



Robust Estimation of the Generalized Loggamma Model: The R Package `robustloggamma`

Claudio Agostinelli
University of Trento

Alfio Marazzi
University of Lausanne

Víctor J. Yohai
University of Buenos Aires

Alex Randriamiharisoa
University of Lausanne

Abstract

`robustloggamma` is an R package for robust estimation and inference in the generalized loggamma model. We briefly introduce the model, the estimation procedures and the computational algorithms. Then, we illustrate the use of the package with the help of a real data set.

Keywords: generalized loggamma model, R, robust estimators, `robustloggamma`, τ estimator, weighted likelihood.

1. Introduction

The generalized loggamma distribution is a flexible three parameter family introduced by Stacy (1962) and further studied by Prentice (1974) and Lawless (1980). This family is used to model highly skewed positive data on a logarithmic scale and it includes several asymmetric families, such as logexponential, logWeibull, and loggamma distributions, and also the normal distribution. In the parametrization given by Prentice (1974) the three parameters are location μ , scale σ , and shape λ . We denote the family by $LG(\mu, \sigma, \lambda)$, $\mu \in \mathbb{R}$, $\sigma > 0$, $\lambda \in \mathbb{R}$. If y is a random variable with distribution $LG(\mu, \sigma, \lambda)$ then y is obtained by location and scale transformation

$$y = \mu + \sigma u$$

of the random variable u with density

$$f_{\lambda}(u) = \begin{cases} \frac{|\lambda|}{\Gamma(\lambda^{-2})} (\lambda^{-2})^{\lambda^{-2}} \exp\left(\lambda^{-2} (\lambda u - e^{\lambda u})\right) & \text{if } \lambda \neq 0, \\ \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{u^2}{2}\right) & \text{if } \lambda = 0, \end{cases}$$

where Γ denotes the gamma function. Hence, the density of y is $f_{\theta}(y) = f_{\lambda}((y - \mu)/\sigma) / \sigma$ where $\theta = (\mu, \sigma, \lambda)$. The normal model ($\lambda = 0$), logWeibull model ($\lambda = 1$), logexponential model ($\lambda = 1$ and $\sigma = 1$), and loggamma model ($\sigma = \lambda$) are special cases. The *generalized gamma* family is obtained by transforming back onto the original scale, i.e., $t = \exp(y)$; in this situation the expectation $\eta = E(t) = \delta \Gamma(\alpha + 1/\gamma) / \Gamma(\alpha)$ where $\alpha = \lambda^{-2}$, $\gamma = \lambda/\sigma$, $\delta = \exp(\mu + 2 \log(\lambda)\sigma/\lambda)$ is an important parameter.

The R (R Core Team 2016) package **robustloggamma** (Agostinelli, Marazzi, Yohai, and Radrilmiharisoa 2016) is available from the Comprehensive R Archive Network (CRAN) at <https://CRAN.R-project.org/package=robustloggamma>. It provides density, distribution function, quantile function and random generation for the loggamma distribution using the common syntax `[dpqr]loggamma`. In Figure 1 we draw the density of some relevant distributions using the following code.

```
R> library("robustloggamma")
R> plot(function(x) dloggamma(x, mu = 0, sigma = 1, lambda = 0),
+       from = -8, to = 4, ylab = "density")
R> plot(function(x) dloggamma(x, mu = 0, sigma = 2, lambda = 1),
+       from = -8, to = 4, add = TRUE, col = 2)
R> plot(function(x) dloggamma(x, mu = 0, sigma = 1, lambda = 1),
+       from = -8, to = 4, add = TRUE, col = 3)
R> plot(function(x) dloggamma(x, mu = 0, sigma = 2, lambda = 2),
+       from = -8, to = 4, add = TRUE, col = 4)
R> legend("topleft", legend = c("normal(0,1)", "logWeibull(0,2)",
+ "logexponential", "loggamma(0,2)"), col = 1:4, lty = 1, inset = 0.01)
```

2. Robust estimation and inference

We consider the three parameter family $LG(\mu, \sigma, \lambda)$ with $\theta = (\mu, \sigma, \lambda)$, and distribution function $F_{\theta}(y) = F^*((y - \mu)/\sigma, \lambda)$. For $0 < u < 1$ let $Q(u, \theta)$ be the u -quantile of $F_{\theta}(y)$. Then, $Q(u, \theta) = \sigma Q^*(u, \lambda) + \mu$, where $Q^*(u, \lambda) = Q(u, (0, 1, \lambda))$. Let $y_{(1)}, \dots, y_{(n)}$ be the order statistics of a sample of size n from $LG(\mu_0, \sigma_0, \lambda_0)$, and $\theta_0 = (\mu_0, \sigma_0, \lambda_0)$ is the unknown vector of parameters to be estimated. Since, $y_{(j)}$ is the quantile $u_{n,j} = (j - 0.5)/n$ of the empirical distribution, it should be close to the corresponding theoretical quantile $\sigma_0 Q^*(u_{n,j}, \lambda_0) + \mu_0$. Hence, the residuals

$$r_{n,j}(\theta) = y_{(j)} - \mu - \sigma Q^*(u_{n,j}, \lambda)$$

are a function of θ and they should be as small as possible. To summarize their size a *scale* s is often used. Given a sample $\mathbf{u} = (u_1, \dots, u_n)$, s is a function of \mathbf{u} with two basic properties: (i) $s(\mathbf{u}) \geq 0$; (ii) for any scalar γ , $s(\gamma \mathbf{u}) = |\gamma|s(\mathbf{u})$. The most common scale is

$$\left(\sum_{j=1}^n u_j^2 / n \right)^{1/2},$$




Figure 1: Density for some relevant members of the generalized loggamma family.

which is clearly non robust. To gain robustness we use a τ scale, $\tau(r_{n,1}(\boldsymbol{\theta}), \dots, r_{n,n}(\boldsymbol{\theta}))$, introduced by [Yohai and Zamar \(1988\)](#). A short review of τ scales can be found in [Appendix A](#). Then, the $Q\tau$ estimator is defined by

$$\tilde{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} \tau(r_{n,1}(\boldsymbol{\theta}), \dots, r_{n,n}(\boldsymbol{\theta})) .$$

We note that, fixing λ , the values of μ and σ minimizing the τ scale are obtained by a simple regression τ estimate for the responses $y_{(j)}$ and the regressors $Q^*(u_{n,j}, \lambda)$. We also note that $n^{1/2}r_{n,j}(\boldsymbol{\theta}_0)$ is approximately distributed according to $N(0, v^2(\boldsymbol{\theta}_0, u_{n,j}))$ ([Serfling 1980](#)), where

$$v^2(\boldsymbol{\theta}_0, u) = \frac{\sigma_0^2 u(1-u)}{f_{\lambda_0}^2(Q^*(u, \lambda_0))} .$$

Then, the variances of the regression errors can be estimated by

$$\tilde{\sigma}_j^2 = v^2(\tilde{\boldsymbol{\theta}}_n, u_{n,j})$$

and the basic estimator can be improved by means of a weighted procedure. More precisely, one defines the *weighted* $Q\tau$ estimator ($WQ\tau$), with the set of weights $1/\tilde{\sigma}_1, \dots, 1/\tilde{\sigma}_n$, by

$$\tilde{\boldsymbol{\theta}}^w = \arg \min_{\boldsymbol{\theta}} \tau \left(\frac{r_{n,1}(\boldsymbol{\theta})}{\tilde{\sigma}_1}, \dots, \frac{r_{n,n}(\boldsymbol{\theta})}{\tilde{\sigma}_n} \right) .$$

Monte Carlo simulations ([Agostinelli, Marazzi, and Yohai 2014](#)) show that both the $Q\tau$ and the $WQ\tau$ estimators perform well in the case where the model is correct and also when the sample contains outliers. These empirical findings are corroborated by a theoretical

result showing that $Q\tau$ and $WQ\tau$ have a 50% breakdown point (BDP; according to a special definition of BDP – the finite sample distribution breakdown point – which is particularly designed to assess the degree of global stability of a distribution estimate).

