QUANTILE REGRESSION
AN INTRODUCTION

ROGER KOENKER AND KEVIN F. HALLOCK

Abstract. Quantile regression, as introduced by Koenker and Bassett (1978), may be viewed as an extension of classical least squares estimation of conditional mean models to the estimation of an ensemble of models for several conditional quantile functions. The central special case is the median regression estimator which minimizes a sum of absolute errors. Other conditional quantile functions are estimated by minimizing an asymmetrically weighted sum of absolute errors. Quantile regression methods are illustrated with applications to models for CEO pay, food expenditure, and infant birthweight.

We say that a student scores at the τth quantile of a standardized exam if he performs better than the proportion τ, of the reference group of students, and worse than the proportion $(1-\tau)$. Thus, half of students perform better than the median student, and half perform worse. Similarly, the quartiles divide the population into four segments with equal proportions of the reference population in each segment. The quintiles divide the population into 5 parts; the deciles into 10 parts. The quantiles, or percentiles, or occasionally fractiles, refer to the general case. Quantile *regression* as introduced by Koenker and Bassett (1978) seeks to extend these ideas to the estimation of conditional quantile *functions* – models in which quantiles of the conditional distribution of the response variable are expressed as functions of observed covariates.

In Figure 1 we illustrate one approach to this task based on Tukey's boxplot. Annual CEO compensation is plotted as a function of firm's market value of equity. A sample of 1660 firms was split into 10 groups of equal size according to their market capitalization. For each group of 166 firms we compute the three quartiles of CEO compensation: salary, bonus, and other compensation including stock options valued by Black-Scholes at the time of the grant. For each group the bowtie-like box represents the middle half of the salary distribution lying between the first and third quartiles. The horizontal line near the middle of each box represents the median compensation for each group of CEOs, and the notches represent an estimated confidence interval for each median estimate.The full range of the observed salaries in each group is represented by the horizontal bars at the end of the dashed "whiskers". In cases where the whiskers would extend more than three times the interquartile range, they are truncated and the remaining outlying points are indicated by open circles. The

Version: May 31, 2001. Roger Koenker is William B. McKinley Professor of Economics and Professor of Statistics and Kevin F. Hallock is Associate Professor of Economics and of Labor and Industrial Relations, University of Illinois at Urbana-Champaign. Their e-mail addresses are roger@ysidro.econ.uiuc.edu and hallock@uiuc.edu.

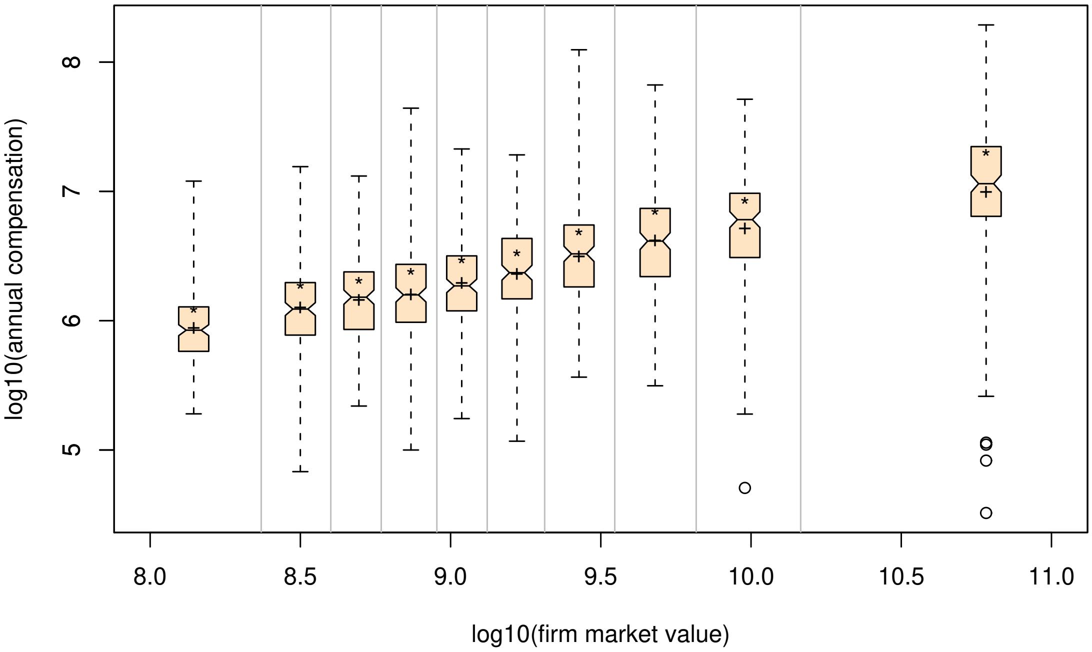

FIGURE 1. CEO Pay by Firm Size: The boxplots provide a summary of the distribution of CEO annual compensation, from EXECUCOMP in 1999, for 10 groupings of firms ranked by market capitalization. The light grey vertical lines demarcate the deciles of the firm size groupings. The upper and lower limits of the boxes represent the first and third quartiles of pay, the median for each group is represented by the horizontal bar in the middle of each box. Both scales are logarithmic base 10, so 6 on the vertical scale represents annual compensation of one million dollars. On the horizontal scale 9 represents market capitalization of one billion dollars.

mean compensation for each group is also plotted: the geometric mean as a "+" and the arithmetic mean as a "$*$".

