Biometrics 000, 000–000 DOI: 000
000 0000

Parametric Quantile Regression Based on the Generalised Gamma Distribution

Angela Noufaily

Department of Mathematics & Statistics, The Open University, Walton Hall, Milton Keynes MK7 6AA, U.K.

email: a.noufaily@open.ac.uk

and

M.C. Jones

Department of Mathematics & Statistics, The Open University, Walton Hall, Milton Keynes MK7 6AA, U.K.

email: m.c.jones@open.ac.uk

Summary: We explore a particular fully parametric approach to quantile regression and show that this approach can be very successful. Motivated by the provision of reference charts, we work in the specific context of a positive response variable, whose conditional distribution is modelled by the generalised gamma distribution, and a single covariate, the dependence of parameters of the generalised gamma distribution on which is through simple linear and log-linear forms. With only 6 parameters at most, such models allow a perhaps surprisingly wide range of distributional shapes that seems adequate for many practical situations. We show that maximum likelihood estimation of the models is computationally quite straightforward; that the estimated quantiles behave well; that use of standard maximum likelihood asymptotics to perform likelihood ratio tests of the number of parameters needed, and to give pointwise confidence bands based on the expected information matrix are reliable in this context; and we more tentatively provide a simple goodness-of-fit test of the whole model. As well as simulations, three data examples from the biometrical sphere are included. We claim that a quite direct parametric maximum likelihood approach like this one — which also obviates the problem of quantile crossing — is adequate for many situations, and there is less need than one might think to resort to more complicated semi- and non-parametric approaches to quantile regression.
Key words: Likelihood inference; Quantile crossing; Reference chart.

1. Introduction and Approach

Quantile regression (QR) is an increasingly popular methodology for understanding the conditional distribution of a response variable given the values of one or more covariates. A vast literature has arisen on the subject since the seminal paper of Koenker and Bassett (1978), including the important book by Koenker (2005). Almost all the work in this literature has involved a non-parametric component either in the functional form of the regression equation or the distribution of the random component of the model, or both. In this paper, we explore a fully parametric approach to QR which treats both the dependence on a single covariate and the random component parametrically. We show that this approach can be very successful, allowing a perhaps surprisingly wide range of distributional shapes that seems adequate for many practical situations.

We work in the specific context of a positive response variable, Y, and a single covariate, X. We were motivated to do so by particular interest in the provision of reference charts (e.g. Rosner et al., 2008, Bagratee et al., 2009, Cole et al., 2009, van Buuren et al., 2009) which typically consist of two or more quantile curves reflecting 'normal' values of a response of interest (positive quantities such as height or weight or certain medical measurements) plotted as a function of an important covariate, often, but far from always, age. As one example of the output of our approach, Figure 1 shows the estimated 10%, 50% (median) and 90% quantiles, each with 95% pointwise confidence bands, for data pertaining to serum immunoglobulin G concentrations in children (this analysis is discussed in Section 5.1).

[Figure 1 about here.]

The response distribution in our model is taken to be the generalised gamma (GG) distribution introduced by Stacy (1962) and Stacy and Mirham (1965) which has density

$$f(y|\theta,\beta,k) = \frac{\beta}{\theta^{k\beta}\Gamma(k)} y^{k\beta-1} e^{-(y/\theta)^{\beta}}, \qquad y > 0, \tag{1}$$

where $\theta > 0$ is the scale parameter and $k > 0$ and $\beta > 0$ are shape parameters. This is

an increasingly popular three-parameter distribution for positive data, since it includes the gamma, Weibull and exponential distributions as special cases and the lognormal distribution as a limiting case. See Section 2.3 of Noufaily (2011) for a review of the literature on the GG distribution. In QR elsewhere, parametric response distributions are relatively rarely used. The (two-parameter) normal distribution, its three-parameter Box-Cox extension (in the LMS method of Cole, 1988, and Cole and Green, 1992) and a scaled exponential transformation thereof (Royston and Wright, 1998), and the asymmetric Laplace distribution used for mathematical convenience, often in Bayesian approaches (e.g. Yu and Moyeed, 2001), rather than as a serious model for the data distribution, are the alternatives of which we are aware. The GG distribution affords a good variety of monotone and unimodal density shapes and has tails ranging from light to quite heavy ('Weibull-like'). As such, it covers many practical situations in a flexible yet estimable way but, of course, is not claimed to be a panacea.

Fully parametric quantile regression (PQR) then assigns parametric functional forms which vary with values of the covariate to some or all of the parameters of the response distribution. These 'parameter curves' are usually not of direct interest but are readily combined to yield the required quantile curves which vary with values of the covariate. After initially reparametrising the GG distribution (Section 2.1), we employ very simple linear and log-linear functions for each of the parameters and find that the appropriate combinations thereof make for a remarkably wide range of simple, but useful, shapes for the quantile functions (Section 2.2). This contrasts with linear quantile regression in which the quantile curves themselves are linear (and the response normal or, more usually, nonparametric), fractional polynomials in Royston and Wright (1998) or else the very widespread use of nonparametric estimates of functional form (in the case of the LMS method in conjunction with parametric response; elsewhere, usually fully nonparametric e.g. Yu and Jones, 1998, Yu, Lu and Stander, 2003, El Ghouch and Van Kielegom, 2009, Chernozhukov, Fernandez-Val and Galichon, 2010,

Dette, Wagener and Volgushev, 2011). The latter can be overly complicated or unnecessary for the relatively simple forms of quantile dependence often observed in real data.

An immediate and important consequence of taking an approach like ours (or LMS) which involves parametric conditional distributions is that the estimated quantile curves can not cross. This seemingly innocuous claim is a major advantage over most nonparametric approaches which are dogged by quantile crossing problems, generating a large literature on avoidance of the problem (e.g. He, 1997, Dette and Volgushev, 2008, Bondell, Reich and Wang, 2010, Chernozhukov et al., 2010). Crossing quantile curves is simply not an issue here.

Model fitting and inference proceed by maximum likelihood (ML) estimation and utilisation of the standard asymptotic properties thereof. Computationally, maximum likelihood estimation is successfully performed by standard numerical optimisation routines run from a few starting values to guard against very occasional failures or local maxima (see Section 3.1). (Our most complex models have only six parameters.) This might be thought to contradict a literature on fitting the GG distribution (with three parameters) outside the QR context which we have argued gives the impression that ML estimation for the GG distribution is fraught with difficulties; see Noufaily and Jones (2011) for references and details.

Pointwise confidence bands are provided for estimated quantile curves, and likelihood ratio tests are performed to select the most appropriate number of parameters to depend on the covariate; see Sections 3.2 to 3.4. We show in Sections 4.2 to 4.4 that these approaches, based on standard asymptotics for ML estimation, work very, perhaps remarkably, well in this context for datasets down to size $n = 200$ (which are around the smallest for which one should contemplate producing reference charts). A simple chi-squared-type test for goodness-of-fit testing is also proposed (Section 3.5), although there is scope for improvement in this.

Three real data examples, two taken from the literature, one entirely new, are analysed in Section 5, and illustrate various aspects of GG QR modelling.

As well as the fundamental works of Roger Koenker and colleagues and the consequent vast literature alluded to above, our work has other close antecedents and relations. Approaches of this kind are nascent in the works of Gilchrist (1997, 2000, 2008). Gilchrist (2008), in particular, makes a strong call for work on fully parametric QR. Gilchrist also reminds us that QR can be traced back to Galton (1875) and that Parzen (1979), like Gilchrist later, propounds the central use of quantile functions in statistics. We must also mention the GAMLSS package (Stasinopoulos and Rigby, 2005, 2007, GAMLSS, 2011) which allows the fitting of regression models with a huge variety of parametric conditional distributions (including GG) and both parametric and non-parametric dependence on covariates. (GAMLSS is not focussed on QR per se but has recently become increasingly used in this context.) This paper, however, constitutes the first in-depth investigation of a specific practical version of PQR beyond the normal/linear case, which confirms its very good properties and practicability.

2. Generalised Gamma Quantile Regression: The Model

In this section, the elements of our approach to parametric quantile regression with one covariate are described. Let the data be (X_i, Y_i) with $Y_i > 0$, $i = 1, ..., n$.

