VAR EXERCISE ANSWER

The parameter values that maximized the posterior density for the factor model were

ρ_{y1}	.9970
ρ_{y2}	.9912
ρ_z	.6462
α_1	.004970
α_2	-.002950
σ_1	.006687
σ_2	.000000021
δ_1	.03451
δ_2	.0005574

Note that σ_2, the standard deviation of the observation error for the second (unemployment) variable emerges as essentially zero. That means that the factor z_t can be recovered without error from the history of the unemployment variable. Buried in the "Remarks" section of the exercise was a request that you plot on the same graph the smoothed and filtered estimates of the state at the posterior-density-maximizing value of the parameters. Since the value of z can be recovered without error from the data, the filtered and smoothed values of the state are identical, so this plot is not interesting. The only interesting result from the smoother is its value for the initial z value, which we treated as a priori $N(0, 400)$. The smoother estimates its value as .4698, with a standard error of .016.

The log posterior density at the peak emerged as 2180.234.

One way to construct the forecasts from the KF model is to first calculate forecasts of z, as $z_T \rho_z^s$, $s = 1 : 8$, where z_T is the end-of-sample (2016:II) value of the Kalman filtered state estimate. Then for each y_i, one can iterate on $y_{it} = \rho_{yi} y_{i,t-1} + \alpha_i \hat{z}_t$ for eight periods past 2016:II. Plots of the KF forecasts of gdp and unemployment are below. They are rather pessimistic, projecting less than one per cent annual growth on average over the next two years and an unemployment rate flattening out near its current level.

Date: October 16, 2016.

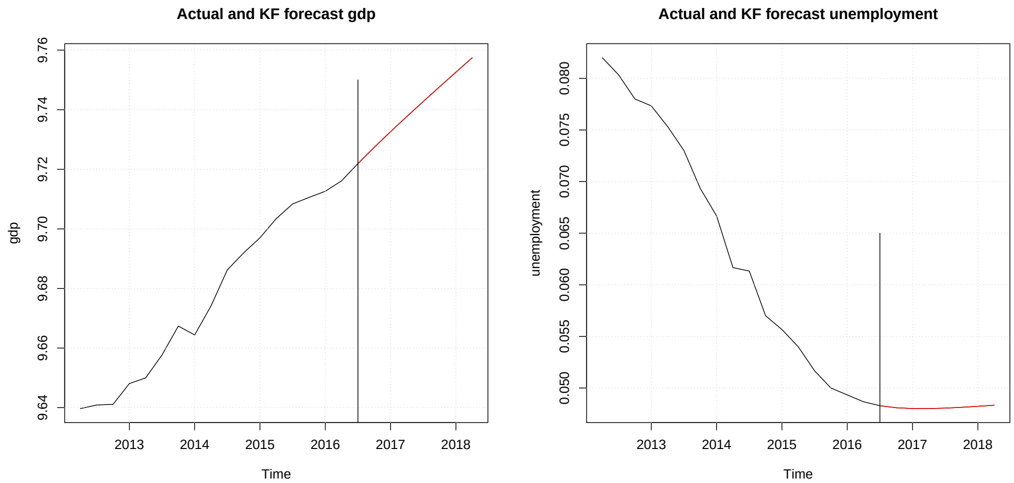

The one-step-ahead forecast error for the y vector consists of the uncorrelated observation errors, plus the contribution of the α vector times the unit-variance z forecast disturbance, plus (in principle) the error from uncertainty about the current z multiplied by ρ_z^2. This last component vanishes for the estimated model, because at the best-fitting parameter values there is no uncertainty about the value of z_t at t, as discussed above. The estimated covariance matrix is

$$\begin{matrix} 0.00006941 & -0.00001466 \\ -0.00001466 & 0.00000870 \end{matrix}$$

This implies a correlation between the innovations of -.6, which makes sense — forecasting too little gdp growth is likely to go with forecasting too high a level of unemployment.

From the plot of standardized forecast errors below, you can see that the largest absolute values are between 3 and 4 standard deviations. The probability, under a normal distribution, of an observation as large as 3.68 is, in a single draw, .000233, i.e. around 2 in 10000. For a sample of size 274, the probability that a draw of this size would occur is .062. So this "test" would not reject normality at the .05 level. However, for both variables there are 5 standardized residuals 2.87 or larger in absolute value. This event, in a sample of size 274, has a probability of .0068 (calculated with R's `pbinom`) function). Of course, because we have estimated the standard deviation from the sample, these tests are at best approximate, but it does seem that the evidence for non-normality due to fat tails is fairly strong. The plots also show clearly the high volatility before 1979 and the "great moderation" period from the mid-eighties to the early 2000's.