2.1. Weighted likelihood estimators

Unfortunately, $Q\tau$ and $WQ\tau$ are not asymptotically normal and therefore inconvenient for inference. Their rates of convergence is however of order $n^{1/2}$ and this makes them a good starting point for a one-step weighted likelihood (WL) procedure which is asymptotically normal and fully efficient at the model. The package **robustloggamma** implements two WL estimators: the fully iterated weighted likelihood and the one step weighted likelihood. Monte Carlo simulations (Agostinelli *et al.* 2014) show that both these estimators maintain the robust properties (BDP) of $Q\tau$ and $WQ\tau$.

In general, a *weighted likelihood estimator* (WLE) as defined in Markatou, Basu, and Lindsay (1998) is a solution of the following estimating equations

$$\frac{1}{n} \sum_{j=1}^n w(y_j, \boldsymbol{\theta}) \mathbf{z}(y_j, \boldsymbol{\theta}) = \mathbf{0} ,$$

where $\mathbf{z}(y, \boldsymbol{\theta})$ is the usual score function vector and $w(y, \boldsymbol{\theta})$ is a weight function defined by

$$w(y, \boldsymbol{\theta}) = \min \left(1, \frac{[A(\delta(y, \boldsymbol{\theta})) + 1]^+}{\delta(y, \boldsymbol{\theta}) + 1} \right) ,$$

where $\delta(y, \boldsymbol{\theta})$ is the *Pearson residual* (Lindsay 1994), measuring the agreement between the distribution of the data and the assumed model. It is defined as $\delta(y, \boldsymbol{\theta}) = [f^*(y) - f_{\boldsymbol{\theta}}^*(y)] / f_{\boldsymbol{\theta}}^*(y)$, where $f^*(y) = \int k(y, t, h) dF_n(t)$ is a kernel density estimate of $f_{\boldsymbol{\theta}}$ (with bandwidth h), $f_{\boldsymbol{\theta}}^*(y) = \int k(y, t, h) f_{\boldsymbol{\theta}}(t) dt$ is the corresponding smoothed model density, F_n is the empirical cumulative distribution function, and $[x]^+ = \max(0, x)$.

The function $A(\cdot)$ is called *residual adjustment function* (RAF). When $A(\delta) = \delta$ the weights $w(y_j, \boldsymbol{\theta}) = 1$ and the WLE equations coincide with the classical MLE equations. Generally, the weight function w uses the RAF that corresponds to minimum disparity problems (Lindsay 1994), see Appendix B for some examples.

The *fully iterated weighted likelihood estimator* (FIWL) is the solution of the weighted equations, while the *one-step weighted likelihood estimator* (1SWL) is defined by

$$\hat{\boldsymbol{\theta}} = \tilde{\boldsymbol{\theta}} - \mathbf{J}^{-1} \sum_{j=1}^n w(y_j, \tilde{\boldsymbol{\theta}}) \mathbf{z}(y_j, \tilde{\boldsymbol{\theta}}) ,$$

where $\mathbf{J} = \int w(y, \tilde{\boldsymbol{\theta}}) \nabla \mathbf{z}(y, \tilde{\boldsymbol{\theta}}) dF_{\tilde{\boldsymbol{\theta}}}(y)$ and ∇ denotes differentiation with respect to $\boldsymbol{\theta}$. This definition is similar to a Fisher scoring step, where an extra term obtained by differentiating the weight with respect to $\boldsymbol{\theta}$ is dropped since, when evaluated at the model, it is equal to zero. Further information on minimum distance methods and weighted likelihood procedures are available in Basu, Shioya, and Park (2011).

3. Algorithms and implementation

In the following sections, we describe the computation of the estimators implemented in the main function `loggammarob`. We first recall its arguments, a reference chart is reported in

Appendix C. The only required argument is \mathbf{x} which contains the data in a numeric vector. The argument `method` allows to choose among the available robust procedures. The default method is "oneWL", a one step weighted likelihood estimator starting from $WQ\tau$. Other alternatives are "QTau" ($Q\tau$), "WQTau" ($WQ\tau$), "WL" (fully iterated weighted likelihood) and "ML" (maximum likelihood). When `method` is not $Q\tau$ an optional numeric vector of length 3 (location, scale, shape) could be supplied in the argument `start` to be used as starting value, otherwise, the default is $WQ\tau$ for the likelihood based methods and $Q\tau$ for $WQ\tau$. By default, weights in the $WQ\tau$ are specified as described in the previous section, if a different set of weights are needed the `weights` argument could be used. Fine tuning parameters are set by the function `loggamma.rob.control` and passed to the main function by the `control` argument.

3.1. Computation of $Q\tau$ and $WQ\tau$

To optimize the τ scale for a given value of λ , `robustloggamma` uses the resampling algorithm described in [Salibian-Barrera, Willems, and Zamar \(2008\)](#). Let $x_j = Q^*(u_{n,j}, \lambda)$ and consider the following steps:

1. Take a random subsample of size 2 made of the pairs $(x_{(j_1)}, y_{(j_1)})$ and $(x_{(j_2)}, y_{(j_2)})$.
2. Compute a preliminary estimate of μ and σ of the form

$$\sigma^{(0)} = \frac{y_{(j_1)} - y_{(j_2)}}{x_{(j_1)} - x_{(j_2)}}, \quad \mu^{(0)} = y_{(j_1)} - \sigma_i^{(0)} x_{(j_1)}.$$

3. Compute the residuals $r_j^{(0)} = y_{(j)} - \sigma^{(0)} x_j - \mu^{(0)}$ for $j = 1, \dots, n$.
4. Compute least squares estimates $\mu^{(1)}, \sigma^{(1)}$ based on the $n/2$ pairs with the smallest absolute residuals $r_j^{(0)}$.
5. Compute the residuals $r_j^{(1)}$ for $j = 1, \dots, n$ and the τ scale $\tau(r_1^{(1)}, \dots, r_n^{(1)})$.

