VAR EXERCISE

The course web site has, in the same directory with this file describing the exercise, monthly data on three components of the produce price index: raw materials, intermediate goods, and final goods. The file pdat.RData has the three series already logged and stored as an R multiple time series object. There are also three .xls files containing the raw data for the three series. The files VARpack2019.zip and optimize.1.zip contain R packages that you can install "locally" in R and then load with the R `library()` function. You don't have to use these packages or R, but if not you may have to skip the part of the exercise that compares marginal data densities or else translate the R code into your working language.

Because of the late posting of the exercise, part

(1) Plot the data. (e.g. `plot(pdat)` in R). Note that the materials index is the least smooth and the final goods index the most smooth of the three series.
(2) Estimate a VAR for these three series. Use 3 lags and 9 lags. This can be done easily with the `mgnldnsty()` function in the VARpack2019 package. That function uses a proper prior, so the "w" component of its output list can be used to compute log odds ratios between models, instead of BIC. Use the results to find the log odds ratio between the 3-lag and 9-lag models. [If not using R, you can substitute by computing both likelihood ratio tests and BIC to decide whether the data favor 3 or 9 lags.] Note that `mgnldnsty()` conditions on initial conditions, and uses all the data, so to make the sample period match between the 3 and 9 lag versions, you need to give the 3 lag version a data matrix with 6 fewer initial observations (e.g., `pdat[-(1:6), ]`).
(3) Check the roots of the system, which is most easily done by setting up the coefficient matrix for the stacked system and calculating its eigenvalues. The `sysmat()` function constructs the stacked-system coefficient matrix from the `mgnldnsty$var()` output. How many roots are there within $1/T$ of 1.0? Few enough to be consisistent with cointegration? Does it look like there could be repeated uni roots (i.e. nearly identical unit roots with nearly the same eigenvectors).
(4) Plot impulse responses for the system, with the default Cholesky ordering, for example with

```
resp <- impulsdtrf(mout$var,nstep=48)
plotir(resp))
```

One might expect that changes in raw materials prices pass through with a delay to intermediate goods prices, and that in turn intermediate goods prices pass through with a delay to final good prices. Do the impulse response estimates support this idea?
(5) How would you formulate and estimate a VECM model that was a restricted version of the 3-lag system you estimated above, under the assumption that there is a single cointegrating vector?

Date: March 5, 2021.

(6) How could you set up the plant and observation equations of a state-space model of these data, in which the 3 prices are all linear functions of an unobserved one-dimensional index, all serially correlated, but with the deviation of the ith price from the index uncorrelated with the deviations of the j'th price for $j \neq i$? Can this model be formulated to be a restricted version of the 3-lag VAR?