Methodological Advances in Robust Small Area Estimation

María Bugallo Porto and Domingo Morales González

Center of Operations Research, Miguel Hernández University of Elche, 03202 Elche, Spain
Correspondence: mbugallo@umh.es

DOI: *https://doi.org/10.17979/spu.23.c10*

Abstract: Small Area Estimation techniques address the growing demand for reliable disaggregated statistics by fitting models to unit-level or area-level data. While recent advances in unit-level mixed models are significant, the presence of outliers has driven the development of robust methods. M-quantile regression offers a promising alternative to mixed models for Small Area Estimation, though some theoretical aspects remain underexplored. We cover the optimal selection of the robustness parameters for bias correction, improving outlier detection. In addition, we propose bootstrap methods to approximate the distribution of area-specific M-quantile coefficients and optimal robustness parameters, enhancing inference and diagnostics.

1 Introduction

Small Area Estimation methods offer model-based solutions to the challenge of producing reliable disaggregated statistics when direct survey estimates are unreliable due to small sample sizes (Morales et al., 2021; Rao and Molina, 2015). While classical Small Area Estimation relies on mixed models and strong distributional assumptions, M-quantile regression (Breckling and Chambers, 1988; Chambers and Tzavidis, 2006) offers a robust alternative by modeling conditional quantiles without including random effects, enhancing robustness to outliers and model misspecification.

In M-quantile based Small Area Estimation, area-specific M-quantile coefficients play a similar role to random effects in mixed models, capturing inter-area variability. This methodology has been successfully applied to predict indicators such as poverty rates and economic variables (e.g., Marchetti et al. (2018); Salvati et al. (2012)). Despite their advantages, robust methods can lead to biased predictions, particularly when the robustness mechanism is not optimally calibrated (Chambers et al., 2014).

A recent research by Bugallo et al. (2025b) proposes a data-driven bias correction through optimal robustness parameters for M-quantile linear prediction in small areas. Building on this, we approximate the distributions of area-specific M-quantile coefficients and optimal robustness parameters to formally identify potential atypical domains—those whose behavior deviates from the overall patterns (Bugallo and Morales, 2025). Our approach leverages parametric bootstrap techniques for M-quantile models (Bugallo et al., 2025a) and complements existing tests such as that by Bianchi et al. (2018), while being specifically tailored to detect extreme deviations in robust Small Area Estimation contexts.

This research contributes to the limited literature on formal diagnosis in M-quantile models by proposing several outlier detection methods, supported by solid theoretical foundations and simulation-based evidence.

2 Basic notation for inference in finite populations

Consider a finite population U of size N, partitioned into D non-overlapping domains U_d, each of size N_d for $d = 1, \ldots, D$. A sample s of size n is drawn, with $s_d \subseteq U_d$ denoting the n_d sampled units in area d and $r_d = U_d \backslash s_d$ the $N_d - n_d$ non-sampled units. Let $\boldsymbol{x}_{dj}$ be a vector of $p \geqslant 1$ auxiliary variables, known for all units $j = 1, \ldots, N_d$, and let y_{dj} be a continuous target variable observed only for $j = 1, \ldots, n_d$.

For the sake of notation, we define the following vectors and matrices with dimensions $N_d \times 1$ and $N_d \times p$, respectively:

$$\boldsymbol{y}_d = \underset{1 \leqslant j \leqslant N_d}{\text{col}} (y_{dj}) \quad \text{and} \quad X_d = \underset{1 \leqslant j \leqslant N_d}{\text{col}} (\boldsymbol{x}_{dj}).$$

3 The Nested Error Regression model for Small Area Estimation

The nested error regression model is widely used in unit-level Small Area Estimation to account for hierarchical structures with area-specific random effects.

The nested error regression model assumes that

$$\boldsymbol{y}_d \overset{ind.}{\sim} \mathcal{N}(X_d \boldsymbol{\beta}, V_d), \quad d = 1, \ldots, D, \tag{1.1}$$

where $\boldsymbol{\beta}$ is a p-dimensional vector of regression coefficients and $V_d = V(\tau^2, \sigma^2)$, with τ^2 and σ^2 the variances of the area effects and the unit-level errors, respectively.

If $\tau^2 = 0$, the model reduces to the synthetic regression model, with no between-area variability. For further details, see Chapter 7 of Morales et al. (2021).

