FURTHER EXPLANATIONS OF THE HIDDEN CHAIN MCMC EXERCISE

- If you've already got things set up with three lags of three right-hand-side variables and no lagged log CPI, that's OK and you can stick with that, but I had meant to have you include three lags of log CPI also. You also should include a constant term.
- You are meant to follow the procedures described in the course notes. That is, you should set up a Gibbs sampling loop, drawing the regression parameters taking the sequence of states and corresponding residual variances and the transition matrix H as given, then drawing the sequence of states taking everything else as given, then drawing the residual variances conditional on everything else, then drawing H conditional on everything else, repeating this for many draws.
- You will need to find starting values for the parameters to get the MCMC going. A natural starting point is to use the full-sample OLS coefficient estimates for the starting coefficients. For the variances, you could look at the sample distribution of squared residuals and start with the three states having variances at the 25th, 50th and 75th percentiles of that distribution (but other starting points should also work).
- For H you will need a prior, which I failed to mention in the problem handout. I suggest independent Dirichlet priors on the columns of H, i.e. a joint pdf

$$\prod_{i=1}^{3}\prod_{j=1}^{3} p_{ij}^{\alpha_{ij}}$$

  with $p_{3j} = 1 - p_{2j} - p_{1j}$ for all j and all p values between 0 and 1. For the α's you can try $\alpha_{ii} = 9.2$, all i, $\alpha_{ij} = .5$ for $|i-j| = 1$, and $\alpha_{ij} = .3$ for $|i-j| = 2$. This implies the mean stay in a given state lasts about 12 months, that it is a little more likely to pass between neighboring states than to jump between states one and 3, and implies substantial uncertainty about the probabilities.
- Because we are doing no model comparison, and because the coefficients are constant across states, it is probably best to use a flat prior on the coefficients — that is, use the likelihood as the posterior. The likelihood, conditional on the state sequence, is just the likelihood of weighted least squares.
- A proper prior on the state variances is needed, though, to avoid the possibility that likelihood goes to infinity when coefficients are chosen to fit one observation perfectly and the variance for the state that appears at that date

Date: December 10, 2009.

is zero. So use an inverse-gamma prior, e.g. for each σ_i^2 the standard inverse-gamma with 1 degree of freedom:

$$\sigma_i^{-2(p+1)} e^{-\frac{1}{\sigma_i^2}} \, d\sigma_i^2 \, .$$

- Let me emphasize again that all the major steps for this exercise are implemented, for a general regression model, in the programs on the course web site. All you have to do is splice them together to do the MCMC iterations, as in the example `mcexlh.m` and `mcexGibbs.m` programs on the web site.