

QUANTILE REGRESSION METHODS FOR REFERENCE GROWTH CHARTS

YING WEI, ANNELI PERE, ROGER KOENKER, AND XUMING HE

ABSTRACT. Estimation of reference growth curves for children's height and weight has traditionally relied on normal theory to construct families of quantile curves based on samples from the reference population. Age-specific parametric transformation has been used to significantly broaden the applicability of these normal theory methods. Nonparametric quantile regression methods offer a complementary strategy for estimating conditional quantile functions. We compare estimated reference curves for height using the penalized likelihood approach of Cole and Green (1992) with quantile regression curves based on data used for modern Finnish reference charts. An advantage of the quantile regression approach is that it is relatively easy to incorporate prior growth and other covariates into the analysis of longitudinal growth data. Quantile specific autoregressive models for unequally spaced measurements are introduced and their application to diagnostic screening is illustrated.

1. INTRODUCTION

Anthropometric methods for constructing reference growth charts were conceived by Quetelet in the 19th century, and have experienced a vigorous subsequent development. Charts describing the dependence of height, weight, head circumference and a variety of other physical characteristics on age are now in widespread use as screening tools for disease and as reference standards for group health and economic status. The typical growth chart depicts a family of curves representing a few selected quantiles of the distribution of some physical characteristic of the reference population as a function of age.

Since Quetelet, reference growth curves have typically been constructed based on the assumption that heights, and other similar measurements, are normally distributed. Age specific mean and standard deviation curves, say $\mu(t)$ and $\sigma(t)$, are estimated and any chosen quantile curve for a $\tau \in [0, 1]$ can then be constructed as,

$$\hat{Q}(\tau|t) = \hat{\mu}(t) + \hat{\sigma}(t)\Phi^{-1}(\tau)$$

Version December 9, 2004. Corresponding author: Roger Koenker, Department of Economics, University of Illinois, Champaign, Illinois, 61820; (email: rkoenker@uiuc.edu). This research was partially supported by NSF grants DMS-01-02411 and SES-02-40781. The authors would like to thank Steve Portnoy for helpful comments.

where $\Phi^{-1}(\tau)$ denotes the inverse of the standard normal distribution function. Provided that the population is normally distributed at each age, these curves should split the population into two parts with the proportion τ lying below the curve, and the proportion $1 - \tau$ above the curve.

Although adult heights in reasonably homogeneous populations are known to be quite close to normal, children's heights can be quite non-normally distributed. Weights and other physical characteristics are potentially even more problematic. To account for this, several proposals have been made for age-specific transformations to normality. The most successful of these proposals has been the LMS, or $\lambda\mu\sigma$, approach of Cole (1988) based on the power transformation of Box and Cox (1964). Cole proposed the model,

$$Q(\tau|t) = \mu(t)(1 + \lambda(t)\sigma(t)\Phi^{-1}(\tau))^{1/\lambda(t)},$$

effectively assuming that after transformation of the measurements, $Y(t)$ to their standardized values,

$$Z(t) = \frac{(Y(t)/\mu(t))^{\lambda(t)} - 1}{\lambda(t)\sigma(t)},$$

the $Z(t)$'s would be normally distributed. The functions $\{\lambda(t), \mu(t), \sigma(t)\}$ were assumed to evolve smoothly with age. To impose this smoothness, Green in his discussion of Cole(1988), proposed estimating the three functions by minimizing the penalized log likelihood,

$$\ell(\lambda, \mu, \sigma) - \nu_\lambda \int (\lambda''(t))^2 dt - \nu_\mu \int (\mu''(t))^2 dt - \nu_\sigma \int (\sigma''(t))^2 dt,$$

where $\ell(\lambda, \mu, \sigma)$ denotes the Box-Cox log-likelihood,

$$\ell(\lambda, \mu, \sigma) = \sum_{i=1}^n [\lambda(t_i) \log(Y(t_i)/\mu(t_i)) - \log \sigma(t_i) - \frac{1}{2} Z^2(t_i)],$$

and the parameters $(\nu_\lambda, \nu_\mu, \nu_\sigma)$ serve to control the degree of smoothness of the three functions. A concise description of the fitting procedure, can be found in Carey et. al. (2004).

Inevitably, doubts may arise about the ability of any one transformation method to achieve its promised normality over the full range of relevant ages. Facing up to such doubts, it would seem desirable to consider alternative methods of estimating quantile reference curves that impose less stringent global hypotheses on the form of the conditional distributions. In the discussion of Cole (1988), both D.R. Cox and M.C. Jones suggested that one way to accomplish this objective would be to estimate a family of conditional quantile functions by solving nonparametric quantile regression problems of the form,

$$\min_{g \in \mathcal{G}} \sum_{i=1}^n \rho_\tau(Y(t_i) - g(t_i)),$$

subject to a smoothness requirement on the domain of the candidate functions, \mathcal{G} . Here the function ρ_τ denotes the simple piecewise linear function,

$$\rho_\tau(u) = u(\tau - I(u < 0)) = \begin{cases} \tau u & u \geq 0 \\ (\tau - 1)u & u < 0 \end{cases}.$$

This piecewise linear form of the objective function has the effect of enforcing a balance between the the number of observations lying above and below the fitted curve. In the simplest instance, when g is required to be a constant function, minimizing this objective requires that $n\tau$ exceeds the number of Y_i 's strictly less than the optimal value \hat{g} , and that $n\tau$ must be less than the number of Y_i 's less than or equal to \hat{g} . This is precisely the condition that, \hat{g} must be a τ th sample quantile of the Y_i 's.

