VAR EXERCISE

Using the same data set that we used for the regression exercise, estimate a four-variable, four-lag VAR. You can use the `mgnldnsty()` function in R or Matlab. You should make draws from the posterior pdf implied by your estimate, and use them to construct impulse responses with error bands. Use all variables in log units, except unemployment rate, which you should just mulitply by .01, so the variables all have roughly the same scale. Then the Minnesota prior's `vprior$sig` parameter can just be `c(.01,.01,.01,.01)`. You can start with defaults for all the other prior parameters, except that you should use `mnprior$tight=1` rather than the default value of 3. Experiment by moving up and down at least two of the other prior parameters as well. (I perturbed only the `mnprior$tight` parameter, which is why I'm telling you to start with one for that.)

The impulse responses should be computed both for the variables ordered as `e`, `p`, `w`, `u`, which is the default ordering if you use the `psdata` matrix as is, but also ordered as `p`, `w`, `u`, `e`. The latter makes more sense if we are thinking of this system as explaining `e`, the employment ratio, and my explorations suggest that the results look different with this ordering.

In R, the default settings for the prior are automatic. In Matlab, you have to set them yourself. Look at the R code for `mgnldnsty` to see what they are. Also, the Matlab `mgnldnsty.m` returns only the marginal data density, not the var that it has estimated, and you need the var results. So you should modify the function so it returns `[w, var]`

The `impulsdtrf` function needs `By` and `smat` arguments. In R, you can handle this by setting the `vout` argument as `vout=list(By=pdout$By, smat=pdout$smat)`. (Of course the `pdout` object itself will have an index, since you will have created an array to hold the MCMC draws.)

To get impulse responses with a different ordering, use the `pchol()` function in R. It computes cholesky decompositions with a different ordering. There is no corresponding function in Matlab, though the comments to `impulsdtrf.m` explain how to do the same thing with a complicated expression.

Think about whether there is a plausible set of restrictions on A_0 or on impulse response signs here that would let us separate labor supply and labor demand effects in the relation between the wage and the employment ratio. I might add an SVAR section of the exercise if I find a set of restrictions that seems to make sense of the data.