Steps 1–5 are repeated a large number N of times and the values $\mu^{(1)}, \sigma^{(1)}$ corresponding to the minimal τ scale are retained. These values are then used as starting values of an IRWLS algorithm, where the weights are defined, at each iteration, as $w_j = W\phi_{j1} + \phi_{j2}$ and

$$W = \frac{\sum_{j=1}^n 2\rho_2(r_j/s) - \psi_2(r_j/s)r_j/s}{\sum_{j=1}^n \psi_1(r_j/s)r_j/s}, \quad \phi_{jk} = \psi_k(r_j/s)/(r_j/s), \quad k = 1, 2.$$

ψ_k is the first derivative of the ρ function ρ_k (see [Appendix A](#)), and s is a scale which is recursively updated as follows

$$s = s \left(\frac{2}{n} \sum_{j=1}^n \rho_1(r_j/s) \right)^{1/2}$$

with initial value $s = \text{median}(|r_j|)/0.6745$.

This algorithm is used to compute $\mu(\lambda)$ and $\sigma(\lambda)$ for all values of λ in a given grid $\lambda_1, \dots, \lambda_k$. The final value of λ is then obtained by minimizing the τ scale over the grid.

The first part of the algorithm is implemented in Fortran while the IRWLS algorithm is implemented in R. The initial random procedure uses the R uniform pseudo random number generator and can be controlled by setting the seed in the usual way. The function `loggammarob.control` is used to set all the other parameters. `tuning.rho` and `tuning.psi` set the constants c_1 and c_2 of the ρ functions ρ_1 and ρ_2 . `nResample` controls the number of subsamples N . `max.it` and `refine.tol` provide the maximum number of iterations and the tolerance of the IRWLS algorithm. An equally spaced grid for λ is defined by the arguments `lower`, `upper` and `n` with obvious meaning. While there are no boundaries on the values of `lower` and `upper` we advice the users to avoid values outside the interval $[-8, 8]$. Values outside this range might lead to numerical problems and to errors in the internal functions; however, we notice that in most applications this is not a restriction. Default values, for the tuning parameters can be inspected using

```
R> loggammarob.control()
```

Once a Q_τ estimate $\tilde{\theta}$ is obtained, the WQ_τ is easily computed by first evaluating a fixed set of scales

$$\tilde{\sigma}_j^2 = v^2(\tilde{\theta}, u_{n,j}),$$

then the IRWLS is used with $\tilde{\theta}$ as starting value and $r_j/\tilde{\sigma}_j$ in place of r_j .

3.2. Computation of FIWL and 1SWL

Weights need to be evaluated for both 1SWL and FIWL. To compute the kernel density estimate $f^*(y)$, **robustloggamma** uses the function `density` with `kernel = "gaussian"`, `cut = 3`, and `n = 512`. The smoothed model $f_{\tilde{\theta}}^*(y)$ is approximated by

$$\frac{1}{K} \sum_{k=1}^K k(y, y_k, h),$$

where y_k is the quantile of order $(k - 0.5)/K$ of $F_{\tilde{\theta}}$. The bandwidth h is adaptively fixed to `bw` times the actual value of σ and K is controlled by the argument `subdivisions`. The RAF is fixed by `raf` among several choices: "NED" (negative exponential disparity), "GKL" (generalized Kullback-Leibler), "PWD" (power divergence measure), "HD" (Hellinger distance), "SCHI2" (symmetric chi-square distance). `tau` selects the particular member of the family in case of "GKL" and "PWD". Finally, weights smaller than `minw` are set to zero.

For 1SWL, $\mathbf{J} = \int w(y, \tilde{\theta}) \nabla_{\mathbf{z}}(y, \tilde{\theta}) dF_{\tilde{\theta}}(y)$ is approximated by

$$\frac{1}{K} \sum_{k=1}^K w(y_k, \tilde{\theta}) \nabla_{\mathbf{z}}(y_k, \tilde{\theta}) .$$

Here K is controlled by `nexp`. Furthermore, the step can be multiplied by the `step` argument (with default 1).

To provide an overview of the efficiency of the algorithms we performed four Monte Carlo experiments. We consider datasets generated from the generalized loggamma model with parameters $\mu = 0$, $\sigma = 1$ and $\lambda = 1$ or $\lambda = 5$ and different sample sizes $n = 50, 100, 500, 1000, 5000, 10000$. For each of these we consider the case of no contamination and the one 10% contamination, where the 10% largest observations are replaced by observations generated by a

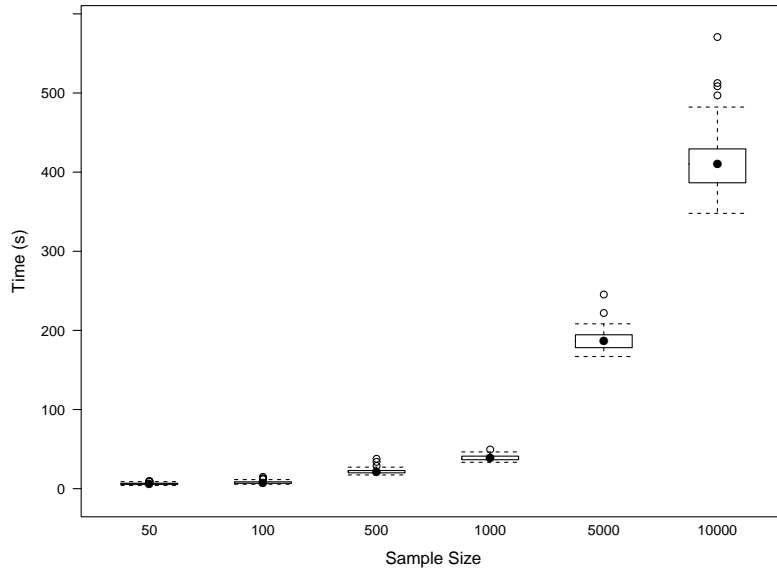


Figure 2: Execution time (seconds) for different sample sizes.

normal distribution with mean 20 and standard deviation 1. We recorded the execution times of the function `loggammarob`, with default parameters, in 100 replications for each sample size, hence $Q\tau$, $WQ\tau$ and $1SWL$ estimators are evaluated at each iteration. The results for the case $\lambda = 1$ without contamination are reported in Figure 2 and they suggest an almost linear relation between execution time and sample size, at least in the range explored. For the other three experiments the results were very similar and hence they are not reported. All the experiments were performed on SCSCF, a multiprocessor cluster system owned by the Ca' Foscari University of Venezia with cores at 2.3Ghz running under GNU/Linux.

4. An illustration

We illustrate the use of `robustloggamma` with the help of the data set `drg2000` included in the package. The data refer to 70323 stays that were observed in year 2000 in a group of Swiss hospitals within a pilot study aimed at the implementation of a diagnosis-related grouping (DRG) system. DRG systems are used in modern hospital management to classify each individual stay into a group according to the patient characteristics. The classification rules are defined so that the groups are as homogeneous as possible with respect to clinical criteria (diagnoses and procedures) and to resource consumption. A mean cost of each group is usually estimated yearly with the help of available data about the observed stays on a national basis. This cost is then assigned to each stay in the same group and used for reimbursement and budgeting.