There is a clear tendency for compensation to rise with firm size, but one can also discern several other features from the plot. Even on the log scale there is a tendency for dispersion, as measured by the interquartile range of log compensation, to increase with firm size. This effect is accentuated if we consider the upper and lower tails of the salary distribution. By characterizing the entire distribution of annual compensation for each group, the plot provides a much more complete picture than would be offered by simply plotting the group means or medians. Here we have the luxury of a moderately large sample size in each group, so we are able to take an essentially nonparametric approach. Had we had several covariates, however, our cell sizes would become rapidly depleted, and we would need to consider estimating parametric models for conditional quantiles, rather than relying on the nonparametric boxplots.

In classical linear regression we also abandon the idea of estimating separate means for grouped data as in Figure 1, and we *assume* that these means fall on a line, or some linear surface, and we estimate instead the parameters of this linear model. Least squares estimation provides a convenient method of estimating such conditional mean models. Quantile regression provides an equally convenient method for estimating models for conditional quantile functions.

Quantiles via Optimization. Quantiles seem inseparably linked to the operations of ordering and sorting the sample observations that are used to define them. So it comes as a mild surprise to observe that we can define the quantiles through a simple alternative expedient – as an optimization problem. Just as we can define the sample mean as the solution to the problem of minimizing a sum of squared residuals, we can define the median as the solution to the problem of minimizing a sum of absolute residuals. What about the other quantiles? If the symmetric absolute value function yields the median, maybe we can simply tilt the absolute value to produce an asymmetric weighting that yields the other quantiles. This "pinball logic" suggests solving

$$\min_{\xi \in \Re} \sum \rho_\tau (y_i - \xi) \tag{1}$$

where the function $\rho_\tau(\cdot)$ is illustrated in Figure 2. To see that this problem yields the sample quantiles as its solutions, it is only necessary to compute the directional derivatives of the objective function with respect to ξ, taken both from the left and from the right.[1]

Having succeeded in defining the unconditional quantiles as an optimization problem, it is easy to define conditional quantiles in an analogous fashion. Least squares regression offers a model for how to proceed. If, presented with a random sample

[1]Consider the median case with $\tau = \frac{1}{2}$. The contributions to the directional derivatives are either $+\frac{1}{2}$ for $y_i > \xi$ or $-\frac{1}{2}$ for $y_i < 0$; zero residuals are counted as $+\frac{1}{2}$ when the derivative is from the right and as $-\frac{1}{2}$ when it is from the left. If $\hat{\xi}$ is taken to be a value such that half the observations lie above $\hat{\xi}$ and half lie below, then both directional derivatives will be positive, and since the objective function is increasing in either direction $\hat{\xi}$ must be a local minimum. The objective function is a sum of convex functions, hence convex, so the local minimum is also a global one. Such a solution is clearly a median by our original definition. The same argument works for the other quantiles, but now we have asymmetric weighting of the number of observations with positive and negative residuals and this leads to solutions $\hat{\xi}(\tau)$ corresponding to the τth quantiles.

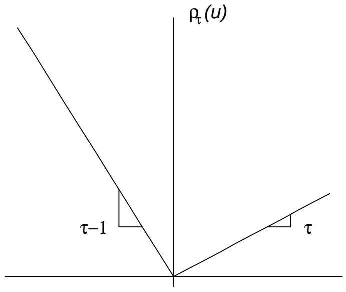

FIGURE 2. Quantile Regression ρ Function

$\{y_1, y_2, \ldots, y_n\}$, we solve

$$\min_{\mu \in \Re} \sum_{i=1}^{n} (y_i - \mu)^2, \tag{2}$$

we obtain the sample mean, an estimate of the unconditional population mean, EY. If we now replace the scalar μ by a parametric function $\mu(x, \beta)$ and solve

$$\min_{\beta \in \Re^p} \sum_{i=1}^{n} (y_i - \mu(x_i, \beta))^2 \tag{3}$$

we obtain an estimate of the conditional expectation *function* $E(Y|x)$.

In quantile regression we proceed in exactly the same way. To obtain an estimate of the conditional median function, we simply replace the scalar ξ in (1) by the parametric function $\xi(x_i, \beta)$ and set τ to $\frac{1}{2}$. Variants of this idea were proposed in the mid 18th century by Boscovich, and subsequently investigated by Laplace and Edgeworth, among others. To obtain estimates of the other conditional quantile functions we simply replace absolute values by $\rho_\tau(\cdot)$, and solve

$$\min_{\beta \in \Re^p} \sum \rho_\tau (y_i - \xi(x_i, \beta)) \tag{4}$$

The resulting minimization problem, when $\xi(x, \hat{\beta}(\tau))$ is formulated as a linear function of parameters, can be solved very efficiently by linear programming methods.

