VAR EXERCISE

The course web site has data on the federal funds rate (ffr), the M1 money stock (m1), the price level (cpi), and industrial production (ip), both as an **R** time series object in `rmpy.R` and as a formatted text file `rmpy.txt`. The data are monthly, 1960:1 to 2005:3. Use these data to estimate a reduced form VAR. Use the default values for the sum of coefficients and cointegration dummy observations in `rfvar3.R` or `rfvar3.m` to do the estimation. If you use other software, you will need to figure out for yourself how to implement the corresponding dummy observations. Log the values of the non-interest-rate variables before you estimate, and it can ease interpretation also to divide the ffr series by 100, so all the residuals are expected to be the same order of magnitude.

(a) Find the eigenvalues of the estimated system. Do they separate cleanly into near-unit roots and others? Do the eigenvectors show evidence of near-repeated-root behavior? [Hint: My own calculations using the default dummy observation prior, with not quite the same dataset, showed four or more roots very close to one, plus clear evidence of repeated-root-like behavior.]

The eigenvalues I found, using 1960:1-2005:3 data, with r divided by 100 and m, p and y logged, were

Date: December 11, 2005.

Root number	Modulus	Period in months
1	0.9999	Inf
2	0.9995	Inf
3	0.9968	Inf
4	0.9826	Inf
5	0.9043	Inf
6	0.8501	Inf
7	0.7068	5.10
8	0.7068	−5.10
9	0.6926	Inf
10	0.6858	8.10
11	0.6858	−8.10
12	0.6314	2.47
13	0.6314	−2.47
14	0.6129	4.38
15	0.6129	−4.38
16	0.5987	2.62
17	0.5987	−2.62
18	0.5929	3.44
19	0.5929	−3.44
20	0.5404	2.00
21	0.5101	3.16
22	0.5101	−3.16
23	0.4385	Inf
24	0.3305	2.00

The **R** commands to arrive at this were:

```
> load("rmpy.Rdata")
> str(rmpy)
 mts [1:1018, 1:4] NA NA NA NA NA NA NA NA NA NA ...
 - attr(*, "dimnames")=List of 2
  ..$ : NULL
  ..$ : chr [1:4] "ffr" "m1" "cpi" "ip"
 - attr(*, "tsp")= num [1:3] 1921 2006   12
 - attr(*, "class")= chr [1:2] "mts" "ts"
> window(rmpy,start=c(1960,1),end=c(1963,1))
          ffr      m1  cpi     ip
Jan 1960 3.99 139.979 29.4 27.026
Feb 1960 3.97 139.867 29.4 26.786
......................
Dec 1962 2.93 147.817 30.4 28.770
Jan 1963 2.92 148.255 30.4 28.981
> ydata <- rmpy
> ydata[,1] <- ydata[,1]/100
> ydata[,2:4] <- log(ydata[,2:4])
> window(ydata,start=c(1960,1),end=c(1963,1))
            ffr       m1       cpi        ip
Jan 1960 0.0399 4.941492 3.380995 3.296799
```

```
Feb 1960 0.0397 4.940692 3.380995 3.287879
.................................
Dec 1962 0.0293 4.995975 3.414443 3.359333
Jan 1963 0.0292 4.998934 3.414443 3.366640
> args("rfvar3")
function (ydata = NA, lags = 6, xdata = NA, breaks = NULL, lambda = 5,
    mu = 2, ic = NULL)
NULL
> ydata <- window(ydata,start=c(1960,1),end=c(2005,3))
> T <- dim(ydata)[1]
> T
[1] 543
> rfout <- rfvar3(ydata=ydata,xdata=matrix(1,T,1))
> Bmat <- rbind(matrix(rfout$By,4,24),diag(1,nrow=20,ncol=24))

> Bev <- eigen(Bmat)
> Bev$values
 [1]  0.9998541+0.0000000i  0.9995128+0.0000000i  0.9967846+0.0000000i
.....................................
22] -0.2070738-0.4661572i  0.4384604+0.0000000i -0.3304822+0.0000000i
> svd4 <- svd(Bev$vectors[,1:4])
> svd4$d
[1] 1.4833842 1.0413492 0.8368099 0.1221155
```

There are four roots very close to one and fairly well separated from the next one, which is .9043. This suggests there is no near-cointegration, unless repeated roots are present. If we consider the first four columns of the eigenvector matrix, corresponding to these four near-unit roots, we can check whether they are nearly singular as a way of checking for repeated-root-like behavior. This was done above with the line that calculates `svd4 <- svd(Bev$vectors[,1:4])`. One of the singular values is one-tenth of the others, but this is not an extreme enough ratio that we would usually think of this as near-singularity. It's enough to tell us, though, that we probably won't get simple exponential decay in the impulse responses.