2.1 *The GG Distribution in the PQR Context*

The data are modelled as

$$Y_i|X_i = x_i \sim f(y|\theta(x_i), \beta(x_i), k(x_i)), \quad i = 1, ..., n, \tag{2}$$

with $Y_1|X_1 = x_1, ..., Y_n|X_n = x_n$ independent. Here, the density f of the conditional distribution of $Y|X = x$ is the GG density given at (1) and the parameters depend on the covariate in as yet undetermined ways (which may include no dependence on the covariate for one or more parameters; see Section 2.2).

We immediately reformulate and reparametrise the model to aid computation and inter-

pretation. Let $Z_i = \log Y_i$, $i = 1, ..., n$ and $z = \log y$ so that $Z_1, ..., Z_n, z \in \mathbb{R}$. Then,

$$Z_i|X_i = x_i \sim f_R(z|\mu(x_i), \sigma(x_i), k(x_i)), \quad i = 1, ..., n, \tag{3}$$

with $Z_1|X_1 = x_1, ..., Z_n|X_n = x_n$ independent. Here, the density f_R of the conditional distribution of $Z|X = x$ is the 'log generalised gamma' distribution (as employed, for example, by Lawless, 1980), with density

$$f_R(z|\mu, \sigma, k) = \frac{k^{k-\frac{1}{2}}}{\sigma\Gamma(k)} \exp\left\{\sqrt{k}w - ke^{w/\sqrt{k}}\right\}, \qquad z \in \mathbb{R}, \tag{4}$$

where $w = (z - \mu)/\sigma$, $\mu \in \mathbb{R}$ is the location parameter, $\sigma > 0$ is the scale parameter and $k > 0$ is the shape parameter. The parameter k is the same in models (1) and (4), while the other parameters are related by

$$\mu = \log\theta + (\log k)/\beta, \quad \sigma = 1/(\beta\sqrt{k}); \tag{5a}$$

$$\theta = e^{\mu}/k^{\sigma\sqrt{k}}, \quad \beta = 1/(\sigma\sqrt{k}). \tag{5b}$$

Let $\gamma(y; k) = \int_0^y v^{k-1}e^{-v}dv$ denote the incomplete gamma function. Then, the cdf associated with density f in (1) is simply

$$F(y|\theta, \beta, k) = \frac{\Gamma((y/\theta)^{\beta}; k)}{\Gamma(k)}.$$

Write $r(q; k)$, $0 < q < 1$, for the inverse function of $\Gamma(y; k)/\Gamma(k)$ (this is the quantile function of the gamma, $G(\theta = 1, k)$, distribution). Then the quantile function associated with density f in (1) can be written, making use of relationships (5), as

$$Q(q|\theta, \beta, k) = Q(q|\mu, \sigma, k) = e^{\mu}\left(\frac{r(q; k)}{k}\right)^{\sigma\sqrt{k}}.$$

In the regression context, we estimate

$$Q(q_j|\mu(x), \sigma(x), k(x)) = e^{\mu(x)}\left(\frac{r(q_j; k(x))}{k(x)}\right)^{\sigma(x)\sqrt{k(x)}} \tag{6}$$

for selected values of $0 < q_j < 1$.

2.2 *Dependence on x in GG QR models*

We consider only simple linear and log-linear parametric forms for the dependence of the parameters of the GG distribution on x:

$$\mu(x) = a + bx, \qquad \sigma(x) = \exp(c + dx), \qquad k(x) = \exp(f + gx), \tag{7}$$

$a, b, c, d, f, g \in \mathbb{R}$. Our GG QR model therefore involves at most six parameters. We also consider the following hierarchy of models:

• first, the shape parameter k does not vary with x ($g = 0$) giving a five-parameter GG QR model;

• second, in addition, the log-scale parameter σ does not vary with x ($d = 0$ as well as $g = 0$) giving a four-parameter GG QR model;;

• finally, in addition, the log-location parameter μ does not vary with x ($b = d = g = 0$) giving the three-parameter GG model with no dependence on x.

Of course, other four- and five-parameter submodels of the GG QR model are possible but are not considered here.

Despite the simplicity of the above, a wide variety of types of dependence of the quantiles on x ensues. The four-parameter GG QR model has $Q(q_j|\mu(x), \sigma, k) = e^{a+bx} C_j(\sigma, k)$ where

$$C_j(\sigma, k) = \left(\frac{r(q_j; k)}{k}\right)^{\sigma\sqrt{k}} \tag{8}$$

are positive constant multipliers differing only with j. In fact, they necessarily increase with the value of q_j. This therefore allows a simple log-linear dependence of the quantiles on the covariate, the quantiles increasing if $b > 0$ and decreasing if $b < 0$ (and constant if $b = 0$). Note that, in this version of the model, all quantiles share the same, monotone, behaviour.

Five- and six-parameter GG QR models allow both more complex shapes and different shapes for different quantiles. The five-parameter GG model does so in an attractively

controlled way. In this case,

$$Q(q_j|\mu(x),\sigma(x),k) = e^{a+bx}\{C_j(\log c, k)\}^{\exp(dx)}.$$

Simple calculus shows these quantiles to be monotone increasing or decreasing, unimodal or uni-antimodal depending on the values of b, d and C_j. The mode or antimode exists and is at $m_j = \log\{-b/(d\log C_j)\}/d$ whenever the argument of the log function is positive. Modes and antimodes increase as q_j decreases. The quantiles are log-concave if $C_j < 1$ or log-convex if $C_j > 1$. Table 1 lists how shapes of *individual* quantiles depend on b, d and C_j.

[Table 1 about here.]

Web Figure 1 shows the quantiles for values

$$q_1 = 0.1, \quad q_2 = 0.25, \quad q_3 = 0.5, \quad q_4 = 0.75, \quad q_5 = 0.9, \tag{9}$$

in an example where $b = -0.1 < 0$, $d = -2 < 0$ (and $a = 1$, $c = -1.5$ and $k = 0.75$), so that quantiles are either monotone decreasing (if $C_j > 1$ which happens in this case for $q_j = 0.75, 0.9$) or unimodal (if $C_j < 1$ which happens in this case for $q_j = 0.1, 0.25, 0.5$) with modes increasing as q_j decreases. It is important, too, to remember that unimodal or uniantimodal quantiles can be monotone over the range of x of interest (see, for example, the 0.1-quantile in Web Figure 1 if x is restricted to $(0, 1)$).

The form of quantile curves in the six-parameter GG model is much more complicated than in its submodels since, now, k also depends on x and enters expression (6) in a complicated way (including within $r(q; k)$). Our experience has been that most quantiles remain (at least over a range of interest) monotone or uni(anti)modal, although we have also observed bimodal behaviour. We demonstrate this fact by a particular example with $\mu(x) = 5 + 3x$, $\sigma(x) = \exp(0.5 - 2x)$ and $k(x) = \exp(9 - 15x)$ in Web Figure 2. Over $x \in (0, 1)$, the 0.9-quantile is bimodal, while the other displayed quantiles are unimodal.

We believe that this subsection illustrates quite well the wide range of realistic collections

of quantile shapes, from the simple to the relatively complex, that result from our GG QR model specification, giving it scope for a wide range of applicability in practice.

3. Generalised Gamma Quantile Regression: Inference

For all its flexibility, our GG QR model is parsimonious, involving a maximum of six parameters. As such, it is amenable to direct maximum likelihood estimation. In this section, we describe and discuss the ingredients of likelihood inference for our model. We follow standard ML procedures, including use of ML asymptotics. In the simulations of Section 4, we show that this is a situation in which standard ML inference can be relied upon to perform very well.

3.1 *Computation*

We work on the log-likelihood function ℓ arising from formulation (3) with $\mu(x)$, $\sigma(x)$ and $k(x)$ given by the simple parametric forms (7), and with none to three of the parameters g, d and b set equal to zero, corresponding to the six- to three-parameter (sub)models laid out at the start of Section 2.2.

