Introduction to the R Statistical Computing Environment
Getting Started With R: Exercises

John Fox
(McMaster University)
ICPSR Summer Program

2010

The two graphs reproduced below are meant to explicate the proportional-odds logistic-regression model, described in Section 5.9 of the *R Companion*. In these graphs, there is a single predictor variable, x, and a four-category ordinal response variable y. Try to duplicate these graphs using R.

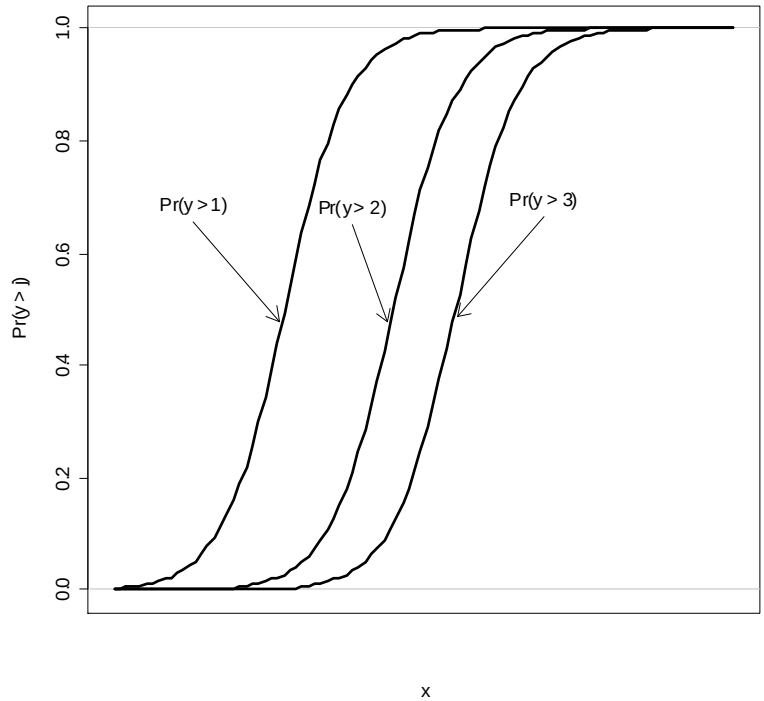

The first graph is similar to Figure 5.8 in the R Companion and is relatively simple to construct. Some hints:

- You can use the `plogis()` or the similar `pnorm()` function to compute cumulative logistic probabilities.

- You can use the mouse to find coordinates for the arrows and the text labeling the curves.

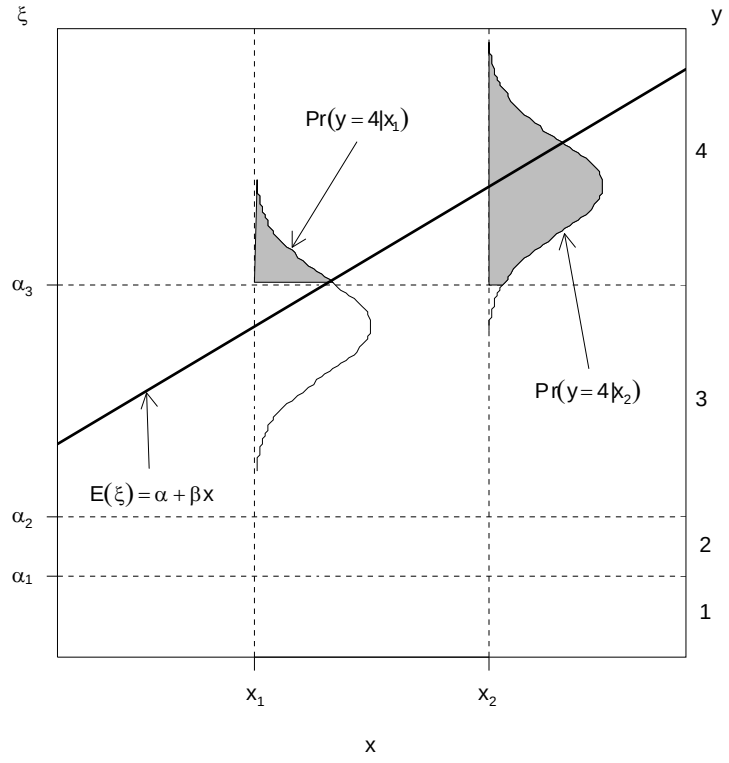

The second diagram is a much more challenging graph, similar to Figure 9.2 in Agresti's *Categorical Data Analysis* (Wiley, 1990), but nicer! The left vertical axis gives the latent continuous response variable ξ, with thresholds at α_1, α_2, and α_3, while the right vertical axis gives the observed ordinal response variable y, with values 1, 2, 3, and 4. The graph shows the regression line, along with the probability that $y = 4$ at two different x-values. Hints:

- All of the techniques required for constructing this graph were covered in the workshop and in Chapter 7 of the *R Companion*.
- I used the normal density function `dnorm()` to draw the curves, figuring that this would be visually indistinguishable from using the logistic density, but you could also use `dlogis()`.
- Most of the text in the graph was positioned with the mouse. Remember that you have to set the argument `xpd=TRUE` in a call to `text()` to write outside of the plotting region. An alternative would be to use `mtext()` to place the text in the plot margins.