4 M-quantile models for Small Area Estimation

This section briefly reviews the M-quantile modeling approach to Small Area Estimation by Chambers and Tzavidis (2006).

For $0 < q < 1$, the two-level M-quantile models are

$$y_{dj} = \boldsymbol{x}'_{dj} \boldsymbol{\beta}_\psi(q) + e_{\psi,dj}(q), \quad d = 1, \ldots, D, \; j = 1, \ldots, N_d, \tag{1.2}$$

where $\boldsymbol{\beta}_\psi(q)$ is a p-dimensional vector of regression coefficients depending on q and $e_{\psi,dj}(q)$ are independent model errors.

The conditional M-quantile function of order q (Breckling and Chambers, 1988) is

$$Q_q(y_{dj}; \sigma_q, \psi \mid \boldsymbol{x}_{dj}) = \boldsymbol{x}'_{dj} \boldsymbol{\beta}_\psi(q),$$

with $Q_q(e_{\psi,dj}(q); \sigma_q, \psi \mid \boldsymbol{x}_{dj}) = 0$ and homoscedastic variance $\sigma_q = \text{var}^{1/2}(e_{\psi,dj}(q))$. Estimation uses a robust loss function ρ, with influence function $\psi(u) = \rho'(u)$, via iteratively reweighted least squares (IRLS) (Bianchi and Salvati, 2015).

In this work, we adopt the Huber function:

$$\psi(u) = u \, \mathbb{I}_{(-c_\psi, c_\psi)}(u) + c_\psi \, \text{sgn}(u) \, \mathbb{I}_{\{|u| \geqslant c_\psi\}}, \quad c_\psi > 0, \tag{1.3}$$

typically with $c_\psi = 1.345$, ensuring 95% efficiency under normality.

Each unit-level M-quantile coefficient q_{dj} is defined as the solution to

$$q_{dj} = \arg \min_{0<q<1} \left\{ Q_q(y_{dj}; \sigma_q, \psi \mid \boldsymbol{x}_{dj}) = y_{dj} \right\},$$

interpreted as the quantile level best representing unit j in area d (Chambers and Tzavidis, 2006; Dawber and Chambers, 2019). For sampled units, it is predicted by

$$\hat{q}_{dj} = \arg \min_{0<q<1} \left\{ \boldsymbol{x}'_{dj} \hat{\boldsymbol{\beta}}_\psi(q) = y_{dj} \right\}. \tag{1.4}$$

Since hierarchical structures often explain part of the variability, units within the same area tend to share similar unit-level M-quantile coefficients. Therefore, the area-specific M-quantile coefficients are defined and predicted as

$$\theta_d = \frac{1}{N_d}\sum_{j=1}^{N_d} q_{dj}, \quad \hat{\theta}_d = \frac{1}{n_d}\sum_{j=1}^{n_d} \hat{q}_{dj},$$

where $\hat{\theta}_d$ serves as a pseudo-random effect summarizing area-level heterogeneity.

4.1 M-quantile linear prediction and error estimation in small areas

This study focuses on the prediction of area-level means, given by:

$$\overline{Y}_d = \frac{1}{N_d}\sum_{j=1}^{N_d} y_{dj} = \frac{1}{N_d}\left(\sum_{j\in s_d} y_{dj} + \sum_{j\in r_d} \boldsymbol{x}'_{dj}\boldsymbol{\beta}_\psi(\theta_d) + \sum_{j\in r_d} e_{\psi,dj}\right).$$

The M-quantile (MQ) predictor replaces unobserved values with model-based estimates:

$$\hat{\overline{Y}}_d^{mq} = \frac{1}{N_d}\left(\sum_{j\in s_d} y_{dj} + \sum_{j\in r_d} \boldsymbol{x}'_{dj}\hat{\boldsymbol{\beta}}_\psi(\hat{\theta}_d)\right).$$