We have found that direct numerical optimisation of these log-likelihoods works perfectly well. We use the Nelder-Mead algorithm (Nelder and Mead, 1965) as implemented in R (R Development Core Team, 2009). We run the algorithm from 10 randomly chosen initial parameter configurations and choose the parameter values giving the maximum log-likelihood achieved in these runs (which typically arise from most if not all runs) in order to guard against possible failures to converge or local maxima.

We report detailed investigations of practical ML estimation for the three- and four-parameter GG models in Noufaily and Jones (2011). There, we also considered other off-the-shelf optimisation algorithms and bespoke iterative algorithms including one of our own. Although there are no relevant theoretical results, the log-likelihood surface appears to be

well-behaved. In Noufaily and Jones (2011), we say: "The GG likelihood surface, while not unimodal, appears to comprise a strong global maximum together with very few local maxima which are much smaller than, and far away from, the global maximum. Also, while the different maximisation methods sometimes fail to find the global maximum on a given single run, they succeed sufficiently often that an approach consisting of running any one (or a mix) of the maximisation methods mentioned above just a few times from different starting points is virtually guaranteed to locate that distinct global maximum."

The continued good performance of this strategy in the five-and six-parameter GG QR model will be clear from the quality of the estimates obtained.

3.2 *Choosing the Best Version of the Model*

Obviously, a more accurate fit to the data is obtained as the number of parameters increases. What remains to check is whether a model with more parameters is significantly more accurate than one with fewer. If not, it is better on grounds of parsimony to retreat to a simpler submodel with fewer parameters.

To test for this significance, we employ the well-known likelihood ratio test (LRT) to compare GG QR models sequentially, using the usual asymptotic chi-squared result. For given data, we denote maximised log-likelihoods of the p-parameter GG models by ℓ_p, $p = 3, ..., 6$. We perform at most three such tests. Write $D_j = 2(\ell_{7-j} - \ell_{6-j})$, $j = 1, 2, 3$. The null distributions of these test statistics are each approximated by χ^2_1, the chi-squared distribution with one degree of freedom, whose 95% quantile is 3.84. So, we make tests $j = 1, 2, 3$ until we meet j such that $D_j > 3.84$ and conclude that a $(7-j)$-parameter GG QR model is most appropriate; if we reach test 3 and $D_3 \leq 3.84$, we use a three-parameter GG model.

3.3 *Asymptotics*

GG PQR is a regular ML problem and so standard asymptotics apply. The ML estimators of the parameters are asymptotically normal, asymptotically unbiased and have asymptotic

variance-covariance matrix given by the inverse of the expected Fisher information matrix, $\mathcal{I}$. The elements of $\mathcal{I}$ are given for the six-parameter GG model in Web Appendix A; the corresponding matrices for three-, four- and five-parameter GG models can be distilled from this matrix. Since we are not especially interested in the parameter estimates per se, the main application of $\mathcal{I}$ is in providing confidence bands for the quantile curves, next.

3.4 *Confidence Bands for the Quantile Curves*

Having obtained ML estimates of the parameters, convert them to estimates of the quantile curves via formula (6) using parametrisation (7); call the resulting estimated quantile curve at the point x, $\hat{t}_{q_j}(x)$. The expected Fisher information matrix, $\mathcal{I}$, for the parameter estimates can be transformed into the asymptotic variance of $\hat{t}_{q_j}(x)$, $\Sigma_j(x)$ say, by the delta method as outlined in the six-parameter case in Web Appendix B. We then use the estimated parameters in $\Sigma_j(x)$ to obtain $\hat{\Sigma}_j(x)$ and obtain $100(1-\alpha)\%$ confidence intervals at each x — and hence $100(1-\alpha)\%$ pointwise confidence bands for the quantile curves — in the usual way through

$$\left(\hat{t}_{q_j}(x) - z_{1-(\alpha/2)}\sqrt{\hat{\Sigma}_j(x)},\ \hat{t}_{q_j}(x) + z_{1-(\alpha/2)}\sqrt{\hat{\Sigma}_j(x)}\right) \tag{10}$$

where $z_{1-(\alpha/2)}$ is the $(1-(\alpha/2))$-quantile of the standard normal distribution.

Note that we apply formula (10) to whichever number of parameters is decided upon by the likelihood ratio testing of Section 3.2, and take no account in our confidence bands of this model selection process. Simulation results to be discussed in Section 4 suggest that this simplifying approach does not lead to underestimation of the variability in our results; the simulation results also give us strong confidence in the classical use of the standard asymptotic approximations in this context.

3.5 *A Chi-Squared Test of Goodness-of-Fit*

Following model identification, parameter estimation and model selection, the final stage in our overall procedure is to test the goodness-of-fit of the fitted model. We propose a simple, global, χ^2-type test to do so.

Ideally, $100(q_{j+1} - q_j)\%$ of the data should lie between the j^{th} and $(j+1)^{th}$ quantile. For our goodness-of-fit test, we consider the five quantiles defined by the probabilities given at (9), which divide the data into six regions. Denote by $\hat{m}_k$ and m_k, $k = 1, ..., 6$, the observed and expected number of data points in each of the regions. If n is the sample size and $q_0 = 0$, $q_6 = 1$, then $m_k = (q_{j+1} - q_j)n$, $k = 1, ..., 6$. We therefore propose the test statistic

$$\tau = \sum_{k=1}^{6} \frac{(\hat{m}_k - m_k)^2}{m_k}. \tag{11}$$

In a simulation study presented in Section 5.2 of Noufaily (2011), it was shown that, approximately,

$$\tau = \tfrac{3}{4}\varrho \quad \text{where } \varrho \sim \chi^2_4, \quad \text{or equivalently, } \tau \sim G(\theta = 1.5, k = 2).$$

Note that we are not able to prove this result theoretically. It is based merely on computational evidence obtained from simulations. We use this result to test the goodness-of-fit of the GG model. For example, the 95% quantile of the $G(\theta = 1.5, k = 2)$ distribution is about 7.12 so that, at a 0.05 level of significance, if $\tau > 7.12$, we reject the fact that the GG is a good fit, otherwise, we conclude that the GG is a good model. We recognise that this test is broad-brush, though it can also be used in a more focussed way to look at model fit over subsets of the range of the covariate. The individual components of τ give information on where any perceived lack-of-fit occurs.

4. Simulation Performance

In this section, we validate, via simulations, some of the theoretical and methodological work described in the paper. We apply the steps described in Section 3 to datasets of size

$n = 200$ and $n = 500$ simulated from each of four parameter configurations from each of the three-, four-, five- and six-parameter GG PQR models (see Section 4.1). 100 replications of each dataset were generated with, in each case, the covariate $X = x$ (when required) generated uniformly over $(0, 1)$ and the $\mathrm{GG}(\theta(x), \beta(x), k(x))$ responses generated by taking the $1/\beta(x)$'th power of a gamma, $G(\theta(x)^{\beta(x)}, k(x))$, random variable. In Section 4.2, we will give evidence that likelihood ratio testing for choice of number of parameters in the GG PQR model (Section 3.2) works well. Then, in Sections 4.3 and 4.4, we present the main results of our simulations. We will compare the practicable case of conducting LRTs to determine the number of parameters to use in the model with the impracticable use of the true number of parameters from which each data set was simulated, and find little or no deterioration in performance. It will also be shown that the quantile curves lie within their confidence intervals (CIs) an appropriate proportion of the time and that the relative biases and mean squared errors of the estimated parameters and quantiles are generally reasonably small. Overall, the simulation results give us considerable confidence in the methodology.

4.1 *The Simulation Scenarios*

Web Table 1 presents the sets of parameters from which each of the three-, four-, five- and six-parameter GG PQR datasets were generated. The 16 sets of parameters were chosen to cover most scenarios a GG PQR model can portray. In the three-parameter case, we tackle situations where $\alpha < 1$ (D32), corresponding to a monotone GG density function, as well as $\alpha > 1$ (D31, D33 and D34), corresponding to unimodal GG density functions; parameters generally have a mix of fairly large and rather smaller values. In the other scenarios, data is dependent on the covariate. With four parameters, we deal with cases where the quantiles are increasing ($b > 0$ in D41 and D43) and decreasing ($b < 0$ in D42 and D44), while covering a large range of parameter values from relatively small (e.g. $b = -0.1$ in D42) to relatively large (e.g. $f = 1.609$ so that $k = 5$ in D43). Similarly, in the five-parameter scenarios, we

cover a reasonably wide range of values for each parameter and take different combinations of positive and negative b and d which result in monotone and unimodal quantiles. Likewise, in the six-parameter case, we ensure that the four scenarios address situations with monotone, unimodal and bimodal quantiles. Quantiles from scenarios D32, D42, D54 and D62, which are the particular focus of Section 4.4 to follow, are shown in Figure 2.

