PRACTICE EXERCISES

(1) In our VAR exercise we found that the four series (output, unemployment rate employment, and employment-population ratio) were best fit by a VAR with two roots very close to one. One possibility is to use this information to simplify the model by putting it into VECM form (while keeping the number of lagged values of y appearing on the right hand side the same), imposing the two unit root assumption, and re-estimating. Another would be to estimate a dynamic factor model of the form

$$y_t = \gamma + \Lambda_0 f_t + \Lambda_1 f_{t-1} + u_t \tag{1}$$
$$f_t = f_{t-1} + \varepsilon_t \tag{2}$$
$$u_t = \delta u_{t-1} + \nu_t \tag{3}$$
$$\varepsilon_t \sim N(0, I)\,, \quad \nu_t \sim N(0, \Sigma)\,, \quad \Sigma \text{ diagonal}\,, \quad \delta \text{ diagonal.}\,, \tag{4}$$
$$\Lambda_i \; 4 \times 2\,. \tag{5}$$

(a) Compare the number of free parameters, after imposing necessary normalizations, in the VECM, the factor model, and the original 5-lag reduced form VAR. Remember to count variance parameters as well as coefficients.
(b) Describe in detail efficient MCMC steps to estimate the VECM and the factor model. Include a discussion of how to handle initial conditions and setting the prior.
(c) Impulse responses for the VECM are easy to calculate using VAR software. (Explain why.) For the factor model, they require substantial computation (I think — I don't see a straightforward analytic approach). Suggest a method for calculating an estimated factor model's impulse responses. [Note: Here we mean by impulse responses, responses of future value forecasts to one-standard deviation innovations (one-step-ahead forecast surprises) in variables. Neither ν, nor ε are innovations in y.]

(2) Construct an example of an ARMA(1,1) process that has a sharp peak at the 12-month seasonal frequency ($\pi/6$).

Date: January 14, 2018.

(3) In Hidden Markov Chain model the transition matrix for the states is usually treated as not varying with the observed data. That is, if the state is S_t, the probability distribution of the observed date y_t conditional on S_{t-1} and lagged observations $\mathcal{I}_{t-1}$ is written as

$$p(y_t \mid \theta(S_t), \mathcal{I}_{t-1})\pi(S_t \mid S_{t-1}) \,. \tag{6}$$

This specification makes calculating the likelihood through a filtering operation straightforward, and makes it possible to make draws from the conditional distribution of the full $\{S_t\}$ sequence given the full data set $\mathcal{I}_t$. In some applications it would be appealing to specify $\pi(S_t \mid S_{t-1}, y_{t-1}, \phi)$, that is to allow the previous period's observed data value to affect the transition probabilities between states.

(a) Is it much harder to evaluate the likelihood with this specification?

(b) Is it much harder to generate draws for $\{S_t\}$ with this specification?

Explain your answers.