Confidence Intervals for Regression Quantiles

Roger Koenker
University of Illinois at Champaign-Urbana

Abstract. Several methods to construct confidence intervals for regression quantile estimators (Koenker and Bassett (1978)) are reviewed. Direct estimation of the asymptotic covariance matrix requires an estimate of the reciprocal of the error density (sparsity function) at the quantile of interest; some recent work on bandwidth selection for this problem will be discussed. Several versions of the bootstrap for quantile regression will be described as well as a recent proposal by Parzen, Wei, and Ying (1992) for resampling from the (approximately pivotal) estimating equation. Finally, we will describe a new approach based on inversion of a rank test suggested by Gutenbrunner, Jurečková, Koenker, and Portnoy (1993) and introduced in Hušková(1994). The latter approach has several advantages: it may be computed relatively efficiently, it is consistent under certain heteroskedastic conditions and it circumvents any explicit estimation of the sparsity function. A small monte-carlo experiment is employed to compare the competing methods.

AMS 1980 Subject Classification: 62G10, 62J10.

Keywords: regression quantiles, confidence intervals, rank tests, bootstrapping, sparsity estimation.

1. Introduction

Quantile regression, as introduced in Koenker and Bassett (1978), is gradually developing into a comprehensive approach to the statistical analysis of linear and nonlinear response models. By supplementing the exclusive focus of least-squares-based methods on the estimation of conditional mean functions with a general technique for estimating conditional quantile functions, it has expanded the flexibility of both parametric and non-parametric statistical methods.

There is already a well-developed theory of asymptotic inference for quantile regression and related L-statistics based on the Bahadur representation of the regression quantile process. See for example Jurečková(1984) and Koenker and Portnoy (1987). However, when inference on discrete quantiles is desired, one is faced with a rather bewildering array of methods based on direct estimation of the asymptotic covariance matrix, an approach which involves estimation of the reciprocal of the error density at the quantile of interest, or some form of the bootstrap. Versions of both approaches are available in existing statistical packages. Recently, however, several alternative approaches to inference have emerged. The objective of this paper is to compare the various methods and offer some guidance from a practical perspective on which methods seem most promising.

Consider the linear model for the τth conditional quantile function of a response variable, Y, given covariates $x \in \mathbf{R}^p$,

$$Q_Y(\tau|x) = x'\beta(\tau)$$

It will be assumed throughout that the first coordinate of x is identically one. We will focus on constructing a confidence interval for the pth coordinate of the parameter vector, β_p. A point estimate $\hat{\beta}(\tau)$, of the parameter $\beta(\tau)$, is obtained by solving

$$\min_{b \in \mathbf{R}^p} \sum_{i=1}^{n} (y_i - x_i'b)$$

where $\rho_\tau(u) = u(\tau - I(u < 0))$. In the simplest case where

$$y_i = x_i'\beta + u_i$$

and the $\{u_i\}$ are iid F with $f = F'$ and $f(F^{-1}(\tau)) > 0$ in a neighborhood of τ. Under mild design conditions we have (Koenker and Bassett(1978))

$$\sqrt{n}(\hat{\beta}(\tau) - \beta(\tau)) \rightsquigarrow \mathfrak{N}(0, \omega^2(\tau, F)D^{-1})$$

where $\beta(\tau) = \beta + F^{-1}(\tau)e_1$, $e_1 = (1, 0, ..., 0)'$, $\omega^2(\tau, F) = \tau(1 - \tau)/f^2(F^{-1}(\tau))$ and $D = \lim\ n^{-1} \sum x_i x_i'$.

It is natural to begin the discussion with direct estimation of the covariance matrix of regression quantiles under iid error conditions. It is an somewhat unhappy fact of life that the asymptotic precision of quantile estimates in general, and quantile regression estimates in particular (Koenker and Bassett(1978)) depend upon the reciprocal of a density function evaluated at the quantile of interest – a quantity Tukey has termed the "sparsity function" and Parzen calls the quantile-density function. It is perfectly natural that the precision of quantile estimates should depend on this quantity since it reflects the density of observations near the quantile of interest. Thus, to estimate the precision of the τth quantile regression estimate directly, the nuisance quantity

$$s(\tau) = [f(F^{-1}(\tau))]^{-1}$$

must be estimated and therefore we must venture into the realm of smoothing. In fact, it may be possible to pull oneself out of this swamp by the bootstraps, or other statistical necromancy, but we defer the exploration of these strategies and explore the direct approach in next section.