The table of absolute values of roots and corresponding periods for the imaginary ones was produced with

```
> BvalMat <- matrix(c(Mod(Bev$values),2*pi/Arg(Bev$values)),ncol=2)
> print(xtable(BvalMat,digits=c(0,4,2)))
............
```

This uses the **R** `xtable` package, which converts dataframes, matrices, or time series into LaTeX (or HTML) tables.

(b) From your estimates, form the sum-of-coefficients matrix that has reduced rank and determines cointegrating vectors in a VECM model. Does the eigenvector-eigenvalue decomposition of this matrix look close to the form expected when there is cointegration?

The matrix `Bsum` is calculated and printed out in the code below:

```
> Bsum <- apply(rfout$By,FUN=sum,MAR=c(1,2))
> Bsum
```

```
                ffr           m1            cpi           ip
ffr   0.976717353  2.209669e-06 -1.093704e-04 2.622418e-05
m1    0.006057106  9.998349e-01 -2.873175e-05 1.256646e-04
cpi   0.009091636 -2.388424e-05  9.998419e-01 2.783903e-04
ip   -0.023948006 -5.680378e-05 -1.337128e-04 9.998220e-01
> svd(Bsum)
$d
[1] 1.0057968 0.9998868 0.9998503 0.9708583
............................
```

Obviously it is very close to an identity matrix, so if we subtract it from the identity we have nearly a zero matrix. One might therefore conclude that there is no evidence here of strongly mean-reverting linear combinations of the data. On the other hand, the half-lives of the four largest roots (the time it takes the root to decay to .5, calculated as `log(.5)/Bev$values[1:4]`, are, in months, 750, 1422, 215, and 39. The first three take well over a decade to decay, but the last decays by one half in about three years. If we calculate the eigenvalues of `diag(4)-Bsum`, we get one of .023 and three that are over 100 times smaller. The linear combination of y's that corresponds to this largest root my therefore be worth looking at as a candidate for a mean-reverting relationship among the variables. It is found by

```
BsumEv <- eigen(Bsum)
BsEvLeft <- solve(BsumEv$vectors)
Re(BsEvLeft[1,]/BsEvLeft[1,2])
 0.5288343  1.0000000 -0.5990081 -0.4859274
```

This looks a bit like a long run money demand or liquidity preference relation. With m1 on the left, it would imply a negative semi-elasticity of .5 with respect to interest rates and a positive elasticity of about .5 with respect to nominal production. It is actually not plausible, in a sample with substantial low frequency movement in inflation, that money demand is less than unit-elastic in the long run with respect to the price level, but the signs and orders of magnitude are are right. One might want to try a VECM here with the coefficients on cpi and m1 in a single cointegrating vector constrained to be 1 and -1 and the coefficients on ffr and ip left free. The constrained model could be estimated as

$$\Delta y_t = C(L)\Delta y_{t-1} + \alpha \begin{bmatrix} \theta_1 & 1 & -1 & \theta_2 \end{bmatrix} y_{t-1} \,.$$

Here α is a 4×1 column vector of coefficients. The model is nonlinear in parameters, so it would have to be estimated by iterative methods based on the full likelihood, but in a model of this size that is quite feasible. To check whether these restrictions are actually easily accepted by the data, we would need to carry out a posterior odds ratio check, which we've not yet discussed.

(c) Calculate forecasts of the four variables from the beginning-of-sample initial conditions. (Here `fcast.R` or `fcast.m` may be useful.) Plot the actual and forecast values for each series, and on each plot show also the unconditional mean (if it exists) of the variable. (The mean will only exist, of course, if your estimates have all roots at least slightly less than one in absolute values.).