To reduce bias, a bias correction based on a robust influence function ϕ is applied:

$$\hat{B}_\phi(\hat{\overline{Y}}_d^{mq}) = -\left(1-\frac{n_d}{N_d}\right)\frac{\hat{\sigma}_{\hat{\theta}_d}}{n_d}\sum_{j\in s_d}\phi(\hat{u}_{\psi,dj}), \quad \hat{u}_{\psi,dj} = \hat{\sigma}_{\hat{\theta}_d}^{-1}(y_{dj} - \boldsymbol{x}'_{dj}\hat{\boldsymbol{\beta}}_\psi(\hat{\theta}_d)), \tag{1.5}$$

which results in the bias-corrected M-quantile (BMQ) predictor developed by Chambers et al. (2014):

$$\hat{\overline{Y}}_d^{bmq} = \hat{\overline{Y}}_d^{mq} + \left(1-\frac{n_d}{N_d}\right)\frac{\hat{\sigma}_{\hat{\theta}_d}}{n_d}\sum_{j\in s_d}\phi(\hat{u}_{\psi,dj}). \tag{1.6}$$

We adopt the Huber function for ϕ, with tuning constant $c_{\phi,d} \geqslant 0$. To balance bias and variance, Bugallo et al. (2025b) propose selecting the optimal $c_{\phi,d}$ by minimizing an estimate of the mean squared error (MSE) of $\hat{\overline{Y}}_d^{bmq}$:

$$\hat{c}_{\phi,d} = \arg\min_{c_{\phi,d}\geqslant 0} \text{mse}_d^{bmq}(c_{\phi,d}) = \arg\min_{c_{\phi,d}\geqslant 0} A_d(c_{\phi,d}), \tag{1.7}$$

where $A_d(c_{\phi,d})$ is the $c_{\phi,d}$-dependent part of the estimated MSE (Chambers et al., 2011, 2014), given in practice by:

$$A_d(c_{\phi,d}) = \left(1-\frac{n_d}{N_d}\right)^2\left(\frac{\hat{\sigma}_{\hat{\theta}_d}}{n_d}\right)^2\left[\sum_{j\in s_d}\phi^2(\hat{u}_{\psi,dj}) + \left(\sum_{j\in s_d}[\phi(\hat{u}_{\psi,dj}) - \hat{u}_{\psi,dj}]\right)^2\right].$$

The existence and uniqueness of $\hat{c}_{\phi,d}$ are established in Bugallo et al. (2025b).

4.2 Bootstrap inference

Several bootstrap methods are proposed to approximate the distributions of $\hat{c}_{\phi,d}$ and $\hat{\theta}_d$ across domains, $d = 1,\ldots,D$, using three approaches (Bugallo and Morales, 2025): a non-parametric (NP) method, a parametric (AP) approach based on approximated residuals and a Naïve (NA) alternative relying on the distribution of the model errors.

Let $b = 1,\ldots,B$ denote the bootstrap replicates and $d = 1,\ldots,D$.

(a1) **Algorithm NP.** Sample n_d standardized residuals $\hat{u}^{*(b)}_{\psi,dj}$ with replacement from the set $\{\hat{u}_{\psi,d1},\ldots,\hat{u}_{\psi,dn_d}\}$.

(a2) **Algorithm AP.** Generate n_d values $\hat{e}^{*(b)}_{\psi,dj} = \left[(\mathbf{I}-\mathbf{H}(\hat{\theta}_d))'\boldsymbol{\xi}^{*(b)}\right]_{dj}$, where $\xi^{*(b)}_{gk} \sim$ GALI$(x'_{gk}\hat{\boldsymbol{\beta}}_\psi(\hat{\theta}_d), \hat{\sigma}_{\hat{\theta}_d}, \hat{\theta}_d)$ and $H(\hat{\theta}_d)$ is the analogue in M-quantile models of the projection matrix in linear models. Then set $\hat{u}^{*(b)}_{\psi,dj} = (\hat{\sigma}^{*(b)}_{\hat{\theta}_d})^{-1}\hat{e}^{*(b)}_{\psi,dj}$.

(a3) **Algorithm NA.** Generate n_d i.i.d. values $\hat{u}^{*(b)}_{\psi,dj} \sim$ GALI$(0,1,\hat{\theta}_d)$.

The term GALI denotes the generalized asymmetric least informative distribution with location $\mu_{\hat{\theta}_d} = Q_{\hat{\theta}_d}(y_{gk};\hat{\sigma}_{\hat{\theta}_d},\psi|x_d) = \boldsymbol{x}'_{gk}\hat{\beta}_\psi(\hat{\theta}_d)$, scale $\hat{\sigma}_{\hat{\theta}_d}$ and probability $q = \hat{\theta}_d$. This distribution was introduced by Bianchi et al. (2018) to provide a working likelihood for inference in M-quantile models, generalizing the asymmetric Laplace distribution commonly associated with quantile regression.