Quantile Engel Curves. To illustrate the basic ideas we briefly reconsider a classical empirical application in economics, Engel's (1857) analysis of the relationship between household food expenditure and household income. In Figure 3 we plot Engel's data taken from 235 European working class households. Superimposed on

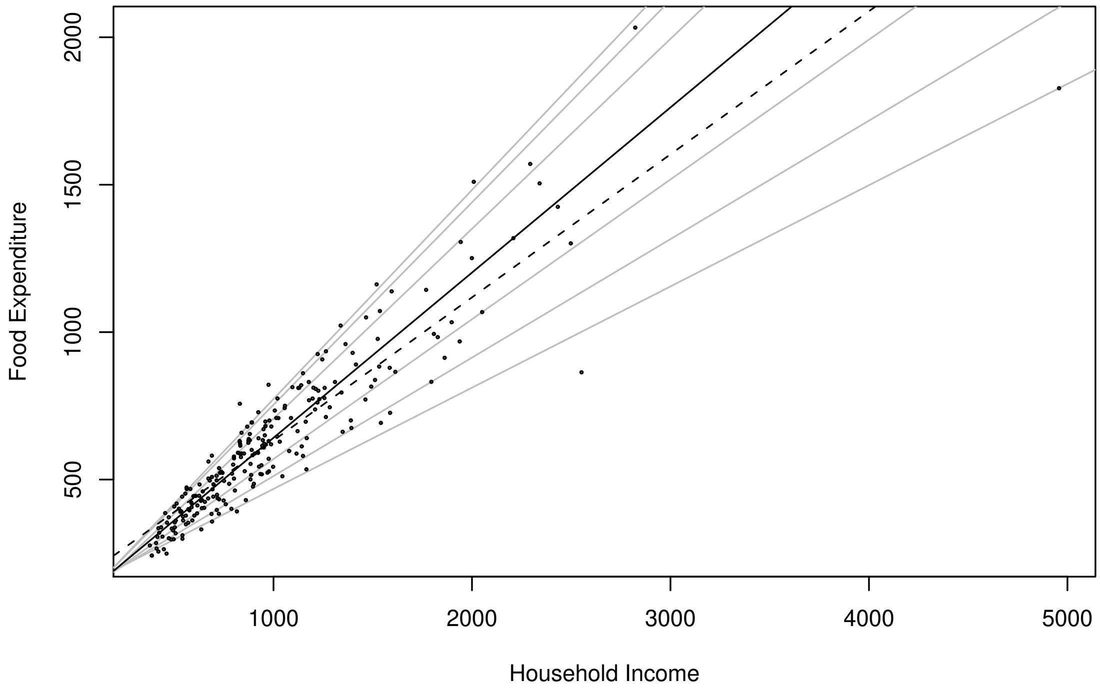

Figure 3. Engel Curves for Food: This figure plots data taken from Engel's (1857) study of the dependence of households' food expenditure on household income. Seven estimated quantile regression lines for $\tau \in \{.05, .1, .25, .5, .75, .9, .95\}$ are superimposed on the scatterplot. The median $\tau = .5$ fit is indicated by the darker solid line; the least squares estimate of the conditional mean function is indicated by the dashed line.

the plot are seven estimated quantile regression lines corresponding to the quantiles $\tau \in \{.05, .1, .25, .5, .75, .9, .95\}$. The median $\tau = .5$ fit is indicated by the darker solid line; the least squares estimate of the conditional mean function is plotted as the dashed line.

The plot clearly reveals the tendency of the *dispersion* of food expenditure to increase along with its level as household income increases. The spacing of the quantile regression lines also reveals that the conditional distribution of food expenditure is skewed to the left: the narrower spacing of the upper quantiles indicating high density and a short upper tail and the wider spacing of the lower quantiles indicating a lower density and longer lower tail.

The conditional median and mean fits are quite different in this example, a fact that is partially explained by the asymmetry of the conditional density and partially by the strong effect exerted on the least squares fit by the two unusual points with high income and low food expenditure. Note that one consequence of this nonrobustness is that the least squares fit provides a rather poor estimate of the conditional mean for the poorest households in the sample. Note that the dashed least squares line passes above *all* of the very low income observations.

We have occasionally encountered the faulty notion that something like quantile regression could be achieved by segmenting the response variable into subsets according to its unconditional distribution and then doing least squares fitting on these subsets. Clearly, this form of "truncation on the dependent variable" would yield disastrous results in the present example. In general, such strategies are doomed to failure for all the reasons so carefully laid out in Heckman's (1979) work on sample selection.

In contrast, segmenting the sample into subsets defined according to the conditioning covariates is always a valid option. Indeed such local fitting underlies all non-parametric quantile regression approaches. In the most extreme cases we have p distinct cells corresponding to different settings of the covariate vector, x, and quantile regression reduces to simply computing ordinary univariate quantiles for each of these cells. In intermediate cases we may wish to project these cell estimates onto a more parsimonious (linear) model. See *e.g.* Chamberlain (1994) and Knight, Bassett, and Tam (2000), for examples of this approach.

