

Privacy-Perturbed Parameter Estimation and Inference for Heavy-Tailed GARCH Models

Lile Zhao*

School of Mathematics, Sichuan University, Chengdu 610065, China

*Corresponding author: zhaolile@stu.scu.edu.cn

Abstract:

Most privacy-aware inference methods are developed for independent data, whereas financial returns combine heavy tails with persistent conditional volatility. This paper examines a Gaussian-perturbed GARCH M-estimation framework and the finite-sample effect of perturbation noise on parameter inference. The idealized large-sample analysis uses stationarity, ergodicity, and a martingale central limit theorem to characterize the sampling and perturbation components. A noise-corrected sandwich covariance estimator and Wald statistic are then constructed. The simulation implementation uses Gaussian QML as the M-estimation center, a projected one-step parameter perturbation, and a working calibration $\Delta = 2B$ with $B = 2$. Across 500 requested replications per design point, the corrected covariance materially improves interval coverage and reduces Wald over-rejection, especially when ϵ is small. The results also reveal finite-sample limitations: projection is non-negligible in the smallest-sample, strongest-perturbation design, and studentized Q-Q plots retain tail departures from normality. Because the initial QML estimate and covariance ingredients use the raw sample and a complete recursive global-sensitivity proof is not provided, the implementation is interpreted as privacy-perturbed and DP-motivated rather than as an end-to-end differentially private release.

Keywords: Financial time series; GARCH; M-estimation; Gaussian perturbation; Noise-corrected inference.

1. Introduction

Financial returns and volatility records can contain commercially or institutionally sensitive information. Their analysis is therefore increasingly conducted

under data-governance constraints that restrict direct access to raw observations. Differential privacy (DP) provides a formal framework for limiting the effect of any single record on a released output. For recursively filtered financial time series, however, translat-

ing this principle into a valid end-to-end release requires particular care because one observation can affect many subsequent conditional variances.

For non-private analysis, M-estimation and related procedures provide established tools for heavy-tailed GARCH models. Mukherjee derived asymptotic normality for a broad class of GARCH M-estimators, while recent rank-based, M-quantile, and bounded M-estimation studies have extended robust volatility inference and forecasting [1-4]. In private statistical inference, Avella-Medina developed privacy-preserving M-estimation and Wald-type procedures, Avella-Medina, Bradshaw and Loh studied inference through noisy optimization, and Jewson, Ghalebikesabi and Holmes proposed a posterior-sampling approach based on beta-divergence [5-7]. These private-inference results primarily address independent or non-recursive settings and therefore do not directly establish a complete privacy guarantee for GARCH filtering.

Several gaps motivate the analysis. First, the recursive GARCH filter complicates sensitivity calibration because changing one observation can affect subsequent conditional variances. Second, conventional Wald covariance formulas omit the additional variability created by parameter perturbation. Third, the finite-sample interaction among heavy tails, parameter constraints, sample size, and perturbation magnitude is not well documented. The paper therefore separates an idealized asymptotic perturbation analysis from the operational Gaussian-QML simulation used to assess these effects.

This paper makes three contributions. It develops an idealized large-sample representation for privacy-perturbed GARCH M-estimation under stationary and ergodic conditions; derives a covariance correction and Wald statistic that include the Gaussian perturbation variance; and evaluates the resulting finite-sample trade-off through Monte Carlo experiments. The simulation evidence is used to assess, rather than presume, the quality of the asymptotic approximation.

This study addresses three research questions:

- (1) How does Gaussian perturbation affect the asymptotic representation and finite-sample accuracy of heavy-tailed GARCH parameter estimators?
- (2) How can the additional perturbation variance be incorporated into a Wald procedure to improve Type I error control?
- (3) How do sample size, perturbation magnitude, heavy tails, and parameter projection interact in the finite-sample accuracy of the estimator?

The analysis addresses these questions through three related contributions.