The plots are at the end. The actual series is a black line, the projected path is a green line, the unconditional mean is a red horizontal line, and two-standard-error-bands around the mean are blue lines. (The exercise did not ask for the error bands, but it should have.) To compute the unconditional mean and the error bands, I used the code below.

```
lyp <- function(A,sig) {
delta <- 1
ly <- rbind(cbind(sig,matrix(0,4,20)),matrix(0,20,24))
while(delta > 1e-10) {
   ly1 <- A %*% ly %*% t(A)
   delta <- sum(abs(ly1))
   A <- A %*% A
   ly <- ly +  ly1
}
return(ly)
}
sig <- crossprod(rfout$u)/T
vcv <- lyp(Bmat,sig)
```

This uses the "doubling algorithm" to calculate the unconditional covariance matrix of the stacked y vector, including all the lags. The unconditional covariance matrix of the 4-dimensional current y vector is the upper left 4×4 submatrix (or any of the other 5 4×4 diagonal blocks below it, since y and its lags all have the same variance-covariance matrix).

I also show the same sort of plot for the first differences of m1, cpi, and ip. Since the estimates make these variables stationary, the unconditional means of their first differences are all zero. If Ω is the unconditional covariance matrix of y, the unconditional covariance matrix of Δy is $\Omega - A\Omega - \Omega A + A\Omega A'$, where A is from $y_t = Ay_{t-1} + \varepsilon_t$ in the stacked system and corresponds to `Bmat` in the code. In the code, Ω is `vcv`.

The ffr series shows some questionable behavior. It suggests that the steady rise in the funds rate from around 2% in the early 60's to around 7% in the early 70's could have been known in advance. It also implies that the actual interest rate was more than three standard deviations above its steady state mean of -9% for the entire sample period. A projection from the end of the sample would probably imply negative interest rates within a few decades.One would not want to use this model, therefore to make multiple-decade projections.

The three levels series all show means *very* far above their observed values in standard deviation units. In each case, this is functioning to provide a near-deterministic linear trend in the projection. Since this behavior would clearly continue in almost the same form in projections out of the sample for many decades, it is not as objectionable as what appears in the interest rate plots, where out of sample behavior would be predicted to be unlike what has been observed historically. It might be tempting to allow a linear trend to enter the system directly. There are two possible problems with doing that: a) unreasonable as it may be to model trending behavior as eventually stopping as some permanent lever is approached, modeling the in-sample trend as persisting forever is arguably equally arbitrary; b) a VAR, which here has used the constant term to generate a near-linear trend, can use a linear trend to generate higher-order polynomials, so the unreasonable behavior of long run projections might well get worse, not better, with a deterministic linear trend in the regressor list.

That the trend-like behavior is not too unreasonable here can be verified from the plots for first-differenced data. They show little implied long-run forecastability and both actual and projected values staying within 2σ bands.

(d) Compare all the results you have obtained above with what emerges when you estimate the VAR with no prior — i.e. with $\lambda = \mu = 0$ in the arguments to `rfvar3`.

Here are the roots.

root #	modulus	period
1	0.9992	Inf
2	0.9849	420.25
3	0.9849	−420.25
4	0.9738	Inf
5	0.8709	270.41
6	0.8709	−270.41
7	0.6974	8.00
8	0.6974	−8.00
9	0.6927	5.09
10	0.6927	−5.09
11	0.6244	Inf
12	0.6164	2.46
13	0.6164	−2.46
14	0.6060	2.62
15	0.6060	−2.62
16	0.5968	4.40
17	0.5968	−4.40
18	0.5956	3.44
19	0.5956	−3.44
20	0.5338	2.00
21	0.5229	3.24
22	0.5229	−3.24
23	0.4924	Inf
24	0.3727	2.00

Now the first four roots have half-lives of 904, 45, 45, and 26 months. Only the first one is extremely persistent. The singular values of the first four eigenvectors are 1.5826, 1.0843, 0.5514, and 0.1248 — not very different from the first estimates. The eigenvalue/vector decomposition of the `Bsum` matrix now has a pair of complex roots as the largest, and the next largest is .3 times the absolute value of the first two. The case for singling out a small number of cointegrating vectors seems weak, and giving an economic interpretation to the complex pair seems difficult.

