BAYESIAN ECONOMETRICS: Tutorial 1 (Numerical Methods)
Professor: Luc Bauwens
Coimbra, April 2013

Exercise 1
This exercise illustrates the methods of deterministic integration, importance sampling, and the Metropolis-Hastings algorithm for Bayesian inference on a scalar parameter.

Consider the independent Student sampling process, for $i = 1, \ldots, n$:
$y_i|\mu \sim I.t(\mu, \nu-2, 1, 5)$, a Student distribution such that $\mathrm{E}(y_i) = \mu$ and $\mathrm{Var}(y_i) = 1$ (for simplicity, we assume ν known and > 2). Write a computer program that:

1. generates a sample of this process (inputs to choose are n, μ, and ν) (see Appendix B of Bauwens, Lubrano and Richard (1999) (BLR) for an algorithm to simulate the Student distribution);[1]

2. computes the ML estimator of μ and its asymptotic standard error (use a library for optimization, such as Maxlik in GAUSS);

3. computes and plots (on the same graph) the posterior densities of μ corresponding to the sampled data and to two prior densities:
   -diffuse,
   -$N(\mu_0, 1/n_0)$ (inputs to choose are μ_0 and n_0);
   NB: To define your prior, you have to choose values of μ_0 and n_0. You can pick arbitrary values and repeat the exercise by changing them to see how they influence the posterior results. An easy way to select μ_0 and n_0 is to sample n_0 initial observations from the DGP and to set μ_0 at the mean of this sample.

4. computes the posterior mean and variance of μ;

5. plots the Student and normal prior densities on the same plot as the corresponding posterior densities.

Use the following methods for doing the computations needed for 3-4-5 above:

DI: Deterministic integration using the trapezoidal or Simpson's rule to compute the posterior moments and the normalizing constant of the posterior of μ for each prior; you must normalize the posterior, i.e. it must integrate to 1, to plot it on the same graph as the other densities.

IS: Importance sampling to compute the posterior moments and marginal density of μ for each prior. Explain how you choose the importance function in each case. Compute the probabilistic relative error bound of the estimate of the normalizing constant of the posterior and of the estimate of the posterior expectation of μ (see formula 3.34 in BLR).

[1] A GAUSS code is available at http://perso.uclouvain.be/luc.bauwens/Bayes/bayesian.htm

MH: the independent MH algorithm to compute the posterior moments and marginal density of μ for each prior. As candidate density, use the importance function you have defined for the IS computations here above. Estimate the rejection probability of the MH sampler by the empirical rejection frequency. Compute the numerical standard error of the posterior mean. Plot the ACF and CUMSUM of the draws and assess the dependence in the generated sample. Apply Geweke's test.

Exercise 2
This exercise illustrates the method of direct sampling with a simple example.

Assume that θ is a parameter in R^2 with as posterior distribution a $N_2(\mu, \Sigma)$ density (see p 48 of the slides). Write a computer program that has as inputs μ, Σ, intervals bounds for θ_1 and θ_2 (see NB below), n and

1. simulates n random draws from the bivariate normal distribution defined above;

2. estimates (in the frequentist sense) μ, Σ, $\Pr(\theta_1 \in A_1)$, and $\Pr(\theta \in A)$ from the draws, and plots these estimates against increasing values of n together with the limits of 95 per cent confidence intervals for them. Take for example $n = 10000$ and plot the estimates for smaller values of n (but use always the same sample, i.e. for $n = 100$, base the estimates on the first 100 values of the sample of size 10 000).

3. estimates the mean of θ_1/θ_2 (; you should see that it is unstable as n changes, reflecting that the expected value does not exist. Estimate also the median of θ_1/θ_2. Is it more stable?

NB: Choose for example $\mu = (1 \; -1)'$, and $\Sigma = \begin{pmatrix} 0.25 & 0.2 \\ 0.2 & 0.64 \end{pmatrix}$, and intervals $A_1 = (0.5, 1.2)$ for θ_1 and $A_2 = (-1.5, 0)$ for θ_2. Define the set A as the Cartesian product of A_1 and A_2).

Exercise 3
This exercise illustrates the Gibbs sampler with a simple example.

Assume that θ is a parameter in R^2 with as posterior distribution a $N_2(0, \Sigma)$ where Σ is a correlation matrix with ρ as the off-diagonal element. Write a computer program that has as inputs ρ, n_0, n and

1. simulates $n_0 + n$ random draws of θ from the bivariate normal distribution defined above, by a Gibbs sampling algorithm that cycles between $\theta_1|\theta_2$ and $\theta_2|\theta_1$;

2. estimates from the last n draws the two means, the two variances and the correlation coefficient from the Gibbs sample;

3. plots the ACF and CUMSUM of the draws of θ_1 and of θ_2 and assess the dependence in each of them;

4. computes the numerical standard error of each estimated posterior mean;

5. applies Geweke's test to the draws of θ_1 and of θ_2;

6. estimate and plots each marginal density using the simulated sample by a kernel method if you know one (otherwise plot a smoothed histogram). Do the same, if you can, for the joint bivariate density.

7. (optionally) does also the computations using Rao-Blackwellisation.

NB: you should execute the program varying
- the initial condition and the length of the warm-up sample (n_0) to assess how they influence (or not) the results (for a given value of ρ);
- the value of ρ (try 0, 0.50, 0.90 and 0.99) to assess the influence of the correlation coefficient between θ_1 and θ_2 on the performance of the Gibbs sampler for a given value of n_0 and n (see p 83 of the slides).