Published in final edited form as:

Stat Interface. 2019 ; 12(1): 71–75. doi:10.4310/SII.2019.v12.n1.a7.

Letter to the Editor

Marco Geraci

Department of Epidemiology and Biostatistics, University of South Carolina, Columbia, SC 29208, USA

Abstract

Galarza, Lachos and Bandyopadhyay (2017) have recently proposed a method of estimating linear quantile mixed models (Geraci and Bottai, 2014) based on a Monte Carlo EM algorithm. They assert that their procedure represents an improvement over the numerical quadrature and non-smooth optimization approach implemented by Geraci (2014). The objective of this note is to demonstrate that this claim is incorrect. We also point out several inaccuracies and shortcomings in their paper which affect other results and conclusions that can be drawn.

1. LINEAR QUANTILE MIXED MODELS

Linear quantile mixed models (LQMMs) were developed by Geraci and Bottai [4] as an extension of the quantile regression model with random intercepts of Geraci and Bottai [3]. We consider data from two-level nested designs in the form $(\mathbf{x}_{ij}^\top, \mathbf{z}_{ij}^\top, y_{ij})$, for $j = 1, \dots, n_i$ and $i = 1, \dots, M$, $N = \sum_i n_i$ where $\mathbf{x}_{ij}^\top$ is the j th row of a known $n_i \times p$ matrix $\mathbf{X}_i$, $\mathbf{z}_{ij}^\top$ is the j th row of a known $n_i \times q$ matrix $\mathbf{Z}_i$ and y_{ij} is the j th observation of the response vector $\mathbf{y}_i = (y_{i1}, \dots, y_{in_i})^\top$ for the i th cluster. The $N \times 1$ vector of responses is denoted by $\mathbf{y} = (\mathbf{y}_1^\top, \dots, \mathbf{y}_M^\top)^\top$. This kind of data arise from longitudinal or panel studies and other cluster sampling designs.

The τ th LQMM is defined as

$$Q_{y_{ij}|\mathbf{u}_i}(\tau) = \mathbf{x}_{ij}^\top \boldsymbol{\beta}_\tau + \mathbf{z}_{ij}^\top \mathbf{u}_i, \quad (1)$$

where $0 < \tau < 1$ is the given quantile level, $\boldsymbol{\beta}_\tau$ is a $p \times 1$ vector of τ -specific coefficients that are common to all clusters, while the $q \times 1$ vector $\mathbf{u}_i$ may vary with cluster. For estimation purposes only, Geraci and Bottai [4] introduced the convenient assumption that the responses y_{ij} , $j = 1, \dots, n_i$, $i = 1, \dots, M$, conditionally on a $q \times 1$ vector of random effects $\mathbf{u}_i$, independently follow the asymmetric Laplace (AL) density

$$p(y_{ij}|\mathbf{u}_i) = \frac{\tau(1-\tau)}{\sigma_\tau} \exp\left\{-\frac{1}{\sigma_\tau} \rho_\tau(y_{ij} - \mu_{\tau,ij})\right\}, \quad (2)$$

where $\rho_\tau(r) = r \{ \tau - I(r < 0) \}$ is the ‘check’ function and I denotes the indicator function, with location and scale parameters given by $\mu_{\tau,ij} = \mathbf{x}_{ij}^\top \boldsymbol{\beta}_\tau + \mathbf{z}_{ij}^\top \mathbf{u}_i$ and σ_τ respectively, which we write as $y_{ij} \sim \mathcal{AL}(\mu_{\tau,ij}, \sigma_\tau)$. (The third parameter of the AL is the skew parameter $\tau \in (0, 1)$ which, in this model, is fixed and defines the quantile level of interest.) Also, they assumed that $\mathbf{u}_i = (u_{i1}, \dots, u_{iq})^\top$, for $i = 1, \dots, M$, is a random vector independent from the model’s error term with mean zero and τ -specific variance-covariance matrix $\boldsymbol{\Sigma}_\tau$ of dimensions $q \times q$. The latter is reparameterized in terms of an m -dimensional vector, $1 \leq m \leq q(q+1)/2$, of non-redundant parameters $\boldsymbol{\theta}_\tau$ i.e. $\boldsymbol{\Sigma}_\tau = \boldsymbol{\Sigma}(\boldsymbol{\theta}_\tau)$.