for this case only the plots for ffr and the growth in cpi are shown. For the levels of the non-ffr variables, the pictures are similar to the case with a prior, but the means are not so far away and the implied trend-like behavior therefore not quite so perfectly linear. For the interest rate, the no-prior estimates show ffr inside the 2σ band for most of the sample, and the mean is only slightly negative, instead of -9%. The rise in inflation from early 60's to early 70's is still fully anticipated, and now some of the post-1980 decline in inflation is also anticipated. For the growth of the cpi, the plot now looks much like the ffr plot, with much of the rise in inflation from early 60's to early 70's anticipated and some of the subsequent decline anticipated. The projection goes outside the 2σ band for several years. The estimates with the prior showed considerably less of this behavior.

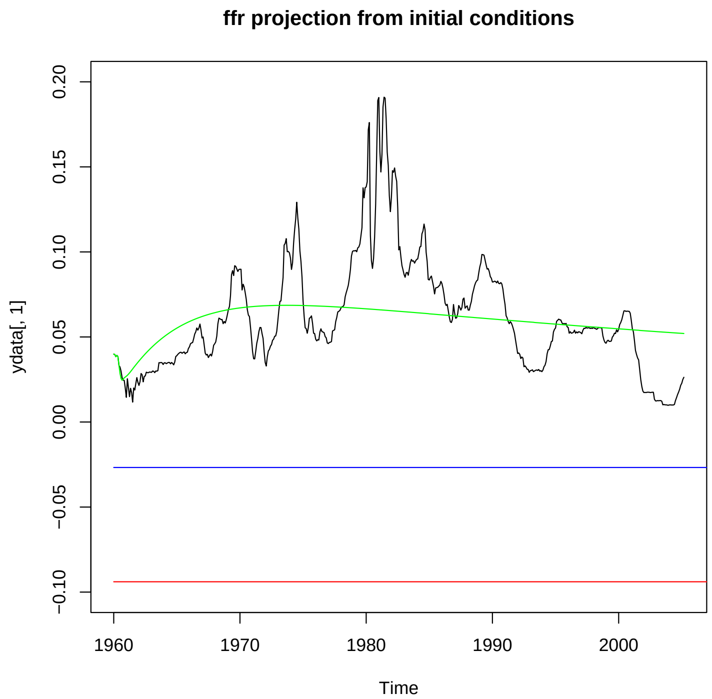

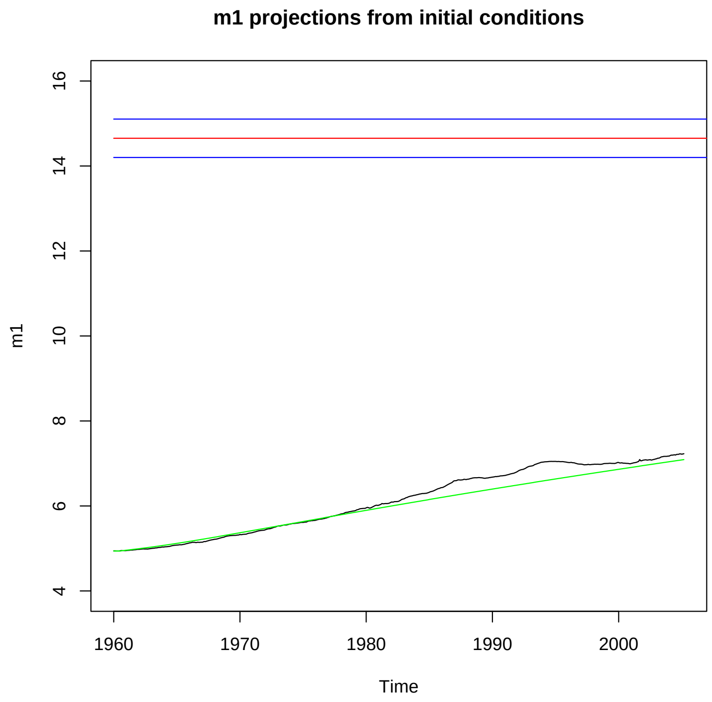

m1 projections from initial conditions
m1
16
14
12
10
8
6
4
1960
1970
1980
1990
2000
Time