Koenker and Bassett (1978) proposed extending this optimization interpretation of the ordinary sample quantiles to the estimation of linear parametric models for conditional quantile *functions*. When we minimize the sum of squared errors $\sum_{i=1}^n (Y_i - \xi)^2$ over $\xi \in \mathbb{R}$ we obtain the sample mean, $\hat{\xi} = \bar{Y}$, as an estimate of the unconditional mean. Minimizing the sum of squares, $\sum_{i=1}^n (Y_i - x_i^\top \beta)^2$ over $\beta \in \mathbb{R}^p$, yields an estimate of the conditional mean function, $g(x) \equiv E(Y|x) = x^\top \beta$. Similarly, minimizing $\sum_{i=1}^n \rho_\tau(Y_i - \xi)$ yields the unconditional τ th sample quantile, and minimizing

$$(1.1) \quad \sum_{i=1}^n \rho_\tau(Y_i - x_i^\top \beta)$$

with respect to the p -dimensional parameter β yields an estimate of the τ th conditional quantile function of Y given the covariate vector, x .

For reference growth charts it is convenient to parameterize the conditional quantile functions as linear combinations of a few fixed basis functions. B-splines are particularly convenient for this purpose. Given a choice of knots for the B-splines, estimation of the growth curves is a straightforward exercise in parametric linear quantile regression with $x_{ij} = \varphi_j(t_i)$ where φ_j is the j th function in the B-spline basis. Solutions to such problems are linear programs and can be computed efficiently, even for very large datasets as described in detail in Koenker (2005).

Recent work on reference growth curves, notably Cole (1994) and Carey et. al. (2004), has emphasized the value of accounting for other covariates in addition to age. Growth history is particularly relevant, but parental characteristics and a variety of other factors may be considered. Another advantage of the quantile regression approach to estimation of growth curves is that it is relatively easy to incorporate new covariates. The primary objective of the present paper is to illustrate how this can be accomplished. Our data, described by Sorva et. al. (1990) and Pere (2000), provides the basis for the modern Finnish reference growth charts.

Before turning to the longitudinal aspects of our analysis, we will briefly introduce the methods and offer some comparisons of their performance with Cole's LMS methods in the context of conventional cross-sectional reference growth charts.

2. DATA

Our data consists of longitudinal measurements on height and weight for 2514 Finnish children. Supine length, rather than height, was measured for infants less than two years of age. As described in greater detail in Pere (2000) the data has been edited to remove a small proportion of children with low or missing birth weight, twins or otherwise suspicious records. After editing, there are 1143 boys and 1162 girls, all full term, healthy, singleton births with between 3 and 44 measurements per child. Infants were measured roughly monthly before the age of two, and annually or biannually thereafter. On average about 20 measurements of height and weight were made between the ages of 0 and 20.

The data was collected retrospectively from health centers and schools. There are two distinct cohorts: one consisting of 1096 children born between 1954 and 1962 (94% between 1959 and 1961), the other of 1209 children born between 1968 and 1972. The former group was followed until the age of 19, the latter group until age 13. The two cohorts constitute more than 0.5 percent of Finns in the respective cohorts.

3. UNCONDITIONAL GROWTH CURVES: A COMPARISON OF METHODS

We will distinguish two general types of reference curves. *Unconditional growth curves* will refer to curves that depend solely on age; *conditional growth curves*, or longitudinal growth curves will connote curves that explicitly account for growth history, and possibly other covariates. In this section we will concentrate on the simpler, more classical problem of estimating unconditional growth curves. Having established a base of comparison in the unconditional setting, we will then turn to the problem of estimating longitudinal curves in the next section. Although we observe multiple measurements $\{Y_i(t_j) : j = 1, 2, \dots, J_i\}$ on each child, we will ignore the longitudinal aspect of the data in this section, treating the sample as if we observed independent measurements on different children.

Our comparison will focus on two methods: the Cole and Green (1992) LMS method as implemented by Carey (2002) in the R (R Development Core Team (2004)) package `lmsqreg`, and a quantile regression method as implemented in the R package `quantreg`, of Koenker (2004) using a linear (B-spline) representation of the curves. R is a public domain language for data analysis sustained by the R Development Core Team (2004). Our comparison complements the recent investigations of Carey et al (2004) and Gaunnoun et al (2002).

3.1. Criteria for Smoothness. Any nonparametric curve estimation method requires some device to control the degree of smoothness of the fitted functions. For the LMS method this control is provided by the parameters $\nu = (\nu_\lambda, \nu_\mu, \nu_\sigma)$. Following recent practice for similar spline smoothing problems, see e.g. Green and Silverman

(1994), it is convenient to represent the degree of smoothness associated with a particular choice of ν in terms of its “effective degrees of freedom.” This quantity can be interpreted as the dimensionality of the fitted function and is measured by computing the trace of the pseudo-projection matrix defining the estimator. Thus, when we report that $\nu = (7, 10, 7)$ for a particular fit, it means that the functions $\hat{\lambda}(t)$, $\hat{\mu}(t)$ and $\hat{\sigma}(t)$ have, respectively, dimension 7, 10 and 7. In contrast to the classical linear regression setting where the trace of least squares projection matrix is an integer equal to the rank of the design matrix, for smoothing splines the trace of the corresponding linear operator is a real number so the dimensionality interpretation should be taken with a grain of salt. Our implementation of the QR method employs the fixed set of B-spline basis functions illustrated in Figure 3.1. Linear combinations of these functions provide a simple and quite flexible model for the entire growth curve from birth to adulthood, as we will see. After inclusion of boundary knots these models have parametric dimension 16.




FIGURE 3.1. Cubic B-Spline Basis functions for the interior knot sequence $\{0.2, 0.5, 1.0, 1.5, 2.0, 5.0, 8.0, 10.0, 11.5, 13.0, 14.5, 16.0\}$. Spacing of the interior knots is dictated by the need for more flexibility during infancy and in the pubertal growth spurt period.

3.2. Comparison of Quantile Growth Curves. To provide a visual comparison of the LMS and QR methods we present in Figures 3.2-5 the results for both methods distinguishing boys from girls and infants from older children. In each plot the central region shows three distinct families of curves: the solid black lines are the quantile regression (QR) curves estimated with the B-spline basis appearing in Figure 3.1.

The solid grey lines represent the LMS estimates with $\nu = (7, 10, 7)$, the dashed lines indicate a higher dimensional LMS fit with $\nu = (22, 25, 22)$. The vertical column of three plots on the right side of the figure shows the fitted λ , μ and σ curves corresponding to the two LMS fits. The more parsimonious LMS fit corresponds to the default choice in the Carey `lmsqreg` package.