2. Sparsity Estimation

Fortunately, there is a large literature on estimating $s(\tau)$ including Siddiqui (1960), Bofinger (1975), Sheather and Maritz (1983), Welsh (1986) and Hall and Sheather (1988). Siddiqui's idea is simplest and has received the most attention in the literature so we will focus on it. Differentiating the identity, $F(F^{-1}(t)) = t$ we have

$$\frac{d}{dt} F^{-1}(t) = s(t)$$

so it is natural to estimate $s(t)$ by using a simple difference quotient of the empirical quantile function, i.e.,

$$\hat{s}_n(t) = [\hat{F}_n^{-1}(t + h_n) \;-\; \hat{F}_n^{-1}(t - h_n)]/2h_n$$

where $\hat{F}^{-1}(\cdot)$ is an estimate of F^{-1} and h_n is a bandwidth which tends to zero as $n \to \infty$. Bofinger (1975) showed that

$$h_n = n^{-1/5}\left(\frac{4.5s^2(t)}{(s''(t))^2}\right)^{1/5}$$

was the optimal (minimum mean squared error) choice of h_n under mild regularity conditions on F. This result follows from standard density estimation asymptotics.

Of course, if we knew $s(t)$ and $s''(t)$ we wouldn't need h_n, but fortunately $s(t)/s''(t)$ is not very sensitive to F so little is lost if we choose h_n for some typical distributional shape - say, the Gaussian. Sheather and Maritz(1983) discuss preliminary estimation of s and s'' as a means of estimating a plug-in h_n. In general,

$$\frac{s(t)}{s''(t)} = \frac{f^2}{2(f'/f)^2 + [(f'/f)^2 - f''/f]}$$

and, for example, if f is Gaussian, $(f'/f)(F^{-1}(t)) = \Phi^{-1}(t)$ and the term in square brackets is 1 , so the optimal bandwidth becomes,

$$h_n = n^{-1/5}\left(\frac{4.5\phi^4(\Phi^{-1}(t))}{(2\Phi^{-1}(t)^2 + 1)^2}\right)^{1/5}.$$

In Figure 1, I have plotted this bandwidth, in the solid lines, as a function of n for three distinct quantiles $t = .50, 75$, and $.95$. For symmetric F the h_n's at t and $1-t$ are obviously the same.

The rule suggested above which is based upon standard density estimation asymptotics has been recently questioned by Hall and Sheather (1988). Based on Edgeworth expansions for studentized quantiles, they suggest

$$h_n = n^{-1/3} z_\alpha^{2/3} [1.5s(t)/s''(t)]^{1/3}$$

where z_α satisfies $\Phi(z_\alpha) = 1 - \alpha/2$ for the construction of $1 - \alpha$ confidence intervals. This bandwidth rule is illustrated in Figure 1 by the dotted curves. It gives somewhat narrower bandwidths for modest to large n. Since the Hall and Sheather rule is explicitly designed for confidence interval construction for the quantiles, rather than simply optimizing mse-performance for the sparsity estimate itself, it seems reasonable to use it for inference. Since the optimal constant in the Hall-Sheather expression depends on the same sparsity functional as the Bofinger bandwidth, the same argument suggests that it may not be unreasonable to use the normal model version. Note that $s(t)/s''(t)$ is location-scale invariant so only the *shape* of the distribution influences this constant.