Another variant is the suggestion that instead of estimating linear conditional quantile models, we could instead estimate a family of binary response models for the probability that the response variable exceeded some prespecified cutoff values. This approach replaces the hypothesis of conditional quantile functions that are linear in parameters with the hypothesis that some transformation of the various probabilities of exceeding the chosen cutoffs, say the logistic, could instead be expressed as linear functions in the observed covariates. In our view the conditional quantile assumption is more natural, if only because it nests within it the iid error location shift model of classical linear regression.

Quantile Regression and Determinants of Infant Birthweight

In this section we reconsider a recent investigation by Abreveya (2001) of the impact of various demographic characteristics and maternal behavior on the birthweight of infants born in the United States. Low birthweight is known to be associated with a wide range of subsequent health problems, and has even been linked to educational attainment and eventual labor market outcomes. Consequently, there has been considerable interest in factors influencing birthweights and public policy initiatives that might prove effective in reducing the incidence of low birthweight infants.

Most of the analysis of birthweights has employed conventional least squares regression methods. However, it has been recognized that the resulting estimates of various effects on the conditional mean of birthweights were not necessarily indicative of the size and nature of these effects on the lower tail of the birthweight distribution. A more complete picture of covariate effects can be provided by estimating a family of conditional quantile functions.[2]

Our analysis is based on the June 1997 Detailed Natality Data published by the National Center for Health Statistics. Like Abreveya (2001), we limit the sample to live, singleton births, with mothers recorded as either black or white, between the ages of 18 and 45, residing in the United States. Observations with missing data for any of the variables described below were dropped from the analysis. This process yielded a sample of 198,377 babies. Birthweight, the response variable, is recorded in grams. Education of the mother is divided into four categories: less than high school, high school, some college, and college graduate. The omitted category is less than high school so coefficients may be interpreted relative to this category. The prenatal medical care of the mother is also divided into 4 categories: those with no prenatal visit, those whose first prenatal visit was in the first trimester of the pregnancy, those with first visit in the second trimester, and those with first visit in the last trimester. The omitted category is the group with a first visit in the first trimester; they constitute almost 85 percent of the sample. An indicator of whether the mother smoked during pregnancy is included in the model, as well as mother's reported average number of cigarettes smoked per day. The mother's reported weight gain during pregnancy (in pounds) is included (as a quadratic effect).

Figure 4 presents a summary of quantile regression results for this example. Unlike the Engel curve example, where we had only one covariate and the entire empirical analysis could be easily superimposed on the bivariate scatter plot of the observations, we now have 15 covariates, plus an intercept. For each of the 16 coefficients we plot the 19 distinct quantile regression estimates for τ ranging from 0.05 to 0.95 as the solid curve with filled dots. For each covariate these point estimates may be interpreted as the impact of a one unit change of the covariate on birthweight holding other covariates fixed. Thus, each of the plots have a horizontal quantile, or τ, scale and the vertical scale in grams indicates the covariate effect. The dashed line in each figure shows the ordinary least squares estimate of the conditional mean effect. The two dotted lines represent conventional 90 percent confidence interval for the least squares estimate. The shaded grey area depicts a 90 percent pointwise confidence band for the quantile regression estimates.

In the first panel of the figure the intercept of the model may be interpreted as the estimated conditional quantile function of the birthweight distribution of a girl born to an unmarried, white mother with less than a high school education, who is 27 years old and had a weight gain of 30 pounds, didn't smoke, and had her first prenatal visit in the first trimester of the pregnancy. The mother's age and weight gain are chosen

[2]In an effort to focus attention more directly on the lower tail, several studies have recently explored binary response (probit) models for the occurrence of low birthweights, conventionally defined by whether infants weigh less than 2500 grams at birth, about 5 pounds 9 ounces.