It extends perturbation-based inference ideas from independent-sample settings to a serially dependent volatility model at the level of asymptotic representation. Its main empirical contribution is to quantify how an explicit

noise-variance correction changes confidence-interval coverage and Wald-test size. A complete end-to-end differential-privacy proof for the operational recursive GARCH procedure is outside the scope of the present paper.

2. Literature Review

2.1 Heavy-Tailed GARCH Estimation and Robust Inference

ARCH and GARCH models provide the classical foundation for modeling conditional heteroskedasticity in financial returns [8, 9]. Heavy-tailed innovations complicate conventional likelihood-based inference because standard asymptotic results may depend on higher-order moments. Hall and Yao developed inference for Gaussian QMLE under heavy-tailed ARCH/GARCH errors, and Mukherjee established asymptotic normality for a broad class of GARCH M-estimators [1, 10]. More recent contributions use residual ranks, M-quantiles, and bounded M-estimation to improve robustness to heavy tails or outliers [2-4]. These studies motivate the M-estimation formulation used in the idealized theory and the Gaussian-QML center used in the simulation.

2.2 Differential Privacy and Private Statistical Inference

Differential privacy was formalized through sensitivity-calibrated randomization mechanisms [11]. Gaussian differential privacy subsequently supplied a hypothesis-testing interpretation and tractable composition framework for Gaussian mechanisms [12]. These results clarify how privacy parameters and sensitivity determine injected noise, but they do not by themselves yield a global-sensitivity bound for a recursive GARCH estimator.

Recent work has expanded private statistical inference beyond basic point estimation. Avella-Medina developed robust private parametric tests, while noisy optimization has been used to construct private M-estimators and confidence regions [5, 6]. Jewson, Ghalebikesabi and Holmes examined private posterior sampling for continuous regression models [7]. Kazan et al. proposed a general framework for converting hypothesis tests into differentially private procedures, and Cai, Wang and Zhang characterized minimax privacy-accuracy trade-offs for fundamental estimation problems [13, 14].

These advances demonstrate that inferential uncertainty introduced by privacy mechanisms must be quantified explicitly. Nevertheless, most available theories are formulated for independent samples, standard optimization objectives, or non-recursive models. A changed return in a GARCH filter can propagate through future conditional variances, so bounding a single contemporaneous score

is insufficient to certify the sensitivity of the complete released procedure.

2.3 Differential Privacy for Time-Series Data

Research linking privacy, robustness, and time-series analysis remains comparatively limited. Dwork and Lei established a foundational connection between differential privacy and robust statistics [15]. Mao et al. reviewed differential-privacy methods for time-series data and later proposed a utility-aware temporal local-DP release mechanism for anomalous time-series observations [16, 17]. These contributions address important sequential and release problems, but they do not provide an end-to-end sensitivity proof or a noise-corrected Wald theory for a recursively filtered heavy-tailed GARCH estimator.

2.4 Research Gap and Study Positioning

The literature therefore identifies a specific methodological gap. Robust GARCH inference is well developed, and private estimation and testing have advanced rapidly [1-7, 10, 12-14]. Time-series privacy research has also developed specialized release mechanisms [16, 17]. However, the joint effect of heavy tails, recursive filtering, parameter perturbation, projection, and covariance correction has not been systematically evaluated within one GARCH simulation framework.

Accordingly, the paper separates two objects: an idealized bounded-score estimating equation used to derive a transparent asymptotic perturbation benchmark, and an operational projected QML simulation used to assess finite-sample inference. This separation prevents the statistical noise experiment from being misinterpreted as a proved end-to-end private release.

The analysis evaluates the following statistical hypotheses:

- (1) Under the stated idealized bounded-score and contraction assumptions, a privacy-perturbed GARCH M-estimator remains consistent and has an asymptotic linear representation;
- (2) In finite samples, a Wald procedure that ignores the Gaussian perturbation variance over-rejects the true null, particularly under stronger perturbation;
- (3) Adding the perturbation-variance component to the covariance estimator materially improves coverage and Wald-test size, although parameter projection and heavy-tailed finite-sample effects may prevent exact nominal performance.