Cost distributions are typically skewed and often contain outliers. When a small number of outliers are observed, the classical estimates of the mean can be much different than when none is observed. And since the values and the frequency of outliers fluctuate from year to year, the classical mean cost is an unreliable estimate. Not surprisingly, since many DRGs must routinely be inspected each year, automatic outlier detection is a recurrent hot topic for discussion among hospital managers.

AP-DRG	Description
185	Dental & oral dis exc extract & restorations, Age > 17
222	Knee procedures w/o cc
237	Sprain, strain, disloc of hip, pelvis, thigh
360	Vagina, Cervix & Vulva Procedures

Table 1: DRG code and description for the available categories in data set `drg2000`.

The data set has four variables: `LOS` length of stay, `Cost` cost of stay in Swiss francs, `APDRG` DRG code (according to the “All Patients DRG” system) and `MDC` major diagnostic category. Packages `xtable` (Dahl 2016) and `lattice` (Sarkar 2008) will be used during the illustration.

```
R> library("xtable")
R> library("lattice")
R> data("drg2000", package = "robustloggamma")
```

We will analyze the variable `Cost` on the logarithmic scale for the four DRGs as described in Table 1.

```
R> APDRG <- c(185, 222, 237, 360)
R> index <- unlist(sapply(APDRG, function(x) which(drg2000$APDRG == x)))
R> DRG <- drg2000[index, ]
```

Figure 3, obtained with the following code, shows a density plot for each selected DRG:

```
R> print(densityplot(~I(log(Cost)) | factor(APDRG), data = DRG,
+   plot.points = "rug", ref = TRUE))
```

Summary statistics are:

```
R> lapply(split(DRG$Cost, DRG$APDRG), summary)
```

```
$`185`
  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
 1228   2645   3462   5059   5047   55770
```

```
$`222`
  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
849.2 4362.0 5288.0 6307.0 6801.0 47240.0
```

```
$`237`
  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
 2038   3144   4100   4987   5169   28780
```

```
$`360`
  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
 1620   2863   3502   4680   4428   51160
```

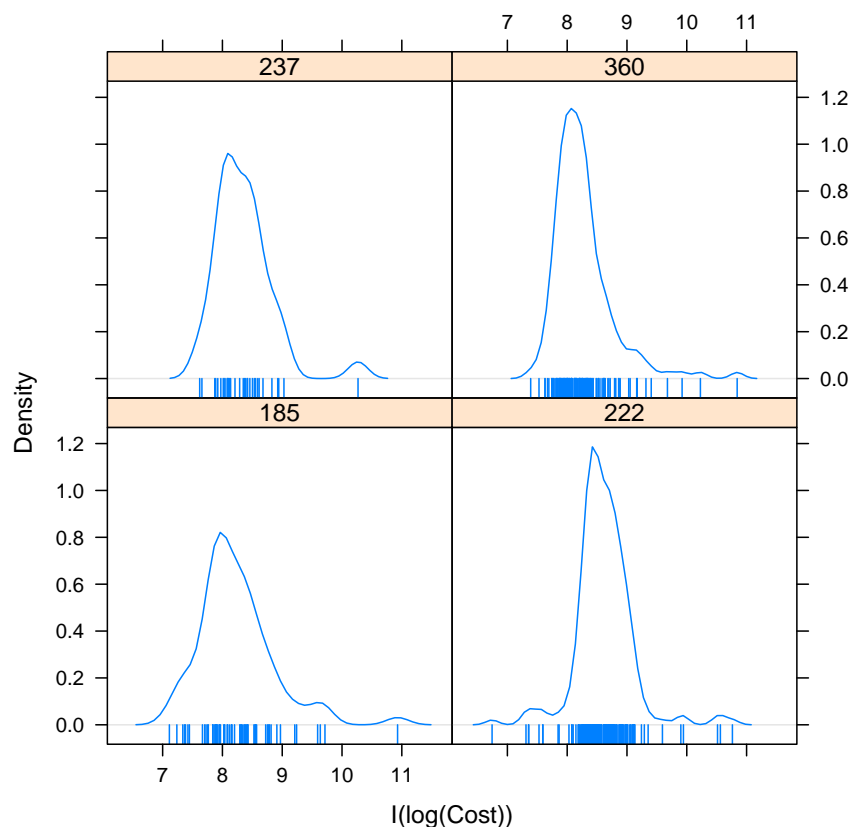



Figure 3: Estimated densities of $\log(\text{Cost})$ for selected DRGs.

A comparison between classical and robust measures of spread, indicates important differences:

```
R> lapply(split(DRG$Cost, DRG$APDRG), function(x) c(sd(x), mad(x)))
```

```
$`185`  
[1] 6894.890 1374.874
```

```
$`222`  
[1] 5095.599 1648.444
```

```
$`237`  
[1] 4490.546 1494.891
```

```
$`360`  
[1] 5066.084 1075.537
```

The differences are due to the presence of outliers. Therefore, it is convenient to analyze the data with the help of robust methods. To begin with, we use the function `loggammarob` to fit a generalized loggamma model to sample `APDRG = 185`. This function provides robust

estimates of the parameters μ (location), σ (scale), and λ (shape) using the default method 1SWL:

```
R> Cost185 <- sort(DRG$Cost[DRG$APDRG == 185])
R> est185 <- loggammарob(log(Cost185))
R> est185
```

```
Call:
loggammарob(x = log(Cost185))
```

```
Location: 8.04 Scale: 0.4944 Shape: -0.6437 E(exp(X)): 4381
```

In addition, a `summary` method is available to calculate confidence intervals (based on the Wald statistics) for the parameters and for selected model quantiles (argument `p`).

```
R> summary(est185, p = c(0.9, 0.95, 0.99))
```

```
Call:
summary.loggammарob(object = est185, p = c(0.9, 0.95, 0.99))
```

```
Location: 8.04 s.e. 0.09841
( 7.847 , 8.233 )
95 percent confidence interval
```

```
Scale: 0.4944 s.e. 0.05071
( 0.395 , 0.5938 )
95 percent confidence interval
```

```
Shape: -0.6437 s.e. 0.3005
( -1.233 , -0.05467 )
95 percent confidence interval
```

```
Mean(exp(X)): 4381 s.e. 426.7
( 3545 , 5218 )
95 percent confidence interval
```

```
Quantile of order 0.9 : 8.932 s.e. 0.2337
( 8.474 , 9.39 )
95 percent confidence interval
```

```
Quantile of order 0.95 : 9.2 s.e. 0.343
( 8.528 , 9.873 )
95 percent confidence interval
```

```
Quantile of order 0.99 : 9.774 s.e. 0.6505
( 8.499 , 11.05 )
95 percent confidence interval
```

Robustness weights:

```
54 weights are ~= 1. The remaining 15 ones are summarized as
  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
0.05591 0.74150 0.88590 0.77530 0.99800 0.99890
```

The function also provides the robust weights that allow outlier identification. For instance:

```
R> which(est185$weights < 0.1)
```

```
[1] 1
```

reports the indices of the observations with weights smaller than 0.1, which in this case is the first one only.

