Supplementary Note S4. Numerical diagnostics for estimation procedures in the Monte Carlo study

This note supplements the Monte Carlo study by documenting the numerical settings and convergence diagnostics for maximum likelihood estimation (MLE), a convergence audit for the method of moments (MoM), and a sensitivity assessment with respect to the number of Monte Carlo replications. It also provides a grid-based examination of the log-likelihood surface for one representative simulated sample in each of Scenarios I and II.

Optimization procedure

The positive parameters were optimized on the log scale, $\eta = \log\lambda$ and $\tau = \log\theta$, using the bounded L-BFGS-B algorithm in R's optim function. The constraints were

$$\log(10^{-4}) \leq \eta, \tau \leq \log(100),$$

so that $10^{-4} \leq \lambda, \theta \leq 100$. The optimizer settings were maxit = 600, factr = 1e7, and pgtol = 1e−8. For each sample, three initial values were used: the method-of-moments estimates and the two fixed starts (1.5, 1.3) and (0.5, 2.0). Only solutions with a finite objective and R convergence code equal to zero were accepted; among accepted solutions, the estimate with the highest numerically evaluated log-likelihood was retained.

The likelihood was evaluated by midpoint quadrature with $K_{MLE} = 2000$ points on the unit interval using the LLU quantile representation; the same implementation was used for optimization and the surface diagnostics. Scenario I used $n = 10$, $(\lambda, \theta) = (1, 0.5)$, and Scenario II used $n = 20$, $(\lambda, \theta) = (0.3, 1)$. Both scenarios considered sample size $N \in \{50, 150, 500, 1500\}$ and 300 replications per configuration.

Convergence and sensitivity to starting values

All 2,400 MLE fits (300 replications for each of four sample sizes in each of two scenarios) converged successfully. Each fit had three successful optimization starts, giving 7,200 successful optimization attempts in total. No accepted solution was flagged at a parameter bound using the implementation's log-scale tolerance of 10^{-5}.

Multi-start agreement was assessed by requiring both (i) a spread in log-likelihood values no greater than 10^{-6} and (ii) a maximum absolute difference between log-parameter estimates no greater than 10^{-3}. Table S4.1 summarizes these results. The agreement criterion was met in 2,375 of the 2,400 replications (98.96%). In the remaining 25 replications, the log-parameter differences still remained below 10^{-3}; the flag was triggered by the log-likelihood-spread criterion. Across all replications, the maximum log-likelihood spread between converged starts was 2.11×10^{-5}, and the maximum absolute log-parameter difference was 2.67×10^{-4}. Thus, the multi-start solutions were numerically very close, although the chosen tolerance flags a small number of cases for which the likelihood values did not meet the stricter 10^{-6} threshold.

Table S4.1. MLE convergence and multi-start agreement across the Monte Carlo configurations.

Scenario	*N*	Successful MLE	Starts per fit	Agreement (%)	Boundary solutions
I	50	300/300	3.00	100.00	0
I	150	300/300	3.00	100.00	0
I	500	300/300	3.00	100.00	0
I	1500	300/300	3.00	98.67	0
II	50	300/300	3.00	100.00	0
II	150	300/300	3.00	99.67	0
II	500	300/300	3.00	99.33	0
II	1500	300/300	3.00	94.00	0

Likelihood-surface diagnostics

For one simulated sample of size $N = 1500$ from each MLE scenario, the numerically evaluated log-likelihood was calculated on a 41 × 41 log-spaced grid in (λ, θ), spanning a multiplicative factor of five on either side of the multi-start estimate (subject to the optimization bounds). In both samples, the grid maximum coincided with the reported multi-start estimate and did not occur on the edge of the inspected grid. The contours in Figures S4.1–S4.2 show relative log-likelihood differences from the grid maximum. They display a single, localized peak within the inspected windows. These plots provide a numerical diagnostic for the representative samples; they do not establish global optimality outside the inspected parameter region.