Generally there is reasonable agreement among the three families. However, especially for infants the more parsimonious LMS curves lack the flexibility of their competitors. Similarly, in the pubertal growth spurt the lower dimensional LMS curves appear to smooth over the curvature seen in the other estimates. There is excellent agreement throughout between the QR curves and the more profligate of the two LMS estimates. Further fine tuning could be expected to reduce the effective dimension of this LMS fit without damaging the fit. Cole, Freeman and Preece (1990) and Pan and Cole (2004) describe a modification of the LMS method in which age is rescaled based on a preliminary estimate of the M -function that leads to somewhat more parsimonious fitting.

Examination of the estimated λ , μ and σ curves for the two LMS fits reveals strong agreement on μ , but substantial differences about λ and σ . The variability of $\hat{\lambda}(t)$ is particularly striking. For infants, $\hat{\lambda}(t)$ takes values around one at birth, rises to two at the age of one, and then falls to nearly zero, indicating a log transformation at age 2.5. For older children $\lambda(t)$ is also somewhat unstable. We find it difficult to interpret the variability in $\hat{\lambda}(t)$ and $\hat{\sigma}(t)$ we see in these plots, particularly in the higher dimensional of the two LMS fits. At the same time it seems evident that the greater flexibility of the larger LMS model is essential to capture important features of the data. This point is reinforced by examining the differences the three estimates of the velocity of growth.

3.3. Comparison of Quantile Velocity Curves. Figure 3.6 illustrates estimated growth velocity curves for our four groups and for five quantiles. These curves are simply the time derivatives of the corresponding quantile functions illustrated in the previous figures. Velocity curves as estimated from cross-sectional data like these are obviously not reasonable substitutes for the longitudinal data analysis to be presented in Section 5, they are introduced here only to provide another perspective on the smoothing choices implicit in the previous figures. Again the solid gray curves represent the QR estimates, the dotted curves are the more parsimonious LMS fits, and the dashed curves are the more profligate LMS fit. In this figure it is even more apparent that the lower dimensional LMS fit is oversmoothing the infant growth experience. For older boys the pubertal growth spurt is significantly attenuated by the lower LMS fit, while the higher LMS and QR estimates match very closely. For girls the agreement among the three fits is somewhat better, but there is still some attenuation in puberty. Although there is excellent agreement between the QR and and high LMS fits in capturing the general shape and level of velocity, the pronounced

oscillation of the LMS curves seems to be an inevitable consequence of increasing the effective dimension of the LMS model.




FIGURE 3.2. Comparison of LMS and QR Growth Curves: The figure illustrates three families of growth curves, two estimated with the LMS methods of Cole and Green, and one using quantile regression methods.

3.4. Comparison of Conditional Densities of Height. Having differentiated with respect to age to obtain the velocity curves, a natural next step is to explore derivatives with respect to the quantile parameter, τ . For univariate quantile functions, differentiating the identity, $F(F^{-1}(t)) = t$, yields

$$\frac{d}{dt}Q(t) = \frac{d}{dt}F^{-1}(t) = 1/f(F^{-1}(t)),$$

that is, the derivative of the quantile function is just the reciprocal of the quantile density function. Rather than examining the reciprocal of the density function, we adopt the more conventional strategy of investigating the age-specific density functions implied by our estimated growth models. For the LMS estimates we have a complete parametric description of the model, so it is straightforward to compute the implied density functions at each age. For the QR estimates we have a nonparametric




FIGURE 3.3. Comparison of LMS and QR Growth Curves: The figure illustrates three families of growth curves, two estimated with the LMS methods of Cole and Green, and one using quantile regression methods.

estimate, and we proceed as follows. The QR model is estimated on a more refined grid of $\tau \in \mathcal{T}$, with spacing roughly 0.01, and then smoothed slightly by regressing the equally spaced τ 's on a B-spline expansion of $\hat{Q}(\tau|t)$ to produce a smoothed conditional distribution function at each age. The estimated density,

$$\hat{f}_{Y|t}(y|t) = \Delta\hat{\tau} / \Delta\hat{Q}(\tau|t)$$

is then computed, where $\Delta\hat{Q}(\tau|t) = \sum_{j=1}^p \varphi_j(t)(\hat{\beta}_j(\tau_k) - \hat{\beta}_j(\tau_{k-1}))$. This estimate is then plotted against $\hat{Q}(\tau|t)$ to obtain a conditional density estimate.

Initially, as an exploratory inquiry estimates were computed for ages 1-18, at annual intervals. For ages greater than two the density estimates produced by the QR and LMS methods corresponded quite closely, however for one year olds a rather surprising discrepancy appeared that we would like to describe in more detail. In Figure 3.7 we illustrate several estimates of the conditional density of height based on the Finnish reference sample at ages .5, 1, and 1.5. In each case we selected a subsample of subjects whose measurement occurred with three days of the target age. Since a