First, Algorithm NP uses the empirical distribution of the residuals, which may be inaccurate when n_d is small. Algorithm AP is more stable for small/moderate sample sizes, as it approximates the residual distribution. Finally, Algorithm NA is the most conservative method and does not depend on $\hat{\sigma}_{\hat{\theta}_d}$ or N_d, only on $\hat{\theta}_d$ and n_d. Though Algorithm NA ignores the variability from estimating $\hat{\boldsymbol{\beta}}_\psi(\hat{\theta}_d)$ and predicting $\hat{\theta}_d$.

For each area $d = 1,\ldots,D$, we approximate the distribution of $\hat{c}_{\phi,d}$ as:

1. For $b = 1,\ldots,B$ (let $j = 1,\ldots,n_d$, $d = 1,\ldots,D$):
   a) Generate $\hat{u}^{*(b)}_{\psi,dj}$ using Algorithms NP, AP or NA.
   GALI random variables are simulated in R using custom code.
   b) Define the bootstrap $c_{\phi,d}$-dependent part of the estimated MSE, $A^{*(b)}_d(c_{\phi,d})$.
   c) Compute $\hat{c}^{*(b)}_{\phi,d} = \arg\min_{c_{\phi,d}\geqslant 0} A^{*(b)}_d(c_{\phi,d})$.
2. Sort $\hat{c}^{*(1)}_{\phi,d},\ldots,\hat{c}^{*(B)}_{\phi,d}$ to estimate the distribution of $\hat{c}_{\phi,d}$ via bootstrap.

A similar procedure is used to approximate the distribution of $\hat{\theta}_d$.

4.3 Hypothesis testing

We consider the hypothesis H_0: area d is not atypical, $d = 1,\ldots,D$, and propose three approaches to test it: (1) based on the optimal robustness parameters; (2) based on the area-specific M-quantile coefficients; and (3) based on the area-specific random effects from nested error regression models (see Section 3).

Let p-value$_d$ be the probability of rejecting H_0 for area d. For non-atypical areas, p-value$_d$ estimates the test size. To evaluate the performance of the different approaches in model-based simulations (see Section 5), we compute the average rejection probability across simulations ($s = 1,\ldots,S$) as follows:

(1) **Optimal robustness parameters.** Define

$$\frac{1}{S}\sum_{s=1}^{S} I\big(\hat{c}^{(s)}_{\phi,d} > \hat{c}^{*(s)}_{\phi,d(\lfloor(1-\alpha)B\rfloor)}(0.5)\big), \quad d = 1,\ldots,D,$$

where $\hat{c}^{*(s)}_{\phi,d(\lfloor(1-\alpha)B\rfloor)}(0.5)$ is computed via parametric bootstrap with B replicates.

(2) **Area-specific M-quantile coefficients.** Define

$$\frac{1}{S}\sum_{s=1}^{S} I\Big(\hat{\theta}^{*(s)}_d(\lfloor(1-\alpha/2)B\rfloor) - \hat{\theta}^{*(s)}_d(\lfloor\alpha/2\cdot B\rfloor) < \ell\Big), \quad d = 1,\ldots,D,$$

where $\ell \in (0,1)$ is a threshold to be selected in a data-driven way:

$$\hat{\ell}^{(s)} = \frac{1}{D}\sum_{d=1}^{D}\left(\hat{\theta}_d^{(s)}\left(\lfloor(1-\alpha/2)\cdot B\rfloor\right) - \hat{\theta}_d^{(s)}\left(\lfloor\alpha/2\cdot B\rfloor\right)\right), \quad s = 1,\dots,S.$$

(3) **Area-specific random effects.** Under nested error regression models, the standardized random effects $\hat{u}_d/\hat{\sigma}_u$ are expected to follow a standard normal distribution, $\mathcal{N}(0,1)$. Define

$$\frac{1}{S}\sum_{s=1}^{S} I\left(\hat{\sigma}_u^{-(s)}\cdot \hat{u}_d^{(s)} \in (z_{\alpha/2}, z_{1-\alpha/2})\right), \quad d = 1,\dots,D,$$

where $z_{\alpha/2}$ and $z_{1-\alpha/2}$ are the standard normal quantiles.