Having chosen a bandwidth h_n the next question is: how should we compute $\hat{F}^{-1}$? The simplest approach is to use the residuals from the quantile regression fit. Let $r_i : i = 1, ..., n$ be these residuals, and $r_{(i)} : i = 1, ..., n$ be the corresponding order statistics. Then define the usual empirical quantile function, $\hat{F}^{-1}(t) = r_{(j)}$ for $t \in$

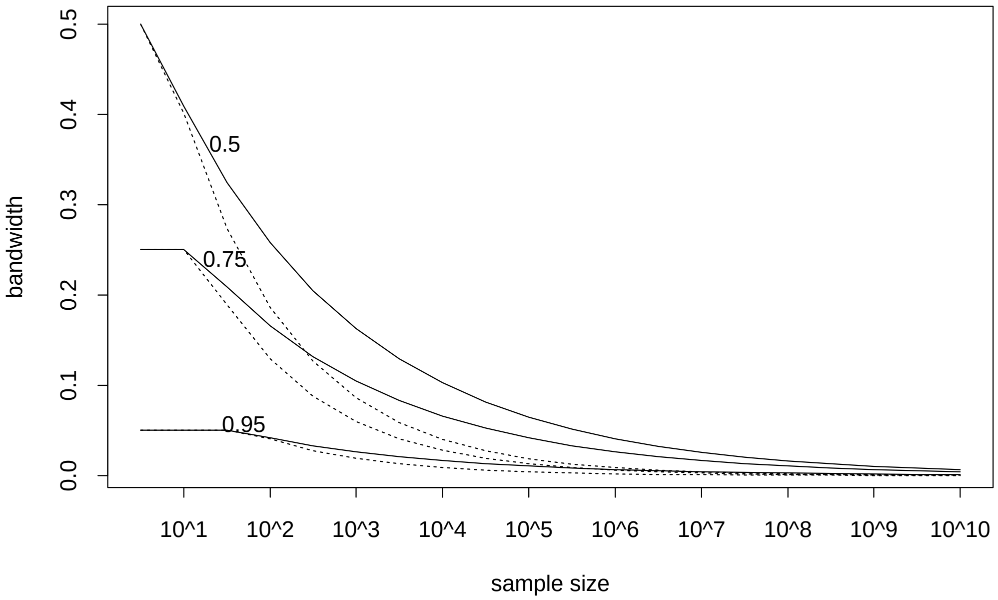

FIGURE 1. Siddiqui Bandwidths for Gaussian Sparsity Estimation

$[(j-1)/n,\ j/n\ )$. Alternatively, one may wish to interpolate to get a piecewise linear version

$$\tilde{F}^{-1}(t) = \begin{cases} r_{(1)} & if\ t \in [0, 1/2n) \\ \lambda r_{(j+1)} + (1-\lambda) r_{(j)} & if\ t \in [(2j-1)/2n, (2j+1)/2n)\ j = 1, \ldots, n-1 \\ r_{(n)} & if\ t \in [(2n-1)/2n, 1] \end{cases}$$

where $\lambda = tn - j + 1/2$. Alternative schemes are obviously possible. A possible pitfall of the residual-based estimates of F^{-1} is that if the number of parameters estimated, say p, is large relative to n, then since there must be p residuals equal to zero at the fitted quantile we must make sure that the bandwidth is large enough to avoid the zero residuals. The simplest approach here seems to be to ignore the zero residuals in the construction of $\hat{F}^{-1}$ and $\tilde{F}^{-1}$ and treat the effective sample size as $n-p$.

An alternative, perhaps less obvious, approach to obtain $\hat{F}^{-1}$ is to employ the empirical quantile function suggested in Bassett and Koenker (1982). In effect this amounts to using $\hat{F}^{-1} = \overline{x}'\hat{\beta}(t)$ where $\hat{\beta}(\cdot)$ is the usual regression quantile process. Like the EQF based on residuals, this is a piecewise constant function, but now the jumps are no longer equally spaced on $[0,1]$. Nevertheless, the same ideas still apply and either the piecewise constant form of the function or the linear interpolant can be used. See Bassett and Koenker(1989) for a detailed treatment of the (strong) consistency of this method.