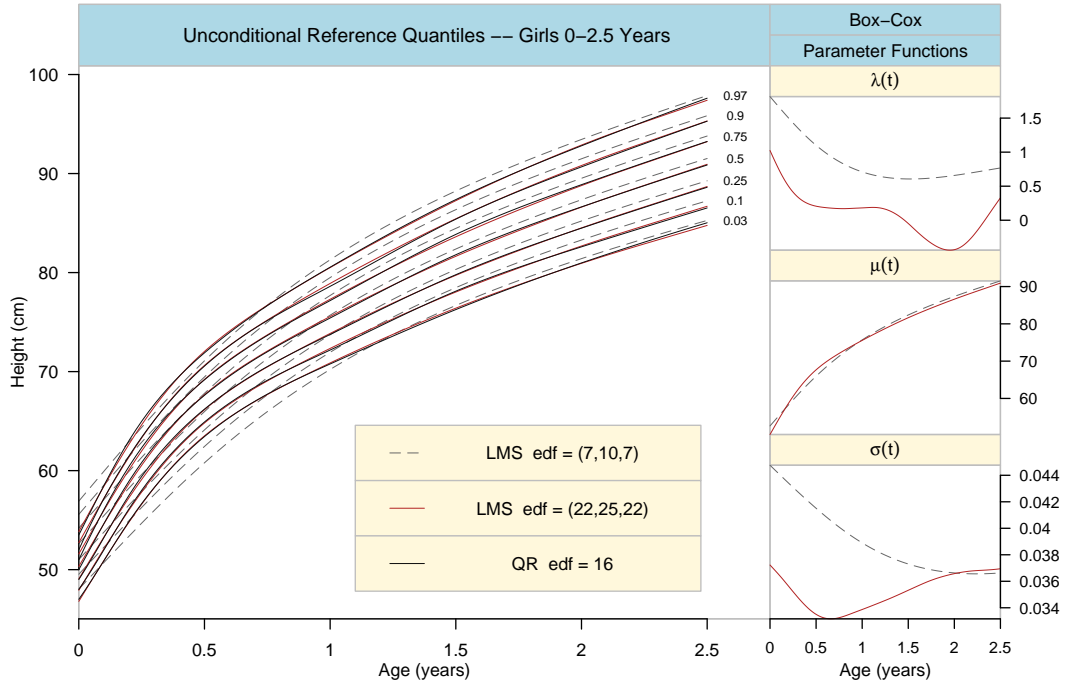


FIGURE 3.4. Comparison of LMS and QR Growth Curves: The figure illustrates three families of growth curves, two estimated with the LMS methods of Cole and Green, and one using quantile regression methods.

large fraction of children in the sample were measured within three days of their first birthday, the sample size for the one year olds is considerably larger than for the adjacent ages.

The striking feature of this figure is the pronounced bimodality in heights of one year olds as revealed by the QR estimates. The LMS estimates *assume* a unimodal form for the conditional density, and therefore are capable only to produce such densities. Is the bimodality of the QR estimates some sort of artifact of the fitting method? We believe this is not the case. To explore this further, we have estimated age specific densities using the original sample observations in narrow (± 3 days) age bands centered at age one. Both histogram estimates and more sophisticated estimates using the adaptive kernel methods of Silverman (1986, p. 101) are shown. Sample sizes for these direct estimates are shown at the top of each of plot panels. For one year olds these estimates still exhibit the same bimodality seen in the QR estimates. For the adjacent age groups, the smaller sample sizes make it more difficult to draw firm conclusions, but the evidence suggests that bimodality is confined quite narrowly to children at age one.

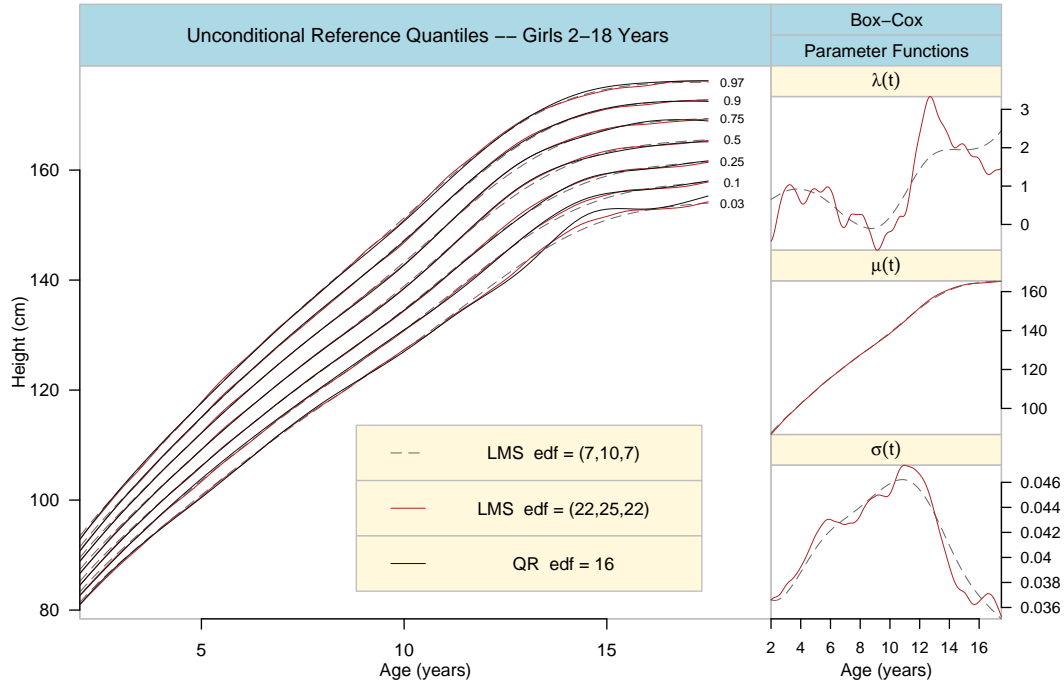


FIGURE 3.5. Comparison of LMS and QR Growth Curves: The figure illustrates three families of growth curves, two estimated with the LMS methods of Cole and Green, and one using quantile regression methods.

What might account for such a surprising violation of the conventional presumption of unimodal distribution of heights? An intriguing feature of the Finnish reference sample is that roughly half of the sample is drawn from children born in 1960, and half from children born around 1970. A natural hypothesis for our observation of bimodality in the full sample is that it may be attributed to this cohort effect. To explore this hypothesis we display in Figure 3.8 separate density estimates, using the adaptive kernel method, for the two cohorts, superimposed on the pooled estimate from the full sample.