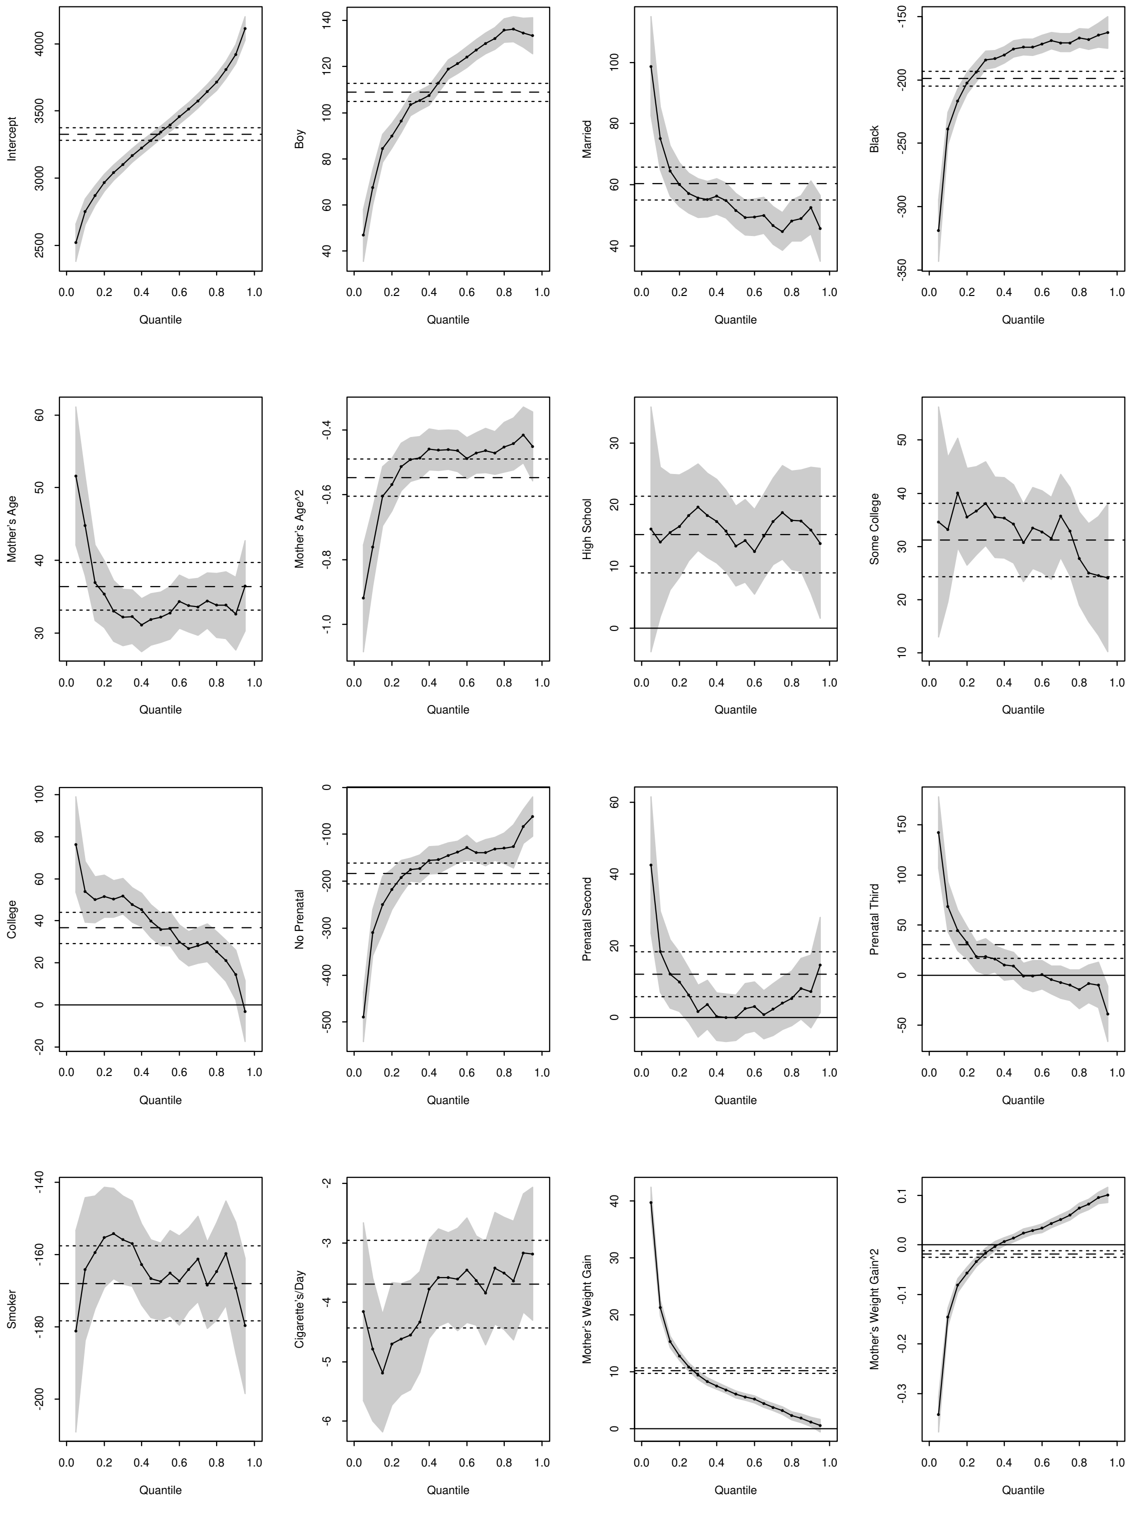

FIGURE 4. OLS and Quantile Regression Estimates for Birthweight Model

to reflect the means of these variables in the sample.[3] Note that the $\tau = .05$ quantile of the distribution for this group is just at the margin of the conventional definition of a low birthweight baby.

We will confine our discussion to only a few of the covariates. At any chosen quantile we can ask, for example, how different are the corresponding weights of boys and girls, given a specification of the other conditioning variables? The second panel answers this question. Boys are obviously larger than girls, by about 100 grams according to the OLS estimates of the mean effect, but as is clear from the quantile regression results the disparity is much smaller in the lower quantiles of the distribution and considerably larger than 100 grams in the upper tail of the distribution. For example, boys are about 45 grams larger at the 0.05 quantile but are about 130 grams larger at the 0.95 quantile. Note that the conventional least squares confidence interval does a poor job of representing this range of disparities.

