UNIVERSITÉ DU QUÉBEC EN OUTAOUAIS

DÉPARTEMENT DES SCIENCES ADMINISTRATIVES

WORKING PAPER NO. 28

Beyond One-Step Distributed Maximum Likelihood: Local Bias Cancellation, Global Synchronization, and the Price of Inexact Curvature

A joint order–round frontier, its sharpness in an exact model, and measured communication costs

Simon-Pierre Boucher

Département des sciences administratives
Université du Québec en Outaouais

simon-pierre.boucher@uqo.ca

Gatineau – Pavillon Alexandre-Taché
283, boulevard Alexandre-Taché
Gatineau, Québec, Canada J9A 1L8

October 2026

Version 1.0

Abstract

Distributed maximum likelihood estimation trades statistical efficiency for communication. We study a two-phase protocol in which each of K workers corrects the bias of its local maximum likelihood estimator by an r-th order split-sample extrapolation over nested sub-blocks, the coordinator averages the corrected estimates, and q globally synchronized Newton corrections follow. For fixed-dimensional regular likelihoods with N observations we prove a joint order–round frontier: first-order equivalence to the centralized estimator holds whenever $K = o(N^{1-1/(2^{q+1}r)})$, so that one extra local extrapolation level and one extra synchronization round are exchangeable in the exponent, and the half-block jackknife followed by one Newton step tolerates $K = o(N^{7/8})$ partitions, against the $K = O(N^{1/2})$ regime stated for the plain-average one-step estimator. In the exponential model we derive exact bias, variance and Newton identities, show that the exponent is sharp for the construction with limiting shift $-\theta_0(\beta_r c^r)^{2^q}$ when $K \sim cN^{\alpha(r,q)}$, and characterize the Newton basin of attraction. A perturbation theorem for inexact curvature recovers the known single-machine-Hessian frontiers as special cases and gives the number of bits to which transmitted derivatives may be quantized at each local order. Monte Carlo experiments with paired replications (exponential, logistic and Poisson likelihoods), an exact-model boundary study, and timed multi-node benchmarks on an Apple-silicon cluster quantify the finite-sample variance cost of higher-order corrections and the wall-clock exchange rate between local fits and synchronization rounds. The components we combine are known; the joint frontier, its sharpness and the bit-precision result are new. No universal communication lower bound is claimed.

Keywords: distributed estimation, maximum likelihood, split-sample jackknife, one-step Newton estimator, communication complexity, asymptotic equivalence, Monte Carlo

JEL Classification: C13, C15, C55, C63

Provenance and AI-use disclosure

This working paper develops an exploratory research note of 28 October 2026 whose theoretical content was drafted with the help of a large language model; the note is frozen with checksums in the replication package (`references/initial_pamle/`). The present version was prepared by the author with Claude (Anthropic) as a research and programming assistant: the literature audit, proofs, code, experiments and text were produced in that collaboration and reviewed by the author, who bears sole responsibility for the content. No funding is claimed for this project.

1. Introduction

When N observations are stored on K machines, the centralized maximum likelihood estimator (MLE) $\widehat{\theta}_N$ is the natural benchmark: under regularity it attains the inverse Fisher information as asymptotic variance, but computing it requires either moving the data or iterating many synchronized passes over the network. The simplest distributed alternative fits the model on each block of $m = N/K$ observations and averages the K local estimates. Averaging reduces variance but not the $O(1/m)$ bias of a nonlinear local estimator, so the average is first-order equivalent to $\widehat{\theta}_N$ only when $K = o(N^{1/2})$ (Zhang et al., 2013; Battey et al., 2018; Rosenblatt and Nadler, 2016). Two ideas relax this barrier. One is *local*: debias each local estimate before averaging, by subsampling (Zhang et al., 2013), by an analytic bias formula (Gu and Chen, 2023; Yin, 2022) or by a jackknife (Wu et al., 2023). The other is *global*: broadcast the average and add one Newton step computed from aggregated scores and Hessians (Huang and Huo, 2019), or several rounds of gradient-only surrogate updates (Jordan et al., 2019; Fan et al., 2023). Each idea has its own cost: a local correction multiplies the number of local fits, a global correction adds a synchronization round and, for dense Hessians, an uplink of order p^2 numbers per worker.

This paper studies the two ideas jointly. Each worker computes the MLE on its block and on nested sub-blocks of sizes $m/s_2, \ldots, m/s_r$, combines them with Richardson-type weights that cancel the first $r-1$ terms of the bias expansion while preserving the first-order influence function exactly, and uploads the corrected estimate. The coordinator averages, broadcasts the average and applies q Newton corrections from aggregated derivatives. Our question is quantitative: for fixed parameter dimension p, how fast may K grow with N while the resulting estimator remains first-order equivalent to $\widehat{\theta}_N$, how does the answer depend on (r,q), when is the answer sharp, and what happens when the curvature used by the coordinator is inexact, stale or quantized?

Contributions.

1. *A joint order–round frontier* (Theorem 1). Under a uniform stochastic expansion with a bias

expansion of order r and local Newton regularity, first-order equivalence to $\widehat{\theta}_N$ holds whenever

$$K = o\big(N^{\alpha(r,q)}\big), \qquad \alpha(r,q) = 1 - \frac{1}{2^{q+1}\, r}. \tag{1}$$

The exponent depends on (r,q) only through the *effective depth* $D = 2^q r$, so that doubling the local order and adding one synchronization round are exchangeable. The half-block jackknife with one Newton step gives $K = o(N^{7/8})$; the plain average with one step gives $3/4$; the quarter-block extrapolation with one step gives $11/12$. The proof keeps the centered random remainder of the local corrections, of order $(m\sqrt{K})^{-1}$, which cannot be dropped even though its mean vanishes.

2. *Exact analysis and sharpness* (Theorem 2). In the exponential model we derive the exact bias and variance of the corrected local estimators, an exact Newton error identity, the basin of attraction of the Newton map, and the boundary behaviour: if $K/N^{\alpha(r,q)} \to c \in (0,\infty)$ then

$$\sqrt{N}\left(\theta_r^{(q)} - \widehat{\theta}_N\right) \xrightarrow{\;p\;} -\theta_0 \left(\beta_r c^r\right)^{2^q}, \qquad q \geq 1, \tag{2}$$

where β_r is an explicit rational constant ($\beta_2 = -2$, so that the half-block jackknife with one step has the limiting shift $-4\theta_0 c^4$), and the estimator is not even $\sqrt{N}$-close to $\widehat{\theta}_N$ when $K/N^{\alpha(r,q)} \to \infty$. The sufficient exponent is therefore necessary for this construction.

3. *The price of inexact curvature* (Theorem 3). If the Newton step uses a Hessian within δ_N of $H_N(\widehat{\theta}_N)$ in operator norm, the admissible exponent changes in a computable way:

$$\alpha = 1 - \frac{1}{2(r+q)} \text{ (stale averaged local Hessians)}, \qquad \alpha = 1 - \frac{1}{2r+q} \text{ (single-machine Hessian)}, \tag{3}$$

the second of which reproduces, for $r = 1$, the $(t+1)/(t+2)$ rounds frontier of Fan et al. (2023) and the $2/3$ exponent of Jordan et al. (2019). Relative quantization of the transmitted score and Hessian to b bits costs nothing asymptotically once

$$b \;\geq\; \left(\tfrac{1}{2} - r(1-\alpha)\right) \log_2 N + \omega(1) \qquad \text{when } K \sim N^{\alpha}, \tag{4}$$

so that higher local order permits coarser derivative uplinks.

4. *Evidence.* Paired-replication Monte Carlo experiments for exponential, logistic and Poisson likelihoods with N up to 2^{22}, an exact-model boundary study, ablations over split patterns, Newton rounds and derivative quantization, and timed benchmarks of the full coordinator–worker protocol on up to 16 Apple-silicon hosts quantify the finite-sample variance cost of higher-order corrections (which can reverse the asymptotic ranking at moderate m), the fraction of local fits that fail, interval coverage, transmitted bytes per round and the measured exchange

rate between local fits and synchronization rounds.

What is not claimed. The two building blocks are known: split-sample extrapolation is the generalized jackknife of Quenouille (1956) and Schucany et al. (1971), applied to panels by Dhaene and Jochmans (2015); one-step distributed Newton estimation is due to Huang and Huo (2019). "Debias locally, then average" exists in several forms (Zhang et al., 2013; Gu and Chen, 2023; Wu et al., 2023; Yin, 2022), and "average, then correct globally" in several others (Huang and Huo, 2019; Zhu et al., 2021; Jordan et al., 2019; Fan et al., 2023). What we add is the joint analysis, the sharpness statement, and the inexact-curvature frontier. We claim no minimax or communication lower bound: in exponential families the pooled sufficient statistic recovers $\widehat{\theta}_N$ exactly with one round of p numbers per worker, and the exponential model is used here as an analytically transparent laboratory, not as an application. All results are for fixed p and i.i.d. data; constants depend on p and the theory does not transfer to growing dimension without new arguments. Wall-clock statements are measurements on specific hardware and are reported as such.

The paper is organized as follows. Section 2 positions the results against the closest theorems in the literature. Section 3 defines the likelihood environment, the protocol and the assumptions. Section 4 develops the general local correction. Section 5 states the frontier theorem and its inferential corollary, Section 6 the exact exponential analysis, Section 7 the inexact-curvature theorem, Section 8 the cost model. Sections 9 and 10 report the Monte Carlo and benchmark evidence, Section 11 the extensions, Section 12 the limitations. Proofs are in Sections A and B; replication material is described in Section D.

2. Related literature and positioning

Table 1 compares the closest results at the level of their theorems; the comparison was made on the full texts, with the caveats listed in the replication package when a published version could not be accessed.

One-shot averaging. Zhang et al. (2013) bound the mean squared error of the average by $O(1/N + 1/m^2)$ (their Theorem 1), which matches the $O(1/N)$ centralized rate only for $K \lesssim N^{1/2}$; Rosenblatt and Nadler (2016) show that the second-order bias of the average equals m times that of $\widehat{\theta}_N$ (their Theorem 2); Battey et al. (2018) obtain the same $\sqrt{N}$ barrier in sparse high dimension; Volgushev et al. (2019) prove for quantile regression that $K = o(N^{1/2})$ is also *necessary* for the average to be first-order equivalent, by exhibiting a model in which the remainder is exactly of order $1/m$ (their Theorem 3.2), and Shang and Cheng (2017) establish a sharp "maximal number of machines" for smoothing splines. Lin and Xi (2011) weight the local estimates by local Jacobians evaluated at the

local estimates; their equivalence theorem requires $K \ll N^{1/4}$.

Local debiasing. The subsampled-average estimator of Zhang et al. (2013) combines each local fit with a fit on a random subsample of size $\lceil \rho m \rceil$ as $(\bar{\theta}_1 - \rho\bar{\theta}_2)/(1-\rho)$, a one-level Richardson extrapolation that removes the $1/m^2$ term of the MSE at the price of a leading constant $(2+3\rho)/(1-\rho)^2$ and a tuning rate $\rho \to 0$ (their Theorem 4); they leave the use of several subsamples open. Gu and Chen (2023) estimate the leading bias analytically on each machine with sample splitting and weight by local information, obtaining the centralized χ^2 limit for $K = o(m^2)$, i.e. $K = o(N^{2/3})$ (their Theorems 4–6), with a smaller constant than the subsampled average. Wu et al. (2023) apply the delete-one jackknife within each subsample and average, with bias $O(m^{-2}) + O(N^{-1})$ for smooth functions of means (their Theorem 2). The split-sample weights themselves go back to Quenouille (1949, 1956) and Schucany et al. (1971); Dhaene and Jochmans (2015) develop the general higher-order split-panel jackknife, with weights solving a Vandermonde-type system (their equations 5.1–5.3), prove that the first-order variance is unchanged and that partitions into $\{2,3,\dots,h+1\}$ parts minimize the inflation of the uncancelled bias terms, and leave the higher-order variance open. Hahn et al. (2024) show that in i.i.d. cross-sections the split-sample bias estimate is inconsistent and that split-sample corrections carry twice the higher-order variance of the delete-one jackknife (their Propositions 5–7), a cost that our exact variance formula and Monte Carlo results make explicit for nested splits.

Global corrections. Huang and Huo (2019) define the distributed one-step estimator: workers send scores and Hessians evaluated at the broadcast average and the coordinator takes one Newton step. Their Theorem 3.4 proves asymptotic equivalence to the full-sample M-estimator when $k = O(\sqrt{N})$; their Theorem 3.5 bounds the MSE by $2\operatorname{tr}\Sigma/N + O(N^{-2}) + O(k^4 N^{-4})$, a bound that is itself $o(N^{-1})$-exact whenever $k = o(N^{3/4})$, although that exponent is not stated. Zhu et al. (2021) prove that one *local* Newton step per machine followed by information-weighted averaging is efficient iff $m/N^{1/4} \to \infty$, "allowing $K = O(N^{3/4})$" (their Theorem 2), and iterate the scheme. Jordan et al. (2019) replace the aggregated Hessian by that of the first machine and send gradients only; their Theorem 3.1 bounds the one-round error by $C(\|\bar{\theta} - \widehat{\theta}_N\|^2 + (N^{-1/2} + m^{-1/2})\|\bar{\theta} - \widehat{\theta}_N\|)$, a contraction by $m^{-1/2}$ per round. Fan et al. (2023) add a proximal term and state the explicit rounds frontier (their remark after Theorem 4.1): from the plain average, t single-machine-Hessian steps match the global minimizer when $m^{t+2} \gg N \cdot \text{polylog}$, i.e. $K = o(N^{(t+1)/(t+2)})$ up to logarithms, and one averaged local-Hessian step in the linear model allows $K = o(N^{3/4})$. Shamir et al. (2014) and Wang et al. (2017) are the optimization-side precursors; Chen et al. (2022) avoid the Hessian entirely by stochastic first-order iterations on one machine. Yin (2022) removes the quadratic terms of the one-step estimator analytically and obtains efficiency iff $K = o(N^{2/3})$. Reviews by Gao et al.

(2022) and Li et al. (2024) confirm that no joint order–round frontier was available at their writing.

Gap. None of these papers combines an extrapolated local start with aggregated-Hessian Newton rounds, states a frontier indexed by both the local order and the number of rounds, proves sharpness for a corrected construction, or expresses the effect of inexact or quantized Hessians as a change in the admissible exponent with a bit count. These are the four results of Sections 5 to 7; the frontiers of Huang and Huo (2019), Zhu et al. (2021), Jordan et al. (2019) and Fan et al. (2023) appear as the cells $(1,1)$, $(1,1)$ with local Hessians, and $(1,t)$ with a single-machine Hessian of our tables.

3. Environment, protocol and assumptions

3.1. Likelihood, partitions and nested splits

Let $X_1,\dots,X_N$ be i.i.d. with density $f(\cdot;\theta_0)$ with respect to a dominating measure, θ_0 interior to an open $\Theta\subset\mathbb{R}^p$ and p fixed. For a sub-sample $S\subset\{1,\dots,N\}$ write $L_S(\theta)=|S|^{-1}\sum_{i\in S}\log f(X_i;\theta)$, $g_S=\nabla L_S$, $H_S=\nabla^2 L_S$ and let $\widehat{\theta}_S$ denote a measurable maximizer of L_S (the MLE on S). The centralized MLE is $\widehat{\theta}_N=\widehat{\theta}_{\{1,\dots,N\}}$. The index set is partitioned, independently of the data, into K blocks $D_1,\dots,D_K$ of equal size $m=N/K$; Section 11 treats unequal sizes. Let $1=s_1<s_2<\cdots<s_r$ be integers, each dividing m. Level j of worker k consists of the s_j disjoint contiguous sub-blocks of D_k of size m/s_j; $T_{s_j,k}$ denotes the average of the s_j MLEs computed on them, so that $T_{1,k}=\widehat{\theta}_{D_k}$. The default pattern is $s=(1,2,\dots,2^{r-1})$; the pattern $(1,2,3)$ of Dhaene and Jochmans (2015) is also used. The collection of all sub-blocks at all levels, all blocks and the full sample is the *sub-sample family* $\mathscr{S}_N$. Asymptotics are along a triangular array $N\to\infty$ with $K=K_N$ and $m\to\infty$.

3.2. The protocol

Algorithm 1 and Figure 1 describe the estimator. Phase 1 is local: worker k fits the $\sum_j s_j$ MLEs of its levels and uploads the p-vector

$$J_{r,k}=\sum_{j=1}^{r} w_j\, T_{s_j,k}, \tag{5}$$

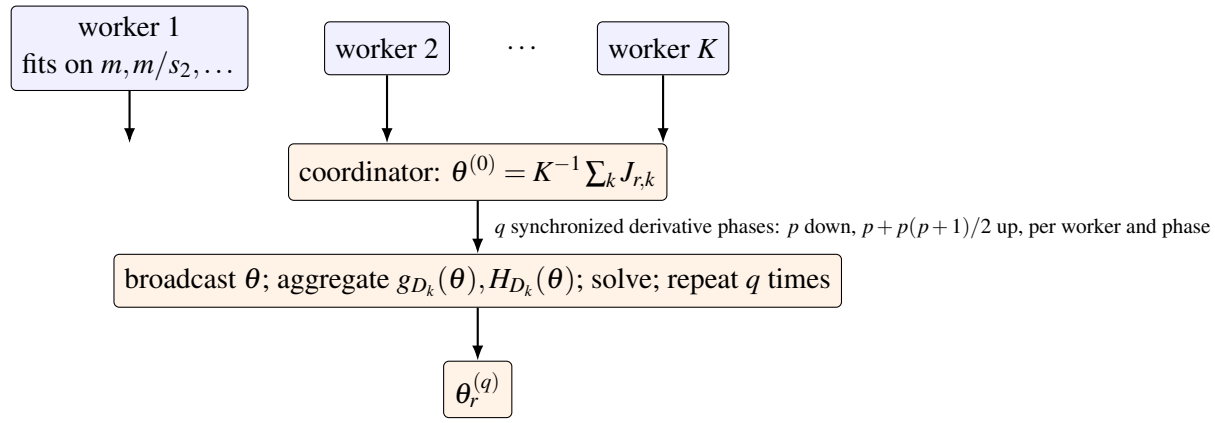
where $w=(w_1,\dots,w_r)$ solves the Vandermonde system of Section 4. The coordinator forms the pilot $\theta_r^{(0)}=K^{-1}\sum_k J_{r,k}$. Each of the q global corrections is one synchronized derivative phase: the coordinator broadcasts the current iterate θ, each worker uploads $g_{D_k}(\theta)$ and the $p(p+1)/2$ distinct entries of $H_{D_k}(\theta)$, the coordinator averages and solves

$$\theta\leftarrow\theta-H_N(\theta)^{-1}g_N(\theta), \tag{6}$$

Algorithm 1 Order-r local correction with q global Newton corrections

Require: blocks $D_1, \ldots, D_K$ on K workers; splits $s_1 < \cdots < s_r$ with weights w; rounds q

1: **Phase 1 (workers, in parallel).** For each level j, fit the MLE on each of the s_j sub-blocks; set $T_{s_j,k}$ to their average; upload $J_{r,k} = \sum_j w_j T_{s_j,k}$ (and a validity flag) ▷ p numbers up
2: **Coordinator.** $\theta \leftarrow K^{-1} \sum_k J_{r,k}$
3: **for** $t = 1, \ldots, q$ **do** ▷ synchronized derivative phase
4: broadcast θ ▷ p numbers down to each worker
5: each worker uploads $g_{D_k}(\theta)$ and the upper triangle of $H_{D_k}(\theta)$ ▷ $p + p(p+1)/2$ numbers up
6: $g \leftarrow K^{-1} \sum_k g_{D_k}(\theta)$, $H \leftarrow K^{-1} \sum_k H_{D_k}(\theta)$; solve $H\Delta = g$; $\theta \leftarrow \theta - \Delta$
7: **end for**
8: **return** $\theta_r^{(q)} = \theta$, with the condition number of $-H$ and $\|g\|$ as diagnostics



Phase 1 uplink: p numbers per worker (plus $p(p+1)/2$ if the local Hessian is shipped for the stale-curvature variant). Arrows are logical exchanges, not measured latency; Section 10 reports measured times and bytes.

Figure 1. Coordinator–worker protocol with exact uplink and downlink accounting.

by a linear solve, never by an explicit inverse. We write $\theta_r^{(q)}$ for the result; $(r,q) = (1,0)$ is the plain average, $(1,1)$ the one-step estimator of Huang and Huo (2019), $(2,1)$ the half-block jackknife with one step. Because the derivatives are evaluated at a common data-dependent point, the q derivative phases cannot be merged with phase 1: the protocol has $1+q$ synchronized exchange phases, $K(1+2q)$ messages, $K(p+q(p+p(p+1)/2))$ uplinked numbers and Kqp downlinked numbers, before headers and serialization (Table 4). Variants in which the coordinator uses a stale, single-machine or quantized Hessian are analysed in Section 7.

3.3. Assumptions

The theorems are conditional on the following assumptions, stated for the sub-sample family $\mathscr{S}_N$. They are high-level: they are implied by standard primitive conditions for regular parametric models (smoothness, moment and identification conditions of the type used by Zhang et al., 2013 and Huang and Huo, 2019, together with the validity of the formal bias expansion, Bhattacharya and

Ghosh, 1978), and they are verified exactly in the exponential model of Section 6.

Assumption 1 (Uniform stochastic expansion)**.** There exist a measurable ψ with $\mathbb{E}\psi(X) = 0$ and $\Sigma_\psi := \operatorname{Var}\psi(X)$ finite and positive definite, an integer n_0 and a constant $C < \infty$ such that for every $S \in \mathscr{S}_N$ with $|S| \geq n_0$,

$$\widehat{\theta}_S - \theta_0 = |S|^{-1} \sum_{i \in S} \psi(X_i) + R_S, \qquad \mathbb{E}\|R_S\|^2 \leq C\,|S|^{-2}. \tag{7}$$

For a correctly specified regular likelihood $\psi = I(\theta_0)^{-1}\nabla_\theta \log f(X;\theta_0)$ and $\Sigma_\psi = I(\theta_0)^{-1}$; for a misspecified model θ_0 is the pseudo-true value, $\psi = -\bar{H}(\theta_0)^{-1}\nabla_\theta \log f(X;\theta_0)$ and Σ_ψ is the sandwich (Section 11). The same ψ serves every nested sub-sample; this is the only property of the influence function used.

Assumption 2 (Bias expansion of order r)**.** There exist vectors $b_1, \ldots, b_{r-1}$ and $C_b < \infty$ with

$$\mathbb{E}R_S = \sum_{a=1}^{r-1} b_a\,|S|^{-a} + \rho_{|S|}, \qquad \|\rho_n\| \leq C_b\, n^{-r}, \qquad S \in \mathscr{S}_N,\ |S| \geq n_0. \tag{8}$$

For $r = 1$ the assumption only asks $\|\mathbb{E}R_S\| \leq C_b|S|^{-1}$, which follows from Assumption 1 by Jensen's inequality. The existence of the expansion for $r \geq 2$ is not implied by consistency or asymptotic normality; it requires smoothness of the likelihood to order $r+1$ and moment conditions, and it fails for estimators with non-smooth criteria.

Assumption 3 (Local Newton regularity)**.** With probability tending to one, $\widehat{\theta}_N$ is an interior zero of g_N; there are a fixed convex neighbourhood U of θ_0 and random variables $M_N, \Lambda_N = O_p(1)$ such that $\sup_{\theta \in U}\|H_N(\theta)^{-1}\| \leq M_N$ and $\|H_N(u) - H_N(v)\| \leq \Lambda_N\|u - v\|$ for $u, v \in U$. Moreover $\sqrt{N}(\widehat{\theta}_N - \theta_0) \Rightarrow \mathscr{N}(0, \Sigma_\psi)$.