For girls the cohort hypothesis is remarkably successful in accounting for the observed bimodality in the full sample. We see that girls born around 1960 have at age one a prominent mode that coincides very closely with the lower peak of the pooled estimate, while girls born around 1970 have a mode that is shifted to the right coinciding with the higher peak of the pooled density. Boys are not quite so cooperative. For both the 1960 and 1970 cohorts of boys we find more pronounced bimodality in the estimated cohort densities. Note, however that the two cohorts exhibit asymmetric bimodality: the 1960 boys have a larger peak coinciding with the lower of the two




FIGURE 3.6. Comparison of LMS and QR Growth Velocity Curves: The figure illustrates three families of growth curves of the prior figures, this times representing the estimated velocity of growth.

peaks of the combined sample, while the 1970 boys have a larger peak that coincides with the higher of the two peaks of the combined sample. Thus, the boys weakly confirm the cohort interpretation suggested by our findings for girls. For the moment we are unable to offer any better explanation of these puzzling findings, however we would like to stress that we would not have noticed this curious bimodality effect had

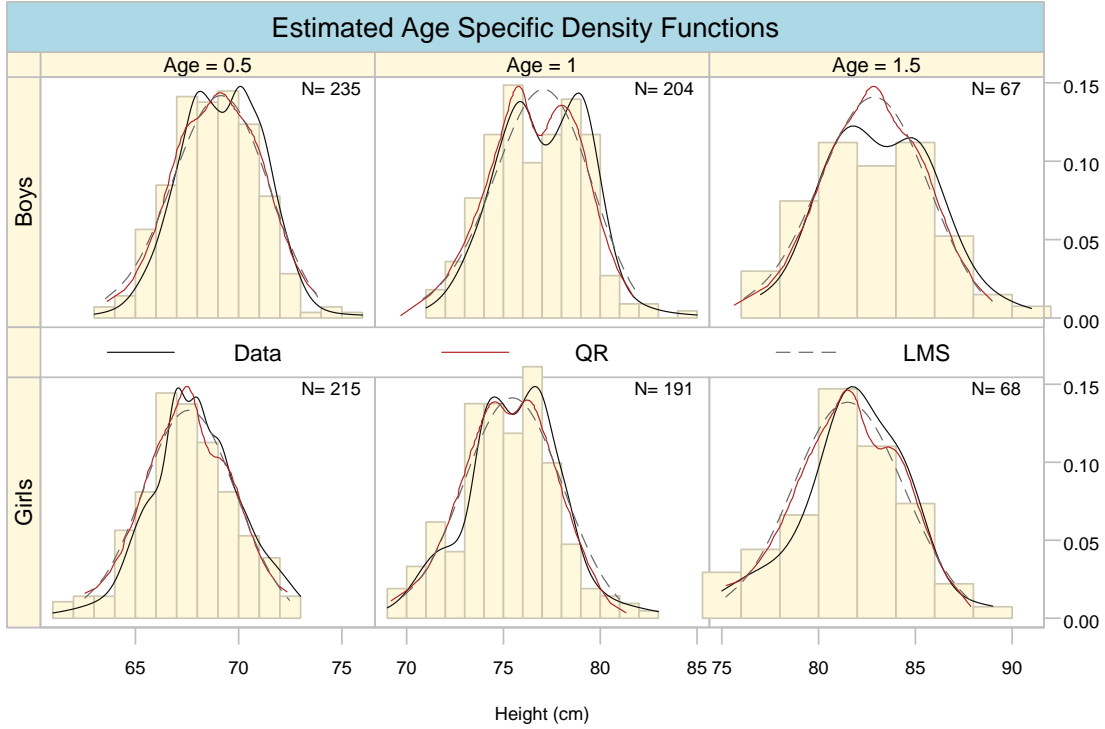


FIGURE 3.7. Comparison of Conditional Density Estimates: Quantile regression estimates of the conditional density of heights at ages .5, 1, and 1.5 years reveals a bimodality in the height distribution of one year olds. This bimodality is not apparent in the LMS estimates, which are unimodal by assumption. Histogram estimates based on the raw data and corresponding adaptive kernel density estimates based on the raw data confirm the bimodality.

we not considered the quantile regression estimates. Transformation models generally impose quite stringent conditions on the shape of the conditional densities; in contrast the quantile regression approach is considerably more flexible. Estimates of the conditional quantile functions at each τ are not constrained by an underlying global model, they are estimated quite independently, consequently they are freer to reflect unusual underlying features of the data.

4. CONDITIONAL GROWTH MODELS BASED ON LONGITUDINAL DATA

Unconditional reference growth curves provide a valuable snapshot of the dispersion of heights at various ages, but for physicians interested in assessing unusual growth, prior growth history and other covariates can offer crucial additional information. One of the main advantages of the QR approach to estimating reference growth

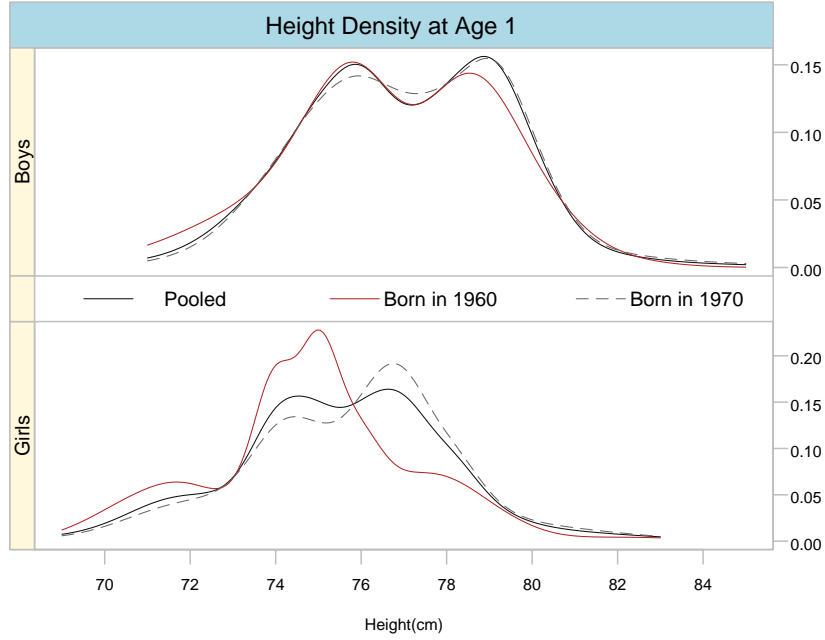


FIGURE 3.8. Cohort Effect in Distribution of Heights of Finnish One Year Olds: The figure compares adaptive kernel estimates of the height density of one year olds in the full reference sample with estimates based on splitting the sample into two cohorts.

curves is that it is relatively easy to incorporate such refinements into the modeling and estimation framework. In this section we will describe some initial steps in this direction using the Finnish height data, and illustrate the role of longitudinal models in screening for unusual growth patterns.