3. Theoretical Basis and Simulation Design

3.1 Model Foundations and Analytical Scope

The study distinguishes a formal privacy benchmark from an operational perturbation experiment. In the idealized benchmark, a bounded estimating equation is perturbed by Gaussian noise whose scale would be valid only if a global-sensitivity bound for the complete recursive procedure were established. The simulation instead applies a projected one-step Gaussian perturbation to a non-private QML estimate. It is designed to evaluate noise-augmented inference and is not represented as an end-to-end DP algorithm.

The financial model is a GARCH process, which captures volatility clustering through a recursive conditional-variance equation. Heavy-tailed Student-t innovations are used in the simulation to examine finite-sample behavior beyond a Gaussian return model. The statistical analysis focuses on parameter bias, dispersion, interval coverage, Wald-test size, and the effect of the perturbation-variance correction; it does not estimate Value-at-Risk or Expected Shortfall.

(1)

The first relation separates the observed return from its conditional volatility, while the second makes current volatility depend on the previous squared return and conditional variance. The coefficients used in the simulation are positive and their persistence sum is 0.95, below one. Standardized Student-t innovations with six degrees of freedom preserve unit innovation variance while producing heavier tails than a Gaussian specification.

3.2 Monte Carlo Design and Data Generation

The Monte Carlo design varies sample size ($T = 500, 1000, \text{ and } 2000$) and perturbation budget ($\epsilon = 0.5, 1, \text{ and } 5$) under a fixed heavy-tailed GARCH (1,1) data-generating process. Each design point requests 500 independent replications. This factorial design permits direct assessment of how sampling variation, perturbation magnitude, and feasibility projection interact.

A burn-in of 500 observations is discarded before each retained sample, reducing dependence on arbitrary initial values. The three sample sizes show how sampling uncertainty changes with information volume, whereas the three perturbation budgets provide strong, intermediate, and weak perturbation settings. Holding the data-generating parameters fixed across the design isolates these two effects.

3.3 Perturbation Procedure and Evaluation Criteria

(2)

The operational procedure first obtains the Gaussian-QML center and local Hessian, then applies the one-step Gaussian perturbation shown above and projects the result onto the admissible parameter set. Consistent with the simulation specification in Section 5, its working calibration uses $\delta = 10^{-5}$, $B = 2$, $\Delta = 2B$, and the stated perturbation budgets. Here B is a working simulation calibration rather than a proved global-sensitivity bound for the complete recursive QML procedure. Projection enforces the model's feasibility restrictions.

Estimation accuracy is summarized by bias, empirical standard deviation, and root mean squared error (RMSE). Inferential performance is assessed through ordinary and noise-corrected 95% coverage and through the rejection rate of a Wald test under a true null hypothesis. Projection rates and valid-replication counts are also reported because numerical failures and boundary projection can affect finite-sample conclusions. Section 4 presents the idealized asymptotic analysis, and Section 5 reports the operational simulation evidence.

(3)

(4)

Here R denotes the number of valid replications, the subscript k indexes a model parameter, the superscript r indexes a replication, and the overbar denotes the mean estimate across valid replications. Corrected coverage is computed from the same indicator expression using the noise-corrected confidence interval. Wald rejection rates, projection rates, and valid-replication counts complement these measures and are defined consistently across all design points.

4. Asymptotic Theory for Privacy-Perturbed GARCH M-Estimation

4.1 Model Setup and Idealized Noise Calibration

This paper studies a strictly stationary GARCH process :

(5)

where the innovations are i.i.d. with mean zero and unit variance, and may exhibit heavy tails—finite fourth moments are not required. The true parameter vector lies in a compact set with .