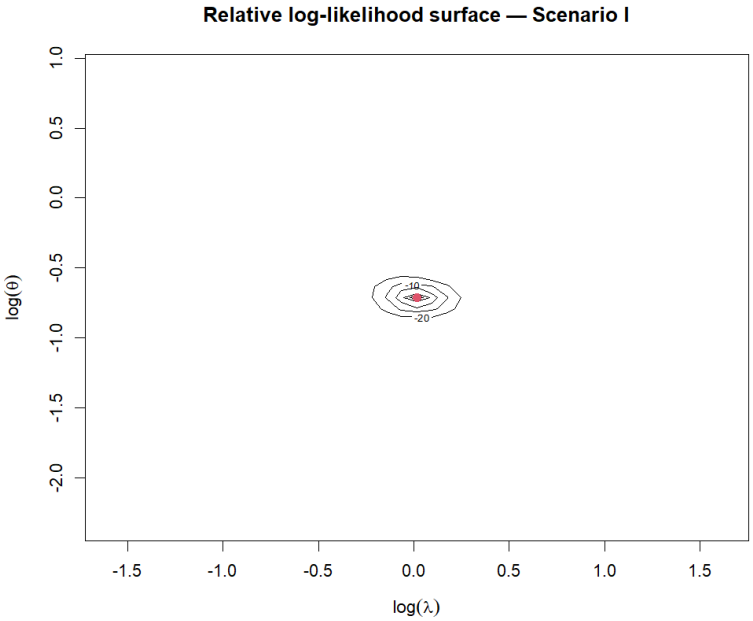


Figure S4.1. Scenario I: relative log-likelihood surface, n = 10, N = 1500.

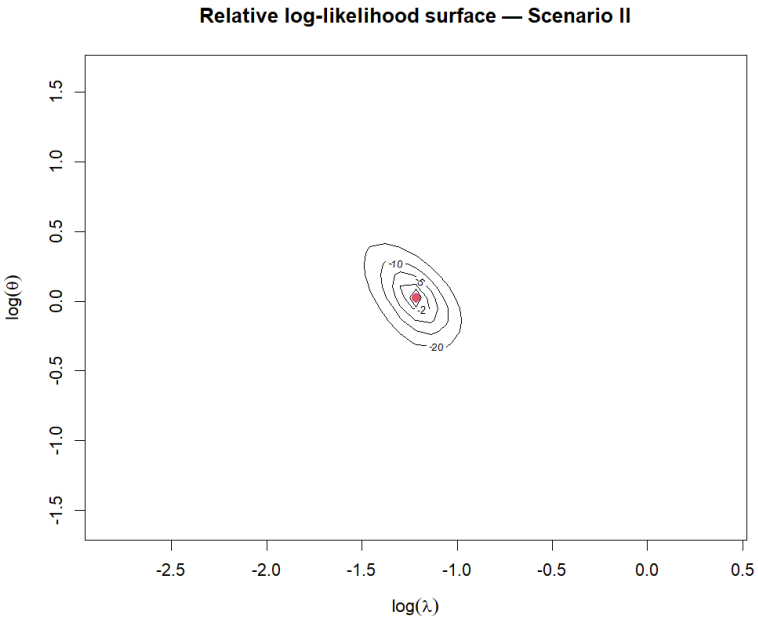


Figure S4.2. Scenario II: relative log-likelihood surface, n = 20, N = 1500.

Additional checks for the main Monte Carlo study

The convergence of the MoM procedure was checked independently on the same simulated samples, regenerated using the original seed and sample-generation order. All 3,600 MoM optimizations (three scenarios, four sample sizes, and 300 replications per configuration) returned a zero optimizer convergence code and finite estimates within the imposed positive parameter bounds. The regenerated estimates matched the stored main-study estimates to a maximum absolute difference of 5.11×10^{-15}. Thus, no MoM optimization failures occurred in the reported Monte Carlo experiments.

The sensitivity of the main Monte Carlo summaries to the number of replications was assessed using cumulative prefixes of M = 100, 200, 250, and 300 replications. For each prefix, the absolute differences between its RMSE values and those obtained with all 300 replications were calculated across the 48 estimator–parameter configurations. Table S4.2 summarizes the median and maximum differences, together with the number of pairwise RMSE rankings that differed from those at M = 300.

Table S4.2. Sensitivity of Monte Carlo RMSE summaries to the number of replications.

M	Median absolute RMSE difference	Maximum absolute RMSE difference	Ranking reversals (out of 24)
100	0.002043	0.036934	3
200	0.001639	0.018695	0
250	0.000773	0.008178	0
300	0	0	0

The pairwise RMSE ordering at M = 200 and M = 250 was identical to that obtained at M = 300; three of the 24 pairwise comparisons differed when only the first 100 replications were used. These calculations are based on nested prefixes of the same ordered set of 300 replications, not on independent Monte Carlo runs. They therefore provide a robustness check for the reported summaries rather than a formal estimate of Monte Carlo sampling uncertainty.