Assumption 4 (Observation-level curvature smoothness; used only in Theorem 3(a))**.** $h(x;\theta) := \nabla^2_\theta \log f(x;\theta)$ is continuously differentiable on U with $\|\nabla h(x;\theta) - \nabla h(x;\theta')\| \leq L(x)\|\theta - \theta'\|$, $\mathbb{E}L(X)^2 < \infty$, $\mathbb{E}\|h(X;\theta_0)\|^2 < \infty$, $\mathbb{E}\|\nabla h(X;\theta_0)\|^2 < \infty$, and $\mathbb{E}\|\widehat{\theta}_{D_k} - \theta_0\|^4 \leq C_4 m^{-2}$.

Remark 1. The assumptions fix p, require independent identically distributed blocks and well-behaved local optimizers. Constants such as C, M_N and Λ_N depend on p and on the parametrization; none of the exponents below should be read as valid when p grows with N. The condition $m \to \infty$ is essential: all statements are about a triangular array in which the number of partitions and the local sample size grow together, not about a fixed data set with a maximal number of machines.

4. The general local correction

4.1. Weights

Let V be the $r \times r$ Vandermonde matrix with entries $V_{aj} = s_j^a$, $a = 0, \dots, r-1$. The weights are the unique solution of

$$\sum_{j=1}^{r} w_j = 1, \qquad \sum_{j=1}^{r} w_j s_j^a = 0, \quad a = 1, \dots, r-1, \tag{9}$$

i.e. $w = V^{-1} e_1$. For $s = (1,2)$, $w = (2,-1)$; for $(1,2,4)$, $w = (8/3, -2, 1/3)$; for $(1,2,4,8)$, $w = (64/21, -8/3, 2/3, -1/21)$; for $(1,2,3)$, $w = (3,-3,1)$. Two constants summarize the cost of a pattern: the *leading bias coefficient* $\beta_r := \sum_j w_j s_j^r$, which multiplies the first uncancelled bias term, and the *remainder amplification* $A_r := \sum_j |w_j| s_j$, which multiplies the L_2 bound of the stochastic remainder. Table 2 lists them; the weights are computed exactly in rational arithmetic in the replication code and the moment constraints are verified to machine precision.

4.2. Influence preservation, bias cancellation and the random remainder

Write $U_S := |S|^{-1} \sum_{i \in S} \psi(X_i)$ and $U_k := U_{D_k}$. The sub-blocks of level j partition D_k exactly, so the average of their influence terms is U_k whatever j; since the weights sum to one, the influence term of $J_{r,k}$ is U_k (Lemma 1). Nothing else about ψ is used, which is why the construction applies verbatim to pseudo-MLEs. Collecting the remainders of the $\sum_j s_j$ fits into

$$R_k^J := \sum_{j=1}^{r} w_j s_j^{-1} \sum_{\ell=1}^{s_j} R_{S_{j\ell,k}}, \tag{10}$$

we have $J_{r,k} - \theta_0 = U_k + R_k^J$. The mean of R_k^J is deterministic and, by (8) and (9),

$$\mathbb{E} R_k^J = \sum_{a=1}^{r-1} \frac{b_a}{m^a} \sum_{j=1}^{r} w_j s_j^a + \sum_{j=1}^{r} w_j \rho_{m/s_j} = O(m^{-r}) \tag{11}$$

(Lemma 2), the first sum vanishing identically. Its L_2 norm is at most $\sqrt{C} A_r / m$ by the triangle inequality (Lemma 3); the nested fits of one worker are dependent, and no independence is used at this step. Across workers the blocks are independent, so the centered average has second moment at most $C A_r^2 / (K m^2)$ (Lemma 4). Finally the full-sample expansion $\widehat{\theta}_N - \theta_0 = U_N + R_N$ with $U_N = K^{-1} \sum_k U_k$ cancels the common influence term exactly:

$$\theta_r^{(0)} - \widehat{\theta}_N = \mathbb{E} R_k^J + K^{-1} \sum_k (R_k^J - \mathbb{E} R_k^J) - R_N = O_p\Big(m^{-r} + \frac{1}{m\sqrt{K}} + \frac{1}{N} \Big), \tag{12}$$

(Lemma 5). The middle term is the centered random part of the local corrections. Its mean is zero but its magnitude, $\sqrt{K}/N$, exceeds N^{-1} by the factor $\sqrt{K}$ and must be carried through every subsequent step; a proof that drops it is incomplete even where the final exponent happens to be unaffected.

4.3. Conditioning and the choice of pattern

Higher order costs conditioning. The weights grow geometrically with r for the doubling pattern ($\max_j|w_j| \approx 3.0, 3.2$ for $r = 4, 5$, with Vandermonde condition numbers 3×10^3 and 5×10^5), and the finest sub-block has $m/2^{r-1}$ observations, which for $p > 1$ quickly reaches the regime where local MLEs fail to exist (separation in logistic regression) or have no moments (sub-blocks of two or three exponential observations). The patterns of Dhaene and Jochmans (2015), $(1,2,3)$ and $(1,2,3,4)$, use coarser finest blocks and smaller $|\beta_r|$ but larger A_r (12 against 8 for $r = 3$): they trade bias constants against remainder amplification. Which trade is better is a finite-sample question answered in Section 9. Adaptive selection of r from the data is not covered by the theory: a data-driven order changes the estimator and its limit, and the theorems below are for fixed (r, q); see Section 12.

5. Main results: the order–round frontier

5.1. One Newton step contracts quadratically

The global correction is a classical one-step estimator (Le Cam, 1956; Bickel, 1975; van der Vaart, 1998), with the difference that the comparison point is $\widehat{\theta}_N$ rather than θ_0. Under Assumption 3, for any $\vartheta_N \xrightarrow{p} \theta_0$ the Newton map $\mathcal{N}_N(\vartheta) := \vartheta - H_N(\vartheta)^{-1} g_N(\vartheta)$ satisfies

$$\|\mathcal{N}_N(\vartheta_N) - \widehat{\theta}_N\| \le \tfrac{1}{2} M_N \Lambda_N \|\vartheta_N - \widehat{\theta}_N\|^2 \tag{13}$$

on an event of probability tending to one (Lemma 6), and q iterates satisfy $\|\theta^{(q)} - \widehat{\theta}_N\| \le (\frac{1}{2} M_N \Lambda_N)^{2^q - 1} \|\theta^{(0)} - \widehat{\theta}_N\|^{2^q}$ (Lemma 7). Applied to the pilot bound (12), repeated squaring acts on all three terms:

$$\|\theta_r^{(q)} - \widehat{\theta}_N\| = O_p\Big(\Big[m^{-r} + \frac{1}{m\sqrt{K}} + \frac{1}{N} \Big]^{2^q} \Big). \tag{14}$$

5.2. The frontier

Theorem 1 (Order–round frontier)**.** *Let Assumptions 1 to 3 hold with order $r \ge 1$ and fixed $q \ge 0$, let $m \to \infty$ and*

$$K = o\big(N^{\alpha(r,q)}\big), \qquad \alpha(r,q) = 1 - \frac{1}{2^{q+1} r}.$$

Then $\sqrt{N}\,(\theta_r^{(q)} - \widehat{\theta}_N) \xrightarrow{p} 0$*, and consequently* $\sqrt{N}(\theta_r^{(q)} - \theta_0) \Rightarrow \mathcal{N}(0, \Sigma_\psi)$.

Proof sketch. By (14), $\sqrt{N}\|\theta_r^{(q)} - \widehat{\theta}_N\|$ is bounded by a sum of three terms, which with $m = N/K$ read

$$\sqrt{N}\,m^{-r2^q} = N^{1/2}\Big(\frac{K}{N}\Big)^{r2^q}, \qquad \sqrt{N}\Big(\frac{\sqrt{K}}{N}\Big)^{2^q} = N^{1/2}K^{2^q-1}N^{-2^q}, \qquad \sqrt{N}N^{-2^q}. \tag{15}$$

The first vanishes if and only if $K^{r2^q} = o(N^{r2^q-1/2})$, that is $K = o(N^{\alpha(r,q)})$. The second equals $\sqrt{K/N}$ for $q = 0$ and vanishes because $\alpha(r,0) < 1$; for $q \geq 1$ it is at most $N^{1/2-2^{q-1}} \leq N^{-1/2}$ whenever $K \leq N$. The third always vanishes. Consistency of the pilot, needed to apply (13), follows from (12) and $m \to \infty$. The limit law follows from Assumption 3 and Slutsky's lemma. Section A gives the details, including the event on which all iterates stay in U. □

Table 3 displays the frontier. Three readings are useful. First, the cells $(1,0)$, $(1,1)$ and $(1,1)$-with-local-curvature are the regimes of Zhang et al. (2013), Huang and Huo (2019) (whose stated condition is the weaker $K = O(N^{1/2})$, while the k^4N^{-4} remainder of their MSE bound already implies $3/4$) and Zhu et al. (2021); the table extends them in both directions. Second, the exponent is symmetric in the exchange "one more round $\leftrightarrow$ double the local order": $(2,1)$, $(1,2)$ and $(4,0)$ all give $7/8$. Which of the three is cheaper is a question about the hardware, not about statistics, and is the subject of Sections 8 and 10. Third, the dominant restriction is always the deterministic bias term; the random remainder and the full-sample term are never binding for $q \geq 1$ and bind only through $K = o(N)$ for $q = 0$. This is a statement about rates, not about finite samples: in Section 9 the random term is visible as variance inflation at moderate m.

5.3. Inference

Corollary 1 (Inferential validity)**.** *Under the conditions of Theorem 1, let* $\widehat{\Sigma}_N \xrightarrow{p} \Sigma_\psi$*. Then for every coordinate* j *and level* $1-\tau$ *the interval*

$$\theta_r^{(q,j)} \pm z_{1-\tau/2}\sqrt{\widehat{\Sigma}_{N,jj}/N} \tag{16}$$

has asymptotic coverage $1-\tau$*. Under correct specification* $\widehat{\Sigma}_N = -H_N(\theta_r^{(q)})^{-1}$*, which is available at no extra communication from the last derivative phase (for* $q \geq 1$*); under misspecification the sandwich* $H_N^{-1}\widehat{J}_N H_N^{-1}$ *with* $\widehat{J}_N = K^{-1}\sum_k |D_k|^{-1}\sum_{i\in D_k}\nabla \log f(X_i;\theta)\nabla \log f(X_i;\theta)^\top$ *requires one additional* $p(p+1)/2$ *upload per worker in that phase.*

The corollary says nothing about estimators outside the frontier: when $K/N^{\alpha(r,q)}$ does not vanish, the point estimator has a non-centered limit (Theorem 2) and the interval under-covers, as the coverage tables of Section 9 show for the plain average and the one-step estimator at small m.

5.4. What the theorem does not say

The theorem is a sufficient statement under high-level assumptions; it gives neither the finite-sample behaviour (the constants C, C_b, M_N, Λ_N, A_r and β_r all matter at moderate m) nor any comparison between protocols at equal cost. The exponent cannot be improved *for this construction* (Theorem 2), but other constructions can do better: in exponential families the pooled sufficient statistic reaches $\widehat{\theta}_N$ exactly for every $K \le N$, and multi-round schemes with $q \to \infty$ remove the restriction on K entirely at the price of unbounded synchronization (Jordan et al., 2019; Fan et al., 2023).

6. Exact analysis in the exponential model and sharpness

Let $X_i \sim \mathrm{Exp}(\theta_0)$ in the rate parametrization. The average log-likelihood, score and Hessian on n observations are

$$L_n(\theta) = \log\theta - \theta\,\bar{X}_n, \qquad g_n(\theta) = \frac{1}{\theta} - \bar{X}_n, \qquad H_n(\theta) = -\frac{1}{\theta^2}, \qquad \widehat{\theta}_n = \frac{1}{\bar{X}_n}. \tag{17}$$

The model is trivial for distributed inference, since $\sum_{i\in D_k} X_i$ is a sufficient statistic that any worker can upload, and the Hessian does not depend on the data, so that single-machine curvature is exact. It is used here because every quantity in Sections 4 and 5 has a closed form, which makes it possible to check the assumptions, the constants and the boundary behaviour rather than assume them. The statements below are proved in Section B from Gamma and Beta moments, verified symbolically, and their algebraic content is machine-checked in Lean 4 (Section D). Write $Y = \theta_0\bar{X}_n$, so that $nY \sim \mathrm{Gamma}(n,1)$, and

$$\psi(x) = \theta_0 - \theta_0^2\,x, \qquad R_n = \widehat{\theta}_n - \theta_0 - \frac{1}{n}\sum_{i=1}^{n}\psi(X_i) = \frac{(1-\theta_0\bar{X}_n)^2}{\bar{X}_n}. \tag{18}$$

Theorem 2 (Exact exponential identities and boundary behaviour)**.** *(i)* Moments. *For $n > 2$,*

$$\mathbb{E}\widehat{\theta}_n = \theta_0\frac{n}{n-1}, \qquad \mathrm{Var}\,\widehat{\theta}_n = \theta_0^2\frac{n^2}{(n-1)^2(n-2)}, \tag{19}$$

$$\mathbb{E}R_n = \frac{\theta_0}{n-1} = \theta_0\sum_{a\ge 1} n^{-a}, \qquad \mathbb{E}R_n^2 = \theta_0^2\frac{3n+2}{n(n-1)(n-2)}, \tag{20}$$

and $\sup_{n\ge 3} n^2\,\mathbb{E}R_n^2/\theta_0^2 = 33/2$. Hence Assumptions 1 and 2 hold for every r with $n_0 = 3$, $C = 33\theta_0^2/2$ and $b_a = \theta_0$ for all a.

(ii) Exact bias of the corrected estimators. *For any pattern with $m/s_r > 1$,*

$$\mathbb{E}J_{r,k} - \theta_0 = \theta_0 \sum_{j=1}^{r} \frac{w_j s_j}{m - s_j} = \beta_r \theta_0 m^{-r}\big(1 + O(m^{-1})\big), \tag{21}$$

in particular

$$\mathbb{E}J_{2,k} - \theta_0 = -\frac{2\theta_0}{(m-1)(m-2)}, \qquad \mathbb{E}J_{3,k} - \theta_0 = \frac{8\theta_0}{(m-1)(m-2)(m-4)}. \tag{22}$$

(iii) Exact variance of the half-block correction. *For $m > 4$,*

$$\operatorname{Var} J_{2,k} = \theta_0^2 \frac{m^2\,(m^2 - 6m + 17)}{(m-4)(m-2)^2(m-1)^2} = \theta_0^2\Big(\frac{1}{m} + \frac{4}{m^2} + O(m^{-3})\Big), \tag{23}$$

the same first two terms as $\operatorname{Var}\widehat{\theta}_m$: the variance cost of the half-block correction is $O(m^{-3})$ in this model.

(iv) Newton identities and basin. *For every $x > 0$,*

$$\begin{aligned} \mathcal{N}_N(x) &= 2x - x^2\bar{X}_N, \qquad \mathcal{N}_N(x) - \widehat{\theta}_N = -\frac{(x - \widehat{\theta}_N)^2}{\widehat{\theta}_N}, \\ \mathcal{N}_N^q(x) - \widehat{\theta}_N &= -\frac{(x - \widehat{\theta}_N)^{2^q}}{\widehat{\theta}_N^{2^q - 1}}. \end{aligned} \tag{24}$$

The iteration converges to $\widehat{\theta}_N$ if and only if $0 < x < 2\widehat{\theta}_N$; it diverges to $-\infty$ when $x < 0$ or $x > 2\widehat{\theta}_N$, and $x = 2\widehat{\theta}_N$ is mapped to 0.

(v) Boundary. *If $m \to \infty$ and $K/N^{\alpha(r,q)} \to c \in (0,\infty)$,*

$$\sqrt{N}\,\big(\theta_r^{(q)} - \widehat{\theta}_N\big) \xrightarrow{p} \begin{cases} -\theta_0\,(\beta_r c^r)^{2^q}, & q \geq 1, \\ \beta_r \theta_0 c^r, & q = 0, \end{cases} \tag{25}$$

so that for $(r,q) = (2,1)$ the limit is $-4\theta_0 c^4$ and $\sqrt{N}(\theta_2^{(1)} - \theta_0) \Rightarrow \mathcal{N}(-4\theta_0 c^4, \theta_0^2)$.

(vi) Necessity. *If $m \to \infty$ and $K/N^{\alpha(r,q)} \to \infty$, then $\sqrt{N}\,|\theta_r^{(q)} - \widehat{\theta}_N| \xrightarrow{p} \infty$.*

Items (v)–(vi) show that the exponent of Theorem 1 is sharp for the construction in this model: the small-o cannot be weakened to O, and beyond the frontier the estimator is not $\sqrt{N}$-consistent for $\widehat{\theta}_N$. Three qualifications matter. The boundary regime $m \asymp N^{1/(2^{q+1}r)}$ is extremely slow to reach—for $(2,1)$, $m \asymp N^{1/8}$, so that $m \leq 16$ for all $N \leq 2^{24}$—and the finite-m bias in (22) exceeds its leading term $-2\theta_0 m^{-2}$ by the factor $m^2/((m-1)(m-2))$, which equals 1.5 at $m = 6$; Section 9 therefore reports both the limit and an exact finite-m prediction. The sharpness is for this estimator

family: it is not a lower bound over all distributed procedures, as pooling $\sum_i X_i$ shows. And the signs alternate, $\beta_2 = -2 < 0 < \beta_3 = 8$, so the higher-order pilot over-corrects where the half-block pilot under-corrects, one reason why finite-sample accuracy is not monotone in r.

Item (iv) also gives an explicit failure gallery. Because the weights have negative entries, $J_{r,k}$ can be negative when a sub-block average is small; the pilot then lies outside the basin $(0, 2\widehat{\theta}_N)$ with positive probability for every finite m, and the Newton iteration diverges rather than converges. The probability is negligible for $m \geq 32$ and visible for $m \leq 8$ in Section 9, where such replications are reported, not discarded. A production implementation would damp or project the step; the theory above describes the undamped map.

7. The price of inexact curvature

The exact protocol uploads $p + p(p+1)/2$ numbers per worker in every derivative phase. Three cheaper variants are common in practice: a *stale* Hessian, equal to the average of the local Hessians evaluated at the local MLEs and shipped once together with the pilot, after which every round costs p numbers; a *single-machine* Hessian evaluated at the current iterate on one worker, as in the surrogate-likelihood approach of Jordan et al. (2019); and a *quantized* exact Hessian. The following perturbation theorem covers all three and inexact scores as well. Let ϑ be the current iterate, $e = \vartheta - \widehat{\theta}_N$, and let the coordinator compute

$$\tilde{\vartheta} = \vartheta - \widehat{H}^{-1}\,\widehat{g}, \qquad \widehat{g} = g_N(\vartheta) + \delta_g. \tag{26}$$

Theorem 3 (Inexact Newton steps)**.** *Let Assumptions 1 to 3 hold,* $\vartheta \xrightarrow{p} \theta_0$, $\|\widehat{H} - H_N(\widehat{\theta}_N)\| = O_p(\delta_N)$ *with* $\delta_N \to 0$, *and* $\|\widehat{H}^{-1}\| = O_p(1)$. *Then*

$$\|\tilde{\vartheta} - \widehat{\theta}_N\| = O_p\big(\|e\|^2 + \delta_N\|e\| + \|\delta_g\|\big). \tag{27}$$

With q such steps from the order-r pilot, each with a Hessian error of order δ_N,

$$\|\theta_r^{(q)} - \widehat{\theta}_N\| = O_p\big(\|e_0\|\,\delta_N^q + \|e_0\|^{2^q}\big), \qquad e_0 = \theta_r^{(0)} - \widehat{\theta}_N, \tag{28}$$

and $\sqrt{N}(\theta_r^{(q)} - \widehat{\theta}_N) \xrightarrow{p} 0$ *provided* $\sqrt{N}\,m^{-r}\delta_N^q \to 0$ *together with the conditions of Theorem 1. In particular:*

(a) Stale averaged local curvature. *Under Assumption 4,* $\widehat{H} = K^{-1}\sum_k H_{D_k}(\widehat{\theta}_{D_k})$ *satisfies* $\delta_N = m^{-1} + N^{-1/2}$, *and for* $q \geq 1$ *the sufficient exponent becomes*

$$\alpha_{\text{stale}}(r,q) = 1 - \frac{1}{2\,(r+q)}. \tag{29}$$

(b) Single-machine curvature. $\widehat{H} = H_{D_1}(\vartheta)$ *satisfies* $\delta_N = m^{-1/2}$ *(no third-derivative condition is needed), and the sufficient exponent becomes*

$$\alpha_{\text{single}}(r,q) = 1 - \frac{1}{2r+q}. \tag{30}$$

(c) Relative quantization. *If every entry of the transmitted score and Hessian is rounded to b significant bits* ($\varepsilon = 2^{-b}$)*, then* (27) *holds with* $\delta_N = \varepsilon$ *and* $\|\delta_g\| = O_p(\varepsilon\|e\|)$*, so that one quantized step from the order-r pilot is first-order equivalent to* $\widehat{\theta}_N$ *whenever* $\sqrt{N}\,\varepsilon\, m^{-r} \to 0$ *in addition to* $K = o(N^{\alpha(r,1)})$*. When* $K \sim N^{\alpha}$ *with* $1 - 1/(2r) \le \alpha < \alpha(r,1)$*, it suffices that*

$$b \;\ge\; \Big(\tfrac{1}{2} - r(1-\alpha)\Big)\log_2 N + \omega(1), \tag{31}$$

and for $\alpha < 1 - 1/(2r)$ *any fixed number of bits suffices, because the pilot alone is already first-order equivalent.*

The proof (Section A) is a matrix-perturbation argument around the integral form of the Newton remainder,

$$\tilde{\vartheta} - \widehat{\theta}_N = \widehat{H}^{-1}\Big[\widehat{H} - \int_0^1 H_N(\widehat{\theta}_N + te)\,\mathrm{d}t\Big]\, e \;-\; \widehat{H}^{-1}\delta_g, \tag{32}$$

whose bracket is bounded by $\|\widehat{H} - H_N(\widehat{\theta}_N)\| + \frac{1}{2}\Lambda_N\|e\|$. The theorem turns the qualitative statements in the literature—the $m^{-1/2}$ factor of Jordan et al. (2019, Theorem 3.1), the $(\beta/\lambda)^2 \asymp m^{-1}$ contraction of averaged local Hessians in Shamir et al. (2014, Theorem 2) and Fan et al. (2023)—into exponents. Table 4 tabulates them with the communication they require. Case (b) with $r = 1$ gives $\alpha = (q+1)/(q+2)$, exactly the rounds frontier that Fan et al. (2023) state for CEASE up to logarithms, and $q = 1$ gives the $2/3$ of Jordan et al. (2019); case (a) with $(r,q) = (1,1)$ gives $3/4$, the regime of Zhu et al. (2021). The stale variant has effective depth $r+q$, linear in q, against $2^q r$ for exact curvature: with stale curvature a second round is worth one more local level; with exact curvature it is worth doubling the local order.

Case (c) is, to our knowledge, new in this form. The precision needed for the derivative uplink is governed by how far the pilot already is from $\widehat{\theta}_N$: a pilot of order r with $K \sim N^{\alpha}$ has error $m^{-r} = N^{-r(1-\alpha)}$, and a relative error ε in the step costs $\varepsilon N^{-r(1-\alpha)}$, which is negligible at the $N^{-1/2}$ scale if and only if (31) holds. For the half-block jackknife at $\alpha = 0.8$ the requirement is $0.1\log_2 N$ bits, two bits of mantissa at $N = 2^{20}$, whereas the plain average at the same α would need $0.3\log_2 N = 6$ bits and is outside its own frontier ($3/4$) in any case. The mantissa bits multiply the $p(p+1)/2$ Hessian entries per worker and round, so the saving is $(64-b)/64$ of the Hessian payload. The statement is asymptotic in the same sense as Theorem 1; the ablation of Section 9 shows at which b the quantization error becomes visible at finite N.

