BAYESIAN ECONOMETRICS: Tutorial 2 (Regression)
Professor: Luc Bauwens
Coimbra, April 2013

Exercise 1
This exercise illustrates Bayesian inference in a linear regression model.

Consider the following wage equation based on the Phillips's curve which relates the growth rate of real wages $\Delta \log(W_t/P_t)$ to the growth rate of the consumption price index $\Delta \log P_t$, the lagged growth rate of labour productivity $\Delta \log Q_{t-1}$, and the rate of unemployment UR_t:

$$\Delta \log(W_t/P_t) = \beta_0 + \beta_1 \Delta \log P_t + \beta_2 \Delta \log Q_{t-1} + \beta_3 UR_t + \epsilon_t$$

where we assume that ϵ_t are independent and $N(0, \sigma^2)$. The data consist of 22 annual observations for Belgium, covering the period 1954-1976. They are available in the data file wage.txt. This file contains indications on which series corresponds to each variable in the above equation.

The prior density for $\beta = (\beta_0, \beta_1, \beta_2, \beta_3)'$ and σ^2 is one of the following:

- Prior 1 (non-informative): $\varphi(\beta, \sigma^2) \propto 1/\sigma^2$.
- Prior 2 (partially non-informative conjugate): you are asked to be informative only on β_1 (the coefficient of inflation) such that there is no monetary illusion on the relation, implying that prior mean of β_1 be equal to 0. You are asked that a priori that (marginally with respect to σ^2) $\Pr(-0.1 < \beta_1 < 0.1) = 0.68$ assuming a normal distribution (actually it is a Student one, but just do like it is a Normal), so that the interval $(-0.1, 0.1)$ corresponds to two standard deviations. For the prior on σ^2, fix s_0 equal to the OLS SSR (the value is ???) and $\nu_0 = 3$ (see slide 114).
  Find the matrix M_0 and the vector β_0 of the prior (see slide p 111).
- Prior 3 (Normal-diffuse): assume a Normal prior for β_1 with mean 0 and variance deduced from $\Pr(-0.1 < \beta_1 < 0.1) = 0.68$.
  Find M_0 and β_0 of this prior.
- Prior 4 (Beta-diffuse prior): replace the prior for β_1 of Prior 3 by a symmetric Beta density on the interval $(-0.3, 0.3)$. NB: β_1 has a symmetric Beta density on $(-c, +c)$ where $c > 0$, with parameter a, if its density is given by $\frac{\Gamma(2a)}{\Gamma(a)\Gamma(a)}(2c)^{-(2a-1)}[(y+c)(c-y)]^{a-1}$. The mean is 0, and the variance is $(2c)^{-2}(4a+2)^{-1}$. Given that $c = 0.3$, choose a such that the variance of β_1 is the same as for Prior 3.

Write a computer program that

1. computes analytically the posterior results for σ^2 and β for Prior 1 and Prior 2;[1]
   NB: These results are known analytically in both cases, so only simple matrix computations are required. Comment the results, in particular how the posterior revises the conjugate prior (compared to the non-informative prior).

2. computes the posterior results (under Prior 2) by the Gibbs sampling algorithm sketched on slide 108 (check that you get results that are close enough to the analytical results.);
   NB: compute the posterior results directly with the sampled draws, and also by using Rao-Blackwellisation.

3. computes the posterior results under Prior 3;
   NB: for this part, use also a Gibbs sampling algorithm (see slide 123).

4. computes the posterior results under Prior 4.
   NB: use a MH step in the previous Gibbs sampling algorithm (see slide 127). Explain what proposal density you use. If possible, try different proposals and discuss how they affect the algorithm.

   NB: for results based on MCMC algorithms (items 2-3-4), do the necessary convergence checks (tests and graphs).

[1] By posterior results here and below we mean: the posterior expectations and standard deviations of β and σ^2, the posterior correlation matrix of β, the posterior probability that $-0.1 < \beta_1 < 0.1$, and a graph of the posterior marginal densities of β_1 and of σ^2.