Finally, we should address the question: what happens if $t \pm h_n$ falls outside $[0,\ 1]$? This can easily happen when n is small. Obviously, some *ad hoc* adjustment is needed in this case, with perhaps a warning to users that the plausibility of the asymptotic theory is strained in such situations. In Figure 2.1 the flattening out of the bandwidth functions for small n reflects a simple rule of this sort.

3. Inversion of Rank Tests

In Gutenbrunner, Jurečková, Koenker and Portnoy (1993) we have developed a new approach to rank-based inference for the linear regression model. The classical theory of rank tests as exposited in the monograph of Hájek and Šidák (1967) begins with the so-called rankscore functions,

$$\hat{a}_{ni}(t) = \begin{cases} 1 & \text{if } t \leq (R_i - 1)/n \\ R_i - tn & \text{if } (R_i - 1)/n < t \leq R_i/n \\ 0 & \text{if } R_i/n < t \end{cases}$$

where R_i is the rank of the i^{th} observation, Y_i, in the sample $\{Y_1, \ldots, Y_n\}$. Integrating $\hat{a}_{ni}(t)$ with respect to various score generating functions φ yields vectors of rank-like statistics which may be used for constructing tests. For example, integrating with respect to Lebesgue measure yields the Wilcoxon scores,

$$b_i = \int_0^1 \hat{a}_{ni}(t)dt = (R_i - 1/2)/n \qquad i = 1, ..., n,$$

while using $\varphi(t) = \text{sgn}(t - 1/2)$ yields the sign scores, $b_i = \hat{a}_{ni}(1/2)$.

How can this idea be extended to regression when, under the null, a nuisance regression parameter is present? This question was answered by Gutenbrunner and Jurečková(1992) who observed that the Hájek-Šidák rankscores may be viewed as a special case of a more general formulation for the linear model in which the functions $\hat{a}_{ni}(t)$ are defined in terms of the linear program

$$\max\{y'a | X'a = (1-t)X'1,\ a \in [0,1]^n\} \tag{3.1}$$

This problem is formally dual to the linear program defining the regression quantiles. Algorithmic details are given in Koenker and d'Orey (1993). Tests of the hypothesis $\beta_2 = 0 \in \mathbf{R}^q$ in the model $y = X_1\beta_1 + X_2\beta_2 + u$ based on the regression rankscore process may be constructed by first computing $\{\hat{a}_{ni}(t)\}$ at the restricted model,

$$y = X_1\beta_1 + u$$

computing the n-vector b with elements $b_i = \int \hat{a}_{ni}(t)d\varphi(t)$, forming the q-vector, $S_n = n^{-1/2}X_2'b$, and noting that, under the null $S_n \rightsquigarrow \mathfrak{N}(0, A^2(\varphi)Q)$ where $A^2(\varphi) = \int_0^1 \varphi^2(t)dt$, $Q = lim_{n\to\infty}Q_n$, $Q_n = (X_2 - \hat{X}_2)'(X_2 - \hat{X}_2)/n$ and $\hat{X}_2 = X_1(X_1'X_1)^{-1}X_1'X_2$. So the test statistic $T_n = S_n'Q^{-1}S_n/A^2(\varphi)$ has an asymptotic χ_q^2 null distribution, and noncentral χ_q^2 distribution under appropriate contiguous alternatives.

In the special case that X_1 is simply a column vector of ones this reduces to the original formulation of Hájek and Šidák. When $\varphi(t) = \text{sgn}(t - 1/2)$ it specializes to the score-test proposed for ℓ_1-regression in Koenker and Bassett (1982). An important feature of these rank tests is that they require no estimation of nuisance parameters, since the functional $A(\varphi)$ depends only on the score function and not on the (unknown) distribution of the vector u. This is familiar from the theory of elementary rank tests, but stands in sharp contrast with other methods of testing in the linear model where, typically, some estimation of a scale parameter, e.g., σ^2, is required to compute the test statistic.

