a linear function of *x,* or at least we attempt to fit a linear function to the conditional expectation, or *regression function.* Instead of the mean, we might choose to work with the median, and to assume that the medians of y conditional on x are linear in *x,* or at least to fit a linear function to the medians. This would be a median regression, or 0.5 quantile regression. In principle, we can do the same for any other quantile of the distribution, thus constructing the p-quantile regression, where *p* is any number between 0 and 1.

Given the idea, why should we be interested, and if we are interested, how can such regressions be calculated? Start with the former. First, by looking at a number of different quantile regressions, we can explore different parts of the conditional distribution. For example, consider the relationship between wages and schooling; at any given number of years of schooling, there is a (conditional) distribution of wages, presumably reflecting unobserved abilities and other labor market skills. In general, thereis no reason to require that the rate of return to an additional year's schooling should be the same at all points in the distribution of abilities conditional on schooling, and quantile regression would pick up the differences, see Chamberlain (1991). Used in this way, quantile regression is essentially a non-parametric technique that describes the shape of the empirical distribution without imposing prior restrictions. As such, it can also provide an indication of heteroskedasticity. If the conditional distribution changes shape with one or more of the explanatory variables, quantile regressions at different quantiles will have different slopes, [see Koenker and Bassett (1982)] for a test that uses this property.

Second, just as the median is Iess sensitive to outliers than is the mean, so are quantile regressions more resistant to outliers than are mean (least-squares) regressions. Median regression is affected by the presence of an outlier, but not by changes in its position, provided of course that it remains above or below the median. As such, quantile regression is one of several regression techniques that have robustness properties superior to OLS [see in particular Huber (1981) and Hampel, Ronchetti, Rousseeuw and Stahel (1986)1. Standard methods of robust regression typically downweight large residuals identified from a previous regression, iterating to convergence. Such procedures require an estimate of the scale of the residuals in order to identify outliers, and thus are sensitive to patterns of heteroskedasticity that are handled naturally by quantile regressions.

Third, (pantiles are not affected by monotonic transformations of the data, so that, for example, the median of the logarithm of y conditional on x is the logarithm of the median of y conditional on *x.* As we shall see in the next subsection, this property has useful consequences.

The estimation of quantile regressions rests on extensions of the well-known result that the median is the point closest to the data in the sense of minimizing the sum of the absolute deviations. Median linear regression parameters are