Robust tests on one or more parameters can be performed by means of the weighted Wald test described in [Agostinelli and Markatou \(2001\)](#). For this purpose, we use the function `loggammarob.test`. For instance, we test the hypothesis that the shape parameter is equal to zero, i.e., that the lognormal model is an acceptable one:

```
R> loggammarob.test(est185, lambda = 0)
```

```
Weighted Wald Test based on oneWL
```

```
data:
```

```
ww = 4.5876, df = 1, p-value = 0.0322
alternative hypothesis: true shape is not equal to 0
95 percent confidence interval:
 -1.23270104 -0.05466989
sample estimates:
[1] -0.6436855
```

To test the hypothesis that the location is zero and the scale is one we use:

```
R> loggammarob.test(est185, mu = 0, sigma = 1)
```

However, in these situations, the confidence intervals are not calculated.

The default estimation method in `loggammarob` is 1SWL (one-step weighted likelihood). However, alternative estimates are made available: $Q\tau$, $WQ\tau$, WL, and ML. $Q\tau$ and $WQ\tau$, typically used as starting values for the weighted likelihood procedures, are obtained as follows:

```
R> qtau185 <- summary(loggammarob(log(Cost185), method = "QTau"))
R> wqtau185 <- summary(loggammarob(log(Cost185), method = "WQTau"))
```

The fully iterated weighted likelihood (FIWL) and the one-step weighted likelihood estimates (1SWL) are obtained as follows:

```
R> fiwl185 <- summary(loggammarob(log(Cost185), method = "WL"))
R> oswl185 <- summary(loggammarob(log(Cost185), method = "oneWL"))
```

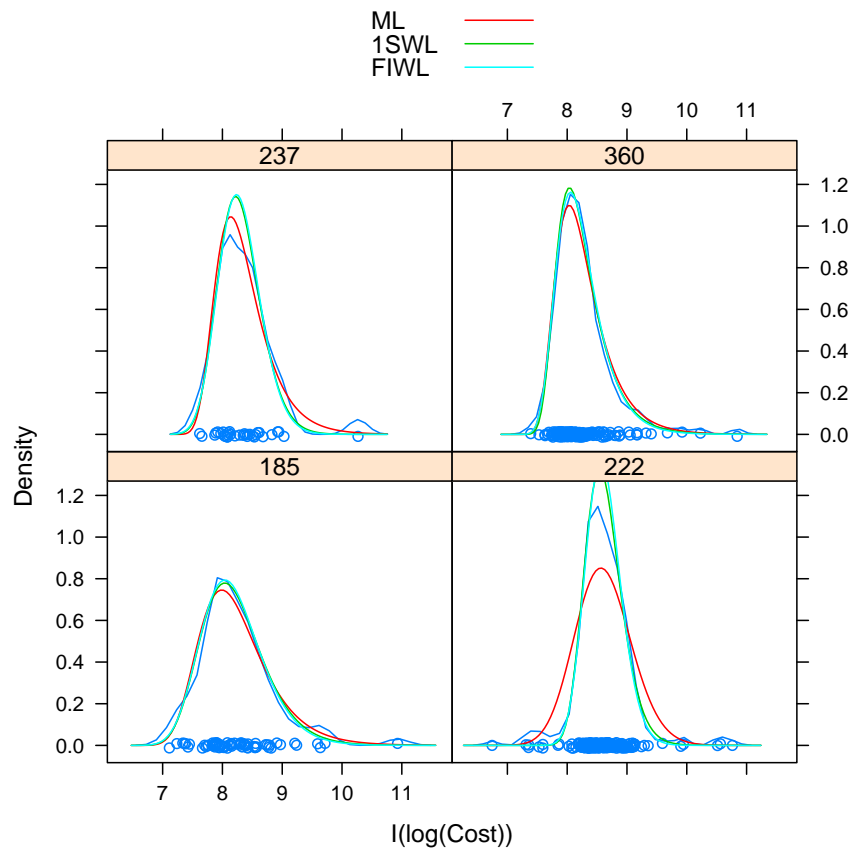


Figure 4: Estimated models provided by ML, 1SWL and FIWL for the four data sets.

The maximum likelihood estimate is also available:

```
R> ml <- summary(loggammarnob(log(Cost185), method = "ML"))
```

We now compare the four samples. For this purpose, the function `analysis` available in the supplementary material must be loaded before running the next command.

```
R> results <- sapply(APDRG, function(x) analysis(APDRG = x, data = DRG),
+   simplify = FALSE)
```

Three estimated densities provided by ML, FIWL, and 1SWL are shown in Figure 4. The robust parameter estimates, their estimated standard errors, and their confidence intervals are shown in Table 2. The table was obtained with the help of the the function `maketable` available in the supplementary material. We note that outliers are present in all samples (Figure 5). However, for $APDRG = 185$ and $APDRG = 360$ they do not have a significant impact on the estimates and the associated inferences. On the contrary, for $APDRG = 222$ the outliers markedly inflate the ML scale estimate and, for $APDRG = 237$ a single outlier has a great impact on the ML estimate of λ . For this sample, the robust parameter estimates and their confidence intervals suggest that the lognormal density is a possible model. To visualize the weights, we use the following code:

DRG	ML				ISWL				FIWL				
	μ	σ	λ	η	μ	σ	λ	η	μ	σ	λ	η	
185	est	7.989	0.501	-0.892	4837	8.040	0.494	-0.644	4381	8.043	0.489	-0.574	4232
	se	0.101	0.056	0.305	657	0.098	0.051	0.301	427	0.095	0.048	0.295	372
	left	7.792	0.392	-1.490	3550	7.847	0.395	-1.233	3545	7.856	0.394	-1.153	3504
	right	8.186	0.610	-0.295	6124	8.233	0.594	-0.055	5218	8.230	0.583	0.005	4961
222	est	8.563	0.467	-0.197	6152	8.544	0.295	-0.493	5836	8.566	0.289	-0.308	5747
	se	0.054	0.025	0.179	237	0.035	0.017	0.183	156	0.035	0.016	0.184	138
	left	8.457	0.419	-0.549	5687	8.475	0.261	-0.851	5531	8.498	0.257	-0.669	5475
	right	8.670	0.515	0.155	6618	8.614	0.329	-0.135	6141	8.634	0.321	0.054	6018
237	est	8.139	0.349	-1.050	4836	8.221	0.344	-0.432	4300	8.236	0.343	-0.329	4266
	se	0.103	0.059	0.455	596	0.095	0.046	0.423	308	0.095	0.045	0.425	290
	left	7.936	0.234	-1.942	3667	8.034	0.254	-1.261	3697	8.050	0.256	-1.162	3697
	right	8.341	0.465	-0.158	6005	8.408	0.434	0.396	4902	8.422	0.431	0.504	4835
360	est	8.034	0.324	-1.191	4444	8.038	0.304	-1.137	4235	8.055	0.318	-0.974	4142
	se	0.049	0.029	0.237	280	0.046	0.026	0.234	228	0.046	0.026	0.221	201
	left	7.938	0.268	-1.655	3896	7.949	0.252	-1.595	3789	7.965	0.267	-1.407	3748
	right	8.129	0.380	-0.727	4993	8.127	0.356	-0.679	4681	8.145	0.368	-0.542	4535

Table 2: Estimates, standard errors and confidence intervals for the four samples.