This raises the question: could we invert a rank test of this form to provide a method of estimating a confidence interval for quantile regression, thus circumventing the problem of estimating $s(t)$. Hušková(1994) considers this problem in considerable generality establishing the validity of sequential fixed-width confidence intervals for general score functions φ. Unfortunately, for general score functions these intervals are difficult to compute. However, in the case of a fixed quantile one particularly natural choice of φ yields extremely tractable computations and we will focus on this case.

Specializing to the scalar β_2 case and using the τ-quantile score function

$$\varphi_\tau(t) = \tau - I(t < \tau)$$

and proceeding as above, we find that

$$\hat{b}_{ni} = -\int_0^1 \varphi_\tau(t) d\hat{a}_{ni}(t) = \hat{a}_{ni}(\tau) - (1-\tau) \tag{3.2}$$

with

$$\bar{\varphi} = \int_0^1 \varphi_\tau(t)dt = 0$$

$$A^2(\varphi_\tau) = \int_0^1 (\varphi_\tau(t) - \bar{\varphi})^2 dt = \tau(1-\tau).$$

Thus, a test of the hypothesis $H_o : \beta_2 = \xi$ may be based on $\hat{a}_n$ from solving,

$$\max\{(y - x_2\xi)'a | X_1'a = (1-\tau)X_1'1, a \in [0,1]^n\} \tag{3.3}$$

and the fact that

$$S_n(\xi) = n^{-1/2} x_2' \hat{b}_n(\xi) \rightsquigarrow \mathfrak{N}(0, A^2(\varphi_\tau)q_n^2) \tag{3.4}$$

under H_o; where $q_n^2 = n^{-1}x_2'(I - X_1(X_1'X_1)^{-1}X_1')x_2$. That is we may compute $T_n(\xi) = S_n(\xi)/(A(\varphi_\tau)q_n)$ and reject H_o if $|T_n(\xi)| > \Phi^{-1}(1-\alpha/2)$. This takes us back to the linear program (3.3) which may now be viewed as a one parameter parametric linear programming problem in ξ. In ξ the dual vector $\hat{a}_n(\xi)$ is piecewise constant; ξ may be altered without compromising the optimality of $\hat{a}_n(\xi)$ as long as the sign of the residuals in the primal quantile regression problem do not change. When ξ gets to such a boundary the solution does change, but may be restored by taking one simplex pivot. The process may continue in this way until $T_n(\xi)$ exceeds the specified critical value. Since $T_n(\xi)$ is piecewise constant we interpolate in ξ to obtain the desired level for the confidence interval. See Beran and Hall (1993) for a detailed analysis of the effect of interpolation like this in the case of confidence intervals for ordinary quantiles. This interval, unlike the Wald type sparsity intervals, is not symmetric; but it *is* centered on the point estimate $\hat{\beta}_2(\tau)$ in the sense that $T_n(\hat{\beta}_2(\tau)) = 0$. This follows immediately from the constraint $X'\hat{a} = (1-\tau)X'1$ in the full problem.

The primary virtue of this approach is that it inherits the scale invariance of the test statistic T_n and therefore circumvents the problem of estimating the sparsity function. Implemented in S, using an adaptation of the algorithm described in Koenker and d'Orey (1993), it has essentially the same computational efficiency as the sparsity methods. More computationally intensive resampling methods offer an alternative route which we explore in the next section.

4. Resampling Methods

There has been considerable recent interest in resampling methods for estimating confidence intervals for quantile-type estimators. However, despite the fact that confidence intervals for quantiles was one of the earliest success stories for the bootstrap (in contrast to the delete-1 jackknife which fails in this case) recent results have been considerably more guarded in their enthusiasm. Hall and Martin (1989) conclude:

> It emerges from these results that the standard bootstrap techniques perform poorly in constructing confidence intervals for quantiles... The percentile method does no more than reproduce a much older method with poor coverage accuracy at a fixed level: bias corrections fail for the same reason; bootstrap iteration fails to improve the order of coverage accuracy; and percentile-t is hardly an efficacious alternative because of non-availability of suitable variance estimates.

