Advanced Programming in Quantitative Economics

Introduction, structure, and advanced programming techniques

Charles S. Bos

VU University Amsterdam
Tinbergen Institute

c.s.bos@vu.nl

15 – 19 August 2011, Aarhus, Denmark

Tutorial Day 3 - Morning

10.30P Tutorial

- Loglikelihood duration model
- Standard deviations
- Analytical score

12.00 Lunch

Duration: The (simplified) model

As a reminder: Durations y_i are assumed to be distributed according to

$$y \sim \text{Weib}(\alpha, \lambda) \qquad f(y; \alpha, \lambda) = \alpha\lambda^{\alpha} y^{\alpha-1} \exp(-\lambda^{\alpha} y^{\alpha})$$

Dependence on personal characteristics can be introduced by taking

$$\lambda_i \equiv \exp(X_i\beta)$$
$$y_i \sim \text{Weib}(\alpha, \lambda_i)$$

Make sure you have data available, e.g. in `data/genrdur.fmt`, or download a file from the web.

Duration: Optimisation

Central to optimisation is the log likelihood function. In this case, it would read

$$\lambda_i \equiv \exp(X_i\beta)$$

$$\begin{aligned}\log \mathcal{L}(y;\theta) &= \sum_i (\log\alpha + \alpha\log\lambda_i + (\alpha-1)\log y_i - \lambda_i^\alpha y_i^\alpha) \\ &= N\log\alpha + \alpha\sum_i \log\lambda_i + (\alpha-1)\sum \log y_i - \sum(\lambda_i y_i)^\alpha\end{aligned}$$

Work on this in steps...

Duration: Steps

Perform, in steps, for instance

1. Get the outline of your loglikelihood function. Call it from main, with a valid vector of parameters, and set the likelihood value equal to the average of your y's.
2. Extract the vector of parameters, into β and α. Print them separately from the loglikelihood function.
3. Check the value of α. If negative, maybe return a zero?
4. Construct a vector of λ's. Does this work?
5. Construct full loglikelihood function. Does the value seem 'logical'?
6. Run `MaxBFGS()`. What return value `ir` do you get, what does it mean? What is `vP`?

Duration: Standard errors

For the standard errors, you had to find

$$\Sigma(\hat{\theta}) = -H(\hat{\theta})^{-1}$$

$$H(\hat{\theta}) = \left. \frac{\delta^2 l(Y;\theta)}{\delta\theta\delta\theta'} \right|_{\theta=\hat{\theta}}$$

Some standard code could look like

```
if (Num2Derivative(AvgLnLiklRegr, avP[0], &mH))
  mS2= invertgen(-mH, 30)/iN,

avS[0]= sqrt(diagonal(mS2)');
```

7. Get the standard errors with it. How do they change if you only use the first 10 observations?
8. Beautify the output: Get a nice print with the maximum likelihood you find, the type of convergence, the parameters, standard errors and t-values

Duration: Analytical score

As an extra: Work out the analytical score for the model

9. Find it on paper
10. In the program, add the necessary code (only compute it if asked for it, if `avScore` is an address!)
11. Check it, contrasting your analytical score to `Num1Derivative`
12. Evaluate the number of function evaluations needed when using numerical scores, and when using analytical scores (hint: Define an extra global variable, `s_iEval`)