2

Worked Examples for Quantile Regression and Deciles

These examples use the ASA software integrated into Excel or SPSS (www.asastat.com). ASA is, in part, a point-and-click interface to R but analyses can be conducted from within SPSS or Excel. All data in the examples are hypothetical. We assume you have read the primer on quantile regression.

The first example analyzes how annual salary changes over time for professors at universities, i.e., how it increases with each additional year of experience. The study is cross-sectional in design. The outcome variable is the annual salary of professors (in the variable called *salary*) and the predictor variable is the number of years since getting a Ph.D. (in the variable called *timephd*). We regress salary onto timephd focusing on three quantiles, 0.25, 0.50, and 0.75. We want to evaluate if the estimated increment in salary as a function of time is different at the lower end of the salary dimension (the first quartile), in the middle of the salary dimension, and at the upper end of the salary dimension (the third quartile). We could include covariates and other predictors in the analysis, but for the sake of pedagogy, we focus on this simple bivariate relationship. For the actual execution of the program, watch the video link on our website. For explication of the statistical approach per se, see the primer on the website. The program is "Quantile regression" in the folder "Robust Statistics > Robust Regression > Robust Regression Models" of the ASA software. It interfaces with the R programs developed by Dr. Rand Wilcox and published on his website at the University of Southern California.

The ASA software routinely reports confidence intervals for key parameters in statistical models. There are different ways of presenting confidence intervals. One strategy is to report them directly. Another strategy is to report them as margins of error, much like the margins of error you see for political polls on television or in print media. In this case, one calculates the half width of the confidence interval and reports it in "plus or minus" format. For example, in a political poll, you might be told that the percent of people endorsing a candidate is 50% ±5%. In this case, the confidence interval is 45% to 55%. This is an efficient way of summarizing the interval. In some cases, confidence intervals are asymmetric. When this occurs, some researchers will report the lower and upper margin of error separately. Alternatively, the researcher might calculate the absolute difference between the lower limit and the parameter estimate as well as the absolute difference between upper limit of the interval minus the parameter estimate and then report whichever difference is larger using the ± format. Some analysts prefer the use of credible intervals in Bayesian analytic frameworks instead of confidence intervals for characterizing margins of error (see Curran, 2005).

QUANTILE REGRESSION

We use the ASA program defaults in our analysis (see the video on our website). The first part of the output provides the results for the 0.25 quantile. Here is the model fit information:

```
MODEL FIT FOR QUANTILE: 0.25

Log likelihood for null model: -8022.7248
Log likelihood for model: -7844.5620
AIC: 15693.1200
BIC: 15702.3600
Pseudo R squared: 0.2114
```

The percent improvement in fit of the model relative to an intercept only model is (0.2114) (100) = 21.1%.

Here is the output for the intercept:

```
Intercept

   Value of intercept: 33909.6177
   Standard error (SE): 1022.6987
   95% confidence interval using above SE: 31901.9164 to 35917.3190
   Lower and upper margin of error: -2007.7013, 2007.7013
   95% rank inverse confidence interval: 32402.2647 to 36517.1013
   Lower and upper margin of error using inverse CI: -1507.3530, 2607.4836
   Critical ratio: 33.1570
   p value: 0.000000
```

The value of the intercept is $33,909 ± $2008, which is the estimated salary for professors at the first quartile of the salary distribution when the time since their doctorate is 0, i.e., it is their starting salary.

Here is the output for the regression coefficient for timephd:

```
Predictor: TIMEPHD

   Coefficient: 1472.2878
   Standard error (SE): 89.1914
   95% confidence interval using above SE: 1297.1925 to 1647.3830
   Lower and upper margin of error: -175.0953, 175.0953
   95% rank inverse confidence interval: 1266.5307 to 1621.2167
   Lower and upper margin of error using rank inverse CI: -205.7571, 148.9290
```

```
   Critical ratio: 16.5071
   p value: 0.000000
```

The coefficient was \$1,473 ± \$175, which was statistically significant (critical ratio (CR) = 16.51, $p < 0.05$). For every additional year of experience, the first quartile for salary is predicted to increase by \$1,473.

Here is the output for the 0.50 quantile or the median:

```
MODEL FIT FOR QUANTILE: 0.5

Log likelihood for null model: -7976.4307
Log likelihood for model: -7810.9510
AIC: 15625.9000
BIC: 15635.1400
Pseudo R squared: 0.1980
```

The percent improvement in fit of the model relative to an intercept only model is (0.1980) (100) = 19.8%.