Nevertheless, there has been considerable recent interest, particularly among econometricians, in using the bootstrap to compute standard errors in quantile regression applications. See Buchinsky(1994), Hahn(1993) and Fitzenberger(1993) for examples.

There are several possible implementations of the bootstrap for quantile regression applications. As in other regression applications we have a choice between the residual bootstrap and the xy-pairs bootstrap. The former resamples with replacement from the residual vector and adds this to the fitted vector $X\hat{\beta}_n(\tau)$ and reestimates, in so doing it assumes that the error process is iid. The latter resamples xy pairs, and therefore is able to accomodate some forms of heteroskedasticity. As in the sparsity estimation approaches we may consider replacing the residual EQF by the EQF obtained directly from the the regression quantile process, but this maintains the iid error assumption. More interesting is the possibility of resampling directly from the full regression quantile process which we will call the Heqf bootstrap. By this I mean for each bootstrap realization of n observations we draw n p-vectors from the estimated process $\hat{\beta}_n(t)$. There are, say, J distinct such realizations

$$\hat{\beta}_n(t) = \hat{\beta}_n(t_j) \quad \text{for} \quad t_j \leq t < t_{j+1}$$

$j = 1, ..., J$ and each is drawn with probability $\pi_j = t_{j+1} - t_j$. For each design row x_i we associate the bootstrapped y observation which is the inner product of that design row and the corresponding ith draw from the regression quantile process. This procedure has the virtue that it is again capable of accommodating certain forms of heteroskedastic regression models, in particular those with linear conditional quantile functions.

Finally, we will describe a new resampling method due to Parzen, Wei and Ying (1993) which is quite distinct from the bootstrap. It arises from the observation that the function

$$S(b) = n^{-1/2} \sum_{i=1}^{n} x_i(\tau - I(y_i \leq x_i'b)) \tag{4.1}$$

which is the estimating equation for the τth regression quantile is a pivotal quantity for the true τth quantile regression parameter $b = \beta_\tau$. That is, its distribution may be generated exactly by a random vector U which is a weighted sum of independent, recentered Bernoulli variables which play the role of the indicator function. They show further that for large n the distribution of $\hat{\beta}_n(\tau) - \beta_\tau$ can be approximated by the conditional distribution of $\hat{\beta}_U - \hat{\beta}_n(\tau)$, where $\hat{\beta}_U$ solves an augmented quantile regression problem with $n+1$ observations and $x_{n+1} = -n^{1/2}u/\tau$ and y_{n+1} is sufficiently large that $I(y_{n+1} \leq z'_{n+1}b)$ is always zero. This is essentially the same as solving $S(b) = -u$ for given realization of u. This approach, by exploiting the asymptotically pivotal role of the quantile regression "gradient condition", also achieves a robustness to certain heteroskedastic quantile regression models. In practice, one might be able to exploit the fact that the solution to the augmented problems is close to the original one, since they differ by only one observation, but we have not tried to do this in our simulation experiments which are reported in the next section.

5. Monte-Carlo Comparison of Methods

In this final section we report on a small monte-carlo experiment to compare the performance of the methods described above. We focus primarily on the computationally less demanding sparsity estimation and inverted rankscore methods, but some results are reported for three of the resampling methods. Preliminary results indicated that the Hall and Sheather bandwidths performed considerably better than the Bofinger choice so we have restricted our reported results mainly to this form of sparsity estimation. We also focus exclusively on the problem of confidence intervals for the median regression parameters, partly because this is the most common practical problem, and also because it restricts the amount of computation and reporting required. In subsequent work, it is hoped to provide a much more exhaustive monte-carlo experiment.

We considered first an iid error model in which both x's and y's were generated from the Student t distribution. The degrees of freedom parameter varies over the set $\{1, 3, 8\}$ for both x's and y's. The first column of the design matrix is ones, all other entries are iid draws from the specified t distribution. For each cell of the experiment the design matrix is drawn once, and 1000 replications of the response vector, y, are associated with this fixed design matrix. Throughout, we have studied only the sample size $n = 50$.

