Predicting the real economy wih financial asset liquidity

Bernt Arne Ødegaard

24 November 2021

1 Predicting the real economy with financial asset liquidity

1.1 Introduction

We look at using financial asset market liquidity as a predictive variable, predicting macroeconomic aggregates. The paper that originally showed that liquidity is a viable variable for predicting the state of the economy was Næs, Skjeltorp, and Ødegaard (2011). We will show some of this style of analysis. We also show some analysis using data for Norway, which is also discussed in Skjeltorp and Ødegaard (2009).

1.2 Descriptives showing that there may be a link between liquidity and the real economy

We use data for Norway, and want to investigate a link between liquidity and real economy, measured by GDP. We will show a plot with liquidity and a measure of the state of the real economy, the output gap, in the same picture. The output gap is a measure of how far the economomy is from full production, usually estimated by the difference between the current GDP and a rolling mean. (Technically, this is calculated using what is called a HP filter.)

Figure 1 shows how liquidity and the economy seems to be moving in opposite directions.

Showing how this picture is generated in R

Reading data (to illustrate)

```
library(zoo)
Liq  <- read.zoo("../data/liq_measures_ose_quarterly.txt",header=TRUE,sep=";",format="%Y%m%
Liq <- as.zooreg(Liq)
RelSpread <- zooreg(as.matrix(Liq$RelSpread),frequency=4,start=c(1980,1))
GDP <- read.table("gdp_mainland_norway.csv",sep=";",skip=5,header=FALSE)
gdp <- zooreg(GDP[,2],frequency=4,start=c(1978,1))
names(gdp)[1] <- "GDP"
```
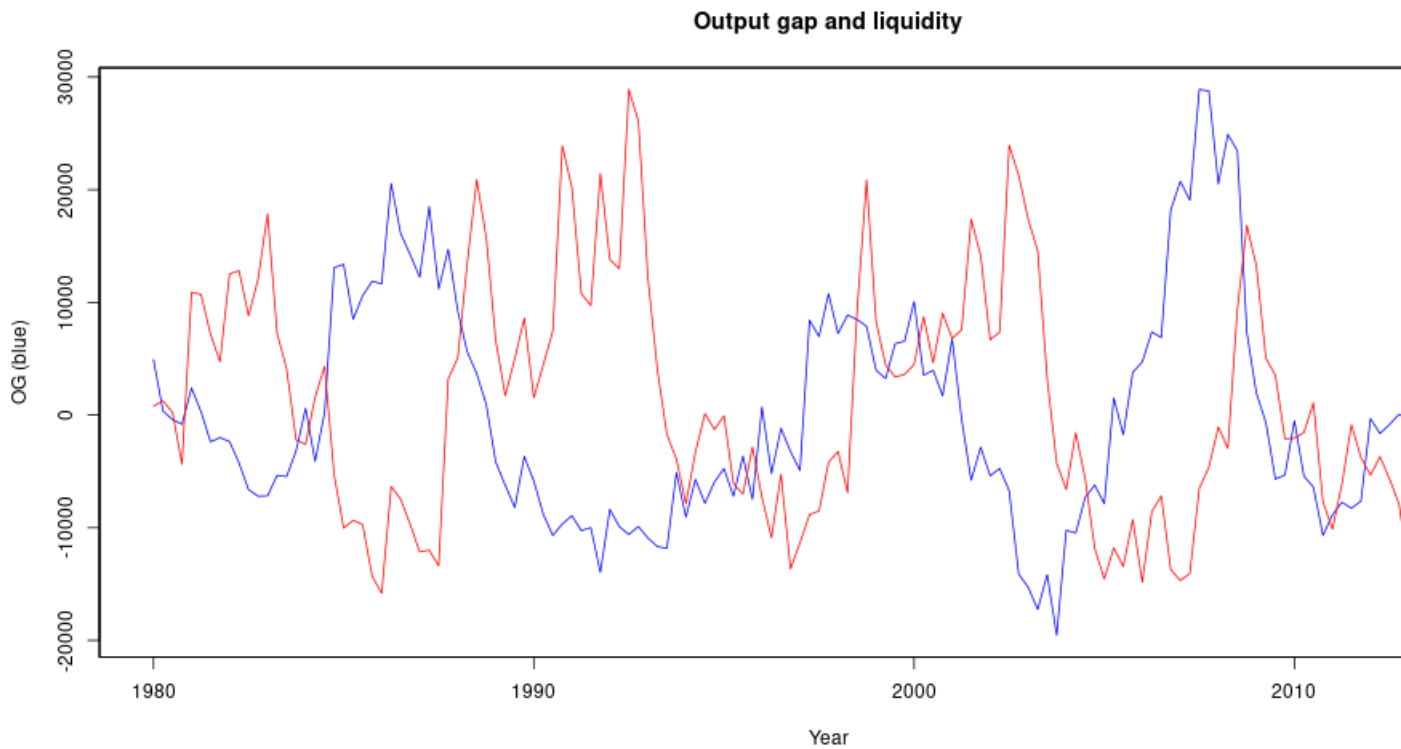
To calculate the output gap, this is done using a HP filter, available in the R library `mFilter`.

```
library(mFilter)
gdp.hp <- hpfilter(gdp,freq=40000)
og <- zooreg(gdp.hp$cycle,frequency=4,start=c(1978,1))
data <- merge(og,RelSpread,all=FALSE)
plot(data$og,main="Output gap and liquidity",xlab="Year",ylab="OG",col="blue")
par(new=T)
plot(RelSpread,col="red",axes=F,ylab="",xlab="")
axis(4)
mtext(side=4,"Relative Spread")
dev.off()
```