[Figure 2 about here.]

4.2 *LRT Performance*

As a precursor to our main simulation results, Web Table 2 displays the number of parameters reported by the LRTs for each of the 16 scenarios described in Web Table 1. Each column presents the numbers of times (percentages) each of the 100 replicated datasets are reported to come from a three-, four-, five- or six-parameter GG PQR model by the LRT. The percentages in the rows associated with the correct numbers of parameters are mostly between 80 and 100, indicating very satisfactory performance in those cases. The few exceptions to this are readily explainable: they correspond to scenarios where one of the true parameters is close to zero, so a model with one fewer parameter than the truth is then selected. (For example, D42 is reported as coming from a three-parameter GG distribution 88% of the time for $n = 200$ and 81% of the time for $n = 500$. It corresponds to $b = -0.1$ which is quite close to zero.) While the difference between a GG PQR model with one small parameter and a model with that parameter set to zero could not be picked up in LRT terms, those two models are very similar and give very similar results, so the LRT performance provides no practical problem. There are no great differences between the $n = 200$ and $n = 500$ cases.

4.3 *Results on Parameter Estimation*

In this analysis, we compare the estimated parameters to the true ones for each of the (practicable) "LRT Version" and the (impracticable) "Known p Version". For each data set

from the 16 scenarios, we compute two statistics, the relative bias and the relative squared error, which are averaged over the 100 data sets from each scenario. We then take the square root of the mean relative squared error. This is straightforward in the "Known p Version" because the number of estimated and true parameters is the same. A few alterations are made in the "LRT Version" when the predicted number of estimated parameters differs from the true number. If the number of predicted parameters is greater than the true number, the actual (not relative) bias and squared error of the extra estimated parameter(s) are given by the averages of their estimated and squared estimated values, the latter square-rooted, respectively. If the true number of parameters is greater than the predicted number, the relative bias and relative squared error for the extra real parameter are given by the averages of the relative quantities where 0 is entered in place of the estimated parameter.

For $n = 200$, Web Tables 3 to 6 report the mean relative biases and the root mean relative squared errors of the parameters for each of the three-, four-, five- and six-parameter GG simulations respectively. In the tables, we use a hopefully obvious terminology for bias and (root) mean squared errors, examples of which are biasmu for the (relative) bias of $\hat{\mu}$ as estimator of μ and rmsek for the root (relative) mean squared error of $\hat{k}$ as estimator of k. Results are given for both the "LRT Version" and the "Known p Version". The analogues of these tables for the $n = 500$ case are given in the Appendix of Noufaily (2011).

Perusing these results, we realise that generally the mean relative biases and the root mean relative squared errors are quite small. For most of the cases, the "Known p Version" is slightly better than the "LRT Version", though the difference is very small. Also, there proves to be very little difference between the $n = 200$ and $n = 500$ cases. Although the $n = 500$ case has slightly smaller biases and errors, our method still performs well in the $n = 200$ case implying that the approach works reasonably well for smaller sized datasets. Of particular interest is that D42, D61 and D62 behave fairly well. Even though the LRT

failed to predict the right number of parameters for those data sets, it still managed to alight on a reasonable model. In rare cases, we notice large values in the tables. For example, "rmsek" and "biask" are quite large in Web Table 3 for D33 and D34. This reflects the failure of the method to accurately estimate k when k is relatively large, as the shape of the distribution only changes slightly when k approaches infinity. We do not observe a noticeably greater/fewer number of positive or negative biases in general which means that there is no systematic over/under-estimation.

4.4 *Results on Quantile Estimation*

Having compared the estimated and true parameters for both versions, we now find out, most importantly, the effect of those results on the estimated quantiles themselves. This was done for all 16 scenarios, but we present results here for the representative set of 4 scenarios, D32, D42, D54 and D62, one with each number of parameters, whose quantiles were shown in Figure 2. In this case we compute, and average over simulations, the averaged relative biases (abias) and root averaged relative squared errors (ramse) of our estimates of the quantiles for each of the quantiles listed at (9). Here, averaging takes place over a fine equi-spaced grid of 1000 x values, to afford a global assessment of the quality of estimation for each quantile. We also report the proportions of gridpoints for which the real quantiles lie within our pointwise confidence bands both for $\alpha = 0.05$ (CI95prop) and $\alpha = 0.01$ (CI99prop).

[Table 2 about here.]

For $n = 200$, Table 2, included here as an example for convenience, and Web Tables 7 to 9 display results for each of scenarios D32, D42, D54 and D62, respectively, when $n = 200$. In Table 2, for the highly dispersed but parallel quantiles of D32, "abias", always positive, and "ramse" are reasonably small. Higher quantiles seem to be a little better estimated than lower quantiles. Also, just under 95% of the time the real quantiles lie within their 95% confidence bands and just under 99% of the time they lie within their 99% confidence bands,

the worst exceptions being when $\alpha = 0.05$ and $q_j = 0.9$. (So it seems that better (worse) "point" estimates of quantiles correspond to worse (better) coverage of confidence bands.) Errors are a bit bigger and coverages a bit smaller for the "LRT Version" compared with the "Known p Version", but coverages, in particular, are not greatly affected by not knowing p.

Similar patterns occur in the other tables. In Web Table 7, corresponding to the gently decreasing quantiles of D42, "abias" and "ramse" are smaller than in Table 2 and coverages of confidence bands are broadly comparable. In Web Table 8, "abias" and "ramse" are higher than in other tables. This seems to be due to the more complex quantile shapes: the 0.9 quantile has a strong minimum while the 0.1, 0.25, 0.5 and 0.75 quantiles are monotone decreasing. "CI95prop" and "CI99prop" are under-estimated a little more. Web Table 9 corresponds to the converging quantiles of D62. Here, the "Known p Version" slightly over-estimates "CI95prop" and "CI99prop", whereas the "LRT Version" slightly under-estimates. Noticeably, "abias" and "ramse" are generally very small, the biases mostly being negative.

It is valuable to note that any biases or large variances in parameter estimation seem not to have a great effect on the quality of quantile estimation. Scenarios D42 and D62 are both amongst those for which the LRT often underestimates the number of parameters (Web Table 2), but no matter since the "missing" parameters are small anyway. In D54, the fairly large root relative mean squared errors of $\hat{k}$, associated with fairly large underlying k, seem not to be the major contributor to the relatively weaker results (the quantile shapes might be more so) since distributions, and quantiles, are similar over a broad range of large k.

5. Examples

In this section, we apply our version of parametric quantile regression to three data sets, concerning: Immunoglobulin Concentration vs Age; Weight vs Height; and Water Table Depth vs Flux. The first two datasets are health-related and the third environment-related.

5.1 *Serum Immunoglobulin G Concentrations in Children*

Isaacs et al. (1983) measured the serum immunoglobulin G concentrations (IgG) of 298 children aged 6 months to 6 years old. Parametric quantile regression models were fitted to these data by Isaacs et al. (1983) and Royston and Wright (1998) (see also Kapitula and Bedrick, 2005); Yu et al. (2003) proposed a kernel-based non-parametric smoothing method to estimate the quantile curves; while, very recently, Jara and Hanson (2011) applied a Bayesian semiparametric approach.