Within the non-private framework, a robust M-estimation framework is adopted. Let be a suitably smooth loss function, and denote its gradient by . The gradient satisfies the moment condition . The classical M-estimator is then defined as the solution to

(6)

For the idealized asymptotic analysis, Gaussian noise is added to a bounded estimating equation. Temporal dependence matters because a changed observation may propagate through the conditional-variance recursion. Accordingly, any formal DP interpretation requires a global-sensitivity bound for the complete released procedure, not merely a bound on one observed score contribution.

Let and denote two neighboring datasets differing only at time index . Exploiting the strict stationarity and contraction properties of the GARCH process (e.g., or a more general contraction condition), the impact of a past shock on both the conditional variance and the gradient decays exponentially. Under the stated uniform boundedness and contraction conditions, suppose that a constant independent of sample size , satisfies the inequality for any ,

(7)

Under additional uniform boundedness and contraction conditions, a representative propagation bound has the form $2M/(1-\rho)$, where ρ is a contraction factor. This expression is an idealized analytic benchmark. It is not identified with the working $\Delta = 2B$ calibration used in Section 5, and the paper does not claim that the operational QML implementation establishes the global sensitivity of the complete recursive GARCH release.

Given a valid sensitivity bound Δ for an idealized bounded estimating equation, a Gaussian perturbation can be calibrated in the usual way. The following display records that benchmark mechanism; its privacy interpretation is conditional on the assumed bound.

(8)

Draw a noise vector independent of the data.

Definition 1 (Idealized Privacy-Perturbed GARCH M-Estimator). The estimator is defined by the perturbed estimating equation

(9)

If the perturbed solution is infeasible, it is projected onto the parameter space; it is never replaced by the non-private estimate. Section 5 uses a related one-step projected QML perturbation. That operational implementation is designed to study noise-augmented inference and is not asserted to inherit an end-to-end DP guarantee from this idealized equation.

4.2 Large-Sample Theory

This section derives large-sample properties for the idealized privacy-perturbed GARCH M-estimator. Consistency is established first, followed by asymptotic normality.

4.2.1 Consistency (Theorem 1)

The analysis relies on the following set of regularity assumptions.

- (A1) The GARCH process is strictly stationary and ergodic; the true parameter lies in the interior of the compact set .

- (A2) The loss function is continuously differentiable in , and there exists a constant such that almost surely for all .
 - (A3) Identification: has a unique zero at and is continuous on .
 - (A4) The perturbation scale is fixed with respect to T, so the additive noise is $O_p(1)$ before division by T.
- Theorem 1 (Consistency of the idealized perturbed estimator). Under conditions (A1)–(A4), as T increases, the perturbation divided by T vanishes and the estimator converges in probability to the identified parameter.

(10)

Proof sketch:

Starting from the perturbed estimating equation (9),

(11)

By (A4), has variance , so . Define . Under (A1)–(A2), ergodicity together with a uniform law of large numbers guarantees that converges uniformly in probability to on . Combined with the identification condition (A3) and the standard M-estimation argument (cf. Avella-Medina, proof of Theorem 1), the resulting estimator must converge in probability to the unique zero [5].

4.2.2 Asymptotic Normality (Theorem 2)

Building on consistency, this subsection derives the asymptotic distribution of the idealized perturbed estimator. In addition to (A1)–(A4), the score process must satisfy a structural condition.

- (A5) (Martingale difference and moment condition) The gradient sequence forms a martingale difference array with respect to the natural filtration); i.e.) for all . There exists such that . The matrix is non-singular, and

(12)

Remark. Gaussian QML scores for correctly specified GARCH models form martingale differences under standard conditions. For other M-scores, the martingale-difference and centering requirements must be verified rather than assumed. The finite-sample projection used in Section 5 can also modify the tails of the operational estimator's distribution.

Theorem 2 (Asymptotic normality) Under conditions (A1)–(A5), as ,

(13)

where and . Under stationarity and the martingale-difference property, .