5 Some outstanding empirical results

The performance of the new methods for detecting atypical areas (Section 4.3) is evaluated through model-based simulations inspired by Chambers et al. (2014). Population and sample sizes are fixed at $N_d = 100$ and $n_d = 5$, respectively. Area-level random effects and unit-level errors are generated for the following three scenarios:

$[0,0]$ — No outliers: $u_d \sim \mathcal{N}(0,3)$, $e_{dj} \sim \mathcal{N}(0,6)$, $d = 1,\dots,D = 40$.

$[e,0]$ — Unit-level outliers: $e_{dj} \sim \delta\,\mathcal{N}(0,6) + (1-\delta)\,\mathcal{N}(20,150)$, $\delta \sim$ Bernoulli(0.97).

$[e,u]$ — Area- and unit-level outliers: same e_{dj} as above, and $u_d \sim \mathcal{N}(9,\sqrt{20})$ for areas $d = 37,\dots,40$.

Each scenario is replicated $S = 500$ times. The target variable is defined as

$$y_{dj}^{(s)} = 100 + 5x_{dj} + u_d^{(s)} + e_{dj}^{(s)}, \quad s = 1,\dots,S,\ j \in U_d,\ d = 1,\dots,D,$$

where $x_{dj} \sim$ LogNormal(1, 0.5) and $u_d^{(s)}, e_{dj}^{(s)}$ are generated according to the scenario.

Table 1 shows the average p-value$_d$ estimates for non-atypical and atypical areas for Methods (1)–(3), based on simulations with nominal level $\alpha = 0.05$.

Method	[0,0]	[e,0]	[e,u] $1 \leqslant d \leqslant 36$	[e,u] $37 \leqslant d \leqslant 40$
(1)	0.118	0.149	0.211	0.745
(2)	0.050	0.050	0.056	0.855
(3)	0.051	0.053	0.012	0.510

Table 1: Average estimates of p-value $_d$ for non-atypical and atypical areas (if applicable) in model-based simulations across different scenarios for $\alpha = 0.05$. The best option is highlighted in light pink.

Columns "[0,0]", "[e,0]" and "[e,u] $1 \leqslant d \leqslant 36$" report results for non-atypical areas, where values near 0.05 indicate good Type I error control; the best performer is Method (2). Column "[e,u] $37 \leqslant d \leqslant 40$" focuses on truly atypical areas, where a larger p-value$_d$ indicate greater power to detect outliers. Among the methods, Approach (2) balances error control and detection power best, yielding p-value$_d$ close to 0.05 in non-atypical areas and high values for the atypical ones. Approach (3), which assumes no contamination, performs poorly in detecting atypical areas.

The performance of Method (2) strongly depends on the choice of $\ell \in (0,1)$, requiring careful selection. Figure 1 summarizes its behavior across scenarios.

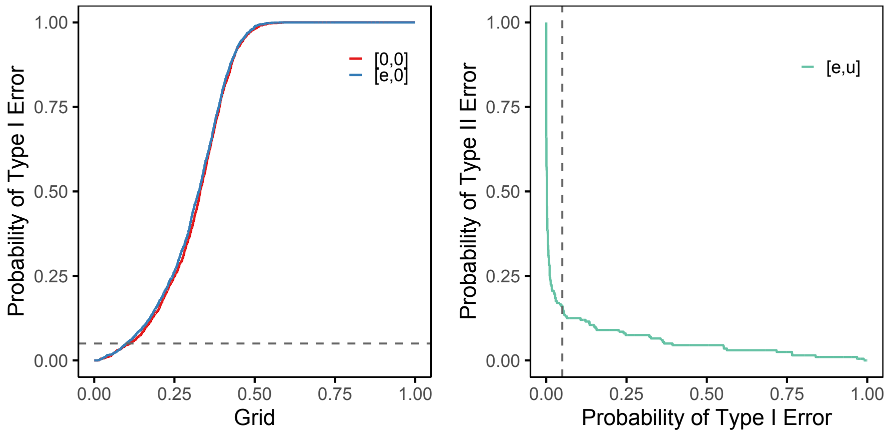

Figure 1: Performance of Method (2) for different settings. Left: Type I error versus $\ell \in (0, 1)$. Right: Type I and Type II errors. The dotted line next to the Type I error is $\alpha = 0.05$.