8. The economics of synchronization

The frontier of Table 3 makes $(2,1)$, $(1,2)$ and $(4,0)$ statistically interchangeable at the level of exponents. Choosing among them is a question of cost, and the cost model below is proposed, calibrated and tested, not proved. Let $T_{\text{fit}}(n,p)$ be the elapsed time of one local MLE fit on n observations, $T_{\text{der}}(m,p)$ that of one local score-and-Hessian evaluation, λ the latency of one synchronized exchange (broadcast plus the slowest worker's reply, excluding payload), B the effective uplink bandwidth at the coordinator, b the bytes per transmitted number and $T_{\text{solve}}(p)$ the coordinator's linear solve. For the doubling pattern the local phase of order r performs 2^r-1 fits touching every local observation r times, so that with a fit cost $T_{\text{fit}}(n,p) \approx c_0(p) + c_1(p)\,n$ (fixed overhead plus a linear pass, the empirical shape of Newton fits with a bounded number of iterations),

$$T(r,q) \;\approx\; \underbrace{(2^r-1)\,c_0 + r\,c_1\,m}_{\text{local phase (slowest worker)}} \;+\; q\,\underbrace{\Big[\lambda + T_{\text{der}}(m,p) + \frac{K\,b\,(p+p(p+1)/2)}{B} + T_{\text{solve}}(p)\Big]}_{\text{one synchronized derivative phase}} \\ +\; T_{\text{sched}}(K) + T_{\text{straggle}}(K), \tag{33}$$

where the last two terms collect worker start-up and the excess of the slowest worker over the mean. Sums are critical-path sums: the local fits of different workers overlap, the q derivative phases do not.

Proposition 1 (Exchange rate between a local level and a round, under (33))**.** *Hold the effective depth $D = 2^q r$ fixed. Moving from (r,q) to $(2r, q-1)$ removes one derivative phase and adds the local levels $r+1,\dots,2r$. Under (33) the round is cheaper iff*

$$\lambda + T_{\text{der}}(m,p) + \frac{Kb(p+p(p+1)/2)}{B} + T_{\text{solve}}(p) \;<\; (2^{2r}-2^r)\,c_0 + r\,c_1 m. \tag{34}$$

The right-hand side grows linearly in m and exponentially in r; the left-hand side grows linearly in K and quadratically in p. For fixed N and small p the rule favours local levels when K is large and m is small, and favours rounds when m is large. Holding statistical accuracy fixed in finite samples requires, in addition, the variance inflation of higher local orders (Section 9) to be accounted for.

The proposition is a tautology given (33); its content is the calibration. Section 10 measures every term of (33) on the cluster—fit and derivative times by m and p, per-phase latency percentiles, bytes per phase, start-up and straggler excess—and reports where (34) lands for the designs run. We do not model overlapping of computation and communication, asynchronous updates (which change the estimator) or energy; these are listed as open in Section 12.

9. Monte Carlo evidence

9.1. Design

All experiments are paired: every estimator in a cell is computed on the same simulated sample and the same contiguous partition (observations are i.i.d., so contiguous and random partitions are equivalent in law). The designs are (i) the exponential model, $N \in \{2^{14}, 2^{16}, 2^{18}, 2^{20}, 2^{22}\}$, $m \in \{8, 16, \ldots, 1024\}$, 2000 replications per cell up to $N = 2^{20}$ and 500 at $N = 2^{22}$ (where the $m = 8$ cell was dropped for cost); (ii) logistic regression with $p \in \{3, 8\}$ (intercept and standard normal covariates, coefficients 0.3 and $\pm 0.5/(1 + 0.15j)$), $N \in \{2^{14}, \ldots, 2^{20}\}$, $m \in \{32, \ldots, 1024\}$ with $m \geq 8p$, 1000 replications, plus a stress design with $p = 20$ and $m \geq 6p$ (500 replications); (iii) Poisson regression with the same structure and half the coefficients; (iv) robustness designs: a misspecified logistic model with an omitted quadratic term $0.5(x_1^2 - 1)$ in the index, an ill-conditioned logistic design (covariate covariance with condition number 100), an exponential sample contaminated by one percent of observations drawn with ten times the scale, and logistic designs with very small blocks. Every cell records, for each estimator, the scaled gap to the centralized MLE $\sqrt{N}\|\tilde{\theta} - \widehat{\theta}_N\|$, the scaled error $\sqrt{N}\|\tilde{\theta} - \theta_0\|$, Wald intervals from the observed information at the estimate, the status of every local fit, the condition number of the aggregated Hessian and the Newton residuals. Replications in which local fits failed are kept, and the tables report both the root-mean-square gap over all replications and its median, together with the fraction of workers whose sub-block fits all converged; rows labelled "valid workers only" drop the workers with a failed fit from the pilot. In the misspecified designs the target of the coverage statistics is the pseudo-true value, estimated by the average of the centralized pseudo-MLE over replications. Monte Carlo standard errors accompany every summary in the processed files and are at most a few percent of the reported gaps. The cells were run on the author's laptop and on 18 nodes of the MacLustr cluster (Section D); the simulations are in memory and measure no communication.

9.2. Rates, the frontier, and the random remainder

Figures 2 and 3 plot the root-mean-square scaled gap against m at the largest N, and Table 11 fits log-log slopes over the failure-free cells. The ordering predicted by Table 3 holds in every design: the plain average decays like m^{-1} (slopes between -1.0 and -1.1), one step from it like m^{-2} (-2.0 to -2.1), two steps like m^{-4} (-4.1 to -4.2). The half-block pilot has slope -2 in the exponential model but only -1.2 to -1.4 in the regressions, and the one-step estimators from the corrected pilots have slopes between -2.3 and -2.9 rather than -4; the analytic correction with one step has slope -1 everywhere. These shallower slopes are not a failure of the bias cancellation but the signature of the centered random remainder of Lemma 4: once the bias of order m^{-r} has been

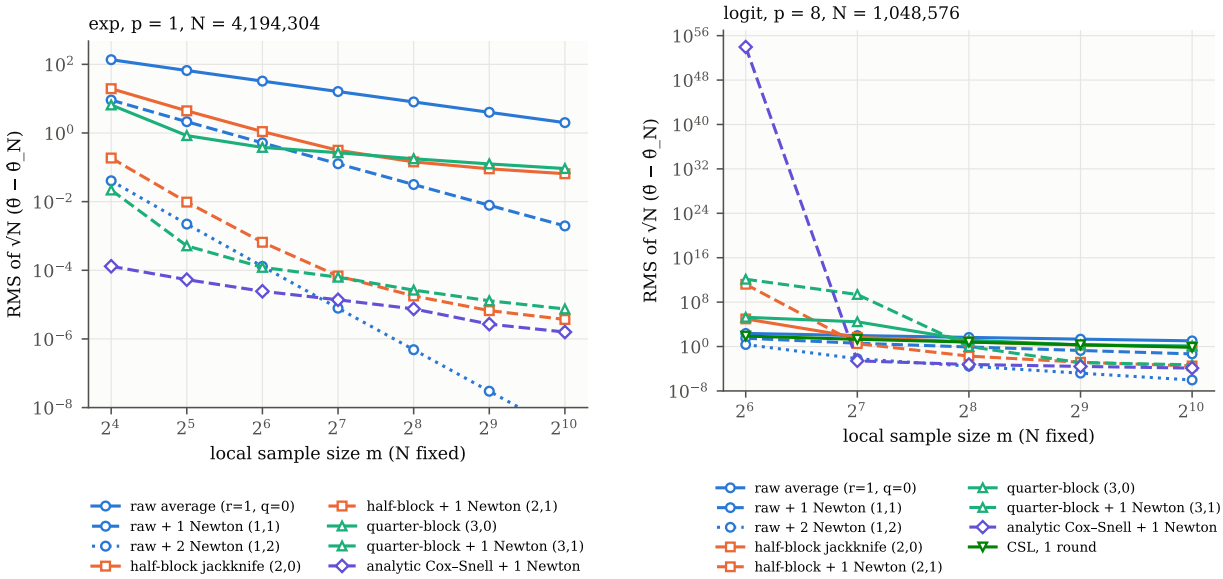


Figure 2. Root-mean-square of $\sqrt{N}\,\|\tilde{\theta} - \widehat{\theta}_N\|$ against the local sample size m at the largest N of the design, exponential model (left, $N = 4,194,304$) and logistic regression with $p = 8$ (right, $N = 2^{20}$). Estimators are labelled by (r, q); the analytic Cox–Snell correction uses one local fit; CSL is the full surrogate maximization of Jordan et al. (2019) with one round. Logarithmic axes; the values at the smallest m of the logistic design include replications with failed sub-block fits (Table 7).

removed, the gap of the one-step estimator is governed by the square of the remainder, $(\sqrt{K}/N)^2$, whose scaled size $\sqrt{N}K/N^2 = (m\sqrt{N})^{-1}$ decays only like m^{-1} at fixed N. The term vanishes asymptotically, as Theorem 1 requires, but at finite N it is the visible floor for every estimator whose bias has been cancelled, and it is the reason the proofs keep it. Figure 4 shows the gaps against $K/N^{\alpha(r,q)}$ for all N: the curves of different N collapse for each (r, q) wherever the bias term dominates.

9.3. Bias, variance, and the cost of higher order

Tables 5 to 9 report gaps and coverage, and Figure 5 the RMSE against θ_0 of the three pilots. Three facts qualify the exponents. First, higher order is not monotone in finite samples. Across all cells of all designs the quarter-block pilot has a smaller gap than the half-block pilot in 62% of the cells with $m = 16$ and in 4% of the cells with $m \geq 256$, where its bias is already negligible and the remainder amplification of the weights $(8/3, -2, 1/3)$ dominates. In logistic regression with $p = 8$ at $m = 256$ the one-step estimator from the quarter-block pilot has gap 0.98 against 0.019 for the half-block pilot, because the quarter-blocks contain 64 observations for 8 parameters; in Poisson regression with $p = 8$ the ranking is reversed at $m = 128$ (0.0025 against 0.0059) and restored at $m = 256$ (8.93×10^{-4} against 7.41×10^{-4}). Second, two Newton steps from the plain average, $(1, 2)$, which share the exponent $7/8$ with $(2, 1)$, have the smaller gap in 100% of the

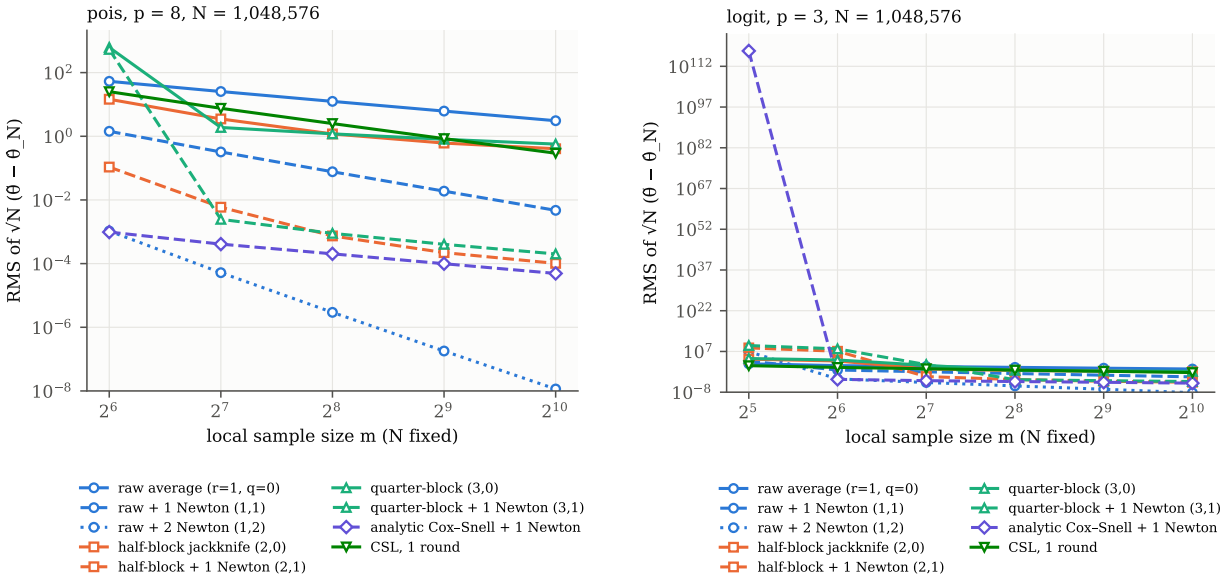


Figure 3. As Figure 2, Poisson regression with $p = 8$ (left) and logistic regression with $p = 3$ (right).

cells (2.92×10^{-4} against 0.019 in the logistic example): the plain pilot's error is dominated by its deterministic bias, which two squarings remove, whereas the corrected pilot's error contains the random remainder, which one squaring turns into the $(m\sqrt{N})^{-1}$ floor. Third, the analytic first-order correction of Cordeiro and McCullagh (1991), which costs one local fit and no sub-blocks, beats the half-block jackknife with one step in 84% of the logistic cells and in every exponential and Poisson cell (5.91×10^{-4} at $m = 256$ in the logistic example), but in none of the contaminated exponential cells, where the bias formula of the assumed model is wrong: the jackknife needs no formula and no third derivatives, which is its advantage under misspecification and outside the canonical GLM class. The stale-curvature variant of $(2, 1)$ sits between the one-step and the pilot (0.21 at $m = 256$), in the order of Theorem 3.

9.4. Failures at small blocks

When the finest sub-blocks are small relative to p, local MLEs fail to exist (separation in logistic regression) or converge to extreme values, and the construction inherits the failure: the tables report the fraction of workers whose sub-block fits all converged. In logistic regression with $p = 3$ and $m = 32$, 96% of the workers have valid half-blocks but only 23% have valid quarter-blocks; with $p = 8$ the quarter-block fits are valid for 0.9% of the workers at $m = 64$ and for 87% at $m = 128$; with $p = 20$ the half-block fits are valid for 82% of the workers at $m = 128$ and the quarter-block fits for 67% at $m = 256$. In those cells the root-mean-square gaps are dominated by a few replications in which an invalid local estimate enters the average and the Newton iteration starts far outside the basin of Theorem 2(iv); the medians are orders of magnitude smaller, and dropping the invalid

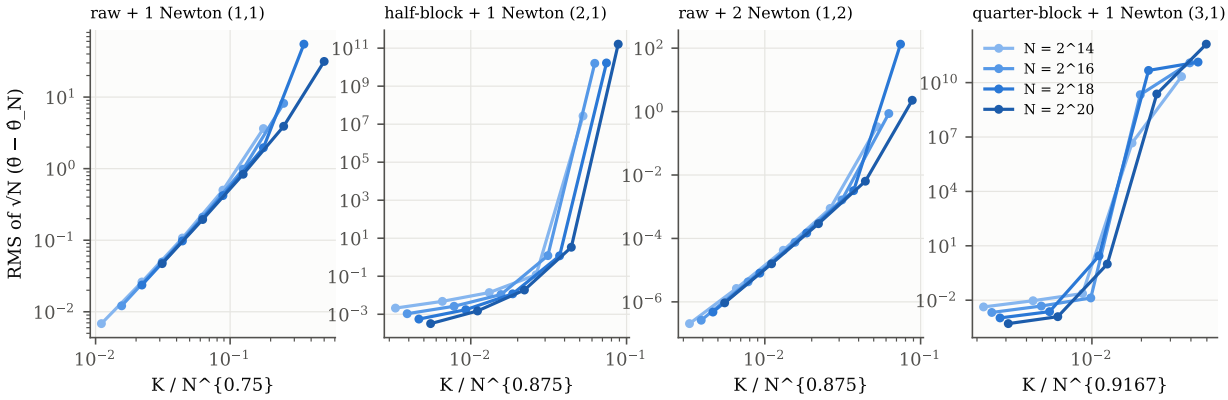


Figure 4. Root-mean-square scaled gap against $K/N^{\alpha(r,q)}$ for several N (logistic regression, $p = 8$). Under Theorem 1 the bias part of the gap is a function of this ratio, so the curves for different N coincide where the bias dominates and separate where the $(m\sqrt{N})^{-1}$ floor takes over.

workers from the pilot (rows "valid workers only") restores the behaviour of the failure-free cells at the price of a small loss of data. The safeguard is a necessity, not an option: the theory assumes that local MLEs exist, and Assumption 1 fails on exactly the sub-blocks that the higher-order patterns create.

9.5. Coverage

Figure 6 and the coverage columns of the tables compare empirical and nominal 95% coverage of intervals built from the observed information at each estimate. Inside the frontier the coverage of $(2,1)$, $(1,2)$, $(3,1)$ and of the analytic correction with one step is indistinguishable from that of the centralized MLE. The plain average under-covers at every m of the grids, since its bias θ_0/m exceeds the standard error $\theta_0/\sqrt{N}$ whenever $m < \sqrt{N}$; the one-step estimator $(1,1)$ and the half-block pilot $(2,0)$, which share the exponent $3/4$, under-cover for $m \lesssim N^{1/4}$ (in the exponential model at $N = 2^{20}$, coverage 0.40 and 0.01 at $m = 32$), and $(2,1)$ is nominal from $m = 16$ on (0.93). Under misspecification the relevant target is the pseudo-true value: with it, all estimators inside the frontier have nominal coverage in the misspecified logistic design with either the observed information or the sandwich, because the information identity nearly holds for a Bernoulli outcome with the fitted mean; in the contaminated exponential design the two differ, the observed-information intervals of every estimator including the centralized pseudo-MLE cover the pseudo-true rate with probability 0.78 at $m = 64$ while the sandwich intervals cover it with probability 0.96 (Table 17), as Corollary 1 and Proposition 2(ii) state.

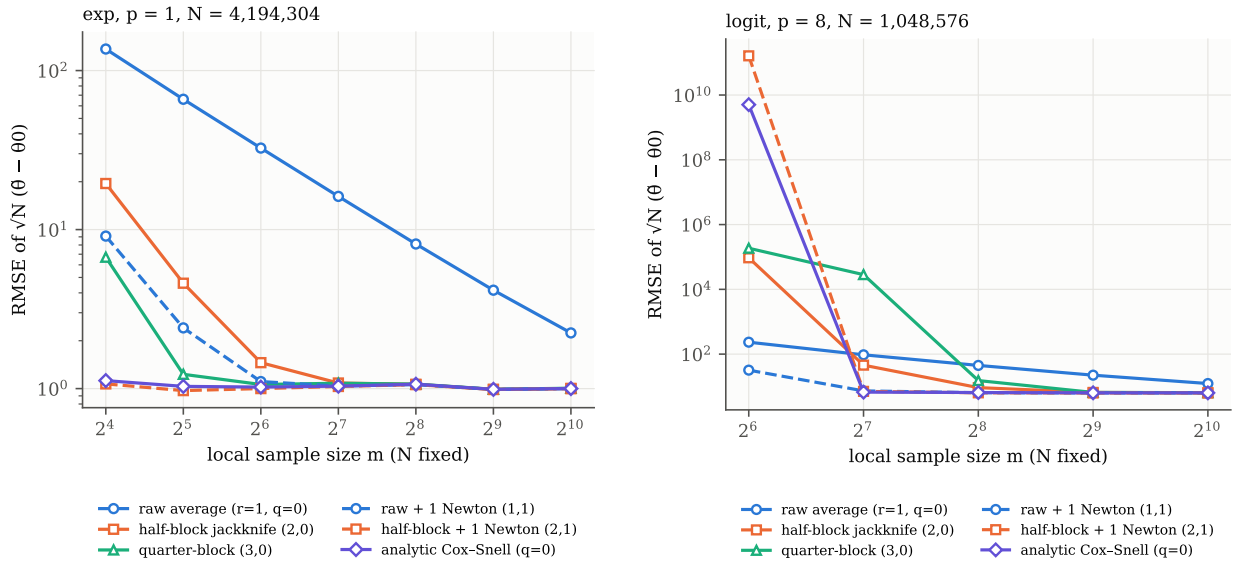


Figure 5. Root-mean-square of $\sqrt{N}\,\|\tilde{\theta} - \theta_0\|$ (error against the truth, not against $\widehat{\theta}_N$) for the pilots of order 1, 2, 3, the one-step estimators from the plain and half-block pilots, and the analytic correction; exponential (left) and logistic $p = 8$ (right). The centralized MLE has RMSE close to $\sqrt{\mathrm{tr}\,\Sigma_\psi}$; a pilot whose RMSE exceeds it pays for its bias (order 1) or for its remainder variance and failed sub-block fits (order 3).

9.6. The boundary

Table 12 and Figure 7 report the exact-model boundary study: K is set to $cN^{\alpha(r,q)}$ (rounded so that m is a multiple of the finest split) and the mean of $\sqrt{N}(\theta_r^{(q)} - \widehat{\theta}_N)$ is compared with the limit of Theorem 2(v) at the realized $c_{\mathrm{eff}} = K/N^{\alpha}$ and with a finite-m prediction, $-\sqrt{N}\,\mathbb{E}[(\theta_r^{(0)} - \widehat{\theta}_N)^2]/\theta_0$ for $q = 1$, computed from the exact bias and the variance of the local remainder (Section B). The finite-m prediction matches the Monte Carlo means to three digits in every cell. The ratio of the mean to the asymptotic limit decreases towards one as m grows: for $(1,1)$ and $(2,0)$, where $m \asymp N^{1/4}$, it is within 10% of one at $N = 2^{24}$; for $(2,1)$, where $m \asymp N^{1/8}$, the local sample sizes at $N \le 2^{24}$ do not exceed 16 and the ratio is still between 1.5 and 2.4. The boundary constants are confirmed as the limit of a regime that is reached only at very large N; at any practical N the finite-m formula is the relevant prediction.

9.7. Ablations: patterns, rounds, curvature and quantization

Tables 13 and 14 vary the split pattern, the number of rounds $q \in \{0,\dots,3\}$, the curvature used by the coordinator and the number of mantissa bits of the transmitted score and Hessian, on paired samples. Three findings. (a) At equal order, the pattern $(1,2,3)$ of Dhaene and Jochmans (2015) and the doubling pattern $(1,2,4)$ are close when all sub-block fits are valid; when the quarter-blocks are too small for the model, $(1,2,3)$, whose finest block is larger, and even the two-level pattern

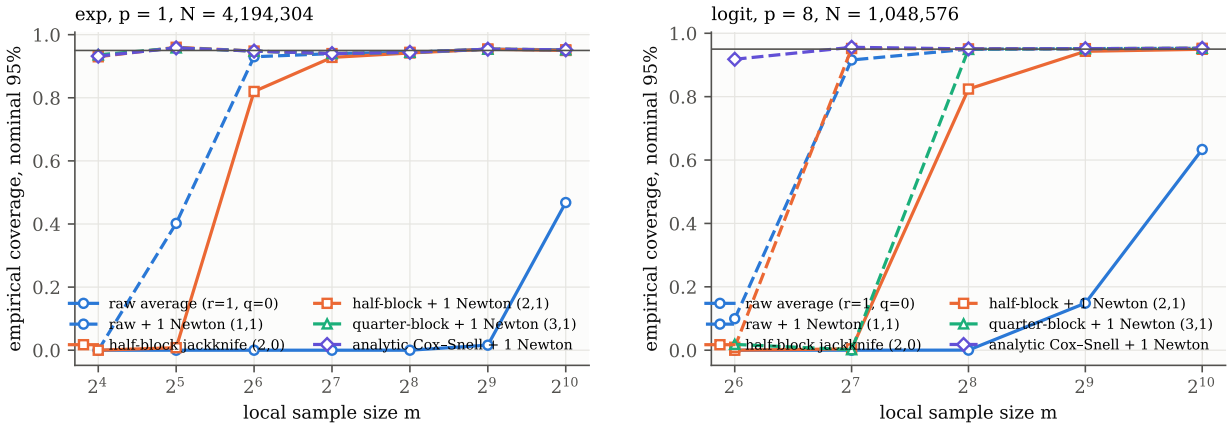


Figure 6. Empirical coverage of nominal 95% Wald intervals (observed information at the estimate) against m, exponential (left) and logistic $p = 8$ (right); Monte Carlo standard error about 0.005 (exponential) and 0.007 (logistic).