Here is the information for the intercept and coefficient:

```
REGRESSION EQUATION FOR QUANTILE: 0.5

Intercept

   Value of intercept: 40269.9772
   Standard error (SE): 1026.1807
   95% confidence interval using above SE: 38255.4401 to 42284.5142
   Lower and upper margin of error: -2014.5370, 2014.5370
   95% rank inverse confidence interval: 38839.4006 to 42686.8762
   Lower and upper margin of error using inverse CI: -1430.5765, 2416.8990
   Critical ratio: 39.2426
   p value: 0.000000

Predictor: TIMEPHD

   Coefficient: 1404.8695
   Standard error (SE): 91.7082
   95% confidence interval using above SE: 1224.8334 to 1584.9056
   Lower and upper margin of error: -180.0361, 180.0361
   95% rank inverse confidence interval: 1257.0693 to 1519.2519
   Lower and upper margin of error using rank inverse CI: -147.8003, 114.3824
   Critical ratio: 15.319
   p value: .000000
```

The median starting salary is \$40,270 ±\$2,015. The coefficient for timephd was \$1,405 ± \$180, which was statistically significant (critical ratio (CR) = 15.32, $p < 0.05$). For every additional year of experience, the median salary is predicted to increase by \$1,405.

Here is the output for the 0.75 quantile:

```
MODEL FIT FOR QUANTILE: 0.75

Log likelihood for null model: -8012.9910
Log likelihood for model: -7851.5730
AIC: 15707.1500
BIC: 15716.3900
Pseudo R squared: 0.2114
```

The percent improvement in fit of the model relative to an intercept only model is (0.2114) (100) = 21.1%.

Here is the information for the intercept and coefficient:

```
REGRESSION EQUATION FOR QUANTILE: 0.75

Intercept

   Value of intercept: 45364.0370
   Standard error (SE): 930.8243
   95% confidence interval using above SE: 43536.6980 to 47191.3760
   Lower and upper margin of error: -1827.3390, 1827.3390
   95% rank inverse confidence interval: 43970.1386 to 46618.7506
   Lower and upper margin of error using inverse CI: -1393.8984, 1254.7136
   Critical ratio: 48.7353
   p value: 0.000000

Predictor: TIMEPHD

   Coefficient: 1423.1542
   Standard error (SE): 90.6498
   95% confidence interval using above SE: 1245.1959 to 1601.1126
   Lower and upper margin of error: -177.9584, 177.9584
   95% rank inverse confidence interval: 1280.6630 to 1565.1999
   Lower and upper margin of error using rank inverse CI: -142.4912, 142.0457
   Critical ratio: 15.6995
   p value: 0.000000
```

The starting salary for those at the third quartile is \$45,364 ±\$1,827. The coefficient for timephd was \$1,423 ± \$178, which was statistically significant (critical ratio (CR) = 15.70, $p < 0.05$). For every additional year of experience, the third quartile for salary is predicted to increase by \$1,423.

The program next provides formal contrasts that evaluate the null hypothesis of equal coefficients for each pair of coefficients at the different quantiles. Here is the output:

```
CONTRASTS COMPARING COEFFICIENTS FOR DIFFERENT QUANTILES

Degrees of freedom for each contrast: 1499

Predictor: TIMEPHD

  .250 vs. .500:   t = .8101       p value = .41803
  .250 vs. .750:   t = .4717       p value = .63718
  .500 vs. .750:   t = .2170       p value = .82821
```

None of the contrasts was statistically significant, so the coefficients at each quartile are roughly comparable (be careful not to "accept" the null hypothesis). If you want, you can control the familywise error rate across the contrasts. This is not necessary in the present case because none of the contrasts were statistically significant without such controls.

Here is how we might write-up these results for a report, assuming we have already explained how we are defining margins of errors (e.g., "Margins of errors (MOEs) are calculated from 95% confidence intervals and are the absolute distance between the lower limit or upper limit of the interval minus the parameter estimate, whichever is larger, unless otherwise noted"):