The left panel in Figure 1 shows the Type I error probability as the parameter $\ell \in (0, 1)$ varies, while the right panel presents an inverted Receiver Operating Characteristic (ROC) curve – plotting Type I error on the x-axis and Type II error on the y-axis– illustrating the trade-off between both errors. Curves closer to the origin in the right panel indicate more powerful and reliable tests. The area under the ROC curve (AUC) value is 0.942 for Scenario [e, u]. Higher AUC values reflect better discrimination, with Method (2) performing quite well for Scenario [e, u].

6 Conclusions

Small Area Estimation is crucial for producing reliable detailed statistics, especially when data are scarce or heterogeneous. This work highlights the importance of robust methods to handle outliers effectively. By focusing on M-quantile regression, we improve the bias correction process through optimal tuning of the robustness parameter, which enhances the accuracy of the estimates. Moreover, our proposed bootstrap techniques provide practical tools to approximate the distributions of area-specific M-quantile coefficients and optimal robustness parameters.

These advances facilitate more reliable inference and improve the detection of atypical areas, making the M-quantile models a competitive and flexible alternative to traditional mixed models in Small Area Estimation.

Future studies should investigate the robustness properties of the predictors derived from M-quantile models and examine how the selection of the optimal robustness parameters influences their performance. For instance, assessing robustness in the presence of outliers or heavily skewed distributions remains an open problem. Additionally, the insights derived from the optimal selection of the robustness parameters offer valuable guidance for enhancing future M-quantile applicationss.

Bibliography

A. Bianchi and N. Salvati. Asymptotic properties and variance estimators of the M-quantile regression coefficients estimators. *Communications in Statistics*, 44:813–827, 2015.

A. Bianchi, E. Fabrizi, N. Salvati, and N. Tzavidis. Estimation and testing in M-quantile regression with applications to Small Area Estimation. *International Statistical Review*, 86(3): 541–570, 2018.

J. Breckling and R. Chambers. M-quantiles. *Biometrika*, 75(4):761–771, 1988.

M. Bugallo and D. Morales. Inference and diagnosis in M-quantile models with applications to Small Area Estimation. *Unpublished manuscript*, 2025.

M. Bugallo, D. Morales, N. Salvati, and F. Schirripa. Bias adjustment for Mean Squared Error in M-quantile models for Small Area Estimation. *Unpublished manuscript*, 2025a.

M. Bugallo, D. Morales, N. Salvati, and F. Schirripa. Temporal M-quantile models and robust bias-corrected small area predictors. *Journal of Survey Statistics and Methodology.*, 2025b.

R. Chambers and N. Tzavidis. M-quantile models for Small Area Estimation. *Biometrika*, 93(2): 255–268, 2006.

R. Chambers, H. Chandra, and N. Tzavidis. On bias-robust Mean Squared Error estimation for pseudo-linear small area estimators. *Survey Methodology*, 37(2):153–170, 2011.

R. Chambers, H. Chandra, N. Salvati, and N. Tzavidis. Outlier robust Small Area Estimation. *Journal of the Royal Statistical Society: Series B*, 76(1):47–69, 2014.

J. Dawber and R. Chambers. Modelling group heterogeneity for Small Area Estimation using M-quantiles. *International Statistical Review*, 87(1):50–63, 2019.

S. Marchetti, M. Beresewicz, N. Salvati, M. Szymkowiak, and L. Wawrowski. The use of a three-level M-quantile model to map poverty at local administrative unit 1 in Poland. *Journal of the Royal Statistical Society, Series A*, 181(4):1077–1104, 2018.

D. Morales, M. D. Esteban, A. Pérez, and T. Hobza. *A Course on Small Area Estimation and Mixed Models: Methods, Theory and Applications in R*. Springer Nature, 2021.

J. N. K. Rao and I. Molina. *Small Area Estimation*. John Wiley & Sons, 2015.

N. Salvati, N. Tzavidis, M. Pratesi, and R. Chambers. Small Area Estimation via M-quantile Geographically Weighted Regression. *TEST*, 21:1–28, 2012.