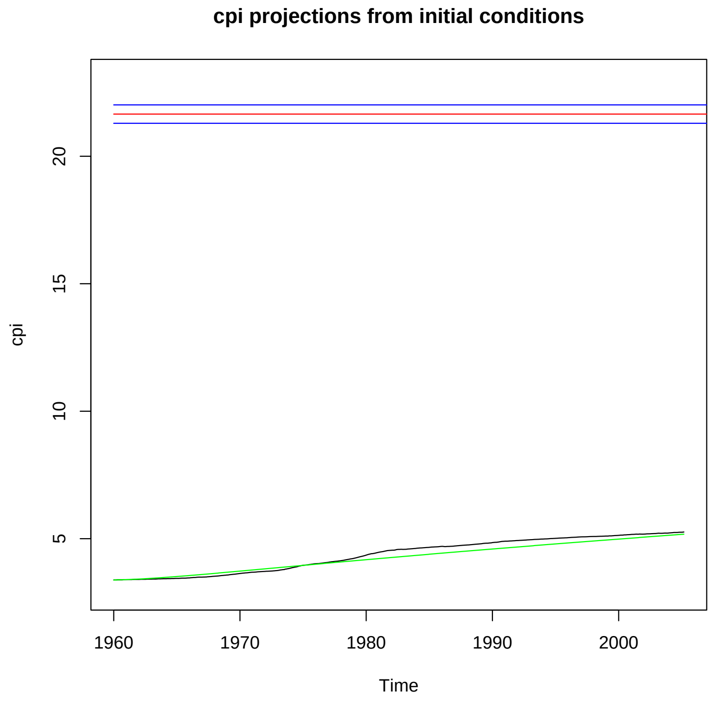

cpi projections from initial conditions
cpi
20
15
10
5
1960
1970
1980
1990
2000
Time

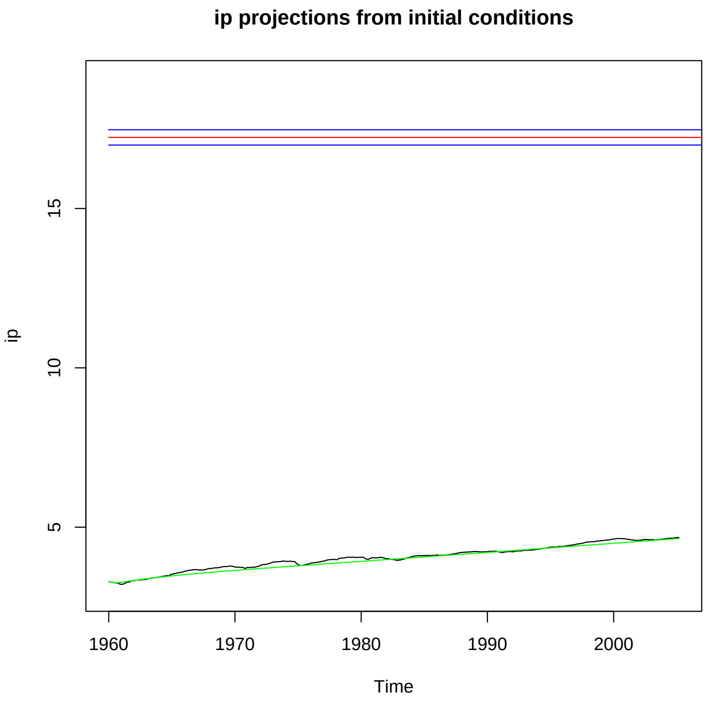
ip projections from initial conditions
ip
15
10
5
1960
1970
1980
1990
2000
Time