Given the different approaches used to fit quantiles to the IgG data, and for the sake of comparison, we apply the methodology developed in this paper to these data. At the first stage, $\ell_6 = -138.54$ and $\ell_5 = -141.89$, so that $D_1 = 6.7$: the significant difference between ℓ_6 and ℓ_5 means that a six-parameter GG QR model is required. (Unusually, we had a little trouble with computational accuracy in the six-parameter model here which did not afflict other versions of the model.) The values of the estimated parameters are $\hat{a} = 1.384$, $\hat{b} = 0.092$, $\hat{c} = -1.021$, $\hat{d} = 0.008$, $\hat{f} = -3.493$ and $\hat{g} = 4.766$. Figure 3 presents a plot of the 10%, 25%, 50%, 75% and 90% estimated quantiles. The quantiles are monotone increasing. They increase sharply up to one year of age and more gradually for individuals older than that. The 95% pointwise confidence bands around the 10%, 50% and 90% estimated quantiles were shown in Figure 1. We observe that near the center (where data are denser), the confidence bands are quite narrow, while they are wide at the edges where there are few data points.

[Figure 3 about here.]

To check the goodness-of-fit of our estimated model, we compute $\hat{m}_k$ and m_k, $k = 1, ..., 6$, respectively the observed and expected number of data points in each of the six regions formed by the 10%, 25%, 50%, 75%, and 90% quantiles, from which we obtain the test statistic τ. It turns out that $\{\hat{m}_l\} = \{25, 55, 63, 81, 55, 19\}$, $\{m_l\} = \{29.8, 44.7, 74.5, 74.5, 44.7, 29.8\}$ and $\tau = 11.78$. Using our suggested approximate $G(\theta = 1.5, k = 2)$ sampling distribution,

the p-value is 0.0034. Despite this (slightly marginal) rejection of the GG QR model, Figure 3 and the $\hat{m}_l$ values suggest that the model fits quite well and reflects the general shape of the data. Several points lie very close to the 10% estimated quantile, so a very slight upward shift of the quantiles would have resolved that part of the lack of fit due to $\hat{m}_1$ and $\hat{m}_2$.

Given the sharper increase of the estimated quantiles in the interval $(0, 1]$, we also tested the goodness-of-fit of the model for two regions separately: $x \leq 1$ and $x > 1$. The results for $x \leq 1$ are $\{\hat{m}_l\} = \{7, 13, 14, 12, 11, 5\}$, $\{m_l\} = \{6.2, 9.3, 15.5, 15.5, 9.3, 6.2\}$ and $\tau = 3.05$, p-value 0.397. This, along with similar behaviour in the fitted models of Royston and Wright (1998), Jara and Hanson (2011) and the nonparametric fit of Yu et al. (2003, Figure 8(b)), contrasting with its non-appearance in simpler models in Isaacs et al. (1983), Yu et al. (2003, Figure 8(a)) and our five-parameter GG fit, give us confidence in the reality of this effect.

5.2 *Weight and Height in the NHANES Study*

Weight problems are becoming more and more common so that some world leaders are considering this issue a threat to societies and are therefore seeking solutions. Since 1999, the National Center for Health Statistics in the United States has conducted the NHANES survey annually. We used the NHANES 2007-2008 survey (CDC, 2007). We considered two important variables, the weights (in kg) and heights (in cm) of individuals across all age ranges, and limited ourselves to the data on 4,448 males with known weight and height.

Figure 4 displays a scatter plot of the data superimposed on which are our estimated 10%, 25%, 50%, 75% and 90% quantiles, together with their pointwise confidence bands. Of course, generally, weight is increasing with height; for smaller heights, the variability of height is quite small, increasing as height increases. The estimated quantile curves reflect the shape of this dataset very nicely, in a smoothly increasing manner. The confidence bands around the estimated quantiles are narrow because the data set is quite large.

[Figure 4 about here.]

As we fitted the GG QR model to the data, results showed that the estimated k was taking very high values. The large k limiting case of the GG distribution is the lognormal distribution, and so the model fitting process, advantageously, led us to a simplified, lognormal-based, QR fit. The estimated quantiles shown in Figure 4 correspond to a 5-parameter GG QR model, and hence a 4-parameter lognormal QR model, in which the usual parameters μ and σ take linear and loglinear forms as functions of the covariate, respectively. The estimated quantiles have the relatively simple formulae

$$\hat{t}_{q_j}(x) = \exp\{2.378 + 2.685x + e^{-2.325+1.101x}\Phi^{-1}(q_j)\} \tag{12}$$

where Φ^{-1} is the quantile function of the standard normal distribution.

The observed and expected counts are $\{\widehat{m_l}\} = \{427, 763, 1086, 1118, 606, 448\}$ and $\{m_l\} = \{444.8, 667.2, 1112, 1112, 667.2, 444.8\}$ respectively, while $\tau = 20.74$ with small p-value. We discount this rejection of the goodness of fit of our model somewhat because of the large sample size (and the possible inadequacy of our χ^2 test): the observed inter-quantile counts seem to us to match their expected values pretty well. With (relative) confidence in our fitted model, another beauty of the parametric approach to quantile regression is that having fitted the model, different quantiles can be estimated trivially without refitting the model (not so for nonparametrically fitted quantile curves). A reference chart with the 5% and 95% quantiles is therefore given in Figure 5.

[Figure 5 about here.]

5.3 *Water Table Depth and Flux in the Cape Floristic Region of South Africa*

In this context, flux is the amount of rainfall that goes into the ground. Since some of this water evaporates as it touches the ground, flux can take negative values. As water filters through pore spaces in the soil, it first passes through soil which is unsaturated. At increasing depths, water fills more spaces, until the soil is saturated. The top surface of this zone of saturation constitutes the water table. Water regimes dictate different vegetation types. The

Cape Floristic Region (Cape Town, South Africa) has about 9,000 plant species most of which are unique to this area. According to changing climate scenarios, annual rainfall in the Cape is likely to decrease. A better understanding of the water distribution will enhance the way this biodiversity spot is managed. Researchers from the Open University, working in collaboration with a South African team and other environmental institutions, established around ten sites in the Cape to monitor the hydrology of the area (Araya and Walker, 2009). Reference charts giving an idea of 'normal' and 'unusual' behaviour are of interest to the researchers to allow them to classify geographical areas as having water-rich soil, dry soil, floods etc., these having a direct effect on the plant population. Here, we study one of the collected "water table depth (WTD, in cm) vs flux (in mm/day)" data sets, with $n = 307$.

On their original scale, the flux values are quite condensed near zero, therefore we aim to take the logarithm of the flux; as the minimum value for flux is -6.8 in this case, we take $\log(\text{flux} + 6.9)$ and then rescale the resulting values to the interval (0,1). A scatter plot of the data is given in Figure 6 along with five estimated quantile curves: the GG QR model required all six parameters in this case. Unsurprisingly, WTD generally decreases as flux increases, but the 6-parameter GG quantile curves are sufficiently flexible to afford some interestingly different shapes for different quantiles. Confidence bands around the estimated quantiles are shown in Web Figures 3 and 4, and display considerable uncertainty in all quantiles for small flux and in the 90% quantile for large flux.

[Figure 6 about here.]

Single, possibly outlying, points, immediately appear to have large influence in the areas of greatest uncertainty and so we considered their effect by removal, of the leftmost point (at $(0.01, 31.3)$) and the point at $(0.96, 69.8)$, both singly and together. Briefly, removal of the leftmost point causes the quantiles to become much more parallel and increasing in much the same way to the left of the plot, while removal of the other point makes little difference.

(See Section 6.1 of Noufaily, 2011, for details.) The remainder of this section concerns the dataset with its original leftmost point removed.

[Figure 7 about here.]

Our goodness-of-fit test rejects the fit of the GG QR model although again we contend that the fitted values seem pretty good: $\{\widehat{m_l}\} = \{22, 51, 75, 103, 27, 28\}$, while $\{m_l\} = \{30.6, 45.9, 76.5, 76.5, 45.9, 30.6\}$. The greatest lack-of-fit concerns the position of the 75% quantile which should be a little lower. This, in turn, appears to be due to the many points with zero flux which appear to require a short-lived dip in the 75% quantile around a rescaled flux value of about 0.57 which our model cannot accommodate. A reference chart using the 10% and 90% estimated quantiles fits well, however, and is given, on the original covariate scale, in Figure 7. Here, unusually low WTD observations are associated only with low flux, while unusually high WTD observations are mostly associated with greater values of flux.