"A quantile regression analysis was performed evaluating the relationship between annual salary and years since one's Ph.D. for the 0.25, 0.50, and 0.75 quantiles of salary. Coefficient estimation used the Barrodale-Roberts method, with standard errors/confidence intervals estimated using a Huber based sandwich approach (Koenker, 2005). The pseudo R squared for the first, second, and third quartiles were 0.21, 0.20 and 0.21, respectively, with associated log-likelihoods of -7,844.6, -7811.0 and -7851.6. The corresponding log likelihoods for the intercept only models were -8,022.8, -7976.4, and -8,013.0, respectively. Table 1 presents the regression coefficients, their associated margins of error, and their critical ratios. The coefficients for the time since the doctorate were similar in magnitude (near $1, 410, all $p < 0.05$) for each quartile. Formal contrasts of the coefficients comparing a coefficient for a given quartile with the coefficient for the other quartiles did not reveal any statistically significant results (for quantiles 0.25 versus 0.50, t ratio = 0.81, $p < 0.42$; for quantiles 0.25 versus 0.75, t ratio = 0.47, $p < 0.64$; for quantiles 0.50 versus 0.75, t ratio = 0.22, $p < 0.83$)."

Table 1: Results of Quantile Regression Analysis

	Coefficient	Critical Ratio
0.25 Quantile		
Time since Ph.D	1,427 ±175	16.51*
Intercept	33,909 ±2,007	-
0.50 Quantile		
Time since Ph.D	1,404 ±180	15.32*
Intercept	40,270 ±2,015	-
0.75 Quantile		
Time since Ph.D	1,423 ±177	15.70*
Intercept	45,364 ±1,827	-

(Notes: N = 750, * $p < 0.05$)

DECILES

If all you want is to compare the distributions of two groups, then a simpler approach than quantile regression is to compare the deciles of the groups. The ASA software has such an option. Our example that illustrates comparing deciles for two independent groups again focuses on the annual salary of starting professors (called *salary* in the data set). We want to test if there are differences in the decile values of salary as a function of gender (the variable called *dfemale*, scored 1 = female, male = 0). We are not concerned about confounds or covariates; we just want to determine if there is a gender difference in the salary deciles ignoring everything else. The program is “Two independent groups: Deciles” in the folder “Robust Statistics > Group Differences in Entire Distributions” of the ASA software. We use the default options associated with the program.

Here is the output for the 0.10 decile:

```
DECILE 0.10

DFEMALE group 0 decile value: 44140.0406
95% confidence interval: 42899.2395 to 45498.1745
Margin of error: -1240.8011, 1358.1338
```

```
DFEMALE group 1 decile value: 40564.0969
95% confidence interval: 39265.4624 to 41744.7322
Lower and upper margin of error: -1298.6345, 1180.6353

Decile difference: 3575.9437
95% confidence interval: 1613.4913 to 5466.7170
Lower and upper margin of error: -1962.4524, 1890.7733
p value: 0.000

95% simultaneous confidence interval for difference: 1023.21 to 6128.68
Simultaneous lower/upper margin of error for difference: -2552.74, 2552.74
```

The estimated salary cut point for the lower 10% of males is $44,140 (lower margin of error (MOE) = -$1,241, upper MOE = $1,358) and for females it is $40,564 (lower MOE = -$1,298, upper MOE = $1,181). The difference in these income levels is $3,576 (lower MOE = -$1,962, upper MOE = $1,891). These confidence intervals are unadjusted for the fact that nine different contrasts were performed. If you want to adjust for this fact, then you would use the simultaneous confidence intervals for the difference. If a 95% simultaneous confidence interval does not contain the value of 0, then the contrast is statistically significant at $p < 0.05$ after making familywise adjustments.