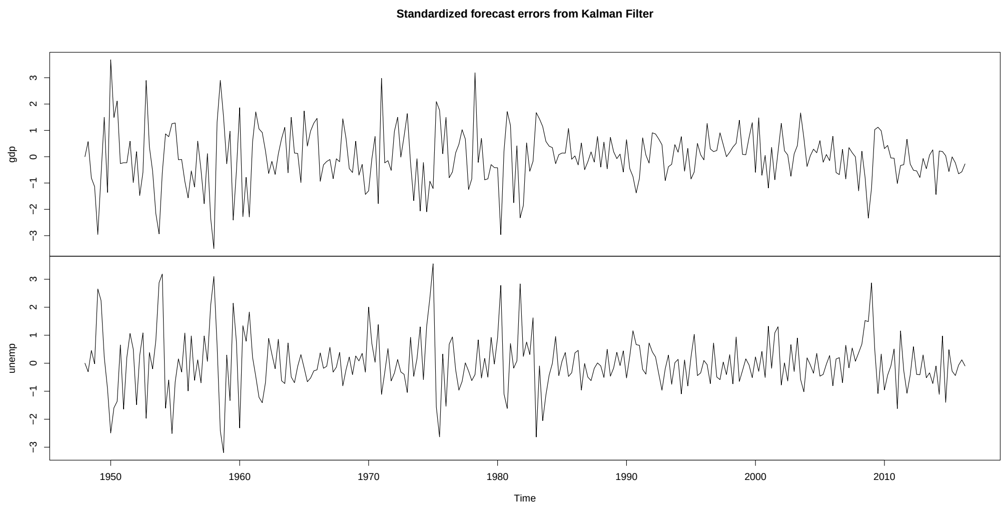

The impulse responses are the effects on predicted future values of one time changes in an observable variable's innovation. For the factor model, it's easiest to "orthogonalize" with unemployment first in the ordering, as unemployment's innovation is estimated to be just α_2 times the innovation in z, due to the near-zero observation error variance in that equation. Then, for gdp, the innovation minus its projection on the unemployment innovation, is just the gdp equation's observation error. The response to a one-standard-deviation gdp shock is then just the standard deviation of ε_1 times ρ_1^s, $s = 0, \ldots, H$, where H is the forecast horizon. These calculations are a bit simplified by using the R `filter` function. The impulse responses for the factor model are below

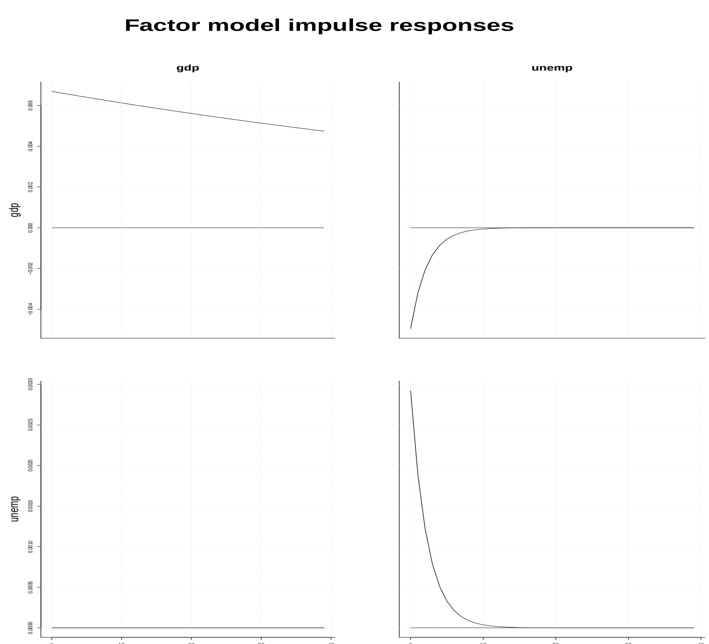

For the VAR model, the log marginal data density emerges as 2141.345. Note that this is the integrated posterior density and thus not directly comparable to what we calculated for the factor model, which was the log of the maximized posterior density.

The VAR forecasts from the end of the sample are shown below.