The algorithm to estimate $(\boldsymbol{\beta}_\tau, \boldsymbol{\theta}_\tau, \sigma_\tau)$ is described in detail by [2, 4]. First, the (quasi) log-likelihood is integrated numerically over the distribution of the random effects, i.e.

$$\ell_{\text{GQ}}(\boldsymbol{\beta}_\tau, \boldsymbol{\theta}_\tau, \sigma_\tau | \mathbf{y}) = \sum_i^M \log \left\{ \sum_{k_1=1}^K \cdots \sum_{k_q=1}^K p(\mathbf{y}_i | \mathbf{v}_{k_1, \dots, k_q}) \prod_{l=1}^q w_{k_l} \right\}, \quad (3)$$

with $\mathbf{v}_{k_1, \dots, k_q} = (v_{k_1}, \dots, v_{k_q})^\top$, where v_{k_l} and w_{k_l} , $k_l = 1, \dots, K$, $l = 1, \dots, q$, denote the abscissas and the weights of the (one-dimensional) Gaussian quadrature. Second, the integrated log-likelihood (3) is maximized via a non-smooth optimizer. In principle, one can consider different distributions for the random effects, which may be naturally linked to different quadrature rules (or penalties). For example, it is immediate to verify that the normal distribution is akin to a Gauss-Hermite quadrature, a special case of LQMM discussed by Geraci and Bottai.

In their paper, Galarza et al. write

[Geraci and Bottai (2014)] extended [Geraci and Bottai’s (2007)] setup to accommodate multiple random effects [...]. Here, we consider a more general correlated random effects framework with general dispersion matrix $\boldsymbol{\Psi} = \boldsymbol{\Psi}(\mathbf{a})$.

This statement is not justified since their model is precisely the LQMM defined in (1)-(2) with normal random effects and general variance-covariance matrix $\boldsymbol{\Sigma}_\tau = \boldsymbol{\Sigma}(\boldsymbol{\theta}_\tau)$.

2. SIMULATION STUDY

To fit the LQMM (1)-(2), Galarza et al. proposed to use a stochastic approximation of the expectation-maximization (SAEM) algorithm. They compared their estimation approach, as implemented in the qrLMM package, to the quadrature-based algorithm implemented in the lqmm package [2]. The description of their simulation study is not clear as it lacks several

details. First of all, it does not specify which versions of lqmm, qrLMM or R were used in their study. Most importantly, there is no indication about lqmm optimization settings and the syntax used for modeling. We speculate that the default options were used. We tried to replicate their design and we believe that the setting described below is quite similar to theirs.

We generated data according to the model

$$y_{ij} = \mathbf{x}_{ij}^T \boldsymbol{\delta} + \mathbf{z}_{ij}^T \mathbf{u}_i + \varepsilon_{\tau, ij}, \quad (4)$$

where $\boldsymbol{\delta} = (0.8, 0.5, 1)^T$, $x_{1,ij} = 1$, $x_{2,ij} \sim \mathcal{N}(0, 1)$, $x_{3,ij} \sim \mathcal{N}(0, 1)$, $z_{1,ij} \sim \mathcal{N}(0, 1)$, $z_{2,ij} \sim \mathcal{N}(0, 1)$, and $\varepsilon_{\tau, ij} \sim \mathcal{AL}(0, 0.2)$. Moreover, $\mathbf{u}_i \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$, where

$$\boldsymbol{\Sigma} = \begin{bmatrix} 0.8 & 0.5 \\ 0.5 & 1 \end{bmatrix}.$$

The number of clusters M varied (50, 100, 200, 300) while the size of the clusters was fixed to $n_i = 3$, $i = 1, \dots, M$, throughout the simulation. We considered five LQMMs (1) for $\tau \in \{0.05, 0.1, 0.5, 0.9, 0.95\}$. (Note that the error $\varepsilon_{\tau, ij}$ in (4) is sampled from an AL with skewness determined by the same τ that defines the quantile to be estimated, therefore $\boldsymbol{\beta}_\tau = \boldsymbol{\delta}$ for all τ .) Data were replicated 100 times for each combination of sample size and quantile level.

