1. EXERCISE DUE MONDAY, 9/28

Use monthly data on total nonfarm payroll employment from 1939 to the present. We will use y_t to refer to the natural log of employment at date t. We want to separate a seasonal from a non-seasonal component of this series. Our model is

$$\begin{aligned} y_t &= s_t + b_t \\ s_t &= s_{t-12} + \varepsilon_t \\ b_t &= 1.9 b_{t-1} - .9 b_{t-2} + \nu_t \,. \end{aligned}$$

Define a state vector that lets you cast this model in the form of the standard Kalman filter state evolution and observation equations. Take the standard deviation of ν_t to be .004 and that of ε_t to be .001. As is standard for the Kalman Filter, treat ε_t, ν_t as i.i.d. across time and uncorrelated with lagged states. Here we assume they are also mutually uncorrelated at each t. You will need 12 values of s_t in the state vector. Take their initial covariance matrix to be the matrix with typical element in the t'th row, j'th column $200 \cdot .96^{|i-j|}$. Take the initial covariance matrix of the two value of b you need in the state vector to be the matrix with 200 on the diagonal and 196 in the off-diagonal position. Take the initial values of s and b to be uncorrelated in the prior.

Apply the Kalman filter. Plot the filtered estimates of b_t over the whole sample. For dates after 1947, plot filtered b_t and y_t on the same graph. (You might want to adjust the level of $\hat{b}_t$ so it matches y_t in January 1948, e.g.) Plot filtered s_t the last two years in the sample and for $1950 - 51$. Does it look like the seasonal pattern has changed much? Plot $\hat{s}_t$ and y_t on the same graph for the last four years of the sample, adjusting the level of $\hat{s}_t$ so it matches y_t in January 2006. How much of the variation in the data is due to the seasonal?

The payroll employment data and code for one iteration of the Kalman filter, in R and in Matlab versions, are on the course website. The data are in both CSV format, which can be read in to excel, Matlab (I think) or R (via `read.csv()`) and `.RData` format, which can be read in to R with the `load()` command. The latter file, when read in, will put an R time series object `y` in your workspace that is logged payroll employment and has dates associated with it.

Date: September 21, 2009.