All of the computations were carried out in the 'S' language of Becker, Chambers, and Wilks(1988) on a Sun workstation. Further details on the experiment are available from the author on request, since space limitations dictate a rather abbreviated treatment here.

In Table 1 we report observed monte-carlo coverage frequencies for nine situations and three non-resampling methods. Confidence intervals are computed for all three slope coefficients for each situation so in each cell we report the number of times the interval covers the true parameter (zero in all cases) in 3000 trials. Throughout the experiment the nominal size is .10. In these iid error situations we see that the size of the rank inversion method is quite accurate throughout, as is the Hall-Sheather sparsity estimate. However the Bofinger results are considerably less satisfactory.

Generally, the rank-inversion intervals are shorter than the sparsity intervals except for the anomalous cases of Cauchy design.

TABLE 1. Confidence Interval Performance – IID Errors

	coverage			length		
	dfy= 1	dfy= 3	dfy= 8	dfy= 1	dfy= 3	dfy= 8
dfx= 1						
rank-inverse	0.892	0.923	0.922	0.320	0.392	0.359
sparsity-HS	0.893	0.907	0.909	0.240	0.142	0.079
sparsity-BS	0.932	0.927	0.931	0.288	0.153	0.083
dfx= 3						
rank-inverse	0.875	0.904	0.890	0.625	0.504	0.501
sparsity-HS	0.923	0.911	0.923	0.614	0.505	0.544
sparsity-BS	0.954	0.932	0.937	0.736	0.544	0.577
dfx= 8						
rank-inverse	0.887	0.885	0.884	0.791	0.617	0.585
sparsity-HS	0.941	0.920	0.919	0.921	0.683	0.640
sparsity-BS	0.968	0.948	0.935	1.107	0.737	0.680

To compare the performance of the resampling methods we report in Table 2 results for 3 iid error situations and 5 methods. Since the resampling methods are quite slow, 500 resamples are done for each of them, we restrict attention to only the diagonal cases of the previous table with the degrees of freedom parameter for x's and y's equal. We are primarily interested in resampling as a means of acheiving consistent confidence intervals in heteroskedastic situations so we restrict attention to the Parzen-Wei-Ying (PWY) approach, the heteroskedastic empirical quantile function bootstrap (Heqf), and the xy-pairs bootstrap. It can be seen from the table that again the rank-inversion method is quite reliable in terms of size, and also performs well with respect to length. The PWY resampling method has empirical size less than half the nominal 10 percent, while the xy-bootstrap is also undersized. The Heqf-bootstrap is accurately sized except for the Cauchy situation. It is obviously difficult to compare the lengths acheived by various methods, given the discrepancies in size, however the rank inversion approach seems to perform reasonably well in this respect.

TABLE 2. Confidence Interval Performance – IID Errors

	coverage			length		
	df= 1	df= 3	df= 8	df= 1	df= 3	df= 8
rank-inverse	0.900	0.893	0.879	0.335	0.427	0.558
sparsity-HS	0.872	0.922	0.915	0.217	0.455	0.613
PWY	0.961	0.957	0.957	0.411	0.520	0.680
Heqf-BS	0.802	0.881	0.895	0.220	0.380	0.512
XY-BS	0.929	0.948	0.945	0.331	0.486	0.640

A more challenging problem for estimation of confidence intervals for quantile regression problems involves heteroskedastic situations. We consider a simple case which bears a close resemblance to the previous iid error situations. Again, we generate 3 columns of the design matrix X as iid draws, this time from the lognormal distribution. The response vectors are then drawn from a Student t distribution with location 0 and scale given by $\sigma_i = \sum_{i=1}^{4} x_i/5$. For all i, $x_{1i} = 1$. Again the design is fixed for a given configuration, and hence scale is fixed. In this model all the conditional quantile functions are linear, so the Heqf-bootstrap is applicable, however the simple sparsity estimation approach is obviously not consistent under these conditions.

