	4,122	1,986	-877	226	0.0452
21	0.9778	856	0.956 (1,138)	5,142	1,601	-1,121	201	0.0520
28	0.9875	642	0.988 (1,263)	6,207	1,560	-1,248	236	0.0556
(b) Truthful-response probability $r = 0.9$								
7	0.6002	2,567	0.972 (291)	3,665	3,908	-96	173	0.0635
14	0.9505	1,284	0.952 (409)	4,184	4,122	-340	177	0.1106
21	0.9778	856	0.936 (425)	3,925	2,648	-370	148	0.1196
28	0.9875	642	0.938 (414)	3,947	3,287	-356	168	0.1243

Notes. The fixed population has $Q_{\mathrm{pop}} = 80{,}000$ USD; results use 500 replications. Entries in the spatial–temporal column are ECP (MAE). MAE, CI length, bias, and Δ_c^{abs} are in USD. Terminal denotes the terminal spatial estimator. Δ_c^{abs} is the median absolute gap between the spatial–temporal and centralized temporal estimates; D_{med} is the median round-averaged RMS disagreement.

Across the four network configurations, spatial–temporal MAE ranges from USD 291 to USD 425 and ECP from 0.936 to 0.972 at $r = 0.9$. Under stronger privacy, $r = 0.6$, MAE ranges from USD 640 to USD 1,263. Temporal averaging reduces MAE relative to the terminal spatial estimator in every configuration and both privacy regimes.

On the degree-2 cycles, communication with only two neighbors per node yields MAE between USD 409 and USD 425 and ECP between 0.936 and 0.952 at $r = 0.9$. The median absolute gaps

from the centralized temporal estimator are between USD 148 and USD 177. The disagreement D_{med} is higher on these cycles, ranging from 0.1106 to 0.1243, compared with 0.0635 for $K = 7$. The results illustrate accurate estimation of the common salary threshold using locally privatized data and limited neighbor communication.