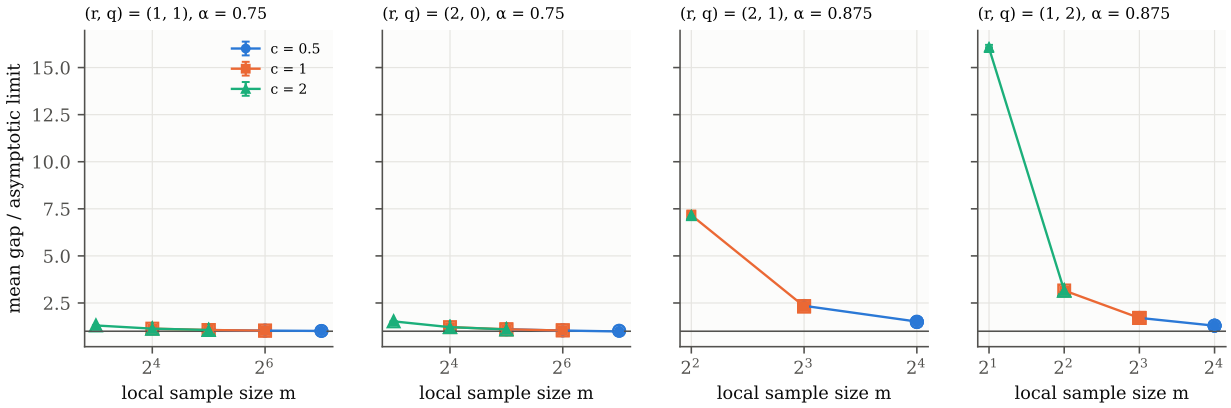


Figure 7. Boundary study in the exponential model: ratio of the Monte Carlo mean of $\sqrt{N}(\theta_r^{(q)} - \widehat{\theta}_N)$ to the asymptotic limit of Theorem 2(v) (filled markers, ± 1.96 standard errors) and of the finite-m prediction to the same limit (hollow markers), against the local sample size m that $K = cN^{\alpha(r,q)}$ implies at each N.

$(1,3)$ are far better, and the order-4 patterns are worse than order 3 at every m of the grids. (b) With stale curvature the first round removes most of the gap but each further round gains far less than with exact curvature, and the single-machine Hessian is worse again, in the order of Theorem 3; in the exponential model single-machine curvature is exact and coincides with the exact variant. (c) In the logistic design with $N = 786,432$, $m = 768$ ($K = 1,024$, so $K = N^{0.51}$), quantizing the Hessian to 12 mantissa bits leaves the one-step gap of the half-block pilot within five percent of its exact value, and the same number of bits suffices when the score is quantized as well; fewer bits degrade the gap smoothly. The requirement (31) is asymptotic and gives a smaller number at this α; the ablation shows where the finite-N constants put the threshold.

9.8. Unequal blocks and contamination

Table 15 reports the unequal-size experiment of Proposition 2: with size ratios 4 and 16 the size-weighted pilots behave as in the equal-size case, the equal-weight pilots have a gap of order one that does not vanish for $q = 0$, and one Newton step removes it. In the logistic design with $K = 1024$ and ratio 16 the smallest blocks contain a few dozen observations for $p = 3$, the half-blocks separate, and both the equal-weight and the size-weighted jackknife pilots fail while the raw pilots survive: the same small-block failure as above, now caused by the size imbalance. The contaminated exponential design (Table 17) is a misspecified exponential model whose pseudo-true rate differs from θ_0; all estimators remain first-order equivalent to the centralized pseudo-MLE, as Proposition 2(ii) states.

Table 1. Closest results, compared at theorem level (our notation: K blocks of size $m = N/K$).

Paper	Local correction	Global correction (curvature)	Uplink per worker	Regime for first-order equivalence
Zhang et al. (2013) Thm 1	none	none	p	$K \lesssim N^{1/2}$ (MSE rate)
Zhang et al. (2013) Thm 4	one random subsample	none	p	$K \lesssim N^{2/3}$ (MSE rate)
Lin and Xi (2011) Thm 4.2	none	local-Jacobian weighting	$p + p^2$	$K \ll N^{1/4}$
Huang and Huo (2019) Thm 3.4/3.5	none	1 step, aggregated Hessian	p; $p + p^2$	$K = O(N^{1/2})$ stated; $N^{3/4}$ implicit
Zhu et al. (2021) Thm 2	none	1 local step + weighting	$p + p^2$	$K = O(N^{3/4})$
Jordan et al. (2019) Cor 3.4	none	t rounds, machine-1 Hessian	p per round	$K = o(N^{2/3})$ for $t = 1$ from the average
Fan et al. (2023) rem. after Thm 4.1	none	t rounds, machine-1 Hessian	p per round	$K = o(N^{(t+1)/(t+2)})$ (polylog)
Gu and Chen (2023) Thm 4–6	analytic bias, split	information weighting	$p + p^2$	$K = o(N^{2/3})$
Wu et al. (2023) Thm 2–4	delete-one jackknife	none	p	$K = o(N^{3/4})$ (functions of means)
Yin (2022) Thm 3.2	none	1 step + analytic adjustment	p + summaries	$K = o(N^{2/3})$
Volgushev et al. (2019) Cor 3.3	none	none	p	$K = o(N^{1/2})$ necessary (quantile regr.)
This paper, Theorem 1	order-r split extrapolation	q steps, aggregated Hessian	p; $q(p + p(p+1)/2)$	$K = o(N^{1-1/(2^{q+1}r)})$, sharp in Theorem 2
This paper, Theorem 3	order r	q steps, stale / machine-1 / b-bit Hessian	see Table 4	$1 - \frac{1}{2(r+q)}$; $1 - \frac{1}{2r+q}$; bits $\geq (\frac{1}{2} - r(1-\alpha))\log_2 N$

Regimes are sufficient conditions unless marked necessary; "MSE rate" means attainment of the $O(1/N)$ rate rather than $\sqrt{N}$-equivalence to $\widehat{\theta}_N$. Theorem numbers refer to the versions read (arXiv or accepted manuscripts where publisher pages were inaccessible; see `literature/PRIOR_ART_MATRIX.csv`).

Table 2. Split patterns, weights and constants.

Pattern s	r	Weights w	β_r	A_r	$\sum_j \lvert w_j \rvert$	$\kappa(V)$	Local fits
(1)	1	1	1	1	1	1	1
$(1,2)$	2	$2, -1$	-2	4	3	6.9	3
$(1,3)$	2	$3/2, -1/2$	-3	3	2	5.8	4
$(1,2,4)$	3	$8/3, -2, 1/3$	8	8	5	80	7
$(1,2,3)$	3	$3, -3, 1$	6	12	7	71	6
$(1,2,4,8)$	4	$64/21, -8/3, 2/3, -1/21$	-64	$80/7$	6.4	3284	15
$(1,2,3,4)$	4	$4, -6, 4, -1$	-24	32	15	1171	10

$\kappa(V)$ is the spectral condition number of the Vandermonde matrix; "local fits" is $\sum_j s_j$ (each level costs one pass over the block's data). Source: `results/tables/weights_analysis.csv`.

Table 3. Sufficient exponents $\alpha(r,q)$ in $K = o(N^{\alpha})$ (Theorem 1), by local order r and number of aggregated Newton corrections q.

	$q = 0$	$q = 1$	$q = 2$	$q = 3$
$r = 1$ (plain local MLE)	$1/2$	$3/4$	$7/8$	$15/16$
$r = 2$ (half-block jackknife)	$3/4$	$7/8$	$15/16$	$31/32$
$r = 3$ (quarter-block, or $(1,2,3)$)	$5/6$	$11/12$	$23/24$	$47/48$
$r = 4$ (eighth-block, or $(1,2,3,4)$)	$7/8$	$15/16$	$31/32$	$63/64$

The exponent depends on (r,q) only through $D = 2^q r$; cells on an anti-diagonal of constant D coincide. These are sufficient conditions under fixed p; they are not machine counts for a given N, not minimax statements and not evidence of speed.

Table 4. Protocol variants: local fits, synchronized phases, uplink per worker, and sufficient exponents.

Variant	Curvature used by the coordinator	Local fits	Uplink per worker (numbers)	Sufficient exponent α
$(r,0)$	none	$\sum_j s_j$	p	$1-1/(2r)$
(r,q) exact	aggregated $H_N(\theta)$ at each iterate	$\sum_j s_j$	$p+q\,(p+p(p+1)/2)$	$1-1/(2^{q+1}r)$
(r,q) stale	$K^{-1}\sum_k H_{D_k}(\widehat{\theta}_{D_k})$, shipped once	$\sum_j s_j$	$p+p(p+1)/2+q\,p$	$1-1/(2(r+q))$
(r,q) single	$H_{D_1}(\theta)$ on worker 1	$\sum_j s_j$	$p+q\,p$	$1-1/(2r+q)$
$(r,1)$ quantized	aggregated, b mantissa bits	$\sum_j s_j$	$p+(p+p(p+1)/2)\,b/64$	$1-1/(4r)$ if (31)
AEE (Lin and Xi, 2011)	local Jacobians at local estimates	1	$p+p(p+1)/2$	$< 1/4$ (their Thm 4.2)
CSL (Jordan et al., 2019), t rounds	machine 1, full surrogate	1	$t\,p$	$1-1/(2+t)$ from the average
exact pooling (exponential families)	sufficient statistic	1	$\dim(T)$	any $K \le N$

Downlink per worker: qp numbers for every variant with $q \ge 1$. Headers, serialization and the pilot's validity flag are excluded; measured bytes appear in Section 10. Exponents are sufficient conditions from Theorems 1 and 3; the CSL and AEE rows are the cited authors' results translated to our notation; the quantized row is in units of 64-bit numbers.

Table 5. Monte Carlo, exp design, $p = 1$, $N = 4,194,304$: scaled gap to $\widehat{\theta}_N$ (root mean square and median) and 95% coverage by local sample size m.

Estimator	$m=16$, $K=262,144$			$m=64$, $K=65,536$			$m=256$, $K=16,384$			$m=1024$, $K=4,096$		
	gap	median	cov	gap	median	cov	gap	median	cov	gap	median	cov
(1,0) plain average	136.531	136.513	0.000	32.502	32.496	0.000	8.031	8.030	0.000	2.003	1.999	0.468
(1,1) average + 1 Newton	9.102	9.100	0.000	0.516	0.516	0.930	0.032	0.031	0.944	0.002	0.002	0.952
(1,2) average + 2 Newton	0.040	0.040	0.936	1.30e-4	1.30e-4	0.948	4.85e-7	4.84e-7	0.944	1.88e-9	1.86e-9	0.952
(2,0) half-block jackknife	19.523	19.497	0.000	1.098	1.053	0.820	0.144	0.096	0.942	0.065	0.046	0.950
(2,1) half-block + 1 Newton	0.186	0.186	0.930	6.55e-4	5.41e-4	0.948	1.79e-5	4.52e-6	0.944	3.69e-6	1.06e-6	0.952
(2,2) half-block + 2 Newton	1.71e-5	1.69e-5	–	2.93e-10	1.43e-10	–	6.25e-13	0	–	2.90e-13	0	–
(3,0) quarter-block	6.557	6.575	–	0.380	0.252	–	0.177	0.123	–	0.092	0.064	–
(3,1) quarter-block + 1 Newton	0.022	0.021	0.936	1.22e-4	3.09e-5	0.948	2.64e-5	7.35e-6	0.944	7.47e-6	2.02e-6	0.952
(2,1) valid workers only	–	–	–	–	–	–	–	–	–	–	–	–
(3,1) valid workers only	–	–	–	–	–	–	–	–	–	–	–	–
analytic Cox–Snell + 1 Newton	1.31e-4	3.02e-5	0.932	2.44e-5	8.12e-6	0.948	7.50e-6	1.88e-6	0.944	1.59e-6	4.66e-7	0.952
CSL, 1 round	2.99e-13	0	–	2.96e-13	0	–	3.05e-13	0	–	2.87e-13	0	–
CSL, 2 rounds	2.88e-13	0	–	2.99e-13	0	–	3.22e-13	0	–	2.86e-13	0	–
AEE (Lin–Xi)	120.465	120.444	–	31.499	31.496	–	7.968	7.965	–	1.999	1.996	–
(2,1) stale Hessian	0.971	0.970	–	0.016	0.016	–	5.43e-4	3.75e-4	–	6.40e-5	4.48e-5	–
centralized MLE	0	–	0.932	0	–	0.948	0	–	0.944	0	–	0.952

500 paired replications per column. gap = root mean square of $\sqrt{N}\,\|\tilde{\theta} - \widehat{\theta}_N\|$ over all replications (including those with failed local fits); median = its median; cov = empirical coverage of nominal 95% Wald intervals (observed information at the estimate; '–' when not computed). Monte Carlo standard errors of the gaps and RMSEs are in `results/processed/summary.csv`; coverage standard error ≈ 0.010 or less.

Fraction of workers with all local fits valid (orders 2 / 3): $m = 16$: 1.000 / 1.000; $m = 32$: 1.000 / 1.000; $m = 64$: 1.000 / 1.000; $m = 128$: 1.000 / 1.000; $m = 256$: 1.000 / 1.000; $m = 512$: 1.000 / 1.000; $m = 1024$: 1.000 / 1.000. Rows 'valid workers only' drop the workers with a failed sub-block fit from the pilot (weights renormalized) and coincide with the corresponding rows when no fit failed.

Table 6. Monte Carlo, logit design, $p = 3$, $N = 1,048,576$: scaled gap to $\widehat{\theta}_N$ (root mean square and median) and 95% coverage by local sample size m.

Estimator	$m = 32$, $K = 32,768$			$m = 128$, $K = 8,192$			$m = 512$, $K = 2,048$			$m = 1024$, $K = 1,024$		
	gap	median	cov	gap	median	cov	gap	median	cov	gap	median	cov
(1,0) plain average	286.566	149.017	0.000	29.019	29.016	0.008	6.969	6.973	0.572	3.463	3.457	0.860
(1,1) average + 1 Newton	2689.961	9.014	0.331	0.318	0.318	0.951	0.018	0.018	0.951	0.004	0.004	0.957
(1,2) average + 2 Newton	8.41e+6	0.030	0.945	3.76e-5	3.75e-5	0.952	1.21e-7	1.20e-7	0.951	6.47e-9	6.22e-9	0.957
(2,0) half-block jackknife	1.45e+4	8938.189	0.000	3.782	3.678	0.829	0.514	0.434	0.949	0.338	0.296	0.958
(2,1) half-block + 1 Newton	2.04e+8	2.36e+5	0.004	0.006	0.005	0.952	1.09e-4	4.71e-5	0.951	4.55e-5	2.19e-5	0.957
(2,2) half-block + 2 Newton	1.18e+24	2.30e+9	–	1.39e-8	7.90e-9	–	1.08e-9	1.07e-9	–	1.08e-9	1.08e-9	–
(3,0) quarter-block	2.50e+4	2910.682	–	79.768	4.556	–	0.678	0.594	–	0.473	0.422	–
(3,1) quarter-block + 1 Newton	1.52e+9	7587.587	0.008	144.811	0.007	0.947	1.73e-4	8.65e-5	0.951	8.65e-5	4.37e-5	0.957
(2,1) valid workers only	–	–	–	–	–	–	–	–	–	–	–	–
(3,1) valid workers only	–	–	–	–	–	–	–	–	–	–	–	–
analytic Cox–Snell + 1 Newton	4.62e+117	2.76e+12	0.457	1.76e-4	9.11e-5	0.952	4.49e-5	2.05e-5	0.951	2.13e-5	9.97e-6	0.957
CSL, 1 round	57.743	39.925	–	4.334	3.654	–	0.510	0.442	–	0.170	0.150	–
CSL, 2 rounds	50.483	15.308	–	1.233	0.691	–	0.064	0.040	–	0.014	0.009	–
AEE (Lin–Xi)	102.585	102.562	–	26.847	26.836	–	6.829	6.827	–	3.423	3.422	–
(2,1) stale Hessian	1.24e+4	5469.786	–	0.102	0.099	–	0.004	0.003	–	0.001	0.001	–
centralized MLE	0	–	0.950	0	–	0.952	0	–	0.951	0	–	0.957

1000 paired replications per column. gap = root mean square of $\sqrt{N}\,\|\tilde{\theta} - \widehat{\theta}_N\|$ over all replications (including those with failed local fits); median = its median; cov = empirical coverage of nominal 95% Wald intervals (observed information at the estimate; '–' when not computed). Monte Carlo standard errors of the gaps and RMSEs are in `results/processed/summary.csv`; coverage standard error ≈ 0.007 or less.

Fraction of workers with all local fits valid (orders 2 / 3): $m = 32$: 0.961 / 0.234; $m = 64$: 1.000 / 0.924; $m = 128$: 1.000 / 1.000; $m = 256$: 1.000 / 1.000; $m = 512$: 1.000 / 1.000; $m = 1024$: 1.000 / 1.000. Rows 'valid workers only' drop the workers with a failed sub-block fit from the pilot (weights renormalized) and coincide with the corresponding rows when no fit failed.

Table 7. Monte Carlo, logit design, $p = 8$, $N = 1,048,576$: scaled gap to $\widehat{\theta}_N$ (root mean square and median) and 95% coverage by local sample size m.

Estimator	$m = 64$, $K = 16,384$			$m = 128$, $K = 8,192$			$m = 512$, $K = 2,048$			$m = 1024$, $K = 1,024$		
	gap	median	cov	gap	median	cov	gap	median	cov	gap	median	cov
(1,0) plain average	233.195	226.356	0.000	94.927	94.925	0.000	21.491	21.476	0.148	10.589	10.579	0.633
(1,1) average + 1 Newton	31.483	23.333	0.099	3.909	3.908	0.916	0.195	0.194	0.952	0.047	0.047	0.953
(1,2) average + 2 Newton	2.275	0.225	0.942	0.006	0.006	0.956	1.59e-5	1.57e-5	0.952	9.30e-7	9.19e-7	0.953
(2,0) half-block jackknife	9.42e+4	1.81e+4	0.000	45.313	36.592	0.004	1.888	1.804	0.943	0.958	0.909	0.950
(2,1) half-block + 1 Newton	1.62e+11	2.23e+6	0.009	3.356	0.550	0.952	0.001	0.001	0.951	3.22e-4	2.26e-4	0.953
(2,2) half-block + 2 Newton	2.38e+24	1.52e+12	–	0.125	1.26e-4	–	1.40e-9	5.28e-10	–	7.37e-11	2.05e-11	–
(3,0) quarter-block	1.86e+5	8047.508	–	2.85e+4	5947.678	–	1.925	1.838	–	1.281	1.218	–
(3,1) quarter-block + 1 Newton	1.32e+12	1.56e+5	0.018	2.36e+9	4.44e+4	0.002	0.001	9.01e-4	0.952	5.11e-4	3.91e-4	0.953
(2,1) valid workers only	–	–	–	–	–	–	–	–	–	–	–	–
(3,1) valid workers only	–	–	–	–	–	–	–	–	–	–	–	–
analytic Cox–Snell + 1 Newton	9.07e+53	0.059	0.918	0.003	0.002	0.956	2.62e-4	1.90e-4	0.951	1.30e-4	9.20e-5	0.953
CSL, 1 round	76.119	66.001	–	20.125	18.523	–	2.097	1.973	–	0.718	0.681	–
CSL, 2 rounds	80.931	47.543	–	10.636	7.749	–	0.431	0.344	–	0.101	0.083	–
AEE (Lin–Xi)	163.596	163.559	–	81.848	81.838	–	20.707	20.711	–	10.393	10.385	–
(2,1) stale Hessian	9.31e+4	1.32e+4	–	7.823	2.738	–	0.031	0.030	–	0.008	0.008	–
centralized MLE	0	–	0.945	0	–	0.956	0	–	0.952	0	–	0.953

1000 paired replications per column. gap = root mean square of $\sqrt{N}\,\|\tilde{\theta} - \widehat{\theta}_N\|$ over all replications (including those with failed local fits); median = its median; cov = empirical coverage of nominal 95% Wald intervals (observed information at the estimate; '–' when not computed). Monte Carlo standard errors of the gaps and RMSEs are in `results/processed/summary.csv`; coverage standard error ≈ 0.007 or less.

Fraction of workers with all local fits valid (orders 2 / 3): $m = 64$: 0.930 / 0.009; $m = 128$: 1.000 / 0.865; $m = 256$: 1.000 / 1.000; $m = 512$: 1.000 / 1.000; $m = 1024$: 1.000 / 1.000. Rows 'valid workers only' drop the workers with a failed sub-block fit from the pilot (weights renormalized) and coincide with the corresponding rows when no fit failed.

Table 8. Monte Carlo, pois design, $p = 3$, $N = 1,048,576$: scaled gap to $\widehat{\theta}_N$ (root mean square and median) and 95% coverage by local sample size m.

Estimator	$m = 32$, $K = 32,768$			$m = 128$, $K = 8,192$			$m = 512$, $K = 2,048$			$m = 1024$, $K = 1,024$		
	gap	median	cov	gap	median	cov	gap	median	cov	gap	median	cov
(1,0) plain average	41.906	41.910	0.145	9.934	9.932	0.630	2.461	2.460	0.727	1.229	1.225	0.877
(1,1) average + 1 Newton	0.882	0.882	0.905	0.049	0.048	0.949	0.003	0.003	0.951	7.49e-4	7.43e-4	0.947
(1,2) average + 2 Newton	3.97e-4	3.97e-4	0.944	1.16e-6	1.16e-6	0.949	3.88e-9	3.81e-9	0.951	3.52e-10	3.51e-10	0.946
(2,0) half-block jackknife	12.288	9.597	0.007	0.721	0.678	0.926	0.221	0.193	0.948	0.154	0.136	0.947
(2,1) half-block + 1 Newton	0.768	0.053	0.941	3.25e-4	2.39e-4	0.949	3.89e-5	2.16e-5	0.951	1.84e-5	1.05e-5	0.946
(2,2) half-block + 2 Newton	0.009	1.65e-6	–	5.70e-10	5.77e-10	–	6.11e-10	6.09e-10	–	6.12e-10	6.12e-10	–
(3,0) quarter-block	1210.274	223.376	–	0.650	0.569	–	0.309	0.269	–	0.218	0.191	–
(3,1) quarter-block + 1 Newton	1192.998	30.713	0.014	3.30e-4	1.89e-4	0.949	7.58e-5	4.16e-5	0.951	3.72e-5	2.09e-5	0.946
(2,1) valid workers only	–	–	–	–	–	–	–	–	–	–	–	–
(3,1) valid workers only	–	–	–	–	–	–	–	–	–	–	–	–
analytic Cox–Snell + 1 Newton	3.65e-4	2.25e-4	0.944	7.83e-5	4.51e-5	0.949	1.81e-5	1.04e-5	0.951	8.83e-6	5.59e-6	0.946
CSL, 1 round	14.351	9.883	–	1.368	1.057	–	0.172	0.139	–	0.060	0.050	–
CSL, 2 rounds	22.792	4.201	–	0.500	0.199	–	0.026	0.014	–	0.006	0.003	–
AEE (Lin–Xi)	40.398	40.398	–	9.860	9.859	–	2.450	2.449	–	1.223	1.223	–
(2,1) stale Hessian	0.803	0.083	–	0.004	0.003	–	3.24e-4	2.46e-4	–	1.16e-4	9.14e-5	–
centralized MLE	0	–	0.944	0	–	0.949	0	–	0.951	0	–	0.946

1000 paired replications per column. gap = root mean square of $\sqrt{N}\,\|\tilde{\theta} - \widehat{\theta}_N\|$ over all replications (including those with failed local fits); median = its median; cov = empirical coverage of nominal 95% Wald intervals (observed information at the estimate; '–' when not computed). Monte Carlo standard errors of the gaps and RMSEs are in `results/processed/summary.csv`; coverage standard error ≈ 0.007 or less.