For this simulation, we used lqmm 1.5.3, which is, at the time of writing, the latest version available on the Comprehensive R Archive Network, and qrLMM 1.3 for R version 3.4.2 [6]. By default, the function QRLMM starts the SAEM algorithm with estimates of $\boldsymbol{\beta}_\tau$ and σ_τ obtained from linear programming (package quantreg). In contrast, lqmm starts by default from ordinary least squares estimates. Therefore, we changed the option lqmmControl(startQR = TRUE) for the sake of comparability. Moreover, we used $K = 9$ quadrature knots instead of the default $K = 7$ to improve accuracy since $q > 1$ (see [4] for details). For the SAEM algorithm we used the same settings as in Galarza et al., namely 20 Monte Carlo simulations, 500 maximum iterations and 0.2 for the cut-point that determines the proportion of initial iterations with no memory. The variance-covariance matrix was specified as a general positive-definite matrix in both estimation procedures. All the other estimation settings in lqmm and QRLMM were left unchanged to their default values. In a preliminary analysis, we assessed the computational time needed to run the full simulation and we estimated it would take approximately two months for QRLMM, but less than half an hour for lqmm. Given the excessive computational time needed for QRLMM, we ran the latter only for selected scenarios, namely $M \in \{50, 300\}$ and $\tau \in \{0.05, 0.5, 0.95\}$.

Table 1 shows a summary of the actual performance of the two algorithms, while Figures 1 and 2 show, respectively, the absolute bias and root mean squared error (RMSE) of the two estimators. For comparison, Figure 2 also shows the RMSE values reported for SAEM by Galarza et al. in their paper [1, Table 2].

The average bias and RMSE calculated for β_τ and σ_τ in selected scenarios were small for both estimators (Table 1), with average RMSE slightly lower for lqmm. Most notably, the time needed by qrLMM to run one model was on average 42.52 minutes (min 5.82, max 132.30 minutes) with a 21% convergence failure rate. In contrast, lqmm took less than 1 second for one replication (min 0.09, max 9.27 seconds) with no convergence failures in any of the selected scenarios (and no convergence failures in any of the other scenarios either). The bias and RMSE for specific sample sizes and quantiles were similar in most cases except for the intercept, in which case lqmm seemed to have an advantage over qrLMM at more extreme quantiles (Figures 1 and 2). Note also that the RMSE results reported by Galarza et al. in their paper are close to those we obtained in our simulation for selected scenarios (Figure 2); thus, it is reasonable to conclude that our simulation design is a faithful reproduction of theirs.

Table 2 shows that the (scaled) average log-likelihood resulting from the two fitting algorithms was comparable in all selected scenarios.

As a side note, we found Galarza et al.'s simulation setting rather unusual since the vector $\mathbf{z}_{ij}$ is typically a subset of $\mathbf{x}_{ij}$, for the random effects are constrained to have zero-mean (see also the discussion in the next section). Moreover, the AL distribution in LQMM provides only a quasi-likelihood for point estimation, i.e. it is not assumed to be the true distribution. Galarza et al.'s simulation is very limited and is more of a sanity check than a validation study. It would be more realistic to simulate errors from a variety of distributions with different shapes, along with heteroscedastic variants of these models. An extensive simulation of this kind is provided by Geraci and Bottai [4].

3. FRAMINGHAM STUDY

Galarza et al. also provided a comparison between the two algorithms using a subset of the cholesterol data from the Framingham study [8]. Once again, the description of their analysis is incomplete. First of all, the authors state

Interestingly, for the extremes quantiles, some warnings messages on convergence were displayed while fitting Geraci's method, even after increasing the number of iterations and reducing [sic] the tolerance, as suggested in the lqmm manual.

The authors do not say which estimation settings were used initially and how these were changed afterwards. Most importantly, they do not say whether the warning messages were obtained during model fitting or bootstrapping, and how many warnings were produced. These warnings may be of little concern (see [2] for a discussion on this point) and, as shown further below, they can be addressed with an appropriate tweaking of optimization parameters [2], along with a thoughtful examination of the data and model.

The authors are also silent on the model for Σ_τ . This is irrelevant for the qrLMM package since it provides only one model (i.e., the general positive-definite matrix). However, the lqmm package provides four different models, including the diagonal variance-covariance structure as the default.