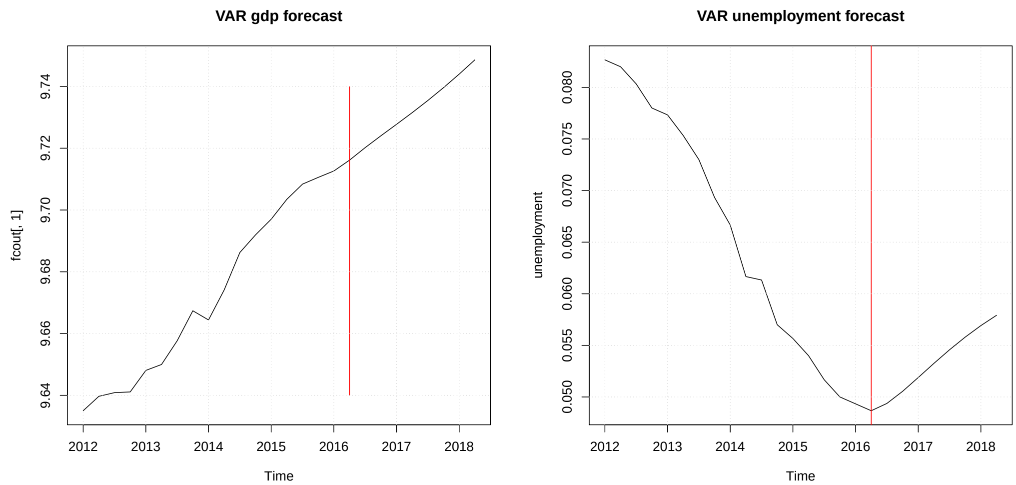

The VAR forecasts are even more pessimistic than the KF forecasts. They predict somewhat slower GDP growth and nearly a percentage point rise in unemployment over the next year.

The covariance matrix of innovations for the VAR model is

	y	u
y	0.00007374	-0.00001396
u	-0.00001396	0.00000809

which implies a correlation of -.57 between the two innovations, very close to the factor model result. The factor model implies slightly larger forecast error variances. This would be unsurprising if the factor z had not emerged as observable, since with z unobservable the Kalman filter would be improving its estimate of the current state through the first part of the sample. But with z observed without error in the unemployment equation, both models are at similar risk of "overfitting", by estimating unrealistically small error variances. The facror model, though, has only 9 parameters free, while the VAR model has 13, and we expect a maximized posterior or likelihood to produce more overfitting for larger models.

The standardized VAR residuals, shown below, include more 4-standard-deviation values than did the Kalman filter residuals, and thus even stronger evidence of non-normality. They also show the same pattern of volatility changes as did the factor model residuals.

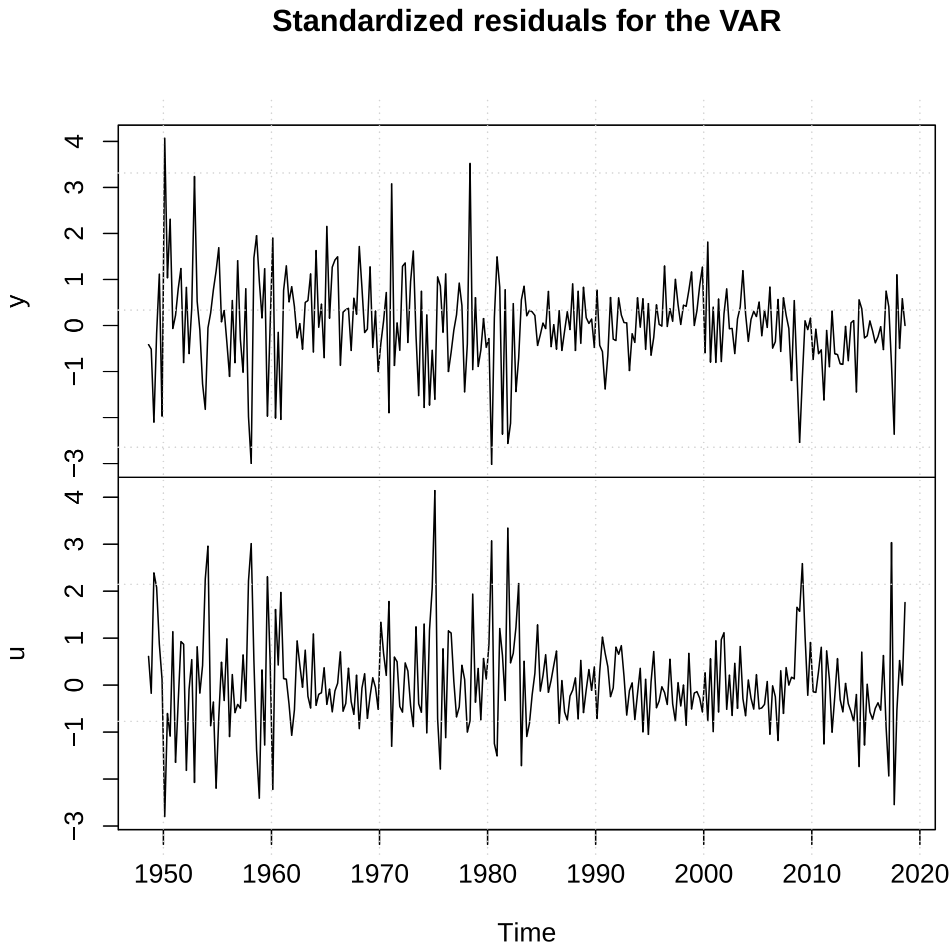

Though you weren't asked to calculate it, it is interesting that the correlations of the residuals for the two models are .93, for both gdp and unemployment.