Fraction of workers with all local fits valid (orders 2 / 3): $m = 32$: 1.000 / 1.000; $m = 64$: 1.000 / 1.000; $m = 128$: 1.000 / 1.000; $m = 256$: 1.000 / 1.000; $m = 512$: 1.000 / 1.000; $m = 1024$: 1.000 / 1.000. Rows 'valid workers only' drop the workers with a failed sub-block fit from the pilot (weights renormalized) and coincide with the corresponding rows when no fit failed.

Table 9. Monte Carlo, pois design, $p = 8$, $N = 1,048,576$: scaled gap to $\widehat{\theta}_N$ (root mean square and median) and 95% coverage by local sample size m.

Estimator	$m = 64$, $K = 16,384$			$m = 128$, $K = 8,192$			$m = 512$, $K = 2,048$			$m = 1024$, $K = 1,024$		
	gap	median	cov	gap	median	cov	gap	median	cov	gap	median	cov
(1,0) plain average	53.236	53.245	0.354	25.472	25.475	0.777	6.187	6.184	0.829	3.082	3.078	0.845
(1,1) average + 1 Newton	1.433	1.433	0.917	0.322	0.322	0.948	0.019	0.019	0.951	0.005	0.005	0.949
(1,2) average + 2 Newton	0.001	0.001	0.952	5.19e-5	5.18e-5	0.948	1.80e-7	1.79e-7	0.951	1.15e-8	1.13e-8	0.949
(2,0) half-block jackknife	14.549	14.547	0.065	3.503	3.484	0.732	0.612	0.587	0.944	0.406	0.393	0.946
(2,1) half-block + 1 Newton	0.108	0.107	0.951	0.006	0.006	0.948	2.22e-4	1.86e-4	0.951	1.02e-4	8.48e-5	0.949
(2,2) half-block + 2 Newton	6.53e-6	6.32e-6	–	2.13e-8	1.76e-8	–	4.59e-11	2.07e-11	–	9.93e-12	4.46e-12	–
(3,0) quarter-block	618.373	304.800	–	1.891	1.828	–	0.805	0.770	–	0.568	0.544	–
(3,1) quarter-block + 1 Newton	534.838	54.517	0.055	0.002	0.002	0.948	4.04e-4	3.27e-4	0.951	2.01e-4	1.61e-4	0.949
(2,1) valid workers only	–	–	–	–	–	–	–	–	–	–	–	–
(3,1) valid workers only	–	–	–	–	–	–	–	–	–	–	–	–
analytic Cox–Snell + 1 Newton	9.78e-4	7.84e-4	0.952	4.10e-4	3.31e-4	0.948	9.93e-5	8.02e-5	0.951	4.91e-5	3.85e-5	0.949
CSL, 1 round	25.196	21.206	–	7.563	6.728	–	0.845	0.791	–	0.294	0.279	–
CSL, 2 rounds	39.205	18.291	–	5.394	3.233	–	0.219	0.156	–	0.048	0.039	–
AEE (Lin–Xi)	51.852	51.847	–	25.195	25.185	–	6.159	6.160	–	3.063	3.062	–
(2,1) stale Hessian	0.063	0.062	–	0.011	0.010	–	9.31e-4	8.25e-4	–	3.32e-4	2.94e-4	–
centralized MLE	0	–	0.952	0	–	0.948	0	–	0.951	0	–	0.949

1000 paired replications per column. gap = root mean square of $\sqrt{N}\,\|\tilde{\theta} - \widehat{\theta}_N\|$ over all replications (including those with failed local fits); median = its median; cov = empirical coverage of nominal 95% Wald intervals (observed information at the estimate; '–' when not computed). Monte Carlo standard errors of the gaps and RMSEs are in `results/processed/summary.csv`; coverage standard error ≈ 0.007 or less.

Fraction of workers with all local fits valid (orders 2 / 3): $m = 64$: 1.000 / 1.000; $m = 128$: 1.000 / 1.000; $m = 256$: 1.000 / 1.000; $m = 512$: 1.000 / 1.000; $m = 1024$: 1.000 / 1.000. Rows 'valid workers only' drop the workers with a failed sub-block fit from the pilot (weights renormalized) and coincide with the corresponding rows when no fit failed.

Table 10. Monte Carlo, logit p20 design, $p = 20$, $N = 1,048,576$: scaled gap to $\widehat{\theta}_N$ (root mean square and median) and 95% coverage by local sample size m.

Estimator	$m = 128$, $K = 8,192$			$m = 256$, $K = 4,096$			$m = 1024$, $K = 1,024$			$m = 2048$, $K = 512$		
	gap	median	cov	gap	median	cov	gap	median	cov	gap	median	cov
(1,0) plain average	323.376	323.451	0.000	131.255	131.250	0.000	29.210	29.203	0.343	14.338	14.326	0.719
(1,1) average + 1 Newton	49.477	49.512	0.075	7.699	7.698	0.889	0.369	0.368	0.953	0.088	0.088	0.954
(1,2) average + 2 Newton	1.031	1.032	0.947	0.025	0.025	0.952	5.84e-5	5.81e-5	0.953	3.37e-6	3.33e-6	0.954
(2,0) half-block jackknife	1.65e+5	5.03e+4	0.000	61.069	61.034	0.033	3.171	3.126	0.941	1.668	1.638	0.947
(2,1) half-block + 1 Newton	2.26e+11	5.46e+7	0.066	1.564	1.560	0.947	0.004	0.004	0.953	9.12e-4	7.57e-4	0.954
(2,2) half-block + 2 Newton	1.29e+131	2.42e+16	–	0.001	0.001	–	7.91e-9	5.18e-9	–	5.12e-10	2.39e-10	–
(3,0) quarter-block	3.25e+5	8.69e+4	–	4.40e+4	1.67e+4	–	3.404	3.342	–	2.272	2.253	–
(3,1) quarter-block + 1 Newton	1.84e+12	2.89e+8	0.899	4.14e+9	6.92e+5	0.008	0.003	0.003	0.953	0.002	0.001	0.954
(2,1) valid workers only	–	–	–	–	–	–	–	–	–	–	–	–
(3,1) valid workers only	–	–	–	–	–	–	–	–	–	–	–	–
analytic Cox–Snell + 1 Newton	0.253	0.249	0.949	0.009	0.008	0.952	7.34e-4	6.60e-4	0.953	3.59e-4	3.16e-4	0.954
CSL, 1 round	112.258	106.215	–	29.269	27.982	–	2.937	2.881	–	1.017	0.989	–
CSL, 2 rounds	141.077	114.985	–	18.314	15.377	–	0.691	0.640	–	0.163	0.150	–
AEE (Lin–Xi)	230.409	230.338	–	112.533	112.500	–	28.049	28.049	–	14.027	14.022	–
(2,1) stale Hessian	1.63e+5	4.41e+4	–	6.341	6.328	–	0.066	0.065	–	0.018	0.018	–
centralized MLE	0	–	0.948	0	–	0.952	0	–	0.953	0	–	0.954

500 paired replications per column. gap = root mean square of $\sqrt{N}\,\|\tilde{\theta} - \widehat{\theta}_N\|$ over all replications (including those with failed local fits); median = its median; cov = empirical coverage of nominal 95% Wald intervals (observed information at the estimate; '–' when not computed). Monte Carlo standard errors of the gaps and RMSEs are in `results/processed/summary.csv`; coverage standard error ≈ 0.010 or less.

Fraction of workers with all local fits valid (orders 2 / 3): $m = 128$: 0.821 / 0.000; $m = 256$: 1.000 / 0.674; $m = 512$: 1.000 / 1.000; $m = 1024$: 1.000 / 1.000; $m = 2048$: 1.000 / 1.000. Rows 'valid workers only' drop the workers with a failed sub-block fit from the pilot (weights renormalized) and coincide with the corresponding rows when no fit failed.

Table 11. Empirical log-log slopes of the RMS scaled gap to $\widehat{\theta}_N$ against m at fixed N (largest N of each design), with the slopes predicted by the leading bias term.

Design	p	N	r1_q0	r1_q1	r1_q2	r2_q0	r2_q1	r3_q0	r3_q1	cs_q1
predicted			-1	-2	-4	-2	-4	-3	-6	-4
exp	1	4,194,304	-1.01	-2.03	-4.05	-1.39	-2.61	-0.89	-1.69	-1.05
exp contam	1	1,048,576	-0.90	-1.81	-3.62	-1.10	-2.12	-1.47	-2.76	-1.50
logit	3	1,048,576	-1.02	-2.05	-4.16	-1.15	-2.29	-2.28	-6.32	-1.01
logit	8	1,048,576	-1.03	-2.07	-4.15	-1.40	-2.94	-1.71	-5.46	-1.09
logit misspec	3	1,048,576	-1.02	-2.06	-4.11	-1.30	-2.62	-2.29	-5.92	-1.97
logit p20	20	1,048,576	-1.04	-2.09	-4.17	-1.36	-2.90	-0.93	-2.11	-1.16
pois	3	1,048,576	-1.02	-2.04	-4.15	-1.22	-2.78	-1.96	-3.98	-1.07
pois	8	1,048,576	-1.03	-2.06	-4.11	-1.28	-2.48	-2.14	-4.53	-1.07

Slopes are least-squares fits over the available m grid; they approach the predicted exponent only where the bias term dominates the random remainder and the sample is in the asymptotic regime. The predicted slope for the Cox–Snell + 1 Newton estimator assumes an exact first-order bias formula (-4). Source: `results/processed/summary.csv`.

Table 12. Boundary study in the exponential model (Theorem 2(v)): mean of $\sqrt{N}(\theta_r^{(q)} - \widehat{\theta}_N)$ when $K \approx cN^{\alpha(r,q)}$, against the limit $-\theta_0(\beta_r c_{\text{eff}}^r)^{2^q}$ ($q \geq 1$) or $\beta_r \theta_0 c_{\text{eff}}^r$ ($q = 0$), with $c_{\text{eff}} = K/N^\alpha$ after rounding m to a multiple of the finest split.

(r,q)	c	N	K	m	reps	mean gap (s.e.)	finite-m prediction	limit at c_{eff}	gap/limit
(1,1)	0.5	2^{16}	2,048	32	2000	-0.266 (0.00039)	-0.2667	-0.25	1.06
(1,1)	0.5	2^{18}	4,096	64	2000	-0.1289 (0.00013)	-0.1291	-0.125	1.03
(1,1)	0.5	2^{20}	16,384	64	2000	-0.2582 (0.00013)	-0.258	-0.25	1.03
(1,1)	0.5	2^{22}	32,768	128	1000	-0.127 (6.4e-05)	-0.127	-0.125	1.02
(1,1)	0.5	2^{24}	131,072	128	400	-0.2541 (0.0001)	-0.254	-0.25	1.02
(1,1)	1	2^{16}	4,096	16	2000	-1.138 (0.0012)	-1.138	-1	1.14
(1,1)	1	2^{18}	8,192	32	2000	-0.5328 (0.00039)	-0.5329	-0.5	1.07
(1,1)	1	2^{20}	32,768	32	2000	-1.065 (0.00038)	-1.066	-1	1.07
(1,1)	1	2^{22}	65,536	64	1000	-0.5158 (0.00019)	-0.516	-0.5	1.03
(1,1)	1	2^{24}	262,144	64	400	-1.032 (0.00029)	-1.032	-1	1.03
(1,1)	2	2^{16}	8,192	8	2000	-5.227 (0.0043)	-5.226	-4	1.31
(1,1)	2	2^{18}	16,384	16	2000	-2.274 (0.0012)	-2.276	-2	1.14
(1,1)	2	2^{20}	65,536	16	2000	-4.55 (0.0012)	-4.551	-4	1.14
(1,1)	2	2^{22}	131,072	32	1000	-2.13 (0.00052)	-2.131	-2	1.07
(1,1)	2	2^{24}	524,288	32	400	-4.262 (0.00088)	-4.262	-4	1.07
(1,2)	0.5	2^{16}	8,192	8	2000	-0.1068 (0.00018)	-0.1068	-0.0625	1.71
(1,2)	0.5	2^{18}	16,384	16	2000	-0.01011 (1.1e-05)	-0.01012	-0.007812	1.29
(1,2)	0.5	2^{20}	65,536	16	2000	-0.02023 (1.1e-05)	-0.02023	-0.01562	1.29
(1,2)	0.5	2^{22}	262,144	16	1000	-0.04046 (1.6e-05)	-0.04046	-0.03125	1.29
(1,2)	0.5	2^{24}	1,048,576	16	400	-0.08092 (2.3e-05)	-0.08091	-0.0625	1.29
(1,2)	1	2^{16}	16,384	4	2000	-3.163 (0.0045)	-3.165	-1	3.16
(1,2)	1	2^{18}	32,768	8	2000	-0.2131 (0.00017)	-0.2134	-0.125	1.71
(1,2)	1	2^{20}	131,072	8	2000	-0.4264 (0.00017)	-0.4265	-0.25	1.71
(1,2)	1	2^{22}	524,288	8	1000	-0.8533 (0.00025)	-0.853	-0.5	1.71
(1,2)	1	2^{24}	2,097,152	8	400	-1.705 (0.00038)	-1.706	-1	1.71
(1,2)	2	2^{16}	32,768	2	2000	-257.5 (0.92)	–	-16	16.09
(1,2)	2	2^{18}	131,072	2	2000	-513.1 (0.74)	–	-32	16.03
(1,2)	2	2^{20}	262,144	4	2000	-12.64 (0.0044)	-12.64	-4	3.16
(1,2)	2	2^{22}	1,048,576	4	1000	-25.29 (0.006)	-25.28	-8	3.16
(1,2)	2	2^{24}	4,194,304	4	400	-50.55 (0.01)	-50.57	-16	3.16
(2,0)	0.5	2^{16}	2,048	32	2000	-0.5522 (0.0085)	-0.5505	-0.5	1.10
(2,0)	0.5	2^{18}	4,096	64	2000	-0.2589 (0.0058)	-0.2622	-0.25	1.04
(2,0)	0.5	2^{20}	16,384	64	2000	-0.5192 (0.0058)	-0.5243	-0.5	1.04
(2,0)	0.5	2^{22}	32,768	128	1000	-0.2469 (0.0056)	-0.256	-0.25	0.99
(2,0)	0.5	2^{24}	131,072	128	400	-0.5122 (0.0089)	-0.5119	-0.5	1.02
(2,0)	1	2^{16}	4,096	16	2000	-2.438 (0.014)	-2.438	-2	1.22
(2,0)	1	2^{18}	8,192	32	2000	-1.102 (0.0086)	-1.101	-1	1.10
(2,0)	1	2^{20}	32,768	32	2000	-2.208 (0.0088)	-2.202	-2	1.10
(2,0)	1	2^{22}	65,536	64	1000	-1.036 (0.0085)	-1.049	-1	1.04
(2,0)	1	2^{24}	262,144	64	400	-2.104 (0.014)	-2.097	-2	1.05
(2,0)	2	2^{16}	8,192	8	2000	-12.19 (0.028)	-12.19	-8	1.52
(2,0)	2	2^{18}	16,384	16	2000	-4.863 (0.014)	-4.876	-4	1.22
(2,0)	2	2^{20}	65,536	16	2000	-9.737 (0.014)	-9.752	-8	1.22
(2,0)	2	2^{22}	131,072	32	1000	-4.391 (0.013)	-4.404	-4	1.10
(2,0)	2	2^{24}	524,288	32	400	-8.8 (0.018)	-8.809	-8	1.10
(2,1)	0.5	2^{16}	8,192	8	2000	-0.5876 (0.0026)	-0.5863	-0.25	2.35
(2,1)	0.5	2^{18}	16,384	16	2000	-0.04726 (0.00026)	-0.0472	-0.03125	1.51
(2,1)	0.5	2^{20}	65,536	16	2000	-0.09306 (0.00027)	-0.09326	-0.0625	1.49
(2,1)	0.5	2^{22}	262,144	16	1000	-0.1866 (0.00038)	-0.1859	-0.125	1.49
(2,1)	0.5	2^{24}	1,048,576	16	400	-0.3725 (0.0006)	-0.3716	-0.25	1.49
(2,1)	1	2^{16}	16,384	4	2000	-28.8 (0.17)	–	-4	7.20
(2,1)	1	2^{18}	65,536	4	2000	-57.3 (0.2)	–	-8	7.16
(2,1)	1	2^{20}	131,072	8	2000	-2.323 (0.0027)	-2.323	-1	2.32
(2,1)	1	2^{22}	524,288	8	1000	-4.637 (0.0036)	-4.645	-2	2.32
(2,1)	1	2^{24}	2,097,152	8	400	-9.287 (0.0057)	-9.288	-4	2.32
(2,1)	2	2^{16}	16,384	4	2000	-28.72 (0.14)	–	-4	7.18
(2,1)	2	2^{18}	65,536	4	2000	-57.07 (0.12)	–	-8	7.13
(2,1)	2	2^{20}	262,144	4	2000	-113.8 (0.12)	–	-16	7.11
(2,1)	2	2^{22}	1,048,576	4	1000	-227.7 (0.18)	–	-32	7.12
(2,1)	2	2^{24}	4,194,304	4	400	-455.1 (0.33)	–	-64	7.11
(3,1)	0.5	2^{16}	8,192	8	2000	-0.7203 (0.1)	–	-0.0625	11.53
(3,1)	0.5	2^{18}	32,768	8	2000	-1.178 (0.0066)	–	-0.125	9.42
(3,1)	0.5	2^{20}	131,072	8	2000	-2.339 (0.0077)	–	-0.25	9.36
(3,1)	0.5	2^{22}	524,288	8	1000	-4.668 (0.012)	–	-0.5	9.34
(3,1)	0.5	2^{24}	2,097,152	8	400	-9.308 (0.016)	–	-1	9.31
(3,1)	1	2^{16}	8,192	8	2000	-0.6026 (0.0082)	–	-0.0625	9.64
(3,1)	1	2^{18}	32,768	8	2000	-1.178 (0.0065)	–	-0.125	9.42

Table 13. Ablation (exponential, $N = 6,291,456$): RMS of $\sqrt{N}\,\|\tilde{\theta} - \widehat{\theta}_N\|$ by split pattern, order r, Newton rounds q and curvature used (exact / stale averaged local / single machine); quantized exact curvature with b mantissa bits for $q = 1$.

Pattern	curvature	q	$m = 24$	$m = 96$
(1)	exact	0	109.045	26.408
(1)	exact	1	4.741	0.278
(1)	exact	2	0.009	3.08e-5
(1)	exact	3	3.20e-8	4.80e-13
(1)	stale	1	8.723	0.545
(1)	stale	2	0.378	0.006
(1)	stale	3	0.015	5.91e-5
(1)	single	1	4.741	0.278
(1)	single	2	0.009	3.08e-5
(1)	single	3	3.20e-8	4.79e-13
(1,2)	exact	0	9.932	0.601
(1,2)	exact	1	0.039	1.77e-4
(1,2)	exact	2	6.32e-7	2.22e-11
(1,2)	exact	3	2.86e-13	3.14e-13
(1,2)	stale	1	0.359	0.006
(1,2)	stale	2	0.014	6.18e-5
(1,2)	stale	3	5.72e-4	6.35e-7
(1,2)	single	1	0.039	1.77e-4
(1,2)	single	2	6.32e-7	2.22e-11
(1,2)	single	3	2.86e-13	3.14e-13
(1,3)	exact	0	15.578	0.870
(1,3)	exact	1	0.097	3.28e-4
(1,3)	exact	2	3.76e-6	5.65e-11
(1,3)	exact	3	2.89e-13	3.11e-13
(1,2,4)	exact	0	2.100	0.312
(1,2,4)	exact	1	0.002	6.55e-5
(1,2,4)	exact	2	3.07e-9	5.44e-12
(1,2,4)	exact	3	2.93e-13	3.14e-13
(1,2,4)	stale	1	0.086	0.003
(1,2,4)	stale	2	0.003	3.32e-5
(1,2,4)	stale	3	1.38e-4	3.43e-7
(1,2,4)	single	1	0.002	6.55e-5
(1,2,4)	single	2	3.07e-9	5.44e-12
(1,2,4)	single	3	2.93e-13	3.14e-13
(1,2,3)	exact	0	1.697	0.429
(1,2,3)	exact	1	0.002	1.30e-4
(1,2,3)	exact	2	3.23e-9	2.25e-11
(1,2,3)	exact	3	3.03e-13	3.10e-13
(1,2,4,8)	exact	0	1.302	0.380
(1,2,4,8)	exact	1	9.81e-4	9.80e-5
(1,2,4,8)	exact	2	8.24e-10	1.20e-11
(1,2,4,8)	exact	3	2.93e-13	3.13e-13
(1,2,3,4)	exact	0	2.411	0.957
(1,2,3,4)	exact	1	0.004	6.46e-4
(1,2,3,4)	exact	2	1.96e-8	6.13e-10
(1,2,3,4)	exact	3	3.01e-13	3.10e-13
(1)	b-bit Hessian, $b = 2$	1	4.546	0.277
(1)	b-bit Hessian, $b = 3$	1	10.384	0.277
(1)	b-bit Hessian, $b = 4$	1	2.423	0.277
(1)	b-bit Hessian, $b = 6$	1	4.312	0.140
(1)	b-bit Hessian, $b = 8$	1	4.778	0.266
(1)	b-bit Hessian, $b = 12$	1	4.741	0.278
(1)	b-bit Hessian, $b = 16$	1	4.741	0.278
(1)	b-bit Hessian, $b = 24$	1	4.741	0.278
(1)	b-bit Hessian, $b = 52$	1	4.741	0.278
(1)	b-bit score, $b = 2$	1	18.981	3.612
(1)	b-bit score, $b = 3$	1	2.378	1.401
(1)	b-bit score, $b = 4$	1	8.316	1.115
(1)	b-bit score, $b = 6$	1	5.450	0.380
(1)	b-bit score, $b = 8$	1	4.729	0.283
(1)	b-bit score, $b = 12$	1	4.740	0.278
(1)	b-bit score, $b = 16$	1	4.741	0.278
(1)	b-bit score, $b = 24$	1	4.741	0.278
(1)	b-bit score, $b = 52$	1	4.741	0.278

Table 14. Ablation (logistic, $N = 786,432$): RMS of $\sqrt{N}\|\tilde{\theta} - \widehat{\theta}_N\|$ by split pattern, order r, Newton rounds q and curvature used (exact / stale averaged local / single machine); quantized exact curvature with b mantissa bits for $q = 1$.