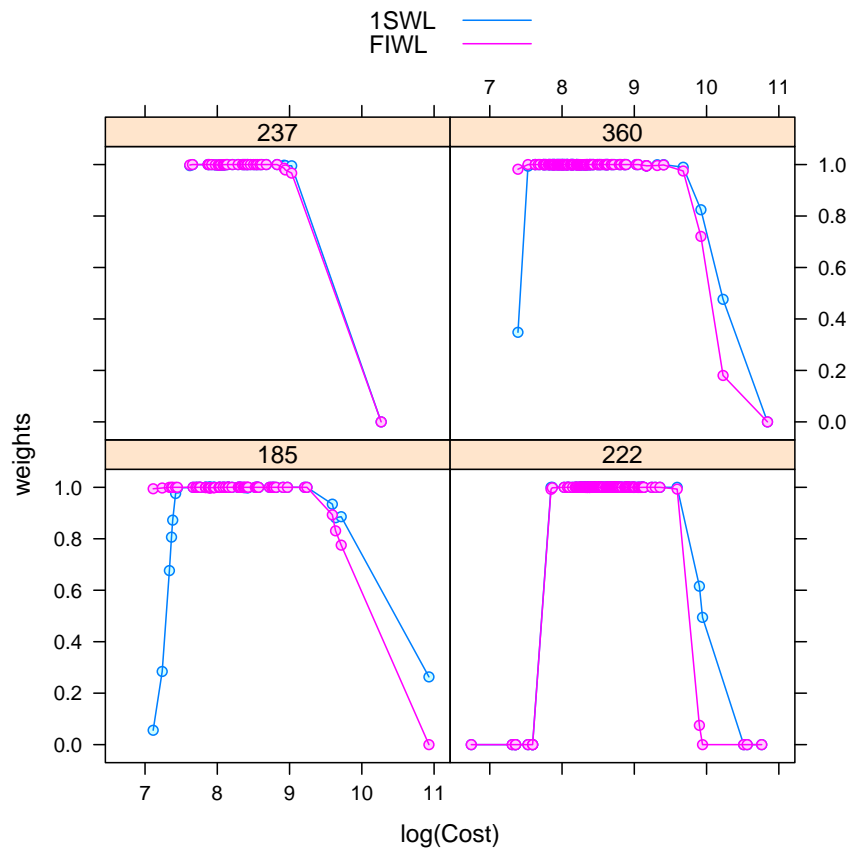


Figure 5: Weights provided by 1SWL and FIWL.

```
R> weights <- function(x, results) {
+   os <- results[[x]]$os
+   wl <- results[[x]]$wl
+   ans <- t(cbind(os$weights, wl$weights, wl$data, x))
+   return(ans)
+ }
R> w <- as.data.frame(matrix(unlist(sapply(1:4,
+   function(x) weights(x, results = results))), ncol = 4, byrow = TRUE))
R> colnames(w) <- c("OSWL", "FIWL", "data", "drg")
R> w$drg <- factor(w$drg, labels = APDRG)
R> lattice.theme <- trellis.par.get()
R> col <- lattice.theme$superpose.symbol$col[1:2]
R> print(xyplot(OSWL+FIWL~data | drg, data = w, type = "b",
+   col = col, pch = 21, key = list(text = list(c("1SWL", "FIWL")),
+   lines = list(col = col)), xlab = "log(Cost)", ylab = "weights"))
```

Is the two-parameter gamma distribution an acceptable model for the four samples? To answer this question, we test the hypothesis $\sigma = \lambda$. We use the function `loggammarob.wilks` with weights provided by ML, 1SWL and, FIWL. The results are summarized in Table 3 provided by the function `extractwilks` included in the supplementary material. The hypothesis is

DRG		ML	1SWL	FIWL
185	statistic	45.882	29.1739	15.3918
	p value	0	0	0.0001
	μ_0	8.5288	8.4277	8.3497
	σ_0	0.7281	0.5949	0.5502
222	statistic	52.3126	24.7611	10.8147
	p value	0	0	0.0010
	μ_0	8.7494	8.6712	8.6563
	σ_0	0.5172	0.3208	0.2984
237	statistic	23.1627	1.7348	1.7679
	p value	0	0.1878	0.1836
	μ_0	8.5146	8.3583	8.3567
	σ_0	0.5548	0.3527	0.3518
360	statistic	123.933	75.6255	62.5466
	p value	0	0	0
	μ_0	8.4512	8.3543	8.3363
	σ_0	0.591	0.4654	0.4481

Table 3: Weighted Wilks test for the hypothesis $\sigma = \lambda$.

always strongly rejected for DRGs 185, 222, 360. For DRG 237, the presence of outliers leads to strongly reject the hypothesis based on ML estimation, while the robust methods accept it (μ_0 and $\sigma_0 = \lambda_0$ are the estimated parameters under the null hypothesis).

```
R> wilks <- extractwilks(results)
R> wilks <- cbind(c("185", rep(" ", 3), "222", rep(" ", 3), "237",
+ rep(" ", 3), "360", rep(" ", 3)), rep(c("statistic", "$p$ value",
+ "$\mu_0$", "$\sigma_0$"),4), wilks)
R> xwilks <- xtable(wilks)
```

Finally, we draw Q-Q plots based on ML, 1SWL and FIWL (Figure 6) for the four data sets. Darker points are associated with smaller weights. 90% confidence bands are provided to check the adequacy of the model to the data.

```
R> quant <- function(x, method, results) {
+   res <- results[[x]][[method]]
+   n <- length(res$data)
+   q <- qloggamma(p = ppoints(n), mu = res$mu, sigma = res$sigma,
+     lambda = res$lambda)
+   qconf <- summary(res, p = ppoints(n), conf.level = 0.90)$qconf.int
+   ans <- t(cbind(q, qconf, res$data, res$weights, x, method))
+   return(ans)
+ }
R> q1 <- matrix(unlist(sapply(1:4, function(x)
+   quant(x, method = 1, results = results))), ncol = 7, byrow = TRUE)
R> q2 <- matrix(unlist(sapply(1:4, function(x)
+   quant(x, method = 2, results = results))), ncol = 7, byrow = TRUE)
```