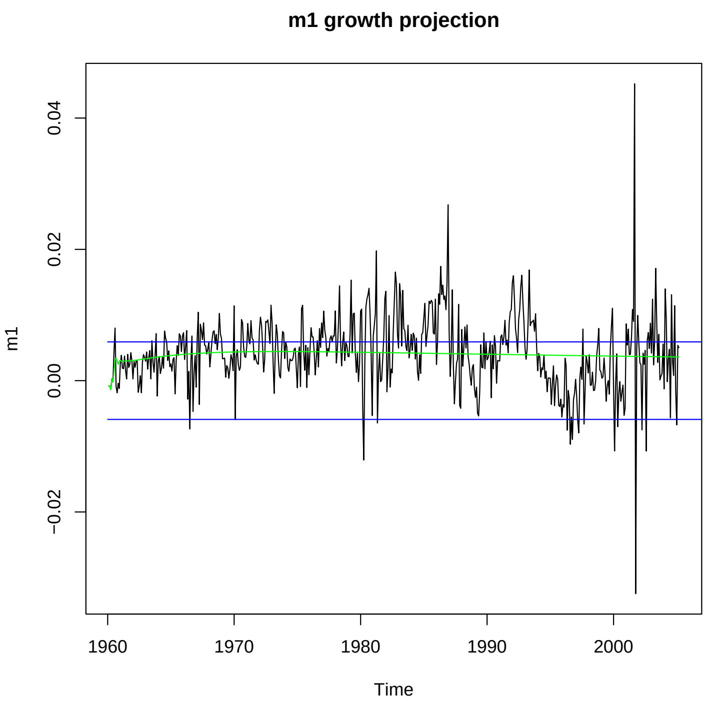
m1 growth projection
m1
0.04
0.02
0.00
−0.02
1960
1970
1980
1990
2000
Time

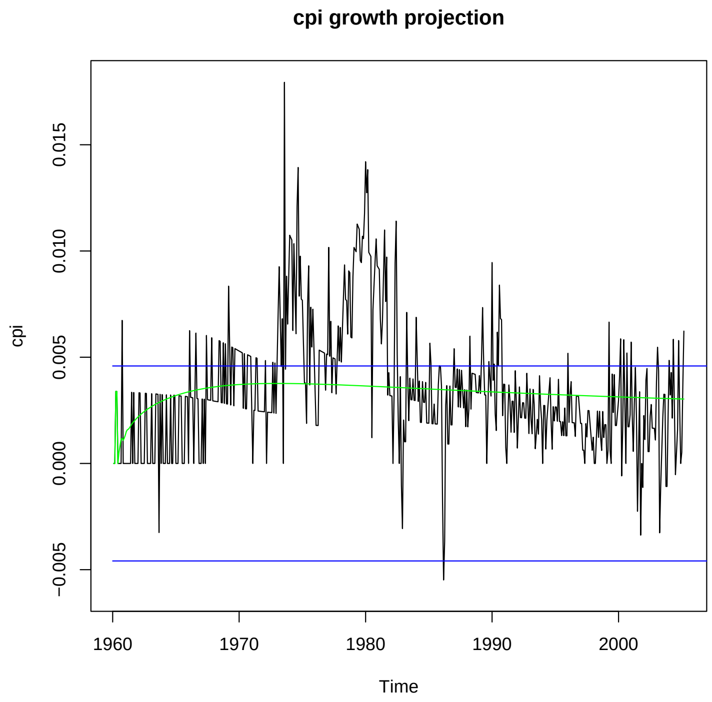
cpi growth projection
cpi
0.015
0.010
0.005
0.000
−0.005
1960
1970
1980
1990
2000
Time

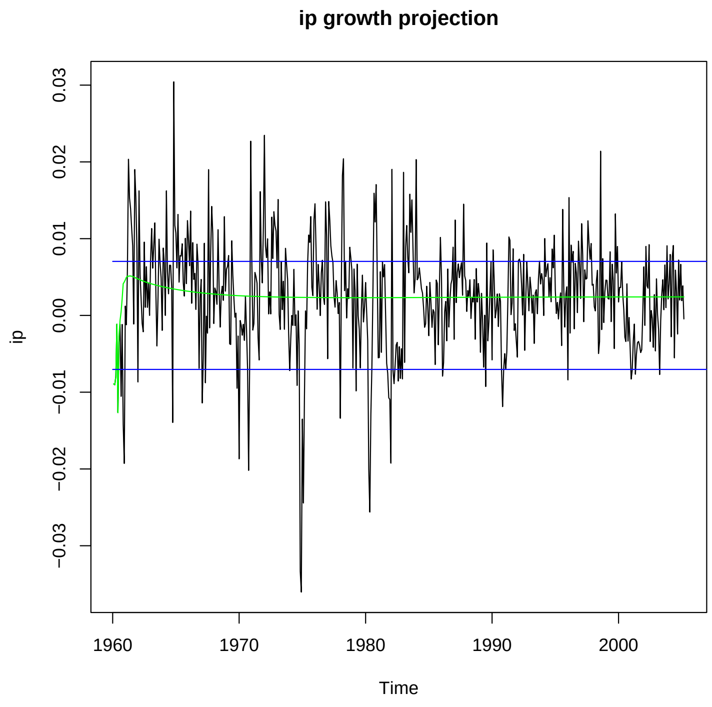
ip growth projection
ip
0.03
0.02
0.01
0.00
−0.01
−0.02
−0.03
1960
1970
1980
1990
2000
Time