A challenging aspect of most longitudinal growth data is the irregular nature of the time series observations. Suppose we observe measurements, $\{Y_i(t_{i,j}) : j = 1, \dots, T_i, i = 1, \dots, n\}$ on n individuals. It would be convenient if the measurements were taken at equally spaced time points, but this is not the case for our Finnish data, nor is it common in other in true clinical settings. To address these difficulties, we adopt a simple first order autoregression model in which the AR(1) parameter is specified as a linear function of the time gap between successive measurements,

$$\begin{aligned}
 Q_{Y_i(t_{i,j})}(\tau | t_{i,j}, Y_i(t_{i,j-1}), x_i) &= g_\tau(t_{i,j}) \\
 (4.1) \quad &+ [\alpha(\tau) + \beta(\tau)(t_{i,j} - t_{i,j-1})]Y_i(t_{i,j-1}) + x_i^\top \gamma(\tau).
 \end{aligned}$$

The τ th conditional quantile function is additively decomposed into a nonparametric trend component, an AR(1) component, g_τ , and a “partially linear” component in the covariate vector x_i . The linearity of the AR(1) coefficient in the measurement gaps is a convenient approximation and could obviously be generalized in a variety

of ways; see Wei (2004) for further details and alternative formulations. We will restrict attention in our application to a single additional covariate, x_i , average parental height.

Following the approach of the previous section we will express the nonparametric growth trend component as a linear expansion in B-splines. This formulation yields a family of linear quantile regression problems that can be easily estimated by standard quantile regression algorithms using linear programming methods.

Rather than assuming that the model (4.1) holds globally across the entire age spectrum, we estimate separate versions of the model for infants, ages 0-2, and young children, ages 6-10. In Tables 4.1 and 4.2 we report estimates of the parametric components of the model (4.1) for these two groups. Estimated standard errors of the parametric estimates obtained by $B = 500$ bootstrap replications reported in parentheses. Bootstrapping is done by sampling the entire longitudinal record of randomly selected children from the sample, a strategy that preserves the dependence structure within individual time-series.

τ	Boys			Girls		
	$\hat{\alpha}(\tau)$	$\hat{\beta}(\tau)$	$\hat{\gamma}(\tau)$	$\hat{\alpha}(\tau)$	$\hat{\beta}(\tau)$	$\hat{\gamma}(\tau)$
0.03	0.845 (0.020)	0.147 (0.011)	0.024 (0.011)	0.809 (0.024)	0.135 (0.011)	0.042 (0.010)
0.1	0.787 (0.020)	0.159 (0.007)	0.036 (0.007)	0.757 (0.022)	0.153 (0.007)	0.054 (0.009)
0.25	0.725 (0.019)	0.170 (0.006)	0.051 (0.009)	0.685 (0.021)	0.163 (0.006)	0.061 (0.008)
0.5	0.635 (0.025)	0.173 (0.009)	0.060 (0.013)	0.612 (0.027)	0.175 (0.008)	0.070 (0.009)
0.75	0.483 (0.029)	0.187 (0.009)	0.063 (0.017)	0.457 (0.027)	0.183 (0.012)	0.094 (0.015)
0.9	0.422 (0.024)	0.213 (0.016)	0.070 (0.017)	0.411 (0.030)	0.201 (0.015)	0.100 (0.018)
0.97	0.383 (0.024)	0.214 (0.016)	0.077 (0.018)	0.400 (0.038)	0.232 (0.024)	0.086 (0.027)

TABLE 4.1. Parametric Components of the Infant Conditional Growth Model: Estimates of the autoregressive parameters α and β , and the parental height effect, γ are given for the seven indicated quantiles. Bootstrapped standard errors are given in parentheses.

The estimated autoregression effect for infants reported in Table 4.1 declines quite dramatically as we move up through the conditional distribution of height. In the lower tail dependence on prior height is quite strong indicating that infants in the lower tail of the height distribution have a steeper growth profile, while infants in the upper tail have a much flatter profile. This is consistent with a “catch-up” hypothesis for smaller infants, and is clearly inconsistent with conventional AR specifications that posit iid errors and a constant AR slope effect. The effect of parental height is weaker

τ	Boys			Girls		
	$\hat{\alpha}(\tau)$	$\hat{\beta}(\tau)$	$\hat{\gamma}(\tau)$	$\hat{\alpha}(\tau)$	$\hat{\beta}(\tau)$	$\hat{\gamma}(\tau)$
0.03	0.976 (0.010)	0.036 (0.002)	0.011 (0.013)	0.993 (0.012)	0.033 (0.002)	0.006 (0.015)
0.1	0.980 (0.005)	0.039 (0.001)	0.022 (0.007)	0.989 (0.006)	0.039 (0.001)	0.008 (0.007)
0.25	0.978 (0.006)	0.042 (0.001)	0.021 (0.006)	0.986 (0.005)	0.042 (0.001)	0.019 (0.006)
0.5	0.984 (0.004)	0.045 (0.001)	0.019 (0.004)	0.984 (0.007)	0.045 (0.001)	0.022 (0.006)
0.75	0.990 (0.004)	0.047 (0.001)	0.014 (0.006)	0.985 (0.007)	0.050 (0.001)	0.016 (0.006)
0.9	0.987 (0.009)	0.049 (0.001)	0.012 (0.009)	0.984 (0.008)	0.052 (0.001)	0.002 (0.012)
0.97	0.980 (0.014)	0.050 (0.002)	0.023 (0.015)	0.982 (0.013)	0.053 (0.001)	0.021 (0.018)

TABLE 4.2. Parametric Components of the Childrens' Conditional Growth Model: Estimates of the autoregressive parameters α and β , and the parental height effect, γ are given for the seven indicated quantiles. Bootstrapped standard errors are given in parentheses.

in the lower tail, but more strongly significant for both boys and girls at the first decile and above.