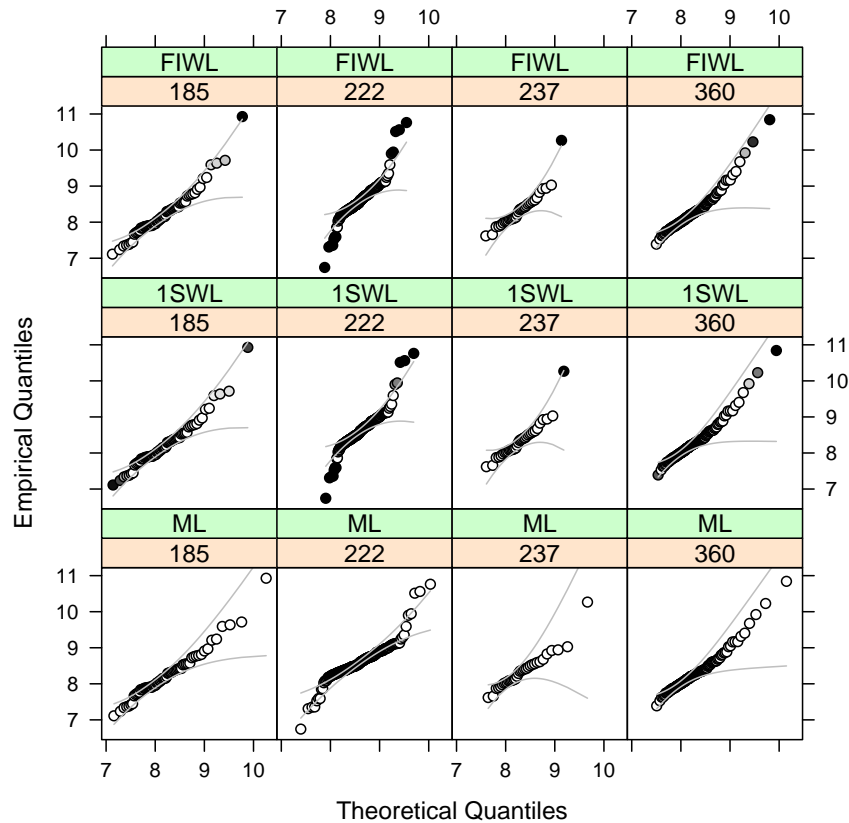


Figure 6: Q-Q plots and 90% confidence bands (grey line). Darker points are associated with smaller weights.

```
R> q3 <- matrix(unlist(sapply(1:4, function(x)
+   quant(x, method = 3, results = results))), ncol = 7, byrow = TRUE)
R> q <- as.data.frame(rbind(q1, q2, q3))
R> colnames(q) <- c("q", "qlower", "qupper", "Cost", "weights", "drg",
+   "method")
R> q$drg <- factor(q$drg, labels = APDRG)
R> q$method <- factor(q$method, labels = c("ML", "1SWL", "FIWL"))
R> print(xyplot(Cost ~ q | drg + method, data = q,
+   xlab = "Theoretical Quantiles", ylab = "Empirical Quantiles",
+   fill.color = grey(q$weights), q = q,
+   panel = function(x, y, fill.color, ..., subscripts, q) {
+     fill <- fill.color[subscripts]
+     q <- q[subscripts, ]
+     panel.xyplot(x, y, pch = 21, fill = fill, col = "black", ...)
+     panel.xyplot(x, y = q$qupper, type = "l", col = "grey75")
+     panel.xyplot(x, y = q$qlower, type = "l", col = "grey75")
+   })
```


Acknowledgments

All statistical analysis were performed on SCSCF (<http://www.dais.unive.it/scscf>), a multiprocessor cluster system owned by Ca' Foscari University of Venice running under GNU/Linux.

This research was partially supported by the Italian-Argentinian project “Metodi robusti per la previsione del costo e della durata della degenza ospedaliera” funded in the collaboration program MINCYT-MAE AR14MO6.

Víctor Yohai research was partially supported by Grants 20020130100279 from Universidad of Buenos Aires, PIP 112-2008-01-00216 and 112-2011-01-00339 from CONICET and PICT 2011-0397 from ANPCYT.

References

- Agostinelli C, Marazzi A, Yohai VJ (2014). “Robust Estimators of the Generalized Log-Gamma Distribution.” *Technometrics*, **56**(1), 92–101. doi:10.1080/00401706.2013.818578.
- Agostinelli C, Marazzi A, Yohai VJ, Randriamiharisoa A (2016). *robustloggamma: Robust Estimation of the Generalized Loggamma Model*. R package version 1.0-1, URL <https://CRAN.R-project.org/package=robustloggamma>.
- Agostinelli C, Markatou M (2001). “Test of Hypotheses Based on the Weighted Likelihood Methodology.” *Statistica Sinica*, **11**(2), 499–514.
- Basu A, Shioya H, Park C (2011). *Statistical Inference: The Minimum Distance Approach*. Chapman & Hall/CRC, Boca Raton.
- Cressie N, Read TRC (1984). “Multinomial Goodness-of-Fit Tests.” *Journal of the Royal Statistical Society B*, **46**, 440–464.
- Cressie N, Read TRC (1988). “Cressie-Read Statistic.” In S Kotz, N Johnson (eds.), *Encyclopedia of Statistical Sciences, Supplementary Volume*, pp. 37–39. John Wiley & Sons.
- Dahl DB (2016). *xtable: Export Tables to L^AT_EX or HTML*. R package version 1.8-2, URL <https://CRAN.R-project.org/package=xtable>.
- Huber PJ (1981). *Robust Statistics*. John Wiley & Sons. doi:10.1002/0471725250.
- Lawless JF (1980). “Inference in the Generalized Gamma and Loggamma Distributions.” *Technometrics*, **22**(3), 409–419. doi:10.1080/00401706.1980.10486173.
- Lindsay BG (1994). “Efficiency versus Robustness: The Case for Minimum Hellinger Distance and Related Methods.” *The Annals of Statistics*, **22**, 1018–1114. doi:10.1214/aos/1176325512.
- Markatou M, Basu A, Lindsay BG (1998). “Weighted Likelihood Equations with Bootstrap Root Search.” *Journal of the American Statistical Association*, **93**, 740–750. doi:10.1080/01621459.1998.10473726.

- Maronna RA, Martin DR, Yohai VJ (2006). *Robust Statistics: Theory and Methods*. John Wiley & Sons. doi:10.1002/0470010940.
- Park C, Basu A (2003). “The Generalized Kullback-Leibler Divergence and Robust Inference.” *Journal of Statistical Computation and Simulation*, **73**(5), 311–332. doi:10.1080/0094965021000033477.
- Prentice RL (1974). “A Loggamma Model and Its Maximum Likelihood Estimation.” *Biometrika*, **61**(3), 539–544. doi:10.1093/biomet/61.3.539.
- R Core Team (2016). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- Salibian-Barrera M, Willems G, Zamar RH (2008). “The Fast- τ Estimator for Regression.” *Journal of Computational and Graphical Statistics*, **17**, 659–682. doi:10.1198/106186008x343785.
- Sarkar D (2008). *lattice: Multivariate Data Visualization with R*. Springer-Verlag, New York.
- Serfling RJ (1980). *Approximation Theorems of Mathematical Statistics*. John Wiley & Sons. doi:10.1002/9780470316481.
- Stacy EW (1962). “A Generalization of the Gamma Distribution.” *The Annals of Mathematical Statistics*, **33**, 1187–1192. doi:10.1214/aoms/1177704481.
- Yohai VJ, Zamar RH (1988). “High Breakdown Estimates of Regression by Means of the Minimization of an Efficient Scale.” *Journal of the American Statistical Association*, **83**, 406–413. doi:10.2307/2288856.

A. τ scale regression

In this section we briefly review the definition of τ scale and τ regression. For a detailed description see [Maronna, Martin, and Yohai \(2006\)](#).