Here are the results for the other deciles:

```
DECILE 0.20

   DFEMALE group 0 decile value: 48250.4145
   95% confidence interval: 46730.3075 to 49795.8278
   Margin of error: -1520.1069, 1545.4133

   DFEMALE group 1 decile value: 44948.1822
   95% confidence interval: 43611.9371 to 45801.2784
   Lower and upper margin of error: -1336.2451, 853.0962

   Decile difference: 3302.2322
   95% confidence interval: 1449.9718 to 5299.2223
   Lower and upper margin of error: -1852.2604, 1996.9900
   p value: 0.002

   95% simultaneous confidence interval for difference: 605.96 to 5998.50
   Simultaneous lower/upper margin of error for difference: -2696.27, 2696.27


DECILE 0.30

   DFEMALE group 0 decile value: 51168.9471
   95% confidence interval: 50115.5061 to 52407.5630
   Margin of error: -1053.4410, 1238.6159

   DFEMALE group 1 decile value: 47985.3952
   95% confidence interval: 46629.6388 to 49223.5114
   Lower and upper margin of error: -1355.7563, 1238.1163
```

```
   Decile difference: 3183.5519
   95% confidence interval: 1349.6708 to 5171.7822
   Lower and upper margin of error: -1833.8811, 1988.2303
   p value: 0.000

   95% simultaneous confidence interval for difference: 475.62 to 5891.48
   Simultaneous lower/upper margin of error for difference: -2707.93, 2707.93


DECILE 0.40

   DFEMALE group 0 decile value: 53790.8700
   95% confidence interval: 52614.0973 to 55357.6387
   Margin of error: -1176.7727, 1566.7687

   DFEMALE group 1 decile value: 50989.2486
   95% confidence interval: 49840.8496 to 51947.7243
   Lower and upper margin of error: -1148.3990, 958.4756

   Decile difference: 2801.6214
   95% confidence interval: 1143.1269 to 4682.0985
   Lower and upper margin of error: -1658.4944, 1880.4771
   p value: 0.000

   95% simultaneous confidence interval for difference: 526.46 to 5076.78
   Simultaneous lower/upper margin of error for difference: -2275.16, 2275.16


DECILE 0.50

   DFEMALE group 0 decile value: 56829.4791
   95% confidence interval: 55223.7910 to 57842.1127
   Margin of error: -1605.6880, 1012.6336

   DFEMALE group 1 decile value: 53293.3518
   95% confidence interval: 52312.5960 to 54254.5053
   Lower and upper margin of error: -980.7558, 961.1535

   Decile difference: 3536.1272
   95% confidence interval: 1799.4979 to 5193.4712
   Lower and upper margin of error: -1736.6293, 1657.3440
   p value: 0.000

   95% simultaneous confidence interval for difference: 1203.59 to 5868.67
   Simultaneous lower/upper margin of error for difference: -2332.54, 2332.54


DECILE 0.60

   DFEMALE group 0 decile value: 59196.1154
   95% confidence interval: 57834.1815 to 60230.9819
   Margin of error: -1361.9339, 1034.8665

   DFEMALE group 1 decile value: 55691.9892
   95% confidence interval: 54385.5250 to 56300.3327
   Lower and upper margin of error: -1306.4642, 608.3434
```

```
   Decile difference: 3504.1261
   95% confidence interval: 1978.4267 to 5010.6689
   Lower and upper margin of error: -1525.6995, 1506.5428
   p value: 0.000

   95% simultaneous confidence interval for difference: 1189.44 to 5818.81
   Simultaneous lower/upper margin of error for difference: -2314.68, 2314.68


DECILE 0.70

   DFEMALE group 0 decile value: 61733.4897
   95% confidence interval: 60281.6051 to 62545.7023
   Margin of error: -1451.8845, 812.2127

   DFEMALE group 1 decile value: 57887.4908
   95% confidence interval: 56548.8187 to 58992.1508
   Lower and upper margin of error: -1338.6721, 1104.6600

   Decile difference: 3845.9989
   95% confidence interval: 1891.4456 to 5366.1195
   Lower and upper margin of error: -1954.5533, 1520.1206
   p value: 0.000

   95% simultaneous confidence interval for difference: 1563.38 to 6128.62
   Simultaneous lower/upper margin of error for difference: -2282.62, 2282.62


DECILE 0.80

   DFEMALE group 0 decile value: 64449.8877
   95% confidence interval: 63026.8626 to 65795.4910
   Margin of error: -1423.0251, 1345.6033

   DFEMALE group 1 decile value: 60815.2419
   95% confidence interval: 59750.8054 to 61737.6840
   Lower and upper margin of error: -1064.4365, 922.4421

   Decile difference: 3634.6459
   95% confidence interval: 1924.6214 to 5452.8686
   Lower and upper margin of error: -1710.0245, 1818.2227
   p value: 0.000

   95% simultaneous confidence interval for difference: 1264.92 to 6004.37
   Simultaneous lower/upper margin of error for difference: -2369.72, 2369.72


DECILE 0.90

   DFEMALE group 0 decile value: 68945.4645
   95% confidence interval: 67201.6875 to 70824.9374
   Margin of error: -1743.7770, 1879.4729

   DFEMALE group 1 decile value: 64567.7126
   95% confidence interval: 63327.3194 to 65847.6735
   Lower and upper margin of error: -1240.3932, 1279.9609
```

```
Decile difference: 4377.7519
95% confidence interval: 2055.5266 to 6657.1927
Lower and upper margin of error: -2322.2253, 2279.4408
p value: 0.000

95% simultaneous confidence interval for difference: 1183.58 to 7571.93
Simultaneous lower/upper margin of error for difference: -3194.17, 3194.17
```

The salary differences are statistically significant at each decile and generally hover around $3,000 to $3,500. Some economists would argue that the differences actually are more exaggerated at the lower end of the distribution because a $3,500 or so difference among individuals who earn about $45,000 per year means more than a $3,500 or so difference among individuals who earn closer to $70,000 per year.