The disparity between birthweights of infants born to black and white mothers is substantial particularly at the left tail of the distribution. At the 5th percentile of the conditional distribution the difference is roughly one third of a kilogram.

Mother's age enters the model as a quadratic effect, shown in the first two figures of the second row. At the lower quantiles the mother's age tends to be more concave, increasing birthweight from age 18 to about age 30, but tending to decrease birthweight when the mother's age is beyond 30. At higher quantiles this "optimal age" becomes gradually older. At the third quantile it is about 36, and at $\tau = .9$ it is almost 40.

Education beyond high school is associated with a modest increase in birthweights. High school graduation has a quite uniform effect over the whole range of the distribution of about 15 grams. This is a rare example of an effect that really does appear to exert a pure location shift effect on the conditional distribution. For this effect, the quantile regression regression results are quite consistent with the least squares results, but this is the exceptional case, not the rule.

Several of the remaining covariates are of substantial public policy interest. These include the effects of prenatal care, marital status, and smoking. However, as in the corresponding least squares analysis, the interpretation of their causal effects may be somewhat controversial. For example, although we find (as expected) that babies born to mothers with no prenatal care are smaller, we also find that babies born to mothers who delayed prenatal visits until the second or third trimester have substantially *higher* birthweights in the lower tail than mothers who had a prenatal visit in the first trimester. This might be interpreted as the self-selection effect of mothers confident about favorable outcomes.

[3]It is conducive for interpretation to center covariates so that the intercept can be interpreted as the conditional quantile function for some representative case – rather than as an extrapolation of the model beyond the convex hull of the data. This may be viewed as adhering to John Tukey's dictum: "Never estimate intercepts, always estimate *centercepts*!"

In almost all of the panels of Figure 4, with the exception of education coefficients, the quantile regression estimates lie at some point outside the confidence intervals for the ordinary least squares regression, suggesting that the effects of these covariates may not be constant across the conditional distribution of the independent variable. Formal testing of this hypothesis is discussed in Koenker and Machado (1999).

SELECTED EMPIRICAL EXAMPLES OF QUANTILE REGRESSION

There is a rapidly expanding *empirical* quantile regression literature in economics which, taken as a whole, makes a persuasive case for the value of "going beyond models for the conditional mean" in empirical economics. Catalysed by Gary Chamberlain's invited address to the 1990 World Congress of the Econometric Society, Chamberlain (1994), there has been considerable work in labor economics: on union wage effects, returns to education, and labor market discrimination. Chamberlain finds, for example, that for manufacturing workers, the union wage premium at the first decile is 28 percent and declines monotonically to a negligible 0.3 percent at the upper decile. The least squares estimate of the mean union premium of 15.8 percent is thus captured mainly by the lower tail of the conditional distribution. The conventional location shift model thus delivers a rather misleading impression of the union effect. Other contributions exploring these issues in the U.S. labor market include the influential work of Buchinsky (1994, 1997). Arias, Hallock, and Sosa-Escudero (2001) using data on identical twins interpret observed heterogeneity in the estimated return to education over quantiles as indicative of an interaction between observed educational attainment and unobserved ability.

There is also a large literature dealing with related issues in labor markets outside the U.S. including Fitzenberger (1999), Machado and Mata (2001) on Portugal; Garcia, Hernandez, and Lopez (2001) on Spain; Schultz and Mwabu (1998) on South Africa; and Kahn (1998) on international comparisons. The work of Machado and Mata (2001) is particularly notable since it introduces a useful way to extend the counterfactual decomposition approach of Oaxaca to quantile regression and suggests a general strategy for simulating marginal distributions from the quantile regression process. Tannuri (2000) has employed this approach in a recent study of assimilation of U.S. immigrants.

In other applied micro areas Eide and Showalter (1998), Knight, Bassett, and Tam (2000), and Levin (2001) have addressed school quality issues. Poterba and Rueben (1995) and Mueller (2000) study public-private wage differentials in the U.S. and Canada. Abadie, Angrist, and Imbens (2001) consider estimation of endogenous treatment effects in program evaluation, and Koenker and Bilias (2001) explore quantile regression models for unemployment duration data. Work by Viscusi and Hamilton (1999) considers public decision making regarding hazardous waste cleanup.

Deaton (1997) offers a nice introduction to quantile regression for demand analysis. In a study of Engel curves for food expenditure in Pakistan he finds that although the median Engel elasticity of 0.906 is similar to the OLS estimate of 0.909, the coefficient at the 10th quantile is 0.879 and the estimate at the 90th percentile is 0.946, indicating a pattern of heteroscedasticity like that of our Figure 3.

In another demand application, Manning, Blumberg, and Moulton (1995) study demand for alcohol using survey data from the National Health Interview Study and find considerable heterogeneity in the price and income elasticities over the full range of the conditional distribution. Hendricks and Koenker (1991) investigate demand for electricity by time of day using a hierarchical model.

Earnings inequality and mobility is a natural arena of applications for quantile regression. Conley and Galenson (1998) explore wealth accumulation in several U.S. cities in the mid-19th century. Gosling, Machin, and Meghir (1996) study the income and wealth distribution in the UK. Trede (1998) and Morillo (2000) compare earnings mobility in the U.S. and Germany.