Pattern	curvature	q	$m = 192$	$m = 768$
(1)	exact	0	52.268	12.296
(1)	exact	1	1.350	0.073
(1)	exact	2	8.80e-4	2.61e-6
(1)	exact	3	3.75e-10	1.27e-13
(1)	stale	1	0.615	0.036
(1)	stale	2	0.024	3.32e-4
(1)	stale	3	9.10e-4	3.14e-6
(1)	single	1	9.113	1.013
(1)	single	2	3.342	0.167
(1)	single	3	1.538	0.030
(1,2)	exact	0	11.228	1.167
(1,2)	exact	1	0.061	5.63e-4
(1,2)	exact	2	1.91e-6	2.41e-10
(1,2)	exact	3	1.30e-13	1.27e-13
(1,2)	stale	1	0.490	0.014
(1,2)	stale	2	0.019	1.62e-4
(1,2)	stale	3	7.41e-4	1.98e-6
(1,2)	single	1	1.900	0.121
(1,2)	single	2	0.710	0.020
(1,2)	single	3	0.336	0.004
(1,3)	exact	0	23.360	1.201
(1,3)	exact	1	0.735	6.83e-4
(1,3)	exact	2	0.004	3.61e-10
(1,3)	exact	3	1.39e-7	1.27e-13
(1,2,4)	exact	0	267.576	1.511
(1,2,4)	exact	1	1410.845	8.22e-4
(1,2,4)	exact	2	5.77e+5	5.37e-10
(1,2,4)	exact	3	1.55e+13	1.27e-13
(1,2,4)	stale	1	167.463	0.018
(1,2,4)	stale	2	112.549	2.21e-4
(1,2,4)	stale	3	64.138	2.76e-6
(1,2,4)	single	1	1415.946	0.167
(1,2,4)	single	2	4.37e+5	0.028
(1,2,4)	single	3	7.63e+94	0.006
(1,2,3)	exact	0	23.384	2.154
(1,2,3)	exact	1	2.991	0.002
(1,2,3)	exact	2	0.073	2.21e-9
(1,2,3)	exact	3	4.99e-5	1.28e-13
(1,2,4,8)	exact	0	3530.386	1.864
(1,2,4,8)	exact	1	8.08e+5	0.001
(1,2,4,8)	exact	2	4.30e+13	1.18e-9
(1,2,4,8)	exact	3	7.62e+39	1.28e-13
(1,2,3,4)	exact	0	792.244	4.986
(1,2,3,4)	exact	1	5.62e+4	0.009
(1,2,3,4)	exact	2	1.55e+10	5.93e-8
(1,2,3,4)	exact	3	1.55e+10	1.27e-13
(1)	b-bit Hessian, $b = 2$	1	3.520	0.759
(1)	b-bit Hessian, $b = 3$	1	3.313	0.682
(1)	b-bit Hessian, $b = 4$	1	1.939	0.306
(1)	b-bit Hessian, $b = 6$	1	1.412	0.142
(1)	b-bit Hessian, $b = 8$	1	1.351	0.077
(1)	b-bit Hessian, $b = 12$	1	1.350	0.073
(1)	b-bit Hessian, $b = 16$	1	1.350	0.073
(1)	b-bit Hessian, $b = 24$	1	1.350	0.073
(1)	b-bit Hessian, $b = 52$	1	1.350	0.073
(1)	b-bit score, $b = 2$	1	3.663	0.861
(1)	b-bit score, $b = 3$	1	2.006	0.446
(1)	b-bit score, $b = 4$	1	1.626	0.241
(1)	b-bit score, $b = 6$	1	1.375	0.095
(1)	b-bit score, $b = 8$	1	1.352	0.075
(1)	b-bit score, $b = 12$	1	1.350	0.073
(1)	b-bit score, $b = 16$	1	1.350	0.073
(1)	b-bit score, $b = 24$	1	1.350	0.073
(1)	b-bit score, $b = 52$	1	1.350	0.073
(1)	b-bit both, $b = 2$	1	4.782	1.133

Table 15. Unequal block sizes (Proposition 2): RMS of $\sqrt{N}\,\|\tilde{\theta} - \widehat{\theta}_N\|$ for size-weighted and equal-weight aggregation of the raw and half-block-jackknife local estimates, with $q = 0, 1, 2$ Newton corrections. Sizes grow geometrically with max/min ratio as shown; $N \approx 2^{18}$.

Model	K	ratio	reps	weights	pilot	$q=0$	$q=1$	$q=2$
exponential	256	1	1000	equal	jack2	0.062	1.24e-5	8.52e-13
exponential	256	1	1000	equal	raw	0.501	4.97e-4	5.13e-10
exponential	256	1	1000	weighted	jack2	0.062	1.24e-5	8.52e-13
exponential	256	1	1000	weighted	raw	0.501	4.97e-4	5.13e-10
exponential	256	4	1000	equal	jack2	0.441	6.71e-4	2.96e-9
exponential	256	4	1000	equal	raw	0.735	0.002	1.03e-8
exponential	256	4	1000	weighted	jack2	0.064	1.37e-5	1.13e-12
exponential	256	4	1000	weighted	raw	0.503	5.02e-4	5.25e-10
exponential	256	16	1000	equal	jack2	0.943	0.003	5.96e-8
exponential	256	16	1000	equal	raw	1.321	0.006	1.78e-7
exponential	256	16	1000	weighted	jack2	0.062	1.28e-5	1.05e-12
exponential	256	16	1000	weighted	raw	0.500	4.95e-4	5.07e-10
exponential	1024	4	1000	equal	jack2	0.447	6.83e-4	3.06e-9
exponential	1024	4	1000	equal	raw	2.415	0.012	3.50e-7
exponential	1024	4	1000	weighted	jack2	0.123	5.26e-5	2.04e-11
exponential	1024	4	1000	weighted	raw	2.009	0.008	1.24e-7
exponential	1024	16	1000	equal	jack2	0.994	0.003	6.41e-8
exponential	1024	16	1000	equal	raw	3.917	0.033	3.05e-6
exponential	1024	16	1000	weighted	jack2	0.133	5.70e-5	1.83e-11
exponential	1024	16	1000	weighted	raw	2.012	0.008	1.25e-7
logistic	256	1	300	equal	jack2	0.508	2.85e-4	1.33e-9
logistic	256	1	300	equal	raw	4.034	0.014	1.77e-7
logistic	256	1	300	weighted	jack2	0.508	2.85e-4	1.33e-9
logistic	256	1	300	weighted	raw	4.034	0.014	1.77e-7
logistic	256	4	300	equal	jack2	2.025	0.004	3.67e-8
logistic	256	4	300	equal	raw	5.096	0.025	8.06e-7
logistic	256	4	300	weighted	jack2	0.522	3.09e-4	1.35e-9
logistic	256	4	300	weighted	raw	3.998	0.014	1.70e-7
logistic	256	16	300	equal	jack2	4.383	0.020	7.36e-7
logistic	256	16	300	equal	raw	8.944	0.083	1.02e-5
logistic	256	16	300	weighted	jack2	0.526	3.29e-4	1.36e-9
logistic	256	16	300	weighted	raw	4.061	0.014	1.84e-7
logistic	1024	4	300	equal	jack2	3.611	0.014	3.54e-7
logistic	1024	4	300	equal	raw	19.797	0.345	1.09e-4
logistic	1024	4	300	weighted	jack2	2.102	0.005	4.00e-8
logistic	1024	4	300	weighted	raw	16.619	0.240	5.00e-5
logistic	1024	16	300	equal	jack2	1821.070	2.09e+6	1.81e+15
logistic	1024	16	300	equal	raw	34.580	1.076	0.001
logistic	1024	16	300	weighted	jack2	372.695	1.35e+4	3.49e+8
logistic	1024	16	300	weighted	raw	17.226	0.258	5.78e-5

Source: `results/processed/skew_summary.csv`.

Table 16. Monte Carlo, logit misspec design, $p = 3$, $N = 1,048,576$: scaled gap to $\widehat{\theta}_N$ (root mean square and median) and 95 coverage by local sample size m.

Estimator	$m = 32$, $K = 32,768$			$m = 64$, $K = 16,384$			$m = 256$, $K = 4,096$			$m = 512$, $K = 2,048$		
	gap	median	cov	gap	median	cov	gap	median	cov	gap	median	c
(1,0) plain average	135.980	120.891	0.232	52.024	51.973	0.309	12.064	12.056	0.339	5.968	5.971	0.6
(1,1) average + 1 Newton	13.101	5.157	0.679	0.900	0.898	0.939	0.047	0.047	0.951	0.011	0.011	0.9
(1,2) average + 2 Newton	0.479	0.009	0.949	2.65e-4	2.63e-4	0.944	7.15e-7	7.08e-7	0.951	4.26e-8	4.20e-8	0.9
(2,0) half-block jackknife	7708.810	6143.458	0.001	41.221	16.880	0.262	0.879	0.765	0.944	0.476	0.416	0.9
(2,1) half-block + 1 Newton	4.73e+6	7.22e+4	0.000	1.800	0.090	0.933	2.93e-4	1.47e-4	0.951	7.54e-5	3.69e-5	0.9
(2,2) half-block + 2 Newton	1.75e+16	7.57e+7	–	0.011	2.65e-6	–	5.96e-11	5.76e-12	–	5.86e-12	5.55e-12	
(3,0) quarter-block	1.32e+4	4098.009	–	2684.344	2013.150	–	0.990	0.856	–	0.643	0.574	
(3,1) quarter-block + 1 Newton	5.83e+7	2.02e+4	0.007	1.55e+5	3111.129	0.002	3.15e-4	1.54e-4	0.951	1.31e-4	6.72e-5	0.9
(2,1) valid workers only	–	–	–	–	–	–	–	–	–	–	–	
(3,1) valid workers only	–	–	–	–	–	–	–	–	–	–	–	
analytic Cox–Snell + 1 Newton	1.75e+18	0.094	0.709	0.021	0.021	0.944	0.001	0.001	0.951	3.26e-4	3.09e-4	0.9
CSL, 1 round	51.104	32.959	–	12.528	10.118	–	1.379	1.221	–	0.464	0.395	
CSL, 2 rounds	54.317	13.218	–	6.106	2.649	–	0.264	0.158	–	0.059	0.035	
AEE (Lin–Xi)	80.818	80.816	–	41.700	41.694	–	10.841	10.840	–	5.470	5.473	
(2,1) stale Hessian	5379.284	3239.275	–	4.858	1.102	–	0.013	0.011	–	0.004	0.003	
centralized MLE	0	–	0.951	0	–	0.944	0	–	0.951	0	–	0.9

1000 paired replications per column. gap = root mean square of $\sqrt{N}\,\|\tilde{\theta} - \widehat{\theta}_N\|$ over all replications (including those with failed loc fits); median = its median; cov = empirical coverage of nominal 95% Wald intervals (observed information at the estimate; '–' when n computed). Monte Carlo standard errors of the gaps and RMSEs are in `results/processed/summary.csv`; coverage standard err ≈ 0.007 or less.

Fraction of workers with all local fits valid (orders 2 / 3): $m = 32$: 0.972 / 0.268; $m = 64$: 1.000 / 0.945; $m = 128$: 1.000 / 1.000; $m = 25$ 1.000 / 1.000; $m = 512$: 1.000 / 1.000. Rows 'valid workers only' drop the workers with a failed sub-block fit from the pilot (weigh renormalized) and coincide with the corresponding rows when no fit failed.

Table 17. Monte Carlo, exp contam design, $p = 1$, $N = 1,048,576$: scaled gap to $\widehat{\theta}_N$ (root mean square and median) and 95% coverage by local sample size m.

Estimator	$m = 16$, $K = 65,536$			$m = 64$, $K = 16,384$			$m = 256$, $K = 4,096$		
	gap	median	cov	gap	median	cov	gap	median	cov
(1,0) plain average	101.734	101.718	0.000	30.226	30.203	0.000	8.279	8.269	0.000
(1,1) average + 1 Newton	11.018	11.014	0.000	0.973	0.971	0.704	0.073	0.073	0.818
(1,2) average + 2 Newton	0.129	0.129	0.812	0.001	0.001	0.786	5.71e-6	5.63e-6	0.814
(2,0) half-block jackknife	12.900	12.860	0.000	4.719	4.712	0.016	0.607	0.560	0.781
(2,1) half-block + 1 Newton	0.179	0.176	0.815	0.024	0.024	0.783	4.98e-4	3.34e-4	0.814
(2,2) half-block + 2 Newton	3.51e-5	3.30e-5	–	6.55e-7	5.94e-7	–	5.03e-10	1.19e-10	–
(3,0) quarter-block	21.466	21.450	–	3.068	3.042	–	0.361	0.253	–
(3,1) quarter-block + 1 Newton	0.493	0.490	0.775	0.011	0.010	0.785	2.37e-4	6.78e-5	0.814
(2,1) valid workers only	–	–	–	–	–	–	–	–	–
(3,1) valid workers only	–	–	–	–	–	–	–	–	–
analytic Cox–Snell + 1 Newton	1.432	1.429	0.602	0.242	0.241	0.778	0.022	0.022	0.816
CSL, 1 round	1.39e-13	0	–	3.13e-13	2.27e-13	–	3.19e-13	2.27e-13	–
CSL, 2 rounds	1.35e-13	0	–	3.12e-13	2.27e-13	–	3.16e-13	2.27e-13	–
AEE (Lin–Xi)	120.265	120.271	–	33.285	33.269	–	8.542	8.534	–
(2,1) stale Hessian	1.806	1.797	–	0.191	0.189	–	0.006	0.005	–
centralized MLE	0	–	0.811	0	–	0.786	0	–	0.814

1000 paired replications per column. gap = root mean square of $\sqrt{N}\,\|\tilde{\theta} - \widehat{\theta}_N\|$ over all replications (including those with failed local fits); median = its median; cov = empirical coverage of nominal 95% Wald intervals (observed information at the estimate; '–' when not computed). Monte Carlo standard errors of the gaps and RMSEs are in `results/processed/summary.csv`; coverage standard error ≈ 0.007 or less.

Fraction of workers with all local fits valid (orders 2 / 3): $m = 16$: 1.000 / 1.000; $m = 64$: 1.000 / 1.000; $m = 256$: 1.000 / 1.000. Rows 'valid workers only' drop the workers with a failed sub-block fit from the pilot (weights renormalized) and coincide with the corresponding rows when no fit failed.

10. Measured distributed benchmarks

10.1. Setup

The protocol of Algorithm 1 was run as written—a coordinator process and K worker processes exchanging length-prefixed messages over TCP—on the MacLustr cluster of Apple-silicon Macs connected by a wired local network (mixed 1 and 10 GbE through one switch), with the coordinator on a Mac Studio (M4 Max, 16 cores, 64 GB) and workers on up to sixteen distinct hosts (Mac Studio M3 Ultra, M2 Ultra, M4 Max, M5 Max; Mac mini M4, M4 Pro, M5 Pro, M6; MacBook Pro M4 Max; the coordinator host also hosts one worker in the sixteen-host designs). Workers run Python 3.9 or 3.14 with NumPy 2.0–2.5; each simulates its own block of a logistic regression with $N = 2^{22}$ observations from a seed, so that no data move. The coordinator logs the wall-clock of each stage, every worker's local fit and derivative times, the reply latency of every worker in every phase (so that the slowest worker and the 99th percentile are known), and the bytes of every message including frame headers and serialization. Every design was repeated three times (51 runs in total); Table 18 reports medians and Figure 8 the decomposition of the elapsed time. Worker start-up over `ssh` and the connection phase (about 1.6 seconds) are excluded from the protocol times and reported separately.

10.2. What the measurements show

Three facts emerge, all specific to this hardware and implementation. First, the synchronized derivative phase is cheap: with $p = 8$ it takes 55 ms end to end on sixteen hosts, of which the slowest worker's score-and-Hessian evaluation accounts for 53 ms and the network round trip for about 0.2 ms; the uplink of 15.9 kB is irrelevant at these sizes. With $p = 64$ the phase takes 1.83 s, almost all of it the slowest worker's computation of the 64×64 Hessian on 2^{18} observations (1.82 s), with an uplink of 292 kB. Second, the local phase dominates and grows with the order: at $p = 8$ the slowest worker needs 0.26 s for one fit, 0.49 s for the three fits of the half-block pattern, 0.72 s for the seven fits of the quarter-block pattern and 0.96 s for the fifteen fits of order four; at $p = 64$ the half-block pattern costs 29.5 s against 14.8 s for one fit. Third, the exchange rate of Proposition 1 therefore favours rounds on this network: at equal exponent $7/8$, $(1,2)$ is the fastest protocol at both dimensions ($0.26 + 0.061$ s against $0.49 + 0.055$ s for $(2,1)$ at $p = 8$; $14.8 + 3.64$ s against $29.5 + 1.83$ s at $p = 64$), and it is also the more accurate one (Section 9). The stale and single-machine variants save nothing here, because the derivative phase they shorten was already short; their statistical cost (Theorem 3) is paid for no gain. In terms of (34), the measured $\lambda + T_{\mathrm{der}} + Kb(p + p(p+1)/2)/B + T_{\mathrm{solve}}$ is two to ten times smaller than $r\,c_1 m$ at every (r,p) tested, so the inequality holds and one more round beats one more level of local fits.

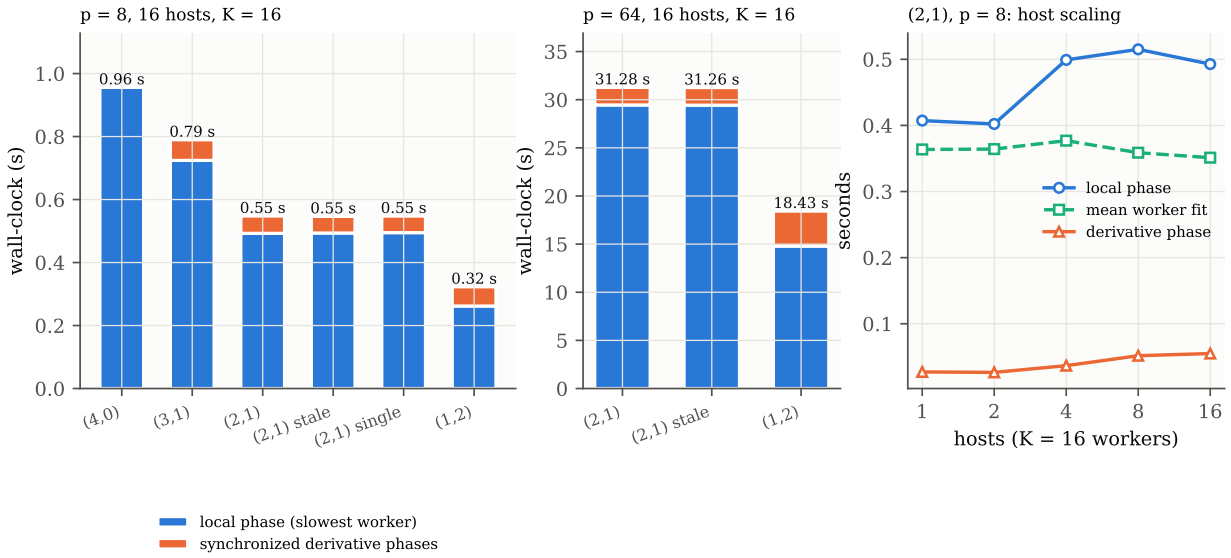


Figure 8. Measured elapsed time of the protocol by variant (medians over three repeats). Left and centre: local phase (slowest worker) and synchronized derivative phases at equal exponent $7/8$ for $(4,0)$, $(3,1)$, $(2,1)$, $(1,2)$ and the stale and single-machine curvature variants of $(2,1)$, logistic regression with $p = 8$ and $p = 64$, $N = 2^{22}$, $K = 16$ workers on 16 hosts. Right: $(2,1)$ with $K = 16$ workers spread over 1, 2, 4, 8 and 16 hosts.

10.3. Hosts, processes and stragglers

Spreading $K = 16$ workers from one host to sixteen (Figure 8, right) does not speed up the local phase: its wall-clock is the slowest worker's, and the slowest of sixteen heterogeneous machines (a Mac mini) is slower than the slowest of sixteen processes on one Mac Studio (0.49 s against 0.41 s), while the mean worker fit time is about the same (0.35 s). The derivative phase roughly doubles from one host to sixteen (27 ms to 55 ms), again through the slowest worker rather than through the network. With $K = 64$ and 128 workers on sixteen hosts (four and eight processes per host, $m = 2^{16}$ and 2^{15}) the local phase falls to 0.10 s for $(2,1)$ and the derivative phase to 7 ms: more partitions mean smaller blocks and faster local fits, but by Theorem 1 they also move the estimator towards the frontier, and the trade is statistical, not computational. The straggler excess, the gap between the slowest and the mean worker, is the quantity that an implementation should attack first; in these runs it comes from hardware heterogeneity rather than from the network.

10.4. Limits of the benchmark

The measurements are for one cluster, one Python/NumPy implementation with pickle framing, one network, and $K \leq 128$; they do not establish a general speed advantage of any protocol. On a network with latency of tens of milliseconds, or with bandwidth-limited uplinks of dense Hessians at large p, the left-hand side of (34) grows and the balance can tip towards local levels; the cost

model is provided so that (34) can be evaluated with the parameters of a given installation. No energy or monetary cost was measured, and none is claimed.

Table 18. Measured multi-node benchmarks of the protocol (logistic regression, $N = 2^{22}$, medians over repeats): wall-clock
local phase (slowest worker) and of the synchronized derivative phases, worker compute times, 99th-percentile reply latenc
bytes moved.

Design	K	hosts	p	r	q	curvature	local (s)	fit max (s)	rounds (s)	per round (s)	der. max (s)	lat. p99 (s)	up (kB)	do
K128 r1q2	128	16	8	1	2	exact	0.059	0.038	0.014	0.007	0.005	0.006	199.7	
K128 r2q1	128	16	8	2	1	exact	0.097	0.076	0.007	0.007	0.005	0.006	127.5	
K64 r1q2	64	16	8	1	2	exact	0.085	0.064	0.019	0.009	0.008	0.009	99.9	
K64 r2q1	64	16	8	2	1	exact	0.148	0.127	0.009	0.009	0.008	0.009	63.8	
cmp p64 r1q2	16	16	64	1	2	exact	14.792	14.682	3.635	1.818	1.819	1.812	569.8	
cmp p64 r2q1	16	16	64	2	1	exact	29.452	29.342	1.826	1.826	1.821	1.818	291.9	
cmp p64 r2q1 stale	16	16	64	2	1	stale	29.431	29.320	1.828	1.828	1.825	1.822	550.1	
cmp r1q2	16	16	8	1	2	exact	0.262	0.231	0.061	0.031	0.030	0.031	25.0	
cmp r2q1 single	16	16	8	2	1	single	0.495	0.464	0.053	0.053	0.051	0.051	15.9	
cmp r2q1 stale	16	16	8	2	1	stale	0.494	0.463	0.053	0.053	0.051	0.051	19.7	
cmp r3q1	16	16	8	3	1	exact	0.724	0.693	0.065	0.065	0.063	0.063	16.0	
cmp r4q0	16	16	8	4	0	exact	0.956	0.925	0.000	0.000	0.000	0.000	7.0	
hosts1	16	1	8	2	1	exact	0.407	0.371	0.027	0.027	0.025	0.025	16.0	
hosts16	16	16	8	2	1	exact	0.493	0.462	0.055	0.055	0.053	0.054	15.9	
hosts2	16	2	8	2	1	exact	0.402	0.369	0.026	0.026	0.024	0.025	16.0	
hosts4	16	4	8	2	1	exact	0.499	0.463	0.037	0.037	0.035	0.036	16.0	
hosts8	16	8	8	2	1	exact	0.515	0.483	0.052	0.052	0.049	0.050	15.9	

Python 3.9/3.14 + NumPy over TCP (length-prefixed pickle frames) on the MacLustr LAN (1 GbE/10 GbE mixed); coordinator on
Studio M4 Max. 'local' is the wall-clock from the first 'fit' broadcast to the last upload, i.e. the slowest worker; 'rounds' sums the synchr
derivative phases. Bytes include frame headers and pickle overhead. Repeats and raw logs: `results/raw/benchmarks.jsonl`.

11. Extensions: unequal blocks and misspecification

Both extensions are known in the literature (Zhu et al., 2021; Gu and Chen, 2023; Jordan et al., 2019) and are stated here only so that the frontier can be used in practice.