Proof sketch:

Consistency of follows from Theorem 1. A first-order Taylor expansion of the perturbed estimating equation (9) around yields

(14)

where stands for the second-order remainder. Combining the consistency result, the convergence of the Hessian in (A5), and the convergence , the following representation can be derived:

(15)

For the normalized score , condition (A5) ensures it is a sum of martingale differences. Under the stated moment bound , the Lindeberg condition holds, which validates the application of martingale central limit theory and delivers the convergence

(16)

Slutsky's theorem and the non-singularity of H establish the asymptotic normality stated in Theorem 2.

4.3 Perturbation-Noise-Corrected Wald Test

The extra variability induced by Gaussian parameter perturbation must be included in finite-sample inference. Otherwise, conventional covariance estimates can understate uncertainty and produce over-rejection. This section derives a noise-corrected Wald statistic for the idealized asymptotic representation.

This analysis focuses on linear null hypotheses formulated as , where the matrix has full row rank . Starting from the asymptotic expansion derived in Equation (15) and the definition of the noise injection mechanism, the privacy-perturbed GARCH M estimator can be written as

(17)

The term has covariance matrix , which is of order and therefore asymptotically negligible relative to the sampling variance. In finite samples, however, especially when the privacy budget is small (and hence is large), this term can substantially elevate the variability of . A naive covariance estimator that ignores the privacy noise, namely , systematically underestimates the true covariance, leading to an over-rejection of the null.

To mitigate this distortion, a privacy-corrected covariance matrix estimator is constructed as

(18)

where , , and is given in (8).

Definition (Noise-corrected Wald statistic).

The corrected Wald test statistic for the hypothesis is

(19)

Theorem 3 (Asymptotic distribution of).

Under conditions (A1)–(A5) and the null hypothesis ,

(20)

Proof sketch:

As established by Theorem 2, with . The consistency of and , combined with the convergence , implies . It follows that shares identical asymptotic properties with the standard Wald statistic, and the chi-squared limiting distribution is established via the continuous mapping theorem.

Implementation note. In the reported simulation, $\hat{H}$ and $\hat{J}$ are evaluated from the raw sample at the perturbed estimate. The corrected covariance and Wald statistic therefore quantify statistical uncertainty caused by Gaussian perturbation; they are not claimed to be privately released outputs. A deployable end-to-end DP procedure would require private or sensitivity-bounded construction

of every released ingredient.

4.4 Perturbation-Accuracy Trade-off

With a fixed perturbation scale, the added parameter noise is of smaller asymptotic order than the sampling term, but it can be important at the sample sizes used in practice. The following expansion separates the idealized sampling and perturbation components.

Starting from the asymptotic expansion (15) employed within the proof of Theorem 2,

(21)

where $\tilde{\theta}$. Under conditions (A1)–(A5) and the additional assumption that $\tilde{\theta}$ is uniformly bounded in a neighbourhood of θ , standard asymptotic reasoning yields $\tilde{\theta}$. Independence between the score term and the Gaussian noise eliminates the cross term in expectation, yielding

(22)

Substituting $\tilde{\theta}$ from (8) delivers the following bound.

Theorem 4 (MSE Upper Bound). Under conditions (A1)–(A5) and with $\tilde{\theta}$ uniformly bounded near θ , there exist positive constants c_1 independent of ϵ and the privacy parameters such that, for all sufficiently large n ,

(23)

where c_1 and c_2 .

This inequality explicitly separates the sampling variation from the privacy-induced variation ϵ . When ϵ is fixed, the latter is asymptotically negligible; however, when the privacy budget shrinks with the sample size n , the two error components scale as ϵ and $n^{-1/2}$, respectively:

- For $\epsilon \gg n^{-1/2}$, the sampling term dominates and the MSE retains the parametric rate n^{-1} .
- For $\epsilon \asymp n^{-1/2}$, the two terms are of the same order; the MSE remains but the leading constant is elevated by the privacy requirement.
- For $\epsilon \ll n^{-1/2}$, the privacy term dominates and the convergence rate degrades to ϵ^2 .