For 6 to 10 year olds the results in Table 4.2 exhibit a much more consistent AR pattern over the quantiles. Both boys and girls have mean AR(1) coefficient, at mean measurement spacing of about 1 year, of about 1.02, implying that the contribution of the AR component to growth is about 2 percent per year. At this age the effect of parental height is weaker, having presumably been already incorporated into the autoregressive effect of prior height. Nonetheless, there is a significant effect of parental height between the first and third quartiles.

5. SCREENING INDIVIDUAL GROWTH PATHS

To illustrate the diagnostic usefulness of the longitudinal form of the reference growth curves we consider screening for unusual individual growth experience. As stressed by Cole (1994) and others, standard reference growth charts are rather unsatisfactory as a screening device since they fail to account for growth history and other potentially relevant covariates. The autoregressive models for unequally spaced measurements provide a flexible way to incorporate this longitudinal information. As we shall illustrate, screening based on unconditional growth curves is prone to both false positive and false negative decisions.

In Figure 5.1 we illustrate the family of unconditional growth curves for boys ages 6 to 10, as estimated by the quantile regression methods described in Section 3. The curves are identical to those appearing in Figure 3.3, restricted to the age range 6




FIGURE 5.1. Screening Individual Growth Paths: The figure illustrates two growth records for individuals between the ages of 6 and 10, superimposed on the estimated unconditional growth curves from the quantile regression approach. Subject 89 is consistently taller than most of his peers and maintains a steady growth pattern over the entire period. Subject 225 is initially somewhat below median height, but a pause in growth between ages 8 and 8.5 shifts his trajectory below the first quartile. The solid points for each subject denote the observation at which we consider the screening decision.

to 10. Superimposed on the figure are the observed height measurements of two of the subjects in the Finnish reference sample. Both subjects are male, subject 89 indicated by the open circles was taller than most of his peers by age 6, and grew steadily thereafter with increments of about 3 cm every 6 months. Subject 225 was only slightly below median height at age 6, again growing quite steadily until age 8, gaining 2.5 to 3 cm between observed measurements at roughly six month intervals. See Table 5.1 for more precise details on the measurements. At age 8.51 subject 225 is reported to have gained only 0.5 cm since his last measurement at age 7.98. After this measurement he falls somewhat below the first quartile reference curve. How unusual are these two cases after consideration of prior growth history and parental height?

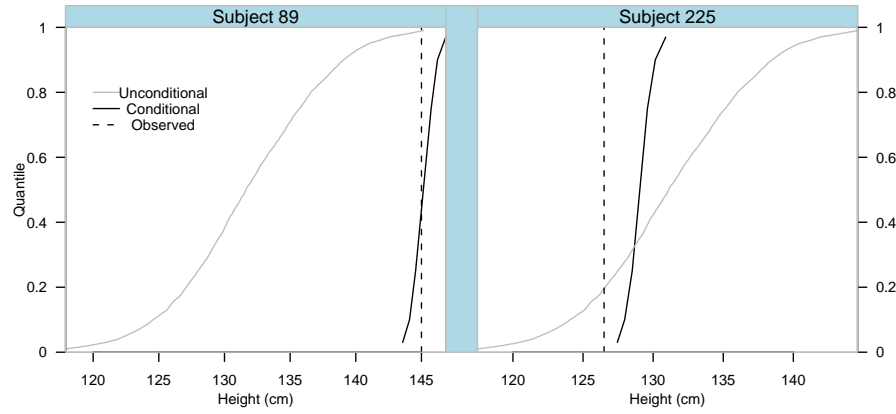


FIGURE 5.2. Conditional vs. Unconditional Screening: The figure illustrates conditional and unconditional predictions of the height distribution for two subjects at roughly age 8.5. Subject 89 is substantially taller than his peers, but his observed height is not at all unusual given his prior growth and parental height. In contrast, subject 225 is not unusual for his age by the standard of the unconditional distribution of heights, but based on the conditional reference standard his growth pause, growth of only 0.5 cm over the six month interval between measurements, is extremely unusual.

	Subject	Measurements							Parental Heights
		1	2	3	4	5	6	7	
Age	89	6.73	7.25	7.61	8.24	8.61	9.09	9.32	166
Height	89	132.0	136.0	139.0	143.0	145.0	148.0	152.0	180
Age	225	6.37	7.02	7.42	7.98	8.51	9.01	9.56	165
Height	225	117.0	120.5	123.5	126.0	126.5	129.0	132.0	177

TABLE 5.1. Measurements of Two Subjects: The table reports observed heights for two subjects between the ages of 6 and 10, as depicted in Figure 5.1. Ages are reported in decimal years, heights in centimeters. Parental heights of the two subjects are given in the last column.

In Figure 5.2 we illustrate the predictive distributions of both the conditional and unconditional growth models for the two subjects. The grey curves representing the estimated unconditional distributions for the two subjects at the respective ages of their next measurement, 8.61 for subject 89, and 8.51 for subject 225, are extremely

dispersed. Conditioning on the prior measurement as well as on parental height dramatically reduces the dispersion of the predictive distributions drawn as the solid black lines in the figure. For subject 89 the observed measurement of 145 cm at age 8.61, is extremely unusual with respect to the unconditional growth curve, representing the 99th percentile of the estimated unconditional distribution. But relative to the estimated conditional model the 145 cm measurement is not at all unusual, indeed it is slightly below the median of the conditional distribution. Given this subjects prior growth experience, it seems reasonable to conclude that although he is exceptionally tall there is nothing unusual about the measurement made at age 8.61. By contrast, subject 225 is quite unexceptional by the standard of the unconditional growth curves. His prior measurement of 126 cm at age 7.98 placed him slightly below the median height for boys of his age. However, conditional on prior height and parental height, his observed measurement of 126.5 cm at age 8.5 – only 0.5 cm taller than his prior measurement six months previously – is extremely unusual, well below the third percentile of the conditional distribution.