We then decided to replicate the analysis of Galarza et al. who considered the model

$$Q_{y_{ij}|\mathbf{u}_t}(\tau) = \beta_{\tau,0} + \beta_{\tau,1}sex_i + \beta_{\tau,2}age_i + u_{1,i} + u_{2,i}T_{ij} \quad (5)$$

where y_{ij} is the ij th measurement of cholesterol (divided by 100) and $T_{ij} = (t_{ij} - 5)/10$, with t_{ij} denoting years since the beginning of the study.

There is clearly something awkward about the specification of model (5) since it does not include a fixed coefficient for T_{ij} . We can only surmise that Galarza et al. misinterpreted equation (9) in [8], where the random slope is actually centered about the fixed slope (not about zero). Whatever the reason, suffice it to say that this oversight not only may introduce bias in the estimates, but it might also render estimation more difficult and prone to failure if the fixed slope is effectively different from zero. Therefore, we proceeded with the following corrected model

$$Q_{y_{ij}|\mathbf{u}_t}(\tau) = \beta_{\tau,0} + \beta_{\tau,1}sex_i + \beta_{\tau,2}age_i + \beta_{\tau,3}T_{ij} + u_{1,i} + u_{2,i}T_{ij} \quad (6)$$

where $(u_{1,i}, u_{2,i})^T \sim \mathcal{N}(\mathbf{0}, \Sigma_\tau)$ and Σ_τ is a general positive-definite matrix. Model fitting and bootstrap standard error estimation were carried out for 19 vigintiles as detailed in Appendix A. The tolerance parameters set for lqmm estimation were less restrictive than the default values but still within reasonable bounds. We obtained only two warning messages of failed convergence during bootstrapping, but not during model estimation. Considering that $19 \times 50 = 950$ models were fitted for the bootstrap, this issue is hardly worthy of note in this case.

Finally, Galarza et al.'s statement

We observe that our SAEM method leads to mostly smaller SEs and AIC compared to the Geraci method. [...] Hence [...] the substantial gain in the AIC criterion and the SEs establish that our SAEM approach provides a much better fit to the dataset

is a hodgepodge of claims that are not substantiated anywhere in their paper. First, the authors found in their simulation study that the asymptotic approximations provide valid standard errors [1, Table 1]. This should be expected since the data were generated under the AL distribution. Theoretical results actually show that outside the AL case it is not appropriate to quantify uncertainty using this distribution as it leads to underestimation of the true variability [7]. This is why lqmm makes use of bootstrap. Moreover, the authors never provided a simulation study to compare SAEM's standard errors with those obtained with lqmm; thus, it is not possible to understand the nature of the differences found for one particular dataset (even if we grant for the sake of argument that the lqmm's estimation settings and modeling syntax were appropriately specified in the Framingham data analysis). Secondly, the comparison between average log-likelihoods in our simulation (Table 2) does not support the superiority of SAEM in terms of goodness of fit (GOF). In summary, the issue of standard error estimation remains to be investigated, while empirical evidence,

although limited, contradicts Galarza et al.'s claim that SAEM gives a better GOF performance.

4. CONCLUSION

Linear quantile mixed models [4] represent a valuable tool available to the scientific community. Computational issues are still an open problem and different approaches have been investigated by several researchers (see [5] for an overview). Galarza et al.'s claim of SAEM's superior performance fails to stand up to closer examination. Our simulation shows that while SAEM produces finite-sample bias and RMSE comparable to those obtained from the quadrature-based algorithm in lqmm, its sluggish convergence and high proportion of convergence failures put Galarza et al.'s proposal at enormous disadvantage.

ACKNOWLEDGEMENTS

This research was partially supported by the National Institutes of Health – National Institute of Child Health and Human Development (Grant Number: 1R03HD084807-01A1).