6. Discussion

When we embarked on this work it was with the expectation of requiring to extend its more straightforward and 'obvious' aspects. It is much to the methodology's practical advantage that we have not found the need to do so. On the one hand, we did not expect that the combination of simple linear and log-linear forms for the dependence of just three parameters on the covariate would lead to such a healthy variety of quantile shapes, as explored in Section 2.2. On the other, we expected to have to adjust maximum likelihood estimates and/or move away from using standard asymptotic results, but neither has proved necessary.

We are not necessarily wedded to always using the generalised gamma as response distribution, however. We see the current study as a torch-bearer on behalf of parametric quantile regression with appropriate few-parameter response distributions — such as those in GAMLSS (2011) — and simple dependence of parameters of response distributions on

the covariate. In particular, if very heavy tails are envisaged, then F or Burr Type XII distributions might be considered, while more heterogeneous responses might require mixture response distributions.

Acknowledgements

We are very grateful to Dr. Yoseph Araya, Department of Life Sciences, The Open University who made the measurements analysed in Section 5.3 for his permission to use them here, and to Professor Frank Critchley and Dr Julian Stander for helpful comments.

Supplementary Materials

Web Appendices A and B, Web Tables 1–9 and Web Figures 1–4 referenced in Sections 2.2, 3.3, 3.4, 4.1, 4.2, 4.3, 4.4 and 5.3 are available with this paper at the Biometrics website on Wiley Online Library.

References

Araya, Y. and Walker, N. (2009). Understanding how water resources shape our flora. *Veld and Flora* **95**, 96–97.

Bagratee, J.S., Regan, L., Khullar, V., Connolly C., and Moodley, J. (2009). Reference intervals of gestational sac, yolk sac and embryo volumes using three-dimensional ultrasound. *Ultrasound in Obstetrics and Gynecology* **34**, 503–509.

Bondell, H.D., Reich, B.J., and Wang, H.X. (2010). Noncrossing quantile regression curve estimation. *Biometrika* **97**, 825–838.

Chernozhukov, V., Fernandez-Val, I., and Galichon, A. (2010). Quantile and probability curves without crossing. *Econometrica* **78**, 1093–1125.

Cole, T.J. (1988). Fitting smoothed centile curves to reference data. *Journal of the Royal Statistical Society Series A* **151**, 385–418.

Cole, T.J. and Green, P.J. (1992). Smoothing reference centile curves: the LMS method and penalized likelihood. *Statistics in Medicine* **11**, 1305–1319.

Cole, T.J., Stanojevic, S., Stocks, J., Coates, A.L., Hankinson, J.L., and Wade, A.M. (2009). Age- and size-related reference ranges: a case study of spirometry through childhood and adulthood. *Statistics in Medicine* **28**, 880–898.

Dette, H. and Volgushev, S. (2008). Non-crossing non-parametric estimates of quantile curves. *Journal of the Royal Statistical Society Series B* **70**, 609–627.

Dette, H., Wagener, J., and Volgushev, S. (2011). Comparing conditional quantile curves. *Scandinavian Journal of Statistics* **38**, 63–88.

El Ghouch, A. and Van Keilegom, I. (2009). Local linear quantile regression with dependent censored data. *Statistica Sinica* **19**, 1621–1640.

Galton, F. (1875). Statistics by intercomparison, with remarks on the law of frequency of error. *Philosophical Magazine* **49** (fourth series), 33–46.

GAMLSS (2011). Generalized Additive Models for Location, Scale and Shape. The GAMLLS Team. URL `http://www.gamlss.org`

Gilchrist, W. (1997). Modelling with quantile distribution functions. *Journal of Applied Statistics* **24**, 113–122.

Gilchrist, W. (2008). Regression revisited. *International Statistical Review* **76**, 401–418.

Gilchrist, W.G. (2000). *Statistical Modelling With Quantile Functions*. Boca Raton, LA: Chapman & Hall/CRC.

He, X.M. (1997). Quantile curves without crossing. *The American Statistician* **51**, 186–192.

Isaacs, D., Altman, D.G., Tidmarsh, C.E., Valman, H.B., and Webster, A.D.B. (1983). Serum immunoglobulin concentration in preschool children measured by laser nephelometry: reference ranges for IgG, IgA, IgM. *Journal of Clinical Pathology* **36**, 1193–1196.

Jara, A. and Hanson, T.E. (2011). A class of mixtures of dependent tail-free processes.

Biometrika **98**, 553–566.

Kapitula, L.R. and Bedrick, E.J. (2005). Diagnostics for the exponential normal growth curve model. *Statistics in Medicine* **24**, 95–108.

Koenker, R. (2005). *Quantile Regression*. Cambridge: Cambridge University Press.

Koenker, R. and Bassett, G. (1978). Regression quantiles. *Econometrica* **46**, 33–50.

Lawless, J.F. (1980). Inference in the generalized gamma and log gamma distributions. *Technometrics* **22**, 409–419.

Nelder, J.A. and Mead, R. (1965). A simplex algorithm for function minimization. *Computer Journal* **7**, 308–313.

CDC (2007). National Health and Nutrition Examination Survey. Centers for Disease Control and Prevention. URL `http://www.cdc.gov/nchs/nhanes/nhanes2007-2008/nhanes07_08.htm`

Noufaily, A. (2011). Parametric quantile regression based on the generalised gamma distribution. PhD thesis, Department of Mathematics & Statistics, The Open University.

Noufaily, A. and Jones, M.C. (2011). On maximisation of the likelihood for the generalized gamma distribution. To appear.

Parzen, E. (1979). Non-parametric statistical data modelling. *Journal of the American Statistical Association* **74**, 105–121.

R Development Core Team (2009). R: a language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL `http://www.R-project.org`

Rosner, B., Cook, N., Portman, R., Daniels, S., and Falkner, B. (2008). Determination of blood pressure percentiles in normal-weight children: some methodological issues. *American Journal of Epidemiology* **167**, 653–666.

Royston, P. and Wright, E.M. (1998). A method for estimating age-specific reference intervals

('normal ranges') based on fractional polynomials and exponential transformation. *Journal of the Royal Statistical Society, Series A* **161**, 79–101.

Stacy, E.W. (1962). A generalization of the gamma distribution. *Annals of Mathematical Statistics* **33**, 1187–1192.

Stacy, E.W. and Mirham, G.A. (1965). Parameter estimation for a generalized gamma distribution. *Technometrics* **7**, 349–358.

Stasinopoulos, D.M. and Rigby, R.A. (2005). Generalized additive models for location scale and shape (with discussion). *Journal of the Royal Statistical Society Series C* **54**, 507–554.

Stasinopoulos, D.M. and Rigby, R.A. (2007). Generalized additive models for location scale and shape (GAMLSS) in R. *Journal of Statistical Software* **23**, Issue 7.

van Buuren, S., Hayes, D.J., Stasinopoulos, D.M., Rigby, R.A., ter Kuile, F.O., and Terlouw, D.J. (2009). Estimating regional centile curves from mixed data sources and countries. *Statistics in Medicine* **28**, 2891–2911.

Yu, K. and Jones, M.C. (1998). Local linear quantile regression. *Journal of the American Statistical Association* **93**, 228–237.

Yu, K., Lu, Z., and Stander, J. (2003). Quantile regression: applications and current research areas. *Journal of the Royal Statistical Society Series D* **52**, 331–350.

Yu, K. and Moyeed, R. (2001). Bayesian quantile regression. *Statistics and Probability Letters* **54**, 437–447.