While the steady growth of subject 89 places him well within the range established by the conditional model, the deceleration in growth experienced by subject 225 is highly unusual by the same standard, and would therefore be a cause of concern and call for closer follow up.. The resumption of normal growth of this subject as indicated in Figure 5.1 by his measurements at ages 9 and 9.5 could be expected. Apparently, the pause was not sustained, nor was it a mere aberration or measurement error, since growth resumes along a significantly lower quantile curve.

6. CONCLUSION

By focusing attention on independent estimation of distinct quantile functions rather than imposing a global distributional hypothesis, quantile regression offers a flexible approach to estimating reference growth curves. For unconditional growth curve estimation based upon cross sectional data the well-established LMS method of Cole and Green (1992) and quantile regression estimates based on linear B-spline expansions both yield growth curves and associated velocity curves that successfully captured the essential features of the data. The two approaches are complementary: LMS imposes more structure and can be more stable particularly in regions where the data is sparse, QR is more flexible and therefore capable of revealing departures from underlying assumptions of parametric models.

Closer inspection of the quantile regression estimates of the unconditional growth curves revealed one surprising feature, not observable from the LMS estimates, for our reference sample of Finnish children. Age specific conditional quantile estimates based on the QR and LMS methods were quite similar at almost all ages, but the QR estimates of age specific conditional densities suggested a distinctly bimodal form for one-year olds. Corresponding LMS estimates are, *by assumption*, unimodal. Adaptive kernel density estimation using subsamples of children near the age of one supported

the plausibility of the bimodality. Further investigation suggested that the bimodality may be attributable to a cohort effect. The full sample is comprised in roughly equal parts of a group of children born around 1960, and a group born around 1970. Splitting the sample into these two groups and separately estimating the density of height for one year olds showed, especially for girls, a larger mode for the later group and a smaller mode for the earlier group. These findings illustrate an advantage that the greater flexibility of the quantile regression approach brings to the reference growth curve problem. This more skeptical attitude toward Gaussian assumptions would be healthy in many biomedical statistical applications – even heights can reveal interesting surprises.

Another compelling motivation for the quantile regression approach to estimating reference growth curves lies in the ability to extend the conventional unconditional models depending only on the age of the subjects to models that incorporate prior growth and other covariates. We have explored one such model, adopting an AR(1) specification in which the autoregressive coefficient is taken as an affine function of the spacing between the irregular successive measurements. Parental height is incorporated into the model as an additional covariate with a linear effect. By conditioning on prior growth and parental height we obtain a more refined diagnostic tool; subjects that are unusual by the standards of the conventional unconditional reference curves can be reassuringly normal, while subjects that seem quite unexceptional conditioning only on age, may appear highly unusual after further conditioning. Estimation of these longitudinal models is quite straightforward given existing quantile regression software.

References

- BOX, G.E.P. AND D.R. COX, (1964) An analysis of Transformations, *J. Royal Stat. Soc. (B)*, 26, 211-252.
- CAREY, V. J. (2002) LMSqreg: An R package for Cole-Green Reference Centile Curves, <http://www.biostat.harvard.edu/~carey>.
- CAREY, V. J., YONG, F. H., FRENKEL, L. M. AND MCKINNEY, R. M. (2004), Growth velocity assessment in paediatric AIDS: smoothing, penalized quantile regression and the definition of growth failure. *Statistics in Medicine*, 23, 509-526.
- COLE, T. J. (1988), Fitting smoothed centile curves to reference data. *Journal of the Royal Statistics Society, Series A, General*, 151, 385-418.
- COLE, T. J. (1994), Growth charts for both cross-sectional and longitudinal data. *Statistics in Medicine*, 13, 2477-2492.
- COLE, T. J. AND P.J. GREEN (1992), Smoothing Reference Centile Curves: The LMS Method and Penalized Likelihood, *Statistics in Medicine*, 11, 1305–1319.

GANNOUN, A., GIRARD, S., CUINOT, C. AND SARACCO, J. (2002), Reference curves based on non-parametric quantile regression. *Statistics in Medicine*, 21, 3119-3135.

GREEN, P. J. AND SILVERMAN, B. W., (1994) *Nonparametric Regression and Generalized Linear Models: A Roughness Penalty Approach*, Chapman Hall: New York.

KOENKER, R. AND BASSETT, G. (1978), Regression quantiles. *Econometrica*, 46, 33-50.

KOENKER, R. (2004), Quantreg: An R package for quantile regression and related methods, <http://cran.r-project.org>.

KOENKER, R. (2005), *Quantile Regression*, Cambridge U. Press.

PERE, A. (2000), Comparison of two methods of transforming height and weight to Normality. *The Annals of Human Biology*, 27, 35-45.

QUETELET, A. (1871) *Anthropométrie*, Muquardt: Brussels.

:q

SILVERMAN, B. (1986) *Density Estimation for Statistics and Data Analysis*, Chapman-Hall: New York.

SORVA, R., LANKINEN, S., TOLPPAEN, E.-M. AND PERHEENTUPA, J. (1990), Variation of growth in height and weight of children. II. After Infancy. *Acta Paediatrica Scandinavica*, 79, 498-506.

WEI, Y. (2004) *Longitudinal Growth Charts based on Semi-parametric Quantile Regression*, Ph.d. Thesis, U. of Illinois, Urbana-Champaign.

DEPARTMENT OF BIostatISTICS, COLUMBIA UNIVERSITY
E-mail address: yw2148@columbia.edu

THE HOSPITAL FOR CHILDREN AND ADOLESCENTS, UNIVERSITY OF HELSINKI
E-mail address: anneli.pere@hus.fi

DEPARTMENT OF ECONOMICS, UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN
E-mail address: rkoenker@uiuc.edu

DEPARTMENT OF STATISTICS, UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN
E-mail address: x-he@uiuc.edu