Let ρ be a function satisfying the following properties (**A**): (i) $\rho(0) = 0$; (ii) ρ is even; (iii) if $|x_1| < |x_2|$, then $\rho(x_1) \leq \rho(x_2)$; (iv) ρ is bounded; (v) ρ is continuous.

Then, an *M scale* ([Huber 1981](#)) based on ρ is defined by the value s satisfying

$$\frac{1}{n} \sum_{j=1}^n \rho\left(\frac{u_j}{s}\right) = b ,$$

where b is a given scalar and $0 < b < a = \sup \rho$. [Yohai and Zamar \(1988\)](#) introduce the family of τ scales. A τ scale is based on two functions ρ_1 and ρ_2 satisfying properties **A** and such that $\rho_2 \leq \rho_1$. To define a τ scale, one considers an *M scale* $s_1^2(\mathbf{u})$ based on ρ_1 . Then, the τ scale is given by

$$\tau^2(\mathbf{u}) = s_1^2(\mathbf{u}) \frac{1}{n} \sum_{j=1}^n \rho_2\left(\frac{u_j}{s_1(\mathbf{u})}\right) .$$

τ scale estimators can be extended easily to the linear regression case. Let us consider the regression model

$$y_j = \boldsymbol{\beta}^\top \mathbf{x}_j + e_j , \quad 1 \leq j \leq n ,$$

where $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)^\top$ and $\mathbf{x}_j = (x_{j1}, \dots, x_{jp})^\top$. For a given $\boldsymbol{\beta}$, let $r_j(\boldsymbol{\beta}) = y_j - \boldsymbol{\beta}^\top \mathbf{x}_j$ be the corresponding residuals. The scale $\tau^2(r_1(\boldsymbol{\beta}), \dots, r_n(\boldsymbol{\beta}))$ may be considered as a measure of goodness-of-fit. Based on this remark, [Yohai and Zamar \(1988\)](#) define robust estimators of the coefficients of a regression model by

$$\hat{\boldsymbol{\beta}} = \arg \min_{\boldsymbol{\beta}} \tau(r_1(\boldsymbol{\beta}), \dots, r_n(\boldsymbol{\beta})) .$$

These estimators are called *τ regression estimators*. If $a/b = 0.5$, the τ estimators have breakdown point (BDP) close to 50% ([Yohai and Zamar 1988](#)). Moreover, we note that, if $\rho_2(u) = u^2$, $\tau^2(u_1, \dots, u_n) = \text{ave}(u_j^2)$ and the regression τ estimator coincides with the least squares estimator. Therefore, taking as ρ_2 a bounded function close to the quadratic function, the regression τ estimators can be made arbitrarily efficient for normal errors. If the errors e_j are heteroscedastic with variances proportional to σ_j^2 , the efficiency of $\hat{\boldsymbol{\beta}}$ can be improved by means of a weighted procedure. A *regression weighted τ estimator* is given by

$$\hat{\boldsymbol{\beta}} = \arg \min_{\boldsymbol{\beta}} \tau(r_1^*(\boldsymbol{\beta}), \dots, r_n^*(\boldsymbol{\beta})) ,$$

where $r_j^*(\boldsymbol{\beta}) = r_j(\boldsymbol{\beta})/\sigma_j$.

Usually, one chooses ρ_1 and ρ_2 in the *Tukey biweight family*

$$\rho(u, c) = \begin{cases} 3(u/c)^2 - 3(u/c)^4 + (u/c)^6 & \text{if } |u| \leq c , \\ 1 & \text{if } |u| > c , \end{cases}$$

using two values c_1 and c_2 of the tuning parameter c . For example, one can take $c_1 = 1.548$ and $c_2 = 6.08$. With $b = 0.5$, these values yield regression estimators with breakdown point 0.5 and normal efficiency of 95%.

B. Residual adjustment functions

The literature provides several proposals for selecting the RAF. In the following, we recall two of them. The RAF based on the *power divergence measure* (PWD; Cressie and Read 1984, 1988), is given by

$$A_{pdm}(\delta, \tau) = \begin{cases} \tau \left((\delta + 1)^{1/\tau} - 1 \right) & \tau < \infty, \\ \log(\delta + 1) & \tau \rightarrow \infty. \end{cases}$$

Special cases are *maximum likelihood* ($\tau = 1$), *Hellinger distance* ($\tau = 2$), *Kullback-Leibler divergence* ($\tau \rightarrow \infty$), and *Neyman's chi-square* ($\tau = -1$). The RAF based on the *generalized Kullback-Leibler divergence* (GKL; Park and Basu 2003) is given by:

$$A_{gkl}(\delta, \tau) = \frac{\log(\tau\delta + 1)}{\tau}, \quad 0 \leq \tau \leq 1.$$

Special cases are maximum likelihood ($\tau \rightarrow 0$) and Kullback-Leibler divergence ($\tau = 1$). This RAF can be interpreted as a linear combination between the likelihood divergence and the Kullback-Leibler divergence. A further example is the RAF corresponding to the *negative exponential disparity* (NED; Lindsay 1994)

$$A(\delta) = 2 - (2 + \delta) \exp(-\delta),$$

which, for discrete models, is second order efficient.

C. Reference chart

Hereafter we provide the reference chart for the main function `loggamma.rob`. The usage has the following form

```
R> loggamma.rob(x, start = NULL, weights = rep(1, length(x)),
+   method = c("oneWL", "WQTau", "WL", "QTau", "ML"), control, ...)
```

where

`x` is a numeric vector, which contains the data.

`start` is `NULL` or a numeric vector containing the starting values of location, scale, and shape to be used when method is "WL", "oneWL" and "ML". Method "QTau" does not require starting values. When start is `NULL`, the methods "QTau" and "WQTau" are called in a series to compute the starting values.

`weights` is a numeric vector containing the weights for method "QTau".

`method` is a character string to select the method. The default is "oneWL" (one step weighted likelihood estimate starting from "WQTau"). Others available methods are "WL" (fully iterated weighted likelihood estimate), "WQTau" (weighted $Q\tau$ estimate), "QTau" ($Q\tau$ estimate), and "ML" (maximum likelihood estimate).

`control` is a list, which contains an object as returned by the function `loggamma.rob.control`. ... further arguments that can be directly passed to the function.

The function returns an object of class ‘`loggammaRob`’. This is a list with the following components:

`mu`: location parameter estimate.
`sigma`: scale parameter estimate.
`lambda`: shape parameter estimate.
`eta`: estimate of $E(\exp(x))$.
`weights`: the final weights.
`iterations`: number of iterations.

Affiliation:

Claudio Agostinelli
Department of Mathematics
University of Trento
Trento, Italy
E-mail: claudio.agostinelli@unitn.it

Alfio Marazzi
University of Lausanne
Lausanne, Switzerland
E-mail: Alfio.Marazzi@chuv.ch

Víctor J. Yohai
University of Buenos Aires
Buenos Aires, Argentina
E-mail: victoryohai@gmail.com

Alex Randriamiharisoa
University of Lausanne
Lausanne, Switzerland
E-mail: Alex.Randriamiharisoa@chuv.ch