These rate comparisons are theoretical benchmarks for the idealized expansion. Because only three sample sizes are used in Section 5, the fitted log-log slopes are reported descriptively and are not treated as a formal empirical verification of the rate bound.

5. Simulation Study

5.1 Experimental Design

The Monte Carlo study evaluates a projected Gaussian-perturbed QML estimator and compares naive infer-

ence, which ignores perturbation variance, with corrected inference based on the same perturbed estimates. Each design point requests 500 independent replications.

Data-generating process. Returns follow a GARCH (1,1) model with $\omega = 0.05$, $\alpha = 0.10$, and $\beta = 0.85$. Innovations are standardized Student-t with six degrees of freedom, representing moderate heavy tails while retaining a finite fourth moment. Sample sizes are $T = 500, 1000$, and 2000 , with a burn-in of 500 observations.

Perturbation mechanism. Gaussian QML provides the center $\hat{\theta}$ and local Hessian $\hat{H}$. The operational estimate is $\tilde{\theta} = \text{projection}[\hat{\theta} - \hat{H}^{-1}\{\epsilon/T\}]$, where ϵ is Gaussian with working scale based on $\Delta = 10^{-5}$, $B = 2$, and $\Delta = 2B$. The budgets $\epsilon = 0.5, 1$, and 5 index stronger to weaker perturbation. B is a working simulation calibration, not a proved global sensitivity bound for the complete recursive QML procedure. $\hat{H}$ and $\hat{J}$ use the raw sample; consequently, the reported implementation is privacy-perturbed/DP-motivated rather than an end-to-end private release.

Comparators. Two inferential treatments are applied to the same perturbed parameter estimate:

- Oracle: the non-perturbed Gaussian QML benchmark;
- Perturbed-naive inference: the projected perturbed estimate with a covariance matrix that omits perturbation variance;
- Proposed corrected inference: the identical perturbed estimate with the additional Gaussian-noise covariance term.

Performance metrics. Bias, SD, RMSE, ordinary 95% coverage (Cov95), corrected 95% coverage (Cov95c), projection rate, and valid replication count are recorded. Across the nine design points, 484-499 of 500 requested replications are valid; failures arise when the numerical Hessian is materially indefinite.

Test evaluation. For $H_0: \alpha + \beta = 0.95$, the naive and noise-corrected Wald rejection rates are calculated at the nominal 5% level using the same perturbed estimates. All simulations are performed in R with fixed seeds; replication-level results and failure counts are retained for auditability.

5.2 Parameter Estimation Results


Table 1 reports the α and β results at $T = 1000$. The Oracle rows are shown separately for each ϵ -specific simulation batch so that all comparisons use the corresponding Monte Carlo draws.

Table 1. Bias, RMSE and empirical coverage for alpha and beta at T = 1000 (500 requested replications per design point)

Method	epsilon	alpha Bias	alpha RMSE	alpha Cov95	alpha Cov95c	beta Bias	beta RMSE	beta Cov95	beta Cov95c
Oracle	0.5	0.0018	0.0359	0.8949	-	-0.0150	0.0581	0.9212	-
Oracle	1.0	0.0021	0.0337	0.9131	-	-0.0169	0.0582	0.9333	-
Oracle	5.0	0.0018	0.0336	0.9038	-	-0.0177	0.0655	0.9018	-
Perturbed	0.5	0.0025	0.0599	0.7535	0.9192	-0.0199	0.1293	0.7434	0.9273
Perturbed	1.0	0.0002	0.0396	0.8828	0.9253	-0.0123	0.0776	0.8949	0.9333
Perturbed	5.0	0.0016	0.0335	0.9058	0.9058	-0.0171	0.0640	0.8978	0.8998


At epsilon = 0.5, ordinary coverage for the perturbed estimator is 0.754 for alpha and 0.743 for beta; adding the perturbation variance raises coverage to 0.919 and 0.927. At epsilon = 1, corrected coverage is 0.925 and 0.933. When epsilon = 5, the perturbation is small and the cor-

rection changes coverage only marginally. Bias remains modest, while RMSE increases as epsilon decreases. At T = 1000 the projection rate is 4.44%, 1.62%, and 0.20% for epsilon = 0.5, 1, and 5, respectively.