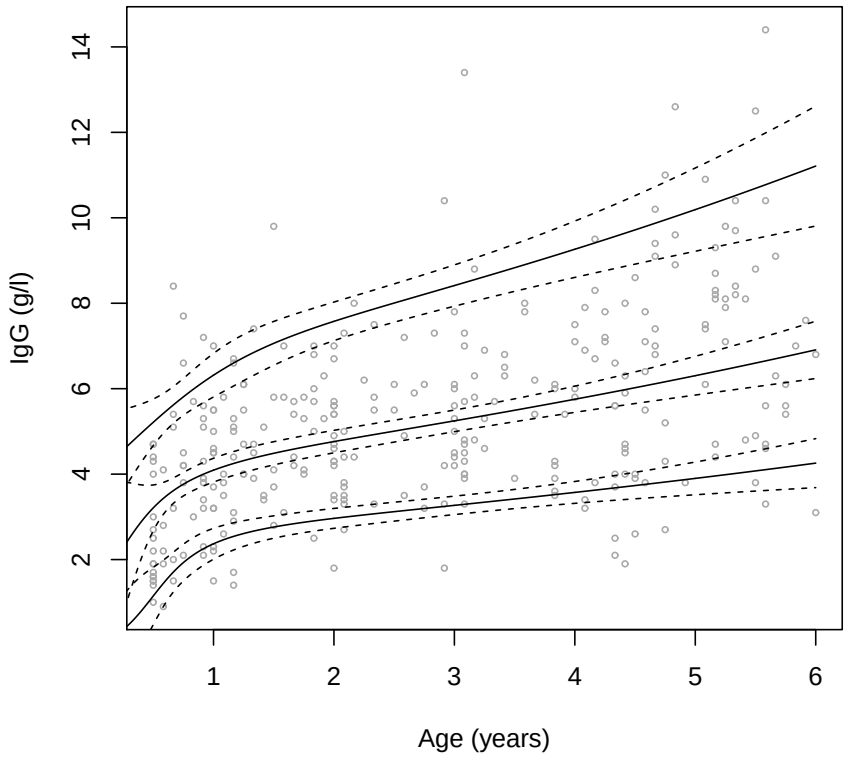

Figure 1. 10%, 50% (median) and 90% estimated quantiles (solid lines), each with 95% pointwise confidence bands (dashed lines), for serum immunoglobulin G concentrations in a sample of 298 children aged between 6 months and 6 years old.

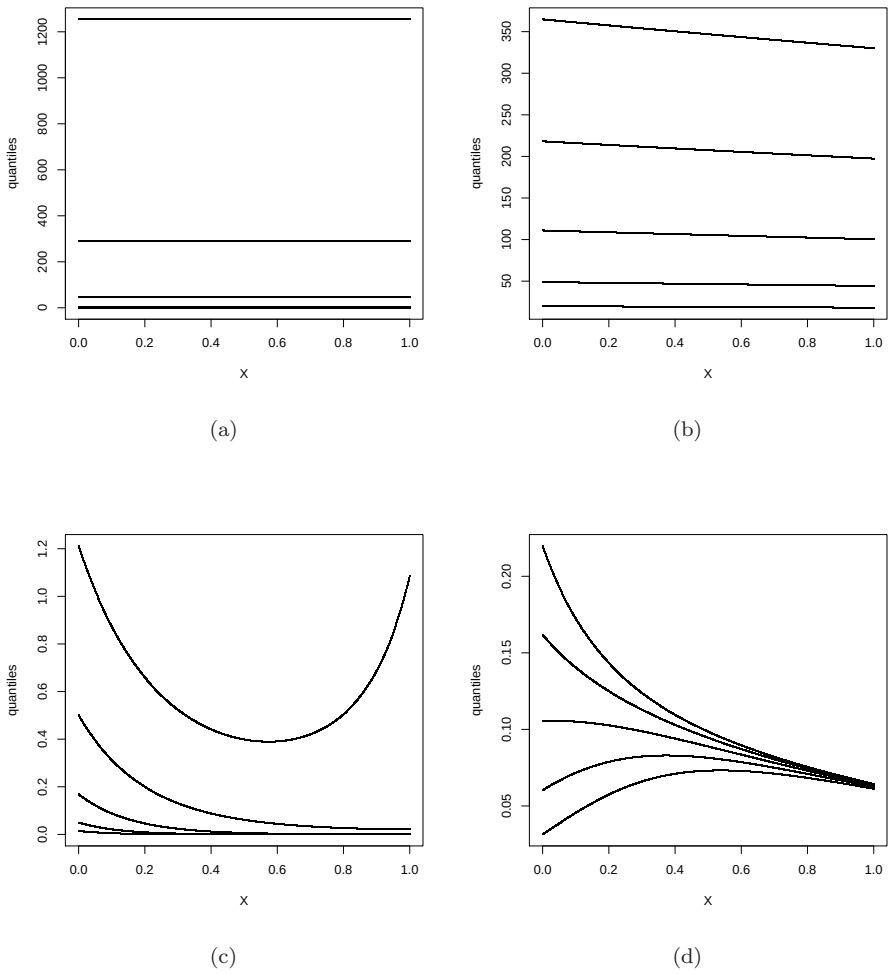

Figure 2. 10%, 25%, 50%, 75% and 90% quantile curves underlying D32 in (a), D42 in (b), D54 in (c) and D62 in (d). In (a), the 10% and 25% quantiles are so close that they produce the illusion of a single, wider, line.

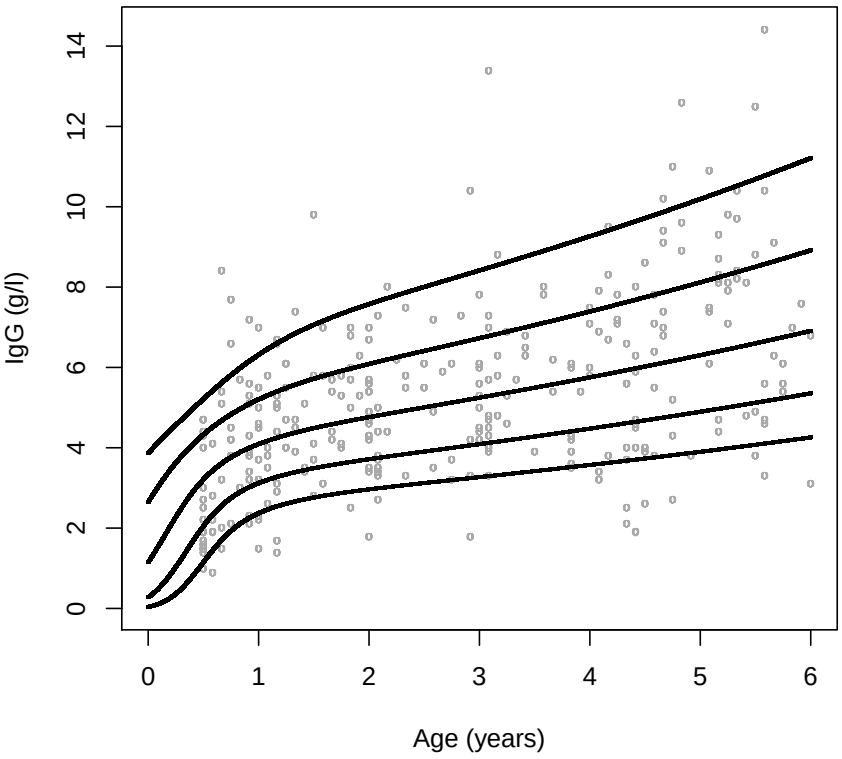

Figure 3. 10%, 25%, 50%, 75%, and 90% estimated quantiles from the six-parameter model for the IgG data.

