POSTERIOR ODDS EXERCISE

Calculate posterior odds on the factor model vs. the VAR in our previous exercise fitting time series on unemployment and real GDP. You will need to formulate a prior on the parameters of the factor model. The log integrated posterior for the VAR is directly available as the `$w` element of the returned value from `mgnldnsty`.

In choosing the prior, you should think of this as a scientific reporting exercise. The prior should not be fully reflecting your personal beliefs, but instead be aimed at being "uninfluential" — trying to rule out only regions of the parameter space that would be uninteresting to most in your audience. For example, you would not want to put much probability on regions of the parameter space that imply explosive behavior or that imply you expect one-step forecast errors for GDP of 20% or for the unemployment rate of 10 percentage points. Note also that the posterior modal parameter values for the factor model with the flat prior imply near-zero variance for the observation error in the unemployment equation. Your prior should not put zero density in the neighborhood of zero for this parameter.

Priors are expected to vary across students in the class, which may give us some insight into how sensitive these results are to priors.

To calculate the integral of the posterior kernel for the factor model, a feasible approach is to use the Kalman filter to evaluate log likelihood at each parameter value and to generate a chain of parameter values (and log likelihood values) using the random walk Metropolis algorithm. Then you can use the modified harmonic mean or bridge sampling to find the integral of the posterior kernel.