There is also a growing literature in empirical finance employing quantile regression methods. One strand of this literature is the rapidly expanding literature on value at risk: this connection is developed in Taylor (1999), Chernozhukov and Umantsev (2001), and Engle and Manganelli (1999). Bassett and Chen (2001) consider quantile regression index models to characterize mutual fund investment styles.

Software and Standard Errors

The diffusion of technological change throughout statistics is closely tied to its embodiment in statistical software. This is particularly true of quantile regression methods since the linear programming algorithms that underlie reliable implementations of the methods appear somewhat esoteric to some users. Since the early 1950's it has been recognized that median regression methods based on minimizing sums of absolute residuals can be formulated as linear programming problems and efficiently solved with some form of the simplex algorithm. The median regression algorithm of Barrodale and Roberts (1974) has proven particularly influential, and can be easily adapted for general quantile regression. Koenker and d'Orey (1987) describe one implementation. For large scale quantile regression problems Portnoy and Koenker (1997) have shown that a combination of interior point methods and effective preprocessing render quantile regression computation competitive with least squares computations for problems of comparable size.

Among commercial programs in common use in econometrics only Stata offers some basic functionality for quantile regression within the central core of the package distributed by the vendor. Since the mid-1980's one of us has maintained a public domain package of quantile regression software designed for the S language of Becker, Chambers, and Wilks (1988), and the related commercial package Splus. Recently, this package has been extended to provide a version for R, the impressive GNU dialect of S. See Koenker (1995). This website also provides software for the Ox and Matlab.

It is a basic principle of sound econometrics that *every serious estimate deserves a reliable assessment of precision.* There is now an extensive literature on the asymptotic behavior of quantile regression estimators and considerable experience with inference methods based on this theory as well as a variety of resampling schemes. We have recently undertaken a comparison of several approaches to the construction of confidence intervals for a problem typical of current applications in labor economics. We find, that the discrepancies between competing methods are slight and inference for quantile regression is, if anything, more robust than most other forms of inference commonly encountered in econometrics. Our initial draft of this paper, Koenker and Hallock (2000), described this exercise in more detail and provides a brief survey of recent work on quantile regression for discrete data models, time series, nonparametric models and a variety of other areas. Many of these developments have been slow to percolate into standard econometric software. With the notable exceptions of Stata and Xplore, see Cizek (2000), and the packages available from the web for `Splus` and `R`, none of the implementations of quantile regression in econometric software packages include any functionality for inference.

CONCLUSION

Much of applied statistics may be viewed as an elaboration of the linear regression model and associated estimation methods of least-squares. In beginning to describe these techniques Mosteller and Tukey (1977) in their influential text remark:

> What the regression curve does is give a grand summary for the averages of the distributions corresponding to the set of x's. We could go further and compute several different regression curves corresponding to the various percentage points of the distributions and thus get a more complete picture of the set. Ordinarily this is not done, and so regression often gives a rather incomplete picture. Just as the mean gives an incomplete picture of a single distribution, so the regression curve gives a correspondingly incomplete picture for a set of distributions.

We would like to think that quantile regression is gradually developing into a comprehensive strategy for completing the regression picture.

□ *We would like to thank the editors as well as Gib Bassett, John DiNardo, Olga Geling, and Steve Portnoy for helpful comments on earlier drafts. The research was partially supported by NSF grant SES-99-11184.*

REFERENCES

ABADIE, A., J. ANGRIST, AND G. IMBENS (2001): "Instrumental Variables Estimation of Quantile Treatment Effects," *Econometrica*, forthcoming.

ABREVEYA, J. (2001): "The effects of demographics and maternal behavior on the distribution of birth outcomes," *Empirical Economics*, 26, 247–257.

ARIAS, O., K. HALLOCK, AND W. SOSA-ESCUDERO (2001): "Indidividual heterogeneity in the returns to schooling: instrumental variables quantile regression using Twins data," *Empirical Economics*, 26, 7–40.

BARRODALE, I., AND F. ROBERTS (1974): "Solution of an overdetermined system of equations in the ℓ_1 norm," *Communications of the ACM*, 17, 319–320.

BASSETT, G., AND H. CHEN (2001): "Quantile Style: Return-Based Attribution Using Regression Quantiles," *Empirical Economics*, 26, 293–305.

BECKER, R., J. CHAMBERS, AND A. WILKS (1988): *The New S Language: A Programming Environment for Data Analysis and Graphics.* Wadsworth: Pacific Grove.

BUCHINSKY, M. (1994): "Changes in US Wage Structure 1963-87: An Application of Quantile Regression," *Econometrica*, 62, 405–458.

——— (1997): "Women's return to education in the U.S.: Exploration by Quantile Regression," *J. Applied Econometrics*, 13, 1–30.

CHAMBERLAIN, G. (1994): "Quantile regression, censoring and the structure of wages," in *Advances in Econometrics*, ed. by C. Sims. Elsevier: New York.