The VAR impulse responses, with unemployment ordered first to match the factor model plot, show a negative response of unemployment to a gdp shock despite the initial response being constrained to zero. ()This was done with `order=c(2,1)` in the call to `impulsdtrf`.)

However if we reverse the ordering, the initial negative response of output to an unemployment innovation is very small, with the overall response being dominated by a long-term positive response of output to the unemployment innovation. This suggests an interpretation with the gdp shock a "labor demand" shock, and the unemployment shock a "labor supply" shock — people enter the labor force, raising unemployment as they look for work, then expand output as they find jobs. Another possibility would be that unexpected increases in unemployment tend to lead to monetary policy easing, and thus increased output. It would take a larger model, in which we could identify a monetary policy equation, to sort this out.

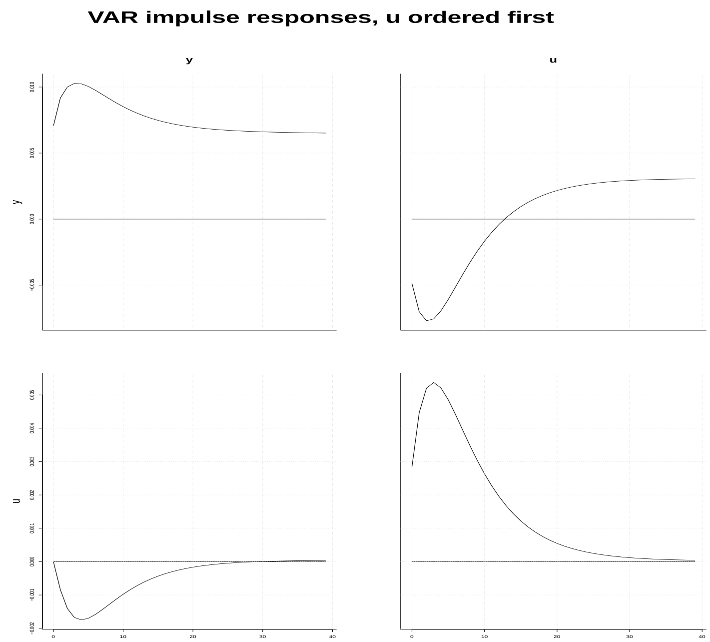
VAR impulse responses, u ordered first
y
u
y
u
0
10
20
30
40

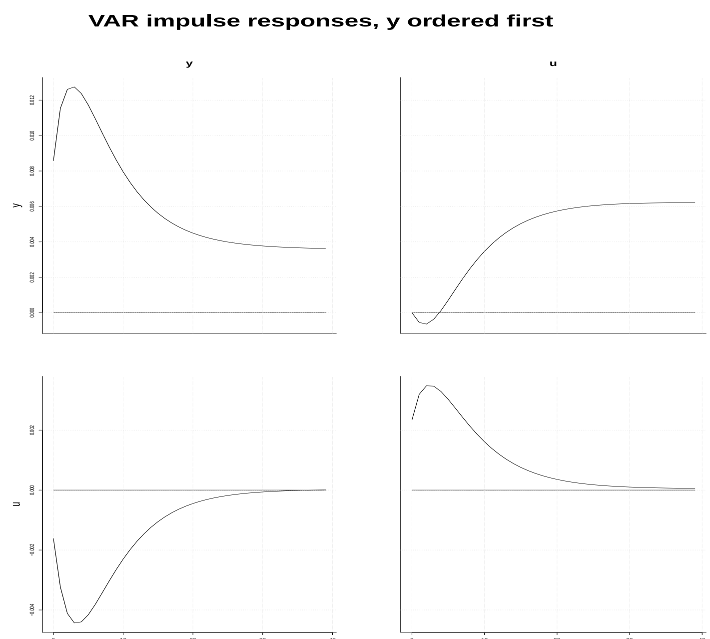
VAR impulse responses, y ordered first
y
u
y
u
0
10
20
30
40