EXERCISE ON AN AR MODEL

From FRED, the St. Louis Fed macro data web site, obtain the data for US real GDP, seasonally adjusted, chain-linked, quarterly, from 1947-2017.

(1) Fit an AR(2), with a constant term, to these data, conditioning on the two initial values, by maximum likelihood assuming normal errors. This is of course just an OLS estimation.
(2) Use these estimates to make 1000 draws from the flat-prior posterior joint distribution of the constant, the two AR coefficients, and the residual variance. No MCMC is needed, since the posterior distribution has a known Normal-Inverse-Gamma form. (Draw the inverse of sigma from the appropriate gamma distribution, then the constant and AR coefficients from their normal distribution conditional on the drawn sigma.) A program that does this in R is below.
(3) Use the draws to generate random draws of the roots of the system (either the roots of the AR polynomial in the lag operator, or their inverses, the eigenvalues of the 2 by 2 system matrix). Use your draws to estimate the probability that there is a root in the unstable region.
(4) Make a scatter plot of the draws of the two AR coefficients.
(5) Generate 1000 draws from the distribution under a flat prior, given the data, of log GDP in the fourth quarter of 2027. You can do this by, for each of your 1000 draws of the model parameters, generating a simulated time path of GDP over the 40 quarters. To do that, recursively forecast one step ahead, adding a randomly drawn normal shock of the appropriate variance at each step. A program that does this in R is below. Display a histogram for the average annual growth rate over the 10 year horizon (i.e., log GDP in 2027:4 - log GDP in 2017:4, divided by 10).
(6) Using the sorted randomly drawn growth rates, estimate the probability of negative average growth over the 10 years and the probability of growth greater than 4 per cent per year over the period. Also find the median of the distribution of growth rates.

Date: April 7, 2018.

(7) Using your ML estimates as if they were non-random, calculate the expected annual growth rate over the next 10 years and its variance. The variance of the 10-year change can be calculated analytically. If we set S to be the 2 by 2 matrix with the variance of the disturbance in the upper left corner and zeros elsewhere, and B to be the system matrix of the stacked first-order system, then repeating $V \leftarrow S + BVB'$ 40 times gives you the answer. Compare this mean forecast and its variance to the distribution you obtained above when taking account of uncertainty in the coefficients.

(8) Write a program to evaluate the likelihood of the first two observations using the unconditional density, assuming stationarity. Your sample of draws from the flat-prior posterior, conditional on initial conditions, can in principle be converted to a weighted sample from the flat-prior posterior for the whole sample, including initial conditions, using your program for the unconditional density. Use your program to calculate the importance weights, and explain why they are useless for practical purposes in this sample. [For numerical stability, you should first compute the logs of the importance weights rather than the weights themselves.]

```
g10 <- function(b,c, sig, yic) {
    shocks <- rnorm(40) * sig
    yf <- vector("numeric", 42)
    yf[1:2] <- yic
    for (it in 1:40) yf[2 + it] <- c + b %*% yf[1:0 + it] + shocks[it]
    return(yf)
}

armcmc <-function(bhat, sshat, xx, T, ndraw) {
    kb <- length(bhat)
    xxifac <- solve(chol(xx))
    sdraw <- sqrt(.5 / rgamma(ndraw, (T-kb)/2, sshat))
    bdraw <-  matrix(rnorm(kb * ndraw), ncol=kb)
    bdraw <- (sdraw * bdraw) %*% t(xxifac) + matrix(bhat, ndraw, kb, byrow=TRUE)
    return(list(bdraw=bdraw, sdraw=sdraw))
}
```