Econ. 556b Spring 1999 Chris Sims

Bayesian Practice

This exercise is meant to give you some experience in Bayesian data analysis, and to illustrate how numerical methods can give conclusions on otherwise intractable problems.

Consider the tobit regression model with zero intercept, in which

$$y_i^* = \beta x_i + \varepsilon_i \, , i = 1, \dots , N \tag{1}$$

$$y_i = \max\{0, y_i^*\} \, . \tag{2}$$

The data are i.i.d. across i, with x_i independent of u_i and $u_i \sim N(0, \sigma^2)$. We observe y_i and x_i for each i, but not y_i^*.

a. Find an expression for the likelihood function for this model.
b. Suppose that we use y and x to refer to the data stacked into $N \times 1$ column vectors, and the data we have are

$$x = \begin{bmatrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \end{bmatrix} \, , \quad \begin{bmatrix} 0 \\ 0 \\ 3.31 \\ 4.72 \\ 2.13 \\ 8.98 \end{bmatrix} \, .$$

Use Matlab's `contour` command or related facilities (if they are available) in Gauss or another program, to display contours of the log likelihood and of the likelihood itself, as a function of β and σ. Also show on your plot the location of the OLS estimate of β, and compare the contour plot of the actual likelihood function with what you would obtain if you ignored the truncation and treated y as if it were y^* (i.e. used the standard OLS assumptions).

Repeat your analysis in (b) for the following two y vectors:

c.

$$\begin{bmatrix} 3.47 \\ 0.91 \\ 2.32 \\ 2.44 \\ 2.03 \\ 4.81 \end{bmatrix}$$

d.

$$\begin{bmatrix} 0.26 \\ 0 \\ 2.24 \\ 0 \\ 4.29 \\ 7.06 \end{bmatrix}$$

e. Show that in the case of (c), the likelihood is just what it would be if we ignored the truncation. Would the Bayesian flat-prior 95% probability interval for β in this case be close to a classical 95% confidence interval computed from the correct model? Why or why not?