Proposition 2 (Unequal block sizes and pseudo-likelihood). *(i) Let the blocks have sizes $m_1, \dots, m_K$, with $\sum_k m_k = N$, with $m_{\min} \to \infty$, and with each s_j dividing every m_k. With the size weights m_k/N in the pilot,*

$$\theta_r^{(0)} - \widehat{\theta}_N = O_p\Big(m_{\text{eff}}^{-r} + \frac{\sqrt{K}}{N} + \frac{1}{N}\Big), \qquad m_{\text{eff}} := \Big(\frac{1}{N}\sum_{k=1}^{K} m_k^{1-r}\Big)^{-1/r} \geq m_{\min}, \tag{35}$$

and Theorem 1 holds with m replaced by m_{eff}; the aggregated derivatives are size-weighted and equal the full-sample derivatives, so the Newton steps are unchanged. With equal weights $1/K$ the influence term is not preserved:

$$\frac{1}{K}\sum_{k=1}^{K} U_k - U_N = \sum_{i=1}^{N} c_i\,\psi(X_i), \qquad c_i = \frac{1}{K\,m_{k(i)}} - \frac{1}{N}, \qquad \sum_{i=1}^{N} c_i^2 = \frac{1}{K^2}\sum_{k=1}^{K}\frac{1}{m_k} - \frac{1}{N} \geq 0, \tag{36}$$

with equality only for equal sizes; the discrepancy is of exact order $N^{-1/2}$ when the sizes are bounded away from equality, so the equal-weight pilot is not first-order equivalent for $q = 0$, while for $q \geq 1$ the first Newton step squares it to $O_p(N^{-1})$.

(ii) If the model is misspecified and θ_0 is the pseudo-true value, all statements hold with

$$\psi(x) = -\bar{H}(\theta_0)^{-1}\nabla_\theta \log f(x;\theta_0), \qquad \Sigma_\psi = \bar{H}(\theta_0)^{-1}\bar{J}(\theta_0)\bar{H}(\theta_0)^{-1}, \tag{37}$$

provided Assumptions 1 to 3 are read for the pseudo-MLE. The estimator is then first-order equivalent to the centralized pseudo-MLE and Wald intervals require the sandwich of Corollary 1; observed-information intervals are not valid.

The proof is in Section A; the Monte Carlo evidence is in Section 9 (skewed partitions with size ratios 4 and 16, and an omitted-nonlinearity logistic design). Heterogeneous distributions across blocks, in which the local pseudo-true values differ, are outside the scope: the common-influence identity of Lemma 1 fails, and the appropriate target and weighting are those of Gu and Chen (2023) and Duan et al. (2022).

12. Discussion and limitations

What is known and what is new. The local correction is the generalized jackknife, the global correction is the distributed one-step estimator, and the idea of debiasing before averaging has analytic, delete-one and subsampling versions in the literature (Section 2). The joint frontier of Theorem 1, its sharpness in Theorem 2 and the inexact-curvature frontier with the bit requirement of Theorem 3 are the new statements; the Monte Carlo and the measured benchmarks are new evidence. We have found no prior statement of the exponents $7/8$ or $11/12$, of a frontier indexed by both r and q, or of a sharpness result for a bias-corrected construction, but we have not been able to check every unpublished thesis (the replication package lists the sources that could not be read in full).

Fixed dimension. All exponents are for fixed p. The constants C, C_b, M_N, Λ_N in Assumptions 1 to 3 depend on p, the Hessian payload grows as p^2, and the number of observations in the finest sub-block is m/s_r, so that for p comparable to m/s_r the local MLEs fail to exist with non-negligible probability (Section 9, $p = 20$). The results do not transfer to $p_N \to \infty$ without a new analysis of how every constant scales, and we make no claim in that direction.

Finite samples versus exponents. The frontier orders estimators by the exponent of the bias term alone. At moderate m the centered random remainder, whose constant grows with A_r, and the alternating signs of β_r make the order-3 pilot worse than the order-2 pilot in RMSE for all but the smallest m in our designs, and the one-step versions inherit part of this ranking. Doubling the local order to save a round is therefore not free statistically, in addition to its computational price; the higher-order variance results of Hahn et al. (2024) describe the same phenomenon for split-sample corrections of cross-section MLEs.

Adaptive order. Choosing r from the data (for instance by comparing successive pilots) changes the estimator, and nothing in Theorem 1 covers a data-dependent r. Sample splitting or a selection event with uniform bounds would be needed to restore a theorem; we leave this open and do not report an adaptive estimator.

Asynchrony, failures and stragglers. The theorems describe synchronized phases: every worker's derivatives are evaluated at the same broadcast point. Asynchronous or stale-iterate updates change the estimator and inherit no guarantee from the present analysis; worker failures were handled in the benchmarks by waiting, not by dropping workers, and the theory for dropped blocks is that of Huang and Huo (2019, Corollary 3.7) rather than ours.

Lower bounds. Nothing here is a minimax or communication lower bound. In exponential families the pooled sufficient statistic recovers $\widehat{\theta}_N$ exactly in one round of $\dim(T)$ numbers per worker, and the exponential model shows that the full surrogate-likelihood maximization is also exact there. The sharpness in Theorem 2 concerns the estimator family $\{\theta_r^{(q)}\}$, not all distributed procedures.

Benchmarks. Measured times are specific to one cluster, one implementation (Python/NumPy over TCP with pickle framing) and the designs run. They calibrate (33) and show which side of (34) the designs fall on; they do not establish a general speed advantage of any protocol, and we report no energy or monetary cost because none was measured.

13. Conclusion

For regular fixed-dimensional likelihoods, local split-sample extrapolation of order r followed by q aggregated Newton corrections is first-order equivalent to the centralized MLE whenever $K = o(N^{1-1/(2^{q+1}r)})$, and this exponent is sharp for the construction in an exactly solvable model, where the limiting shift at the boundary is $-\theta_0(\beta_r c^r)^{2^q}$. One synchronization round and a doubling of the local order are exchangeable in the exponent, which turns the choice between them into a question about hardware that a calibrated cost model and measured benchmarks can answer for a given installation. Inexact curvature changes the exponent in a computable way, recovering the known single-machine-Hessian frontiers and showing that higher local order tolerates coarser quantization of the transmitted derivatives. The finite-sample evidence tempers the exponents: higher-order corrections inflate variance and can fail to exist at small block sizes, and the boundary regime is reached only at very large N. The components combined here are known; the joint frontier, its sharpness and the bit-precision result are the contribution.

References

Battey, Heather, Jianqing Fan, Han Liu, Junwei Lu, and Ziwei Zhu, "Distributed Testing and Estimation under Sparse High Dimensional Models," *Annals of Statistics*, 2018, *46* (3), 1352–1382.

Bhattacharya, R. N. and J. K. Ghosh, "On the Validity of the Formal Edgeworth Expansion," *Annals of Statistics*, 1978, *6* (2), 434–451.

Bickel, P. J., "One-Step Huber Estimates in the Linear Model," *Journal of the American Statistical Association*, 1975, *70* (350), 428–434.

Cam, Lucien Le, "On the Asymptotic Theory of Estimation and Testing Hypotheses," in "Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability," Vol. 1, Berkeley: University of California Press, 1956, pp. 129–156.

Chen, Xi, Weidong Liu, and Yichen Zhang, "First-Order Newton-Type Estimator for Distributed Estimation and Inference," *Journal of the American Statistical Association*, 2022, *117* (540), 1858–1874.

Cordeiro, Gauss M. and Peter McCullagh, "Bias Correction in Generalized Linear Models," *Journal of the Royal Statistical Society: Series B*, 1991, *53* (3), 629–643.

Cox, D. R. and E. J. Snell, "A General Definition of Residuals," *Journal of the Royal Statistical Society: Series B*, 1968, *30* (2), 248–275.

Dhaene, Geert and Koen Jochmans, "Split-Panel Jackknife Estimation of Fixed-Effect Models," *Review of Economic Studies*, 2015, *82* (3), 991–1030.

Duan, Rui, Yang Ning, and Yong Chen, "Heterogeneity-Aware and Communication-Efficient Distributed Statistical Inference," *Biometrika*, 2022, *109* (1), 67–83.

Fan, Jianqing, Yongyi Guo, and Kaizheng Wang, "Communication-Efficient Accurate Statistical Estimation," *Journal of the American Statistical Association*, 2023, *118* (542), 1000–1010.

Gao, Yuan, Weidong Liu, Hansheng Wang, Xiaozhou Wang, Yibo Yan, and Riquan Zhang, "A Review of Distributed Statistical Inference," *Statistical Theory and Related Fields*, 2022, *6* (2), 89–99.

Gu, Jia and Songxi Chen, "Distributed Statistical Inference under Heterogeneity," *Journal of Machine Learning Research*, 2023, *24* (387), 1–56.

Hahn, Jinyong, David W. Hughes, Guido Kuersteiner, and Whitney K. Newey, "Efficient Bias Correction for Cross-Section and Panel Data," *Quantitative Economics*, 2024, *15* (3), 783–816.

Huang, Cheng and Xiaoming Huo, "A Distributed One-Step Estimator," *Mathematical Programming*, 2019, *174* (1–2), 41–76.

Jordan, Michael I., Jason D. Lee, and Yun Yang, "Communication-Efficient Distributed Statistical Inference," *Journal of the American Statistical Association*, 2019, *114* (526), 668–681.

Li, Xuetong, Yuan Gao, Hong Chang, Danyang Huang, Yingying Ma, Rui Pan, Haobo Qi, Feifei Wang, Shuyuan Wu, Ke Xu, Jing Zhou, Xuening Zhu, Yingqiu Zhu, and Hansheng Wang, "A Selective Review on Statistical Methods for Massive Data Computation: Distributed Computing, Subsampling, and Minibatch Techniques," 2024. arXiv:2403.11163.

Lin, Nan and Ruibin Xi, "Aggregated Estimating Equation Estimation," *Statistics and Its Interface*, 2011, *4* (1), 73–83.

Quenouille, M. H., "Approximate Tests of Correlation in Time-Series," *Journal of the Royal Statistical Society: Series B*, 1949, *11* (1), 68–84.

_ , "Notes on Bias in Estimation," *Biometrika*, 1956, *43* (3/4), 353–360.

Rosenblatt, Jonathan D. and Boaz Nadler, "On the Optimality of Averaging in Distributed Statistical Learning," *Information and Inference: A Journal of the IMA*, 2016, *5* (4), 379–404.

Schucany, W. R., H. L. Gray, and D. B. Owen, "On Bias Reduction in Estimation," *Journal of the American Statistical Association*, 1971, *66* (335), 524–533.

Shamir, Ohad, Nathan Srebro, and Tong Zhang, "Communication-Efficient Distributed Optimization Using an Approximate Newton-Type Method," in "Proceedings of the 31st International Conference on Machine Learning," Vol. 32 of *Proceedings of Machine Learning Research* PMLR 2014, pp. 1000–1008.

Shang, Zuofeng and Guang Cheng, "Computational Limits of a Distributed Algorithm for Smoothing Spline," *Journal of Machine Learning Research*, 2017, *18* (108), 1–37.

van der Vaart, A. W., *Asymptotic Statistics*, Cambridge: Cambridge University Press, 1998.

Volgushev, Stanislav, Shih-Kang Chao, and Guang Cheng, "Distributed Inference for Quantile Regression Processes," *Annals of Statistics*, 2019, *47* (3), 1634–1662.

Wang, Jialei, Mladen Kolar, Nathan Srebro, and Tong Zhang, "Efficient Distributed Learning with Sparsity," in "Proceedings of the 34th International Conference on Machine Learning," Vol. 70 of *Proceedings of Machine Learning Research* PMLR 2017, pp. 3636–3645.

Wu, Shuyuan, Xuening Zhu, and Hansheng Wang, "Subsampling and Jackknifing: A Practically Convenient Solution for Large Data Analysis with Limited Computational Resources," *Statistica Sinica*, 2023, *33*, 2041–2064.

Yin, Zhuoran, "Communication-Efficient Distributed Newton-Like Optimization with Gradients and M-Estimators," 2022. arXiv:2207.06253.

Zhang, Yuchen, John C. Duchi, and Martin J. Wainwright, "Communication-Efficient Algorithms for Statistical Optimization," *Journal of Machine Learning Research*, 2013, *14*, 3321–3363.

Zhu, Xuening, Feng Li, and Hansheng Wang, "Least-Square Approximation for a Distributed System," *Journal of Computational and Graphical Statistics*, 2021, *30* (4), 1004–1018.

Appendices

A. Proofs of the general results

Throughout, $S_{j\ell,k}$ ($\ell = 1,\dots,s_j$) denotes the ℓ-th sub-block of level j of worker k, of size m/s_j; $U_S = |S|^{-1}\sum_{i\in S}\psi(X_i)$; w solves (9).

Lemma 1 (Influence preservation)**.** *For every k and j,*

$$\frac{1}{s_j}\sum_{\ell=1}^{s_j} U_{S_{j\ell,k}} = U_k, \qquad \textit{and hence} \qquad \sum_{j=1}^{r} w_j \frac{1}{s_j}\sum_{\ell=1}^{s_j} U_{S_{j\ell,k}} = U_k.$$

Proof. The sub-blocks of level j partition D_k into s_j sets of equal size m/s_j, so

$$\frac{1}{s_j}\sum_{\ell=1}^{s_j} U_{S_{j\ell,k}} = \frac{1}{s_j}\sum_{\ell=1}^{s_j}\frac{s_j}{m}\sum_{i\in S_{j\ell,k}}\psi(X_i) = \frac{1}{m}\sum_{i\in D_k}\psi(X_i) = U_k.$$

The second claim uses $\sum_j w_j = 1$. □

Lemma 2 (Bias cancellation)**.** *Under Assumptions 1 and 2 with $m/s_r \geq n_0$, $J_{r,k} - \theta_0 = U_k + R_k^J$ with R_k^J as in (10), and*

$$\mathbb{E}R_k^J = \sum_{j=1}^{r} w_j \rho_{m/s_j}, \qquad \|\mathbb{E}R_k^J\| \leq C_b\, m^{-r}\sum_{j=1}^{r}|w_j|\, s_j^r.$$

In particular $\mathbb{E}R_k^J$ is deterministic and does not depend on k.

Proof. Applying (7) to each sub-block and Lemma 1,

$$J_{r,k} - \theta_0 = \sum_{j=1}^{r} w_j \frac{1}{s_j} \sum_{\ell=1}^{s_j} \left(U_{S_{j\ell,k}} + R_{S_{j\ell,k}} \right) = U_k + R_k^J.$$

Taking expectations with (8) at $n = m/s_j$,

$$\mathbb{E}R_k^J = \sum_{j=1}^{r} w_j \Big[\sum_{a=1}^{r-1} b_a \Big(\frac{s_j}{m} \Big)^a + \rho_{m/s_j} \Big] = \sum_{a=1}^{r-1} \frac{b_a}{m^a} \underbrace{\sum_{j=1}^{r} w_j s_j^a}_{=0} + \sum_{j=1}^{r} w_j \rho_{m/s_j},$$

and the inner sums vanish by (9). The bound follows from $\|\rho_{m/s_j}\| \le C_b (s_j/m)^r$. □

Lemma 3 (Local remainder)**.** *Under Assumption 1, $(\mathbb{E}\|R_k^J\|^2)^{1/2} \le \sqrt{C} A_r/m$ with $A_r = \sum_j |w_j| s_j$.*

Proof. By Minkowski's inequality in L_2,

$$\left(\mathbb{E}\|R_k^J\|^2 \right)^{1/2} \le \sum_{j=1}^{r} \frac{|w_j|}{s_j} \sum_{\ell=1}^{s_j} \left(\mathbb{E}\|R_{S_{j\ell,k}}\|^2 \right)^{1/2} \le \sum_{j=1}^{r} \frac{|w_j|}{s_j} \cdot s_j \cdot \sqrt{C} \frac{s_j}{m} = \frac{\sqrt{C} A_r}{m}.$$

No independence between the nested sub-samples of a worker is used. □

Lemma 4 (Averaged centered remainder)**.** *Under Assumption 1, $\mathbb{E}\left\| K^{-1} \sum_{k=1}^{K} (R_k^J - \mathbb{E}R_k^J) \right\|^2 \le C A_r^2/(Km^2)$.*

Proof. The blocks are disjoint, so the vectors $R_k^J - \mathbb{E}R_k^J$ are independent and centered; the second moment of the average is $K^{-2} \sum_k \mathbb{E}\|R_k^J - \mathbb{E}R_k^J\|^2 \le K^{-1} \sup_k \mathbb{E}\|R_k^J\|^2 \le C A_r^2/(Km^2)$ by Lemma 3. □

Lemma 5 (Pilot versus centralized MLE)**.** *Under Assumptions 1 and 2, $\theta_r^{(0)} - \widehat{\theta}_N = \mathbb{E}R_1^J + K^{-1} \sum_k (R_k^J - \mathbb{E}R_k^J) - R_N = O_p(m^{-r} + (m\sqrt{K})^{-1} + N^{-1})$.*

Proof. Averaging Lemma 2 over k and using $K^{-1} \sum_k U_k = N^{-1} \sum_{i=1}^{N} \psi(X_i) = U_N$ (equal block sizes),

$$\theta_r^{(0)} - \theta_0 = U_N + \mathbb{E}R_1^J + \frac{1}{K} \sum_{k=1}^{K} \left(R_k^J - \mathbb{E}R_k^J \right), \qquad \widehat{\theta}_N - \theta_0 = U_N + R_N,$$

the second identity being (7) for the full sample. Subtracting, the influence terms cancel exactly. The orders follow from Lemmas 2 and 4 and $\mathbb{E}\|R_N\|^2 \le C N^{-2}$ via Markov's inequality. □

Lemma 6 (Quadratic contraction of one Newton step)**.** *Under Assumption 3, let E_N be the event (of probability tending to one) on which $\widehat{\theta}_N \in U$ is an interior zero of g_N and the bounds M_N, Λ_N hold. On $E_N \cap \{\vartheta \in U\}$, $\|\mathcal{N}_N(\vartheta) - \widehat{\theta}_N\| \le \frac{1}{2} M_N \Lambda_N \|\vartheta - \widehat{\theta}_N\|^2$.*

Proof. Let $e = \vartheta - \widehat{\theta}_N$. Since U is convex, $\widehat{\theta}_N + te \in U$ for $t \in [0,1]$, and because $g_N(\widehat{\theta}_N) = 0$,

$$g_N(\vartheta) = \int_0^1 H_N(\widehat{\theta}_N + te)\, e \, \mathrm{d}t.$$

Hence

$$\begin{aligned} \mathcal{N}_N(\vartheta) - \widehat{\theta}_N &= e - H_N(\vartheta)^{-1} g_N(\vartheta) \\ &= H_N(\vartheta)^{-1} \int_0^1 \big[H_N(\vartheta) - H_N(\widehat{\theta}_N + te)\big]\, e \, \mathrm{d}t. \end{aligned}$$

The Lipschitz bound gives $\|H_N(\vartheta) - H_N(\widehat{\theta}_N + te)\| \le \Lambda_N (1-t)\|e\|$, and $\int_0^1 (1-t)\,\mathrm{d}t = 1/2$. □

Lemma 7 (q steps)**.** *Let $\theta^{(0)} \xrightarrow{p} \theta_0$ and define $\theta^{(j+1)} = \mathcal{N}_N(\theta^{(j)})$. Under Assumption 3, with probability tending to one all iterates $\theta^{(0)}, \dots, \theta^{(q)}$ lie in U and $\|\theta^{(q)} - \widehat{\theta}_N\| \le (\frac{1}{2} M_N \Lambda_N)^{2^q - 1} \|\theta^{(0)} - \widehat{\theta}_N\|^{2^q}$.*

Proof. Let $\kappa_N = \frac{1}{2} M_N \Lambda_N = O_p(1)$ and $e_j = \|\theta^{(j)} - \widehat{\theta}_N\|$. Since $e_0 \xrightarrow{p} 0$ and $\widehat{\theta}_N \xrightarrow{p} \theta_0$, there is a ball $B \subset U$ around θ_0 such that $P(\theta^{(0)} \in B, \widehat{\theta}_N \in B) \to 1$. On E_N, induction with Lemma 6 gives $e_{j+1} \le \kappa_N e_j^2$ and, because $\kappa_N e_0 \xrightarrow{p} 0$, $e_{j+1} \le e_j$ eventually, so that each iterate remains in the ball and the recursion applies q times: $e_q \le \kappa_N^{1+2+\cdots+2^{q-1}} e_0^{2^q} = \kappa_N^{2^q - 1} e_0^{2^q}$. □

Proof of Theorem 1. By Lemma 5 and $m \to \infty$ (which forces $K = o(N)$), $\theta_r^{(0)} - \theta_0 = U_N + O_p(1) \xrightarrow{p} 0$. Lemma 7 applies and, with $\kappa_N = O_p(1)$,

$$\sqrt{N}\|\theta_r^{(q)} - \widehat{\theta}_N\| = O_p\Big(\sqrt{N}\big[m^{-r} + (m\sqrt{K})^{-1} + N^{-1}\big]^{2^q}\Big) \le O_p\Big(3^{2^q}\sqrt{N}\big[m^{-r2^q} + (m\sqrt{K})^{-2^q} + N^{-2^q}\big]\Big),$$

using $(a+b+c)^{2^q} \le 3^{2^q} \max(a,b,c)^{2^q}$. With $m = N/K$ the three terms are

$$\sqrt{N} m^{-r2^q} = N^{1/2}\Big(\frac{K}{N}\Big)^{r2^q} \to 0 \quad \Longleftrightarrow \quad K^{r2^q} = o\big(N^{r2^q - 1/2}\big) \quad \Longleftrightarrow \quad K = o\big(N^{1 - 1/(2^{q+1} r)}\big);$$

$$\sqrt{N}\,(m\sqrt{K})^{-2^q} = N^{1/2} K^{2^{q-1}} N^{-2^q} = \begin{cases} (K/N)^{1/2} \to 0, & q = 0 \text{ (since } K = o(N)), \\ \le N^{1/2 - 2^{q-1}} \le N^{-1/2} \to 0, & q \ge 1,\ K \le N; \end{cases}$$

$$\sqrt{N} N^{-2^q} \to 0.$$

The limit law follows from $\sqrt{N}(\widehat{\theta}_N - \theta_0) \Rightarrow \mathcal{N}(0, \Sigma_\psi)$ and Slutsky's lemma. □

Proof of Corollary 1. $\sqrt{N}(\theta_r^{(q)} - \theta_0) \Rightarrow \mathcal{N}(0, \Sigma_\psi)$ and $\widehat{\Sigma}_N \xrightarrow{p} \Sigma_\psi$ give $\sqrt{N}\widehat{\Sigma}_{N,jj}^{-1/2}(\theta_r^{(q,j)} - \theta_{0,j}) \Rightarrow \mathcal{N}(0,1)$ by Slutsky's lemma. Consistency of $-H_N(\theta_r^{(q)})^{-1}$ under correct specification follows from Assumption 3 (Lipschitz continuity of H_N on U, $\theta_r^{(q)} \xrightarrow{p} \theta_0$) and $-H_N(\theta_0) \xrightarrow{p} I(\theta_0)$; the sandwich case is analogous with a law of large numbers for the aggregated score outer products. □

Proof of Theorem 3. Let $e = \vartheta - \widehat{\theta}_N$ and $\bar{H} := \int_0^1 H_N(\widehat{\theta}_N + te)\,\mathrm{d}t$, so that $g_N(\vartheta) = \bar{H}e$ on $E_N \cap \{\vartheta \in U\}$. Then

$$\tilde{\vartheta} - \widehat{\theta}_N = e - \widehat{H}^{-1}(\bar{H}e + \delta_g) = \widehat{H}^{-1}\big(\widehat{H} - \bar{H}\big)\,e - \widehat{H}^{-1}\delta_g,$$

and

$$\|\widehat{H} - \bar{H}\| \le \|\widehat{H} - H_N(\widehat{\theta}_N)\| + \int_0^1 \|H_N(\widehat{\theta}_N) - H_N(\widehat{\theta}_N + te)\|\,\mathrm{d}t \le \|\widehat{H} - H_N(\widehat{\theta}_N)\| + \tfrac{1}{2}\Lambda_N\|e\|,$$

which gives (27). Iterating with a fixed $\widehat{H}$ (or with a fresh $\widehat{H}$ of the same accuracy at each step) yields $e_{j+1} = O_p(\delta_N e_j + e_j^2)$ and, by induction, $e_q = O_p(e_0\delta_N^q + \sum_{\text{mixed}} \cdots + e_0^{2^q})$ in which every mixed term is dominated by $e_0\delta_N^q$ or $e_0^{2^q}$ once $e_0 \to 0$ and $\delta_N \to 0$; the iterates remain in U as in Lemma 7 because $\delta_N + \kappa_N e_0 \xrightarrow{p} 0$. With $e_0 = O_p(m^{-r} + \sqrt{K}/N + N^{-1})$ the binding condition is $\sqrt{N}\,m^{-r}\delta_N^q \to 0$ (the terms $\sqrt{N}(\sqrt{K}/N)\delta_N^q$ and $\sqrt{N}N^{-1}\delta_N^q$ vanish for any $\delta_N \le 1$ and $K \le N$), together with the exact-curvature condition $\sqrt{N}e_0^{2^q} \to 0$ of Theorem 1.