APPENDIX A.: R CODE

In this Appendix, we provide the R code for the analysis of the Framingham cholesterol data using the lqmm package.

```
library(lqmm)
data(Cholesterol, package = "qrLMM")
Cholesterol$year.c <- (Cholesterol$year - 5)/10
Cholesterol$sex <- as.factor(Cholesterol$sex)
# Set optimization parameters
ctrl <- lqmmControl(method = "df", LP_tol_ll = 1e-3, LP_max_iter = 2000,
startQR = TRUE)
# Fit model for tau = 0.05, 0.1, ..., 0.95
fit <- lqmm(I(cholst/100) ~ year.c + sex + age, random = ~ year.c, group =
ID, data = Cholesterol, tau = 1:19/20,
covariance = "pdSymm", control = ctrl)
# Bootstrap (50 replicates)
fit.s <- summary(fit, R = 50, seed = 178)
```

REFERENCES

- [1]. Galarza CE, Lachos VH, and Bandyopadhyay D (2017). Quantile regression in linear mixed models: A stochastic approximation EM approach. *Statistics and Its Interface* 10 471–482. MR3608555 [PubMed: 29104713]
- [2]. Geraci M (2014). Linear quantile mixed models: The lqmm package for Laplace quantile regression. *Journal of Statistical Software* 57 1–29. [PubMed: 25400517]
- [3]. Geraci M and Bottai M (2007). Quantile regression for longitudinal data using the asymmetric Laplace distribution. *Biostatistics* 8 140–154. [PubMed: 16636139]
- [4]. Geraci M and Bottai M (2014). Linear quantile mixed models. *Statistics and Computing* 24 461–479. MR3192268

- [5]. Marino MF and Farcomeni A (2015). Linear quantile regression models for longitudinal experiments: An overview. *Metron* 73 229–247. MR3386219
- [6]. R Core Team (2017). R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria URL <https://www.R-project.org/>.
- [7]. Yang Y, Wang HJ, and He X (2016). Posterior inference in bayesian quantile regression with asymmetric Laplace likelihood. *International Statistical Review* 84 327–344. MR3580414
- [8]. Zhang D and Davidian M (2001). Linear mixed models with flexible distributions of random effects for longitudinal data. *Biometrics* 57 795–802. MR1859815 [PubMed: 11550930]

Figure 1.

Absolute values of bias for linear quantile mixed effects estimators based on quadrature & non-smooth optimization (filled circles) and approximated EM (filled triangles): $\beta_{\tau,0}$ (upper left panels), $\beta_{\tau,1}$ (upper right panels), $\beta_{\tau,2}$ (lower left panels), and σ_{τ} (lower right panels).

Figure 2.

Root mean squared error (RMSE) for linear quantile mixed effects estimators based on quadrature & non-smooth optimization (filled circles) and approximated EM (filled triangles): $\beta_{\tau,0}$ (upper left panels), $\beta_{\tau,1}$ (upper right panels), $\beta_{\tau,2}$ (lower left panels), and σ_{τ} (lower right panels). The RMSE values reported by Galarza et al. [1, Table 2] for the approximated EM are marked with empty triangles.

Performance summary for the quadrature & non-smooth optimization algorithm (lqmm) and the approximated EM algorithm (qrLMM) run on a 64-bit operating system machine with 32 Gb of RAM and 3.60 GHz clock-rate processor. All figures refer to the same subset of scenarios, namely $M \in \{50, 300\}$ and $\tau \in \{0.05, 0.5, 0.95\}$. Averages are calculated over $2 \times 3 \times 100 = 600$ replicated datasets

Table 1.

Algorithm	Average bias	Average root mean squared error	Total elapsed time	Average elapsed time	Percentage of convergence failures
lqmm	0.0019	0.0852	7.5 (min)	0.7 (s)	0%
qrLMM	0.0017	0.0910	24066.0 (min)	2551.2 (s)	21%

Table 2.

Average log-likelihood at convergence (scaled by n) for the quadrature & non-smooth optimization algorithm (lqmm) and the approximated EM algorithm (qrLMM) run on a 64-bit operating system machine with 32 Gb of RAM and 3.60 GHz clock-rate processor. All figures refer to the same subset of scenarios, namely $M \in \{50, 300\}$ and $\tau \in \{0.05, 0.5, 0.95\}$. Averages are calculated over $2 \times 3 \times 100 = 600$ replicated datasets

Algorithm	Sample size	$\tau = 0.05$	$\tau = 0.5$	$\tau = 0.95$
lqmm	50	-7.78	-4.14	-7.77
	300	-7.83	-4.28	-7.83
qrLMM	50	-7.81	-4.11	-7.80
	300	-7.86	-4.22	-7.86