TABLE 3. Confidence Interval Performance – Heteroskastic Errors

	coverage			**length**		
	df= 1	df= 3	df= 8	df= 1	df= 3	df= 8
rank-inverse	0.887	0.902	0.878	1.196	0.793	0.621
sparsity-HS	0.763	0.717	0.656	0.702	0.552	0.357
PWY	0.953	0.950	0.946	1.557	0.907	0.715
Heqf-BS	0.754	0.813	0.804	0.971	0.655	0.486
XY-BS	0.907	0.911	0.897	1.332	0.799	0.612

Again the rank-inversion approach seems to perform well. As expected the sparsity approach fails miserably. The Parzen-Wei-Ying resampler is again substantially undersized – a rather puzzling result. The xy bootstrap also performs very well, but the Heqf version of the bootstrap has very poor coverage frequencies suggesting that this approach is probably not reliable. Since the rank-inversion method is on the order of 10 times faster than any of the bootstrap methods even for moderate sized problems it appears to have a substantial advantage.

References

Bassett, G.W. and Koenker, R.W. (1982). An empirical quantile function for linear models with iid errors. *Journal of the American Statistical Association.* **77**, 407-415.

Bassett, G.W. and Koenker, R.W. (1986). Strong consistency of regression quantiles and related empirical processes *Econometric Theory*, **2**, 191-201.

Beran, R. and Hall, P. (1993). Interpolated nonparametric prediction intervals and confidence intervals, J. Royal Stat. Soc. (B), **55**, 643-652.

Bofinger, E. (1975). Estimation of a density function using order statistics, *Australian J. of Statistics*, **17**, 1-7.

Buchinsky, M. (1991) The Theory and Practice of Quantile Regression, Ph.d. Thesis, Department of Economics, Harvard University.

Gutenbrunner, C. and Jurečková, J. (1992) Regression quantile and regression rank score process in the linear model and derived statistics, *Ann. Statist.* **20**, 305-330.

Gutenbrunner, C., Jurečková, J., Koenker, R. and Portnoy, S. (1993). Tests of linear hypotheses based on regression rank scores, *J. of Nonparametric Statistics*, **2**, 307-331.

Hahn, J. (1991) Bootstrapping Quantile Regression Models, preprint, Department of Economics, Harvard University.

Hušková, M. (1994) Some sequential procedures based on regression rank scores, *J. of Nonparametric Statistics*, forthcoming.

Fitzenberger, B. (1993) The moving blocks bootstrap and robust inference in linear regression, Ph.d. Thesis, Department of Economics, Stanford University.

Hájek, J. and Šidák, Z. (1967). *Theory of Rank Tests*, Academia, Prague.

Hall, P. and Martin, M.A. (1989) A note on the accuracy of bootstrap percentile method confidence intervals for a quantile. *Stat. and Prob. Letters.* **8**, 197-200.

Hall, P. and Sheather, S.J. (1988). On the distribution of a studentised quantile. *J. R. Statist. Soc.*B, **50**, 381-391.

Jurečková, J. (1984). Regression quantiles and trimmed least squares estimator under a general design, *Kybernetika*, **20**, 345-357.

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

Koenker, R. and Bassett, G (1982). Tests of linear hypotheses and l_1 estimation, *Econometrica*, **50**, 1577-1584.

Koenker, R. and d'Orey, V. (1993) A Remark on Computing Regression Quantiles, *Applied Statistics*, **36**, 383-393.

Koenker, R.W. and Portnoy, S. (1987). L-Estimation for the Linear Model., *Journal of the American Statistical Association*, **82**, 851-857.

Parzen, M. I. Wei, L.J., and Ying, Z. (1992) A resampling method based on pivotal estimating functions, preprint.

Siddiqui, M.M. (1960). Distribution of quantiles in samples from a bivariate population. *J. Res. Nat. Bur. Standards*, **64B**, 145-150.

Sheather, S.J. and J.S. Maritz (1983) An estimate of the asymptotic standard error of the sample median, *Australian J. of Statistics*, **25**, 109-122.

Welsh, A.H. (1988). Asymptotically efficient estimation of the sparsity function at a point. *Stat and Prob Letters*, **6**, 427-432.