KALMAN FILTER AND MCMC EXERCISE

In this exercise you use a simplified version of Stock and Watson's univariate model of the inflation process, estimating using the Kalman filter and MCMC. (They allow time-varying variances, which we do not, for now.)

You will be fitting to data on inflation, defined as $\pi_t = \log(p_t/p_{t-1})$, where p_t is the personal consumption expenditure deflator (PCECTPI on the Fred database at the St. Louis Fed web site). Use quarterly data from 1983 to the present. (Adding earlier data would make the failure to allow for time varying variances more important.)

The model is

$$\pi_t = \tau_t + \eta_t$$
$$\tau_t = \tau_{t-1} + \varepsilon_t$$
$$\begin{bmatrix} \eta_t \\ \varepsilon_t \end{bmatrix} \sim N\left(0, \begin{bmatrix} \sigma_\eta^2 & 0 \\ 0 & \sigma_\varepsilon^2 \end{bmatrix}\right) .$$

η_t and ε_t are independent of τ_{t-1} and are independent across time.

With σ_η^2 and σ_ε^2 given, this model's likelihood can be evaluated with the Kalman filter. You will need an initial distribution for the state τ_t. Make it $N(.005, .0004)$. (Remember these are one-quarter inflation rates, in logs, not per cent, so this is a very dispersed prior — standard error eight per cent at annual rates.) You will also need a prior on the two σ^2's. Make them independent exponentials with parameter 500, i.e. $500\exp(-500\sigma^2)$. This makes the prior mean .002, which is big, but should make results insensitive to the prior.

(i) Find the peak of the posterior density for the two σ^2 parameters by using the Kalman filter together with a hill-climbing unconstrained optimizer. Since the two parameters have to be positive, you might make the hill-climber use their logs as its parameters. Otherwise the function evaluation routine you supply will have to return a penalty value (not just quit with an error message) if given negative values for either σ^2.

(ii) Plot on the same graph the filtered and smoothed values for τ_t at the posterior modal values you have found for the σ^2's.

(iii) Use MCMC — probably straight Metropolis is best — together with the Kalman filter to generate a sample from the posterior distribution of the two σ^2's. Use it to calculate posterior means for the two parameters, and compare them to the posterior modes. Plot the marginal densities of the two parameters. Calculate their posterior correlation. (Be sure to check convergence of your MCMC iterations.)

(iv) As Stock and Watson point out, the best forecast of π_{t+s} from this model, based on information at time t and for any $s > 0$, is the current filtered estimate of τ_t. Calculate the forecast and actual average inflation rates over the year following each of these dates: 2007:III, 2008:III, 2009:II. Also plot the one-step-ahead forecast

Date: October 10, 2010.

errors and one-step-ahead forecast errors divided by their standard errors, for the period 2008:III to the present. Use either the mean or the modal values of the variance parameters, plugging them in as if known exactly. (Note that the Kalman filter output gives you forecast errors and, in the lh matrix, the previous period's variance of the state (from which you can calculate the standard error of the one-step-ahead forecast).)