Exercises for Lecture 1

Exercise 1: Consider the following AR(1) process, initialized in the infinite past:

$$y_t = \phi y_{t-1} + u_t, \tag{1}$$

where $u_t \sim iidN(0,1)$.

(i) Suppose you have a sample of observations $Y_{1:T} = \{y_0, y_1, \dots, y_T\}$. Derive the conditional likelihood function $p(Y_{1:T}|\phi, y_0)$ for θ.

(ii) Consider the following prior for ϕ: $\phi|y_0 \sim N(0, \tau^2)$. Show that the posterior distribution of ϕ is of the form

$$\phi|Y_{0:T} \sim N(\bar{\phi}_T, \bar{V}_T), \tag{2}$$

and provide expressions for $\bar{\phi}_T$ and $\bar{V}_T$.

(iii) Simulate a data set from the AR(1) model. Let $y_0 = 0$, $\phi = 0.95$, and $T = 100$. Plot your simulated data set.

(iv) Compute the posterior mean of ϕ and ϕ^2 using the analytical formulas derived in step (ii).

(v) Now code up a direct sampler that generates N draws from $N(\bar{\phi}_T, \bar{V}_T)$. Based on these draws, generate a Monte Carlo approximation of $\mathbb{E}[\phi|Y]$ and $\mathbb{E}[\phi^2|Y]$. Run your sampler $N_{run} = 100$ times and compare the sampling variation of the Monte Carlo approximation to the asymptotic variance derived in class. Do so for different choices of N.

(vi) Generate draws from the predictive distribution for y_{T+h} given $Y_{1:T}$. You can do so using Algorithm 2 in Del Negro and Schorfheide (2013) "DSGE Model-Based Forecasting." Plot a fan chart for future y_t's as in, e.g., the Bank of England inflation reports:
http://www.bankofengland.co.uk/publications/Pages/inflationreport/irfanch.aspx

Exercise 2: Importance sampling

Suppose the posterior density $\pi(\theta)$ takes the form of a $N(0,1)$. Rather than using direct sampling, now use importance sampling to generate draws from the posterior. Your importance sampling density $g(\theta)$ is the density of a student t distribution with $\nu = 4$ degrees of freedom.

(i) Suppose you are interested in the posterior mean $\mathbb{E}_\pi[h(\theta)]$ for $h(\theta) = \theta$ and $h(\theta) = \theta^2$. Write a program that implements the importance sampling algorithm and can generate Monte Carlo approximations of $\mathbb{E}_\pi[h(\theta)]$. You can let

$$f(\theta) = \exp\left\{-\frac{1}{2}\theta^2\right\}.$$

(ii) Following the steps outlined on Page 35 of HS, derive the asymptotic covariance matrix $\Omega(h)$ in Equation (3.32).

(iii) Choose $N = 100$. Run your importance sampler $N_{run} = 100$ times and compute the bias and the variance of the Monte Carlo approximation of $\mathbb{E}_\pi[h(\theta)]$. Now vary N, compute the (small sample) inefficiency factors described in HS, and generate your version of HS's Figure 3.2. For one of the runs plot the distribution of the importance weights.

(iv) Now reverse the roles of the normal and the t distribution, that is, assume that the posterior distribution corresponds to a t distribution with $\nu = 4$ degrees of freedom. You can let

$$f(\theta) = \left(1 + \frac{\theta^2}{\nu}\right)^{-(\nu+1)/2}.$$

Moreover, let $g(\theta)$ correspond to the $N(0,1)$ distribution. Repeat the experiment in (iii) and discuss the results.