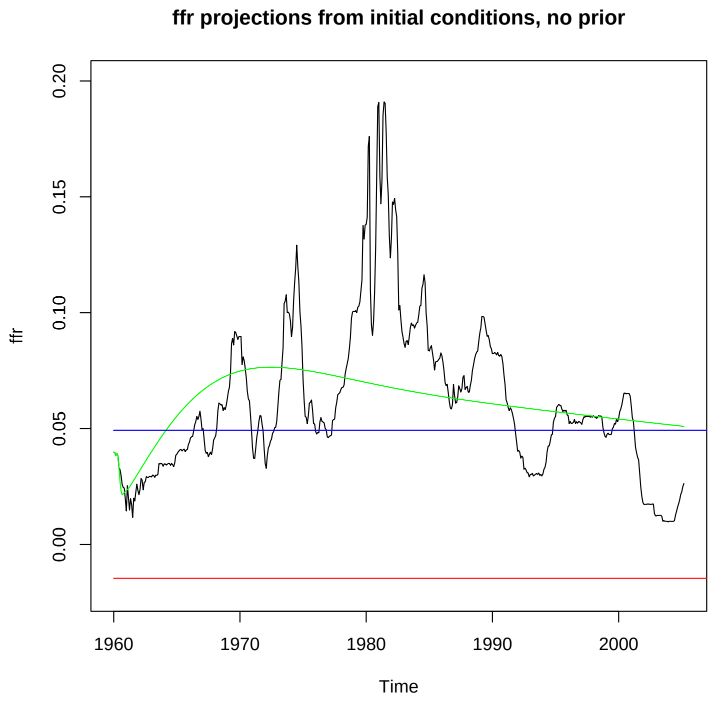
ffr projections from initial conditions, no prior
ffr
0.00
0.05
0.10
0.15
0.20
1960
1970
1980
1990
2000
Time

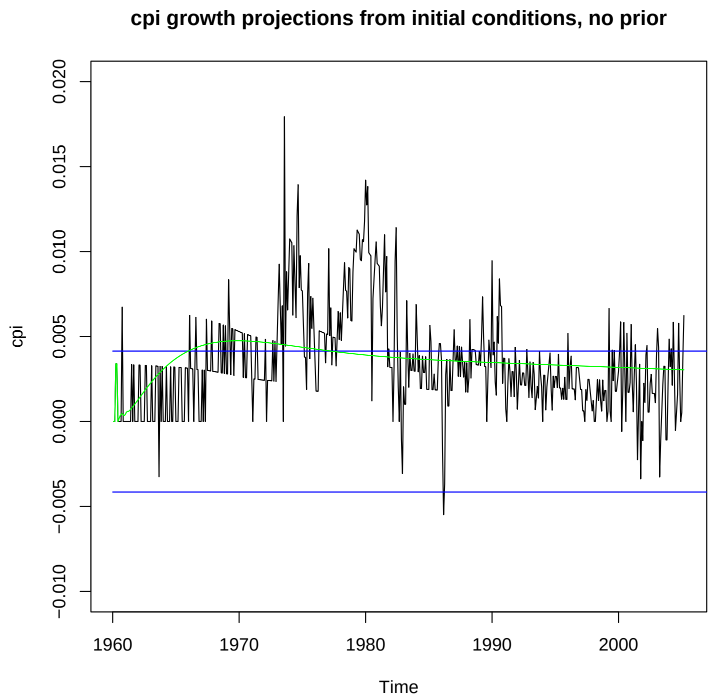
cpi growth projections from initial conditions, no prior
cpi
−0.010
−0.005
0.000
0.005
0.010
0.015
0.020
1960
1970
1980
1990
2000
Time