Here is how we might write-up the results, assuming we have already explained how we are defining margins of errors:

“The deciles for annual salary for males and females were compared using the percentile bootstrap method from Wilcox (2017). The p values and confidence intervals are per comparison rather than familywise adjusted. Table 1 presents the estimated male income, the estimated female income and the difference between them for each decile. In each case, the male salary was statistically significantly greater than the female salary.”

Table 1: Salary Deciles for Males and Females

Decile	Male Income	Female Income	Difference
0.10	44,140 (-1240, 1358)	40,564 (-1298, 1181)	3,575 (-1962, 1891)*
0.20	48,250 (-1520, 1545)	44,948 (-1336, 853)	3,302 (-1852, 1997)*
0.30	51,169 (-1053, 1238)	47,985 (-1356, 1238)	3,184 (-1834, 1988)*
0.40	53,791 (-1178, 1567)	50,989 (-1148, 958)	2,802 (-1658, 1800)*
0.50	56,829 (-1606, 1013)	53,293 (-981, 961)	3,536 (-1737, 1657)*
0.60	59,196 (-1362, 1035)	55,692 (-1306, 608)	3,504 (-1526, 1507)*
0.70	61,733 (-1452, 812)	57,887 (-1339, 1105)	3,845 (-1955, 1520)*
0.80	64,450 (-1423, 1345)	60,815 (-1064, 922)	3,635 (-1710, 1818)*
0.90	68,945 (-1744, 1879)	64,568 (-1240, 1280)	4,378 (-2322, 2279)*

(Note: Estimates are Harrell-Davis; Lower and upper margins of error are in parentheses are percentile bootstrap estimated with 599 replicates – see Wilcox, 2017; * $p < 0.05$)

A NOTE ON ANALYZING RATIOS

As noted, when documenting economic disparities, some economists do not examine the absolute difference in incomes but instead focus on ratios as a form of standardization. For example, at the 0.10 decile, the ratio of male salaries to female salaries is $44,140 / $40,564 = 1.09, so males earn about 1.09 times more than females. At the 0.90 decile, the ratio of male salaries to female salaries is $68,945 / $64,568 = 1.07, or about 7% more. Note that one needs ratio level measures for such statements to be meaningful. To focus the analysis on such ratios, one strategy is to analyze the log salary of each group because log(A/B) = log(A) - log(B). One would then calculate the anti-log of the reported decile difference to obtain the ratio. For example, here is the output for the 0.90 decile when I re-ran the above using the natural log of salary instead of salary as the outcome variable:

```
DECILE 0.90

   DFEMALE group 0 decile value: 11.1410
   95% confidence interval: 11.1153 to 11.1653
   Margin of error: -.0257, .0243

   DFEMALE group 1 decile value: 11.0754
   95% confidence interval: 11.0522 to 11.1005
   Lower and upper margin of error: -.0232, .0251

   Decile difference: .0655
   95% confidence interval: .0298 to .1050
   Lower and upper margin of error: -.0357, .0394
   p value: 0.000

   95% simultaneous confidence interval for difference: .0180 to .1131
   Simultaneous lower and upper margin of error for difference: -.0475, .0475
```

The decile difference is 0.0655 and the exponent (anti-log) of 0.0655 is 1.07. The exponents of the 95% confidence interval limits are exp(0.0298) = 1.03 and exp(0.1050) = 1.11, so the lower margin of error for the ratio is 1.03 - 1.07 = -0.04 and the upper margin of error is 1.11 - 1.07 = 0.04. Thus, at the 0.90 quantile, males are paid 1.07 ±0.04 times more than females, or about 7% more. We could obtain a significance test of the male and female salary difference expressed as ratios between any two deciles (e.g., that for the 0.10 decile as tested against the 0.90 decile) using the ASA quantile regression program with logged salaries as the outcome.

REFERENCES

Curran, J. M. (2005). An introduction to Bayesian credible intervals for sampling error in DNA profiles. *Law, Probability and Risk*, 4, 115-126.

Koenker, R. (2005). *Quantile regression*. New York: Cambridge University Press.

Wilcox, R. (2017). *Introduction to robust estimation and hypothesis testing*. San Diego: Academic Press (Fourth edition).
.