**Fig. 1 RMSE of the perturbed estimator across sample sizes and perturbation budgets.**

As shown in Fig. 1, the log-log RMSE slopes are descriptive because each regression uses only three sample sizes. For alpha the fitted slopes are -0.806, -0.782, and -0.523 as epsilon increases from 0.5 to 5; for beta they are -0.601,

-0.525, and -0.510. These estimates are consistent with declining estimation error as T increases, but they should not be treated as precise estimates of an asymptotic rate.

**Fig. 2 Studentized Q-Q plots for alpha and beta at T = 1000 and epsilon = 0.5**

As shown in Fig. 2, the Q-Q plots align reasonably well through the central quantiles but show visible tail departures. The diagnostic KS p-values are 0.009 for alpha and 0.144 for beta. Thus, the finite-sample normal approximation is imperfect for alpha in the most private design, whereas beta is closer to the reference distribution.

5.3 Privacy-Corrected Wald Test Results

Table 2 reports Wald rejection rates for $H_0: \alpha + \beta = 0.95$. Naive and corrected columns use the same perturbed estimates, so the comparison isolates the effect of adding the perturbation variance to the covariance matrix.

Table 2. Rejection rates for $H_0: \alpha + \beta = 0.95$ using the perturbed estimator

T	epsilon	Wald-Rej-Naive	Wald-Rej-Corrected
500	0.5	0.2686	0.1033
500	1.0	0.1732	0.0680
500	5.0	0.1283	0.1161
1000	0.5	0.1616	0.0687
1000	1.0	0.0768	0.0465
1000	5.0	0.0681	0.0661
2000	0.5	0.1242	0.0641
2000	1.0	0.0743	0.0643
2000	5.0	0.0542	0.0522

The naive Wald test over-rejects most strongly when $\epsilon = 0.5$: rejection rates are 0.269, 0.162, and 0.124 for $T = 500, 1000$, and 2000 . This pattern is consistent with standard errors that omit the added perturbation variance. The corrected rejection rates for $\epsilon = 0.5$ fall to 0.103, 0.069, and 0.064. Across all nine design points they range from 0.046 to 0.116. The correction therefore materially improves size, especially for $T \geq 1000$, but it does not deliver exact 5% control in every finite-sample setting. The largest residual distortions occur at $T = 500$, where the non-private QML benchmark itself has sub-nominal coverage under heavy tails and parameter projection is more frequent.

Taken together, the coverage and Wald results support the direction of the covariance correction: ignoring perturbation variance can materially understate uncertainty, while the corrected procedure moves inference toward the nominal target. The evidence should be interpreted as finite-sample support with remaining heavy-tail and boundary effects, not as exact verification in every scenario.

6. Conclusion

This study analyzes privacy-perturbed M-estimation for heavy-tailed GARCH models and separates an idealized asymptotic perturbation theory from the operational Gaussian-QML simulation.

The theoretical analysis shows how sampling variation and an independent Gaussian perturbation enter an asymptotic linear representation under bounded-score, stationarity, contraction, and identification conditions. The simulation uses a working perturbation calibration and parameter

projection. Because the initial QML estimate and H/J matrices use raw data and the complete recursive global sensitivity is not established, the operational algorithm is not claimed to provide an end-to-end differential-privacy guarantee.

The main inferential contribution is a covariance correction for the added perturbation variance. Across 500 requested replications per design point, the correction substantially improves confidence-interval coverage and reduces Wald over-rejection. Improvements are clearest under stronger perturbation, although exact nominal behavior is not obtained in every small-sample design.