CHERNOZHUKOV, V., AND L. UMANTSEV (2001): "Conditional Value at Risk: Aspects of Modeling and Estimation," *Empirical Economics*, 26, 271–292.

CIZEK, P. (2000): "Quantile regression," in *XploRe Application Guide*, ed. by W. Härdle , Z. Hlavka, and S. Klinke. Springer-Verlag: Heidelberg.

CONLEY, T., AND D. GALENSON (1998): "Nativity and Wealth in Mid-Nineteenth-Century Cities," *J. of Economic History*, 58, 468–93.

DEATON, A. (1997): *The Analysis of Household Surveys.* Johns Hopkins: Baltimore.

EIDE, E., AND M. SHOWALTER (1998): "The Effect of School Quality on Student Performance: A Quantile Regression Approach," *Economics Letters*, 58, 345–50.

ENGLE, R., AND S. MANGANELLI (1999): "CaViaR: Conditional Autoregressive Value at Risk by Regression Quantiles," preprint.

FITZENBERGER, B. (1999): *Wages and employment across skill groups.* Physica-Verlag: Heidleberg.

GARCIA, J., P. HERNANDEZ, AND A. LOPEZ (2001): "How wide is the gap? An investigation of gender wage differences using quantile regression," *Empirical Economics*, 26, 149–167.

GOSLING, A., S. MACHIN, AND C. MEGHIR (1996): "What Has Happened to the Wages of Men since 1966?," in *New inequalities: The changing distribution of income and wealth in the United Kingdom*, ed. by J. Hills.

HECKMAN, J. J. (1979): "Sample Selection Bias As a Specification Error," *Econometrica*, 47, 153–161.

HENDRICKS, W., AND R. KOENKER (1991): "Hierarchical spline models for conditional quantiles and the demand for electricity," *J. of Am. Stat. Assoc.*, 87, 58–68.

KAHN, L. (1998): "Collective Bargaining and the Interindustry Wage Structure: International Evidence," *Economica*, 65, 507–34.

KNIGHT, K., G. BASSETT, AND M. TAM (2000): "Comparing Quantile Estimators for the linear model," preprint.

KOENKER, R. (1995): "Quantile Regression Software," `www.econ.uiuc.edu/~roger/research/rq/rq.html`.

KOENKER, R., AND G. BASSETT (1978): "Regression Quantiles," *Econometrica*, 46, 33–50.

KOENKER, R., AND Y. BILIAS (2001): "Quantile Regression for Duration Data: A Reappraisal of the Pennsylvania Reemployment Bonus Experiments," *Empirical Economics*, 26, 199–220.

KOENKER, R., AND V. D'OREY (1987): "Computing Regression Quantiles," *Applied Statistics*, 36, 383–393.

KOENKER, R., AND K. HALLOCK (2000): "Quantile Regression: An Introduction," `www.econ.uiuc.edu/~roger/research/intro/intro.html`.

KOENKER, R., AND J. MACHADO (1999): "Goodness of fit and related inference processes for quantile regression," *J. of Am. Stat. Assoc.*, 94, 1296–1310.

LEVIN, J. (2001): "Where the reductions count: A quantile regression analysis of effects of class size on scholastic acheivement," *Empirical Economics*, 26, 221–246.

MACHADO, J., AND J. MATA (2001): "Counterfactual decomposition of changes in wage distributions using quantile regression," *Empirical Economics*, 26, 115–134.

MANNING, W., L. BLUMBERG, AND L. MOULTON (1995): "The demand for alcohol: the differential response to price," *Journal of Health Economics*, 14, 123–148.

MORILLO, D. (2000): "Income mobility with nonparametric quantiles: A comparison of the U.S. and Germany," preprint.

MOSTELLER, F., AND J. TUKEY (1977): *Data Analysis and Regression: A Second Course in Statistics*. Addison-Wesley, Reading Mass.

MUELLER, R. (2000): "Public and Private Wage Differentials in Canada Revisited," *Industrial Relations*, 39, 375–400.

PORTNOY, S., AND R. KOENKER (1997): "The Gaussian Hare and the Laplacian Tortoise: Computability of squared-error versus absolute-error estimators, with discusssion," *Stat. Science*, 12, 279–300.

POTERBA, J., AND K. RUEBEN (1995): "The Distribution of Public Sector Wage Premia: New Evidence Using Quantile Regression Methods," NBER Working paper #4734.

SCHULTZ, T., AND G. MWABU (1998): "Labor unions and the distribution of wages and employment in South Africa," *Ind. and Labor Relations Rev.*, 51, 680–703.

TANNURI, M. (2000): "Are the assimilation processes of low and high income immigrants distinct? A Quantile Regression Approach," preprint.

TAYLOR, J. (1999): "A Quantile Regression Approach to Estimating the Distribution of Multiperiod Returns," *Journal of Derivatives*, pp. 64–78.

TREDE, M. (1998): "Making Mobility Visible: A Graphical Device," *Economics Letters*, 59, 77–82.

VISCUSI, W., AND J. HAMILTON (1999): "Are Risk Regulators Rational? Evidence from Hazardous Waste Cleanup Decisions," *American Economic Review*, 89, 1010–27.