(a) Write

$$\widehat{H} - H_N(\widehat{\theta}_N) = \frac{1}{K}\sum_{k=1}^{K}\big[H_{D_k}(\widehat{\theta}_{D_k}) - H_{D_k}(\theta_0)\big] + \big[H_N(\theta_0) - H_N(\widehat{\theta}_N)\big].$$

The second bracket is $O_p(\|\widehat{\theta}_N - \theta_0\|) = O_p(N^{-1/2})$ by Assumption 3. For the first, Assumption 4 gives, for each k and $\Delta_k := \widehat{\theta}_{D_k} - \theta_0$,

$$H_{D_k}(\widehat{\theta}_{D_k}) - H_{D_k}(\theta_0) = \nabla H_{D_k}(\theta_0)[\Delta_k] + r_k, \qquad \|r_k\| \le \tfrac{1}{2}\bar{L}_k\|\Delta_k\|^2,$$

with $\bar{L}_k$ the block average of $L(X_i)$. By Cauchy–Schwarz and Assumption 4,

$$\mathbb{E}\|r_k\| \le \tfrac{1}{2}\big(\mathbb{E}\bar{L}_k^2\big)^{1/2}\big(\mathbb{E}\|\Delta_k\|^4\big)^{1/2} = O(m^{-1}).$$

Write $\nabla H_{D_k}(\theta_0) = T + V_k$ with $T = \mathbb{E}\nabla h(X;\theta_0)$ and V_k a centered block average with $\mathbb{E}\|V_k\|^2 = O(m^{-1})$, and $\Delta_k = U_k + R_{D_k}$. Then $K^{-1}\sum_k T[U_k] = T[U_N] = O_p(N^{-1/2})$; $K^{-1}\sum_k T[R_{D_k}] = O_p(m^{-1})$ by Assumption 1; and $\mathbb{E}\|V_k[U_k]\| \le (\mathbb{E}\|V_k\|^2)^{1/2}(\mathbb{E}\|U_k\|^2)^{1/2} = O(m^{-1})$. Collecting terms, $\delta_N = m^{-1} + N^{-1/2}$, and the condition $\sqrt{N}m^{-r}m^{-q} \to 0$ reads $K = o(N^{1-1/(2(r+q))})$; the $N^{-1/2}$ part of δ_N never binds. Since $r + q \le 2^q r$ for $q \ge 1$, this condition implies that of Theorem 1.

(b) Decompose

$$H_{D_1}(\vartheta) - H_N(\widehat{\theta}_N) = \big[H_{D_1}(\vartheta) - H_{D_1}(\theta_0)\big] + \big[H_{D_1}(\theta_0) - \bar{H}(\theta_0)\big] + \big[\bar{H}(\theta_0) - H_N(\theta_0)\big] + \big[H_N(\theta_0) - H_N(\widehat{\theta}_N)\big],$$

where $\bar{H}(\theta_0) = \mathbb{E}h(X;\theta_0)$. The first and last brackets are $O_p(\|\vartheta - \theta_0\| + N^{-1/2})$ by the Lipschitz property (applied to H_{D_1}, for which the same bound is assumed as for H_N on U, with probability tending to one); the second is $O_p(m^{-1/2})$ and the third $O_p(N^{-1/2})$ by the central limit theorem for block averages of $h(X_i;\theta_0)$. Hence $\delta_N = m^{-1/2} + \|\vartheta - \theta_0\|$, and along the iteration $\|\theta^{(j)} - \theta_0\| = O_p(N^{-1/2} + e_j)$, so the binding condition is $\sqrt{N}m^{-r}m^{-q/2} \to 0$, i.e. $K = o(N^{1-1/(2r+q)})$.

(c) Rounding each entry of $H_N(\vartheta)$ to b significant bits yields $\widehat{H}$ with $|\widehat{H}_{ab} - H_{N,ab}| \le 2^{-b}|H_{N,ab}|$, hence $\|\widehat{H} - H_N(\vartheta)\| \le p\,2^{-b}\|H_N(\vartheta)\|_{\max}$, and $\|H_N(\vartheta) - H_N(\widehat{\theta}_N)\| \le \Lambda_N\|e\|$; so $\delta_N = O_p(\varepsilon) + O_p(\|e\|)$, the second part being absorbed in the $\|e\|^2$ term. Likewise $\|\delta_g\| \le p\,\varepsilon\,\|g_N(\vartheta)\|_{\max} \le p\,\varepsilon\,\|\bar{H}\|\|e\| = O_p(\varepsilon\|e\|)$. Inserting into (27) with $\|e_0\| = O_p(m^{-r} + \sqrt{K}/N + N^{-1})$, first-order equivalence requires $\sqrt{N}\varepsilon m^{-r} \to 0$ (the other two terms vanish for $\varepsilon \le 1$). With $K \sim N^{\alpha}$, $m^{-r} = N^{-r(1-\alpha)}$ up to constants, so $\sqrt{N}\varepsilon m^{-r} \to 0$ iff $\log_2(1/\varepsilon) = b \ge (\frac{1}{2} - r(1-\alpha))\log_2 N + \omega(1)$ when $\frac{1}{2} - r(1-\alpha) \ge 0$, i.e. $\alpha \ge 1 - 1/(2r)$; otherwise $\sqrt{N}m^{-r} \to 0$ already and any $b \ge 0$ works. □

Proof of Proposition 2. (i) With weights m_k/N, $\sum_k (m_k/N)U_k = N^{-1}\sum_i \psi(X_i) = U_N$, so Lemma 5 holds with $\mathbb{E}R^J$ replaced by $\sum_k (m_k/N)\mathbb{E}R_k^J$, whose norm and whose centered second moment satisfy

$$\Big\|\sum_k \frac{m_k}{N}\mathbb{E}R_k^J\Big\| \le C_b\Big(\sum_j |w_j|s_j^r\Big)\frac{1}{N}\sum_{k=1}^K m_k^{1-r} = C_b\Big(\sum_j |w_j|s_j^r\Big)m_{\text{eff}}^{-r}, \qquad \sum_{k=1}^K \Big(\frac{m_k}{N}\Big)^2 \frac{CA_r^2}{m_k^2} = \frac{CA_r^2 K}{N^2}.$$

The Newton steps use $g_N = \sum_k (m_k/N)g_{D_k}$ and $H_N = \sum_k (m_k/N)H_{D_k}$, which are the full-sample derivatives, so Lemmas 6 and 7 are unchanged. With equal weights, $K^{-1}\sum_k U_k - U_N = \sum_i c_i\psi(X_i)$ with $c_i = (Km_{k(i)})^{-1} - N^{-1}$; its covariance is $(\sum_i c_i^2)\Sigma_\psi$ and $\sum_i c_i^2 = \sum_k m_k[(Km_k)^{-1} - N^{-1}]^2 = K^{-2}\sum_k m_k^{-1} - N^{-1}$, which is ≥ 0 with equality iff all m_k are equal (Jensen); when the sizes are bounded away from equality in the sense $K^{-1}\sum_k (N/(Km_k)) \ge 1 + \eta$, this is $\ge \eta/N$, of exact order N^{-1}, so the discrepancy is of exact order $N^{-1/2}$ and $\sqrt{N}(\theta^{(0)} - \widehat{\theta}_N)$ does not vanish. One Newton step squares it to $O_p(N^{-1})$ by Lemma 6. (ii) The only properties of ψ used in Lemmas 1 to 5 are $\mathbb{E}\psi = 0$ and that the same ψ appears in (7) for every sub-sample; Assumption 3 is stated for the pseudo-MLE with the sandwich limit. The proofs go through verbatim. □

B. Proofs for the exponential model

Let $Y = \theta_0 \bar{X}_n$, so that $nY \sim \text{Gamma}(n,1)$, $\widehat{\theta}_n = \theta_0/Y$ and $R_n = \theta_0(1-Y)^2/Y = \theta_0(Y^{-1}-2+Y)$. The Gamma moments $\mathbb{E}(nY)^{-k} = \Gamma(n-k)/\Gamma(n)$ $(k<n)$ give $\mathbb{E}Y = 1$, $\mathbb{E}Y^2 = 1+1/n$, $\mathbb{E}Y^{-1} = n/(n-1)$ and $\mathbb{E}Y^{-2} = n^2/((n-1)(n-2))$.

Proof of Theorem 2(i). From the Gamma moments,

$$\begin{aligned}
\mathbb{E}\widehat{\theta}_n &= \theta_0\,\mathbb{E}Y^{-1} = \theta_0\,\frac{n}{n-1},\\
\operatorname{Var}\widehat{\theta}_n &= \theta_0^2\big(\mathbb{E}Y^{-2} - (\mathbb{E}Y^{-1})^2\big) = \theta_0^2\,\frac{n^2}{(n-1)^2(n-2)},\\
\mathbb{E}R_n &= \theta_0\Big(\frac{n}{n-1} - 2 + 1\Big) = \frac{\theta_0}{n-1} = \theta_0\sum_{a\ge 1} n^{-a},
\end{aligned}$$

so $b_a = \theta_0$ for all a and the remainder $\rho_n = \theta_0\sum_{a\ge r} n^{-a}$ satisfies $\|\rho_n\| \le 2\theta_0 n^{-r}$ for $n \ge 2$. Expanding $(Y^{-1}-2+Y)^2 = Y^{-2} - 4Y^{-1} + 6 - 4Y + Y^2$,

$$\mathbb{E}R_n^2 = \theta_0^2\Big[\frac{n^2}{(n-1)(n-2)} - \frac{4n}{n-1} + 3 + \frac{1}{n}\Big] = \theta_0^2\,\frac{3n+2}{n(n-1)(n-2)},$$

and $n^2\mathbb{E}R_n^2/\theta_0^2 = n(3n+2)/((n-1)(n-2))$ decreases from $33/2$ at $n=3$ to 3, which gives $C = 33\theta_0^2/2$ with $n_0 = 3$. The influence function $\psi(x) = \theta_0 - \theta_0^2 x$ is the same for all sub-samples and $\Sigma_\psi = \theta_0^2 = I(\theta_0)^{-1}$. □

Proof of (ii). A sub-block of level j has $n = m/s_j$ observations, so its MLE has mean $\theta_0 n/(n-1) = \theta_0 m/(m-s_j)$. Using $\sum_j w_j = 1$,

$$\mathbb{E}J_{r,k} - \theta_0 = \theta_0\sum_{j=1}^{r} w_j\Big[\frac{m}{m-s_j} - 1\Big] = \theta_0\sum_{j=1}^{r}\frac{w_j s_j}{m-s_j} = \theta_0\sum_{a\ge 1}\frac{1}{m^a}\sum_{j=1}^{r} w_j s_j^a,$$

where the terms $a < r$ vanish by (9) and the leading term is $\beta_r\theta_0 m^{-r}$. For the two default patterns,

$$\begin{aligned}
s=(1,2):&\qquad \theta_0\Big[\frac{2}{m-1} - \frac{2}{m-2}\Big] = -\frac{2\theta_0}{(m-1)(m-2)},\\
s=(1,2,4):&\qquad \theta_0\Big[\frac{8/3}{m-1} - \frac{4}{m-2} + \frac{4/3}{m-4}\Big] = \frac{8\theta_0}{(m-1)(m-2)(m-4)}.
\end{aligned}$$

Both identities, and the constants $\beta_2 = -2$, $\beta_3 = 8$ and $\beta_4 = -64$, are verified in rational arithmetic in `theory/symbolic_checks.py`. □

Proof of (iii). Let $h = m/2$ and let S_A, S_B be the sums over the two halves, $S = S_A + S_B \sim \text{Gamma}(m,\theta_0)$

and $W = S_A/S \sim \mathrm{Beta}(h,h)$, independent of S. Then

$$\widehat{\theta}_{D_k} = \frac{m}{S}, \qquad \widehat{\theta}_A = \frac{h}{SW}, \qquad \widehat{\theta}_B = \frac{h}{S(1-W)}, \qquad J_{2,k} = \frac{G(W)}{S}, \quad G(W) = 2m - \frac{h}{2}\Big[\frac{1}{W} + \frac{1}{1-W}\Big].$$

Independence gives $\mathbb{E}J_{2,k} = \mathbb{E}S^{-1}\,\mathbb{E}G(W)$ and $\mathbb{E}J_{2,k}^2 = \mathbb{E}S^{-2}\,\mathbb{E}G(W)^2$, with

$$\mathbb{E}S^{-k} = \theta_0^k \frac{\Gamma(m-k)}{\Gamma(m)}, \qquad \mathbb{E}\big[W^{-a}(1-W)^{-c}\big] = \frac{\Gamma(h-a)\Gamma(h-c)\Gamma(2h)}{\Gamma(h)^2\Gamma(2h-a-c)}, \qquad a,c < h.$$

Simplifying the resulting rational functions (SymPy, check S6) gives the stated $\mathrm{Var}\,J_{2,k}$ and the expansion $\theta_0^2(m^{-1} + 4m^{-2} + O(m^{-3}))$; the same expansion holds for $\mathrm{Var}\,\widehat{\theta}_m = \theta_0^2 m^2/((m-1)^2(m-2))$. The condition $m > 4$ ensures $h > 2$ so that all inverse moments exist. □

Proof of (iv). With $g_N(x) = 1/x - \bar{X}_N$ and $H_N(x) = -1/x^2$,

$$\mathscr{N}_N(x) = x + x^2\Big(\frac{1}{x} - \bar{X}_N\Big) = 2x - x^2\bar{X}_N, \qquad \mathscr{N}_N(x) - \widehat{\theta}_N = -\bar{X}_N x^2 + 2x - \frac{1}{\bar{X}_N} = -\bar{X}_N\Big(x - \frac{1}{\bar{X}_N}\Big)^2 = -\frac{(x-\widehat{\theta}_N)^2}{\widehat{\theta}_N},$$

since $\widehat{\theta}_N = 1/\bar{X}_N$. Writing $e_j = \theta^{(j)} - \widehat{\theta}_N$ gives $e_{j+1} = -e_j^2/\widehat{\theta}_N$ and, by induction, $e_q = -e_0^{2^q}/\widehat{\theta}_N^{2^q-1}$. In units of $\widehat{\theta}_N$, $u_{j+1} = -u_j^2$ with $u_j = e_j/\widehat{\theta}_N$, so $|u_j| \to 0$ iff $|u_0| < 1$, i.e. $0 < x < 2\widehat{\theta}_N$; $|u_0| > 1$ gives $|u_j| \to \infty$ with $u_j < 0$ for $j \geq 1$; $u_0 = \pm 1$ gives $u_1 = -1$ and then $u_j = -1$, i.e. the iterate 0, where the log-likelihood is undefined. □

Proof of (v)–(vi). Let $\alpha = \alpha(r,q)$ and $K/N^\alpha \to c \in (0,\infty]$, with $m \to \infty$. By Lemma 5 and (ii), $\theta_r^{(0)} - \widehat{\theta}_N = \mathbb{E}R_1^J + \bar{\xi}_N - R_N$ with $\mathbb{E}R_1^J = \beta_r\theta_0 m^{-r}(1 + O(m^{-1}))$, $\bar{\xi}_N = O_p(\sqrt{K}/N)$ and $R_N = O_p(N^{-1})$; part (i) supplies the constants. Since $m^{-r} = (K/N)^r$ and $r(1-\alpha) = 2^{-(q+1)}$,

$$N^{1/2^{q+1}}\, m^{-r} = N^{1/2^{q+1}}\,(K/N^\alpha)^r N^{-r(1-\alpha)} = (K/N^\alpha)^r \to c^r,$$

while $N^{1/2^{q+1}}\sqrt{K}/N \leq N^{1/2^{q+1}+\alpha/2-1} \to 0$ because $\alpha < 1$ and $1/2^{q+1} \leq 1/2$ give $1/2^{q+1} + \alpha/2 - 1 < 0$; likewise $N^{1/2^{q+1}}N^{-1} \to 0$. Hence $N^{1/2^{q+1}}(\theta_r^{(0)} - \widehat{\theta}_N) \xrightarrow{p} \beta_r\theta_0 c^r$ when $c < \infty$ and $\xrightarrow{p} \pm\infty$ when $c = \infty$. For $q = 0$, $N^{1/2^{q+1}} = \sqrt{N}$ and the claim follows. For $q \geq 1$, (iv) gives

$$\sqrt{N}\big(\theta_r^{(q)} - \widehat{\theta}_N\big) = -\frac{\sqrt{N}\,(\theta_r^{(0)} - \widehat{\theta}_N)^{2^q}}{\widehat{\theta}_N^{2^q-1}} = -\frac{\big[N^{1/2^{q+1}}(\theta_r^{(0)} - \widehat{\theta}_N)\big]^{2^q}}{\widehat{\theta}_N^{2^q-1}},$$

and $\widehat{\theta}_N \xrightarrow{p} \theta_0$ with the continuous mapping theorem yields $-\theta_0(\beta_r c^r)^{2^q}$ when $c < \infty$ and divergence to $-\infty$ in probability when $c = \infty$. Combining with $\sqrt{N}(\widehat{\theta}_N - \theta_0) \Rightarrow \mathscr{N}(0,\theta_0^2)$ proves the displayed limit law for $(2,1)$. Strictly, (iv) requires the pilot to be positive for the iteration to be meaningful;

$P(\theta_r^{(0)} \le 0) \to 0$ since $\theta_r^{(0)} \xrightarrow{p} \theta_0 > 0$. □

C. Reference implementations

The replication package (Section D) contains the package `pamle` with the following components, all exercised by the test suite.

- `jackknife.py`: exact rational Vandermonde weights for any pattern, the constants β_r and A_r, condition numbers, and the combination of level averages.
- `likelihoods.py`: exponential, logistic and Poisson models with batched Newton fitting over blocks and sub-blocks (step halving, convergence on the score, status codes for non-convergence, singular curvature and divergence, and a separation safeguard for logistic regression that refuses a stationary point with fitted $|\eta| > 30$ or a numerically singular information matrix). Derivatives are checked against finite differences.
- `algorithms.py`: local split fits, corrected worker estimates, size-weighted pilots, global Newton corrections with exact, stale, single-machine and quantized curvature (linear solves, condition numbers and step norms as diagnostics; no explicit inverse in the Newton step), the aggregated estimating-equation estimator of Lin and Xi (2011), the surrogate-likelihood estimator of Jordan et al. (2019) with its full local maximization, and the analytic first-order bias correction of canonical GLMs (Cox and Snell, 1968; Cordeiro and McCullagh, 1991).
- `communication.py`: rounds, messages, numbers and bytes for every variant of Table 4.
- `protocol.py`, `worker.py`: the coordinator–worker protocol over TCP with length-prefixed messages, per-stage wall-clock, per-worker compute times, latency percentiles and byte counts (Algorithm 1).

Analytic bias correction used as a baseline. For a canonical-link GLM with per-observation score $\psi_i(\beta) = x_i(y_i - \mu_i)$, the $O(1/n)$ bias of the MLE conditional on the design is b_1/n with

$$b_1 = \tfrac{1}{2} I^{-1} \left[n^{-1} \textstyle\sum_i \nabla^2 \psi_i [I^{-1}] \right] = -\tfrac{1}{2} I^{-1} n^{-1} \textstyle\sum_i x_i \mu''(\eta_i) x_i^\top I^{-1} x_i,$$

where $I = n^{-1} \sum_i x_i x_i^\top \mu'(\eta_i)$; this is the form of Cordeiro and McCullagh (1991) and reduces to $b_1 = \theta_0$ in the exponential model. The corrected local estimate $\widehat{\theta}_{D_k} - \hat{b}_1/m$ needs one local fit and no sub-blocks; it is the local analogue of the debiased weighted estimator of Gu and Chen (2023) and provides, in Section 9, a baseline that attains the order-2 exponent at one third of the local fits of the half-block jackknife whenever the model's bias formula is available and correct.

D. Replication

Everything reported is regenerated from committed raw outputs by the scripts in `experiments/`; no number in the tables or figures is typed by hand. The directory `references/initial_pamle/` contains the prior draft, its archive and SHA-256 checksums; its Monte Carlo (seed 20261028, $N = 2^{20}$, $m \in \{256, 64, 16\}$, 500 replications) was re-run and matched to 5×10^{-13} before any change was made, and the test `tests/test_reproducibility.py` reproduces its summary table from the new code.

Machine-checked algebra. The deterministic identities used by the proofs—the Vandermonde constraints and the constants of Table 2, the exact biases (22), the Newton identities (24) with the q-fold recursion and the basin of attraction, the boundary constant of Theorem 2(v), the exponent algebra behind Theorem 1, the Gamma and Beta inverse moments and the exact variance (23)—are formalized and checked in Lean 4 with Mathlib (`theory/lean/`, 27 theorems, no `sorry`; `scripts/check_lean.sh`). The probabilistic steps (O_p bounds, central limit theorem, Slutsky) are not formalized.

Commands.

```
python3 -m pip install -r requirements.txt
pytest -q                                    # unit tests, exact identities, protocol
python3 theory/symbolic_checks.py            # SymPy verification of all identities
python3 experiments/run_monte_carlo.py --config experiments/configs/main_exponential.yaml
python3 experiments/scatter_cluster.py --launch \
   --configs experiments/configs/main_logistic.yaml \
             experiments/configs/main_poisson.yaml \
             experiments/configs/robustness.yaml         # cluster
python3 experiments/run_boundary.py
python3 experiments/run_ablations.py --model logistic
python3 experiments/run_skewed.py
python3 experiments/run_benchmarks.py \
   --config experiments/configs/benchmarks.yaml \
   --out results/raw/benchmarks.jsonl                  # cluster
python3 experiments/process_results.py --figures
python3 experiments/make_tables.py
bash scripts/compile_paper.sh
bash scripts/validate_pdf.sh
```

Seeds are derived from `seed_base = 20261028` and a hash of the cell description; every raw file has a JSON manifest with git commit, Python and NumPy versions, host, wall time and the SHA-256 of the output (`experiments/EXPERIMENT_MANIFEST.csv`). Monte Carlo cells were

run on the laptop (Apple M5 Max, Python 3.14.4, NumPy 2.4.4) and on 18 cluster nodes (Python 3.9.6/3.14, NumPy 2.0.2/2.5.4); floating-point results differ across these environments at the level of 10^{-12} relative, far below the Monte Carlo standard errors reported.