RMSE declines as sample size grows, illustrating the perturbation-accuracy trade-off. The Q-Q evidence is approximately linear in the center but retains tail deviations, and the fitted RMSE slopes are based on only three sample sizes. These diagnostics are therefore descriptive rather than formal proofs of normality or convergence rates.

The simulation supports a narrower but practically important conclusion: conventional inference that ignores perturbation noise can be seriously misleading, and explicitly adding the noise covariance provides a substantial correction. It does not establish strict end-to-end privacy or uniform finite-sample superiority of the perturbed estimator.

Three limitations remain. First, an end-to-end DP release would require a complete global-sensitivity analysis and private construction of the initial estimate and covariance ingredients. Second, the projection rate reaches 19.4% at $T = 500$ and $\epsilon = 0.5$, so boundary-aware inference warrants further study. Third, the analysis is limited to a univariate simulated GARCH model; multivariate exten-

sions and applications to real financial data are left for future research.

Code and Data Availability

The R simulation code, design-point summaries, replication-level outputs, and failure diagnostics used to produce the reported tables and figures are available from the corresponding author upon reasonable request. The study uses simulated data and does not rely on a confidential empirical dataset.

References

- [1] Mukherjee K. M-estimation in GARCH models. *Econometric Theory*, 2008, 24(6): 1530-1553.
- [2] Liu H., Mukherjee K. R-estimators in GARCH models: Asymptotics and applications. *The Econometrics Journal*, 2022, 25(1): 98-113.
- [3] Patrocinio P. F., Reisen V. A., Bondon P., Monte E. Z., Danilevicius I. M. M-quantile estimation for GARCH models. *Computational Economics*, 2024, 63(6): 2175-2192.
- [4] Fiszeder P., Małacka M., Molnár P. Robust estimation of the range-based GARCH model: Forecasting volatility, value at risk and expected shortfall of cryptocurrencies. *Economic Modelling*, 2024, 141, 106887.
- [5] Avella-Medina M. Privacy-preserving parametric inference: A case for robust statistics. *Journal of the American Statistical Association*, 2021, 116(534): 969-983.
- [6] Avella-Medina M., Bradshaw C., Loh P.-L. Differentially private inference via noisy optimization. *The Annals of Statistics*, 2023, 51(5): 2067-2092.
- [7] Jewson J. E., Ghalebikesabi S., Holmes C. C. Differentially private statistical inference through beta-divergence one posterior sampling. *Advances in Neural Information Processing*

Systems, 2023, 36.

- [8] Engle R. F. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 1982, 50(4): 987-1007.
- [9] Bollerslev T. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 1986, 31(3): 307-327.
- [10] Hall P., Yao Q. Inference in ARCH and GARCH models with heavy-tailed errors. *Econometrica*, 2003, 71(1): 285-317.
- [11] Dwork C., McSherry F., Nissim K., Smith A. Calibrating noise to sensitivity in private data analysis. *Theory of Cryptography Conference (TCC 2006)*: Springer, 2006, 265-284.
- [12] Dong J., Roth A., Su W. J. Gaussian differential privacy. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 2022, 84(1): 3-37.
- [13] Kazan Z., Shi K., Groce A., Bray A. P. The test of tests: A framework for differentially private hypothesis testing. *Proceedings of the 40th International Conference on Machine Learning*, 2023, PMLR 202: 16131-16151.
- [14] Cai T. T., Wang Y., Zhang L. The cost of privacy: Optimal rates of convergence for parameter estimation with differential privacy. *The Annals of Statistics*, 2021, 49(5): 2825-2850.
- [15] Dwork C., Lei J. Differential privacy and robust statistics. *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*, 2009, 371-380.
- [16] Mao Y., Ye Q., Wang Q., Hu H. Differential privacy for time series: A survey. *IEEE Data Engineering Bulletin*, 2024, 47(2): 67-92.
- [17] Mao Y., Ye Q., Wang Q., Hu H. Utility-aware time series data release with anomalies under TLDP. *IEEE Transactions on Mobile Computing*, 2024, 23(6): 7135-7147.