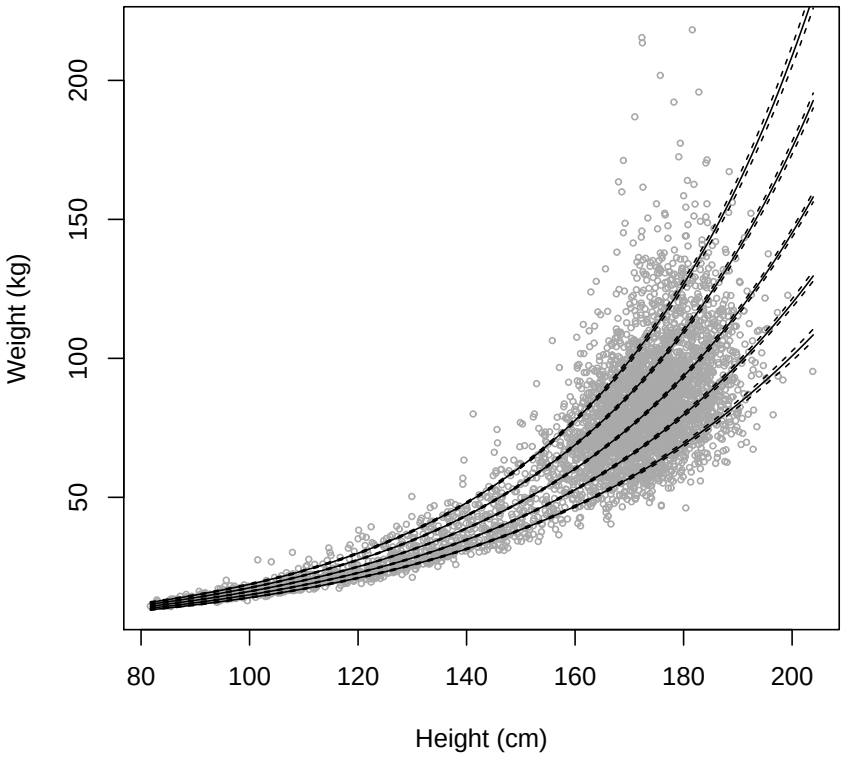

Figure 4. 10%, 25%, 50%, 75%, and 90% estimated quantiles from the four-parameter lognormal model for the NHANES data (solid lines), together with pointwise confidence bands (dashed lines).

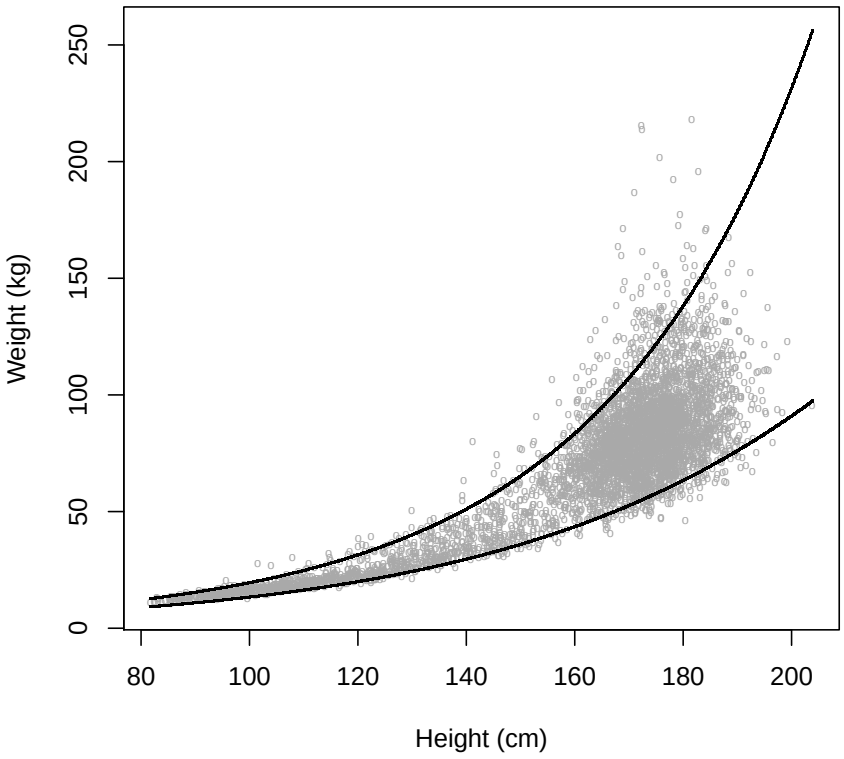

Figure 5. Reference chart of weight against height for males based on the NHANES data shown in Figure 4 using the 5% and 95% estimated quantiles.

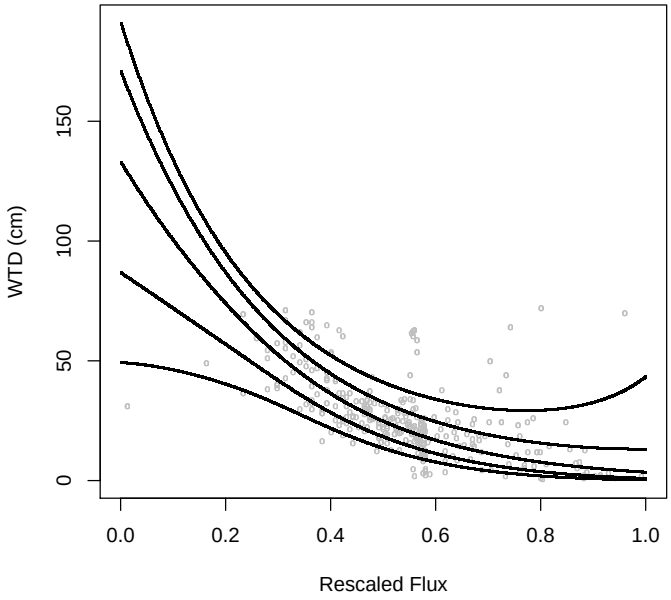

Figure 6. 10%, 25%, 50%, 75%, and 90% estimated quantiles for the WTD vs flux data .

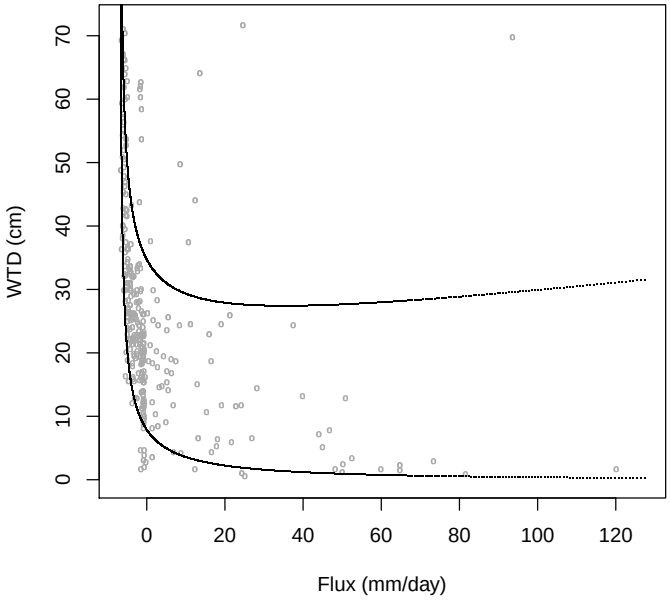

Figure 7. Reference chart of WTD vs flux (where point $(-6.8, 31.1)$ is removed from the original data set) using the 10% and 90% estimated quantiles.

Table 1
Shapes of the five-parameter GG QR model quantile curves.

cases	shape
$b > 0$; $d > 0$; $C_j > 1$	monotone increasing
$b > 0$; $d < 0$; $C_j < 1$	monotone increasing
$b < 0$; $d < 0$; $C_j > 1$	monotone decreasing
$b < 0$; $d > 0$; $C_j < 1$	monotone decreasing
$b > 0$; $d > 0$; $C_j < 1$	unimodal with maximum at m_j
$b < 0$; $d < 0$; $C_j < 1$	unimodal with maximum at m_j
$b > 0$; $d < 0$; $C_j > 1$	uniantimodal with minimum at m_j
$b < 0$; $d > 0$; $C_j > 1$	uniantimodal with minimum at m_j

Table 2

Averaged relative bias (abias) and root averaged mean relative squared error (ramse) of quantiles and proportions of the estimated quantiles inside the confidence bands for scenario D32 when $n = 200$.

Value of q_j	abias	ramse	CI95prop	CI99prop
Known p Version				
0.1	0.099	0.463	0.93	0.99
0.25	0.051	0.293	0.94	0.97
0.5	0.026	0.236	0.95	0.98
0.75	0.007	0.215	0.91	0.98
0.9	0.001	0.236	0.90	0.95
LRT Version				
0.1	0.131	0.592	0.930	0.984
0.25	0.067	0.367	0.937	0.970
0.5	0.033	0.281	0.941	0.978
0.75	0.011	0.243	0.909	0.969
0.9	0.003	0.260	0.878	0.936