Figure 1 Output Gap and Liquidity, Norway

1.3 Predicting GDP Norway with spread – VAR analysis

We look at an example from Norway. Predict GDP for mainlaind Norway, using the relative spread as the explanatory variable.

The following shows the R code for running a VAR - Vector Autoregression analysis.

This is an in-sample analysis.

Reading in the data. Some details hidden in separate input routines (typical way of doing it, a small file with R code massaging the input into an R object.

```
> library(vars)
> outdir <- "../../results/2016_02_var_analysis/"
>
> source ("/home/bernt/data/2016/ssb/read_quarterly_gdp.R")
> head(qGDP)
1978(1) 1978(2) 1978(3) 1978(4) 1979(1) 1979(2)
 247094  249853  250733  255300  257066  257132
>
> source("../../data/read_quarterly_norwegian_liquidity_data.R")
> head(RelSpread)
  1980(1)   1980(2)   1980(3)   1980(4)   1981(1)   1981(2)
0.0391982 0.0396895 0.0386271 0.0338552 0.0496760 0.0494737
>
> dGDP <- diff(log(qGDP))
> names(dGDP) <- "dGDP"
```

2

The following is the data we work with

```
> data <- merge(dGDP,RelSpread,all=FALSE)
> head(data)
                 dGDP RelSpread
1980(1)   0.013019691 0.0391982
1980(2)  -0.011668339 0.0396895
1980(3)   0.002693454 0.0386271
1980(4)   0.004024743 0.0338552
1981(1)   0.017495227 0.0496760
1981(2)  -0.002339087 0.0494737
```

Running the VAR is now just

```
> reg <- VAR(data)
> summary(reg)

VAR Estimation Results:
=========================
Endogenous variables: dGDP, RelSpread
Deterministic variables: const
Sample size: 142
Log Likelihood: 975.925
Roots of the characteristic polynomial:
0.8623 0.2435
Call:
VAR(y = data)


Estimation results for equation dGDP:
====================================
dGDP = dGDP.l1 + RelSpread.l1 + const

             Estimate Std. Error t value Pr(>|t|)
dGDP.l1      -0.224744   0.082735  -2.716  0.00744 **
RelSpread.l1 -0.240555   0.079313  -3.033  0.00289 **
const         0.016686   0.003324   5.019 1.56e-06 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1


Residual standard error: 0.01045 on 139 degrees of freedom
Multiple R-Squared: 0.08967,Adjusted R-squared: 0.07658
F-statistic: 6.846 on 2 and 139 DF,  p-value: 0.001459
```

See here the relevant coefficient, that the coefficient on the relative spread is -0.24, negative and significant at the 1% level. The remaining output is not relevant for asking the question we want, but note that the real variable does not predict the spread.

```
Estimation results for equation RelSpread:
=========================================
```

```
RelSpread = dGDP.l1 + RelSpread.l1 + const

             Estimate Std. Error t value Pr(>|t|)
dGDP.l1      -0.084815   0.047154  -1.799 0.074239 .
RelSpread.l1  0.843497   0.045204  18.660  < 2e-16 ***
const         0.006424   0.001895   3.390 0.000909 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1


Residual standard error: 0.005958 on 139 degrees of freedom
Multiple R-Squared: 0.7337,Adjusted R-squared: 0.7299
F-statistic: 191.5 on 2 and 139 DF,  p-value: < 2.2e-16


Covariance matrix of residuals:
              dGDP  RelSpread
dGDP      1.093e-04 -6.487e-06
RelSpread -6.487e-06  3.550e-05

Correlation matrix of residuals:
           dGDP RelSpread
dGDP      1.0000   -0.1041
RelSpread -0.1041    1.0000
```

Run the test for causality

```
> causality(reg,cause="RelSpread")
$Granger

Granger causality H0: RelSpread do not Granger-cause dGDP

data:  VAR object reg
F-Test = 9.199, df1 = 1, df2 = 278, p-value = 0.00265


$Instant

H0: No instantaneous causality between: RelSpread and dGDP

data:  VAR object reg
Chi-squared = 1.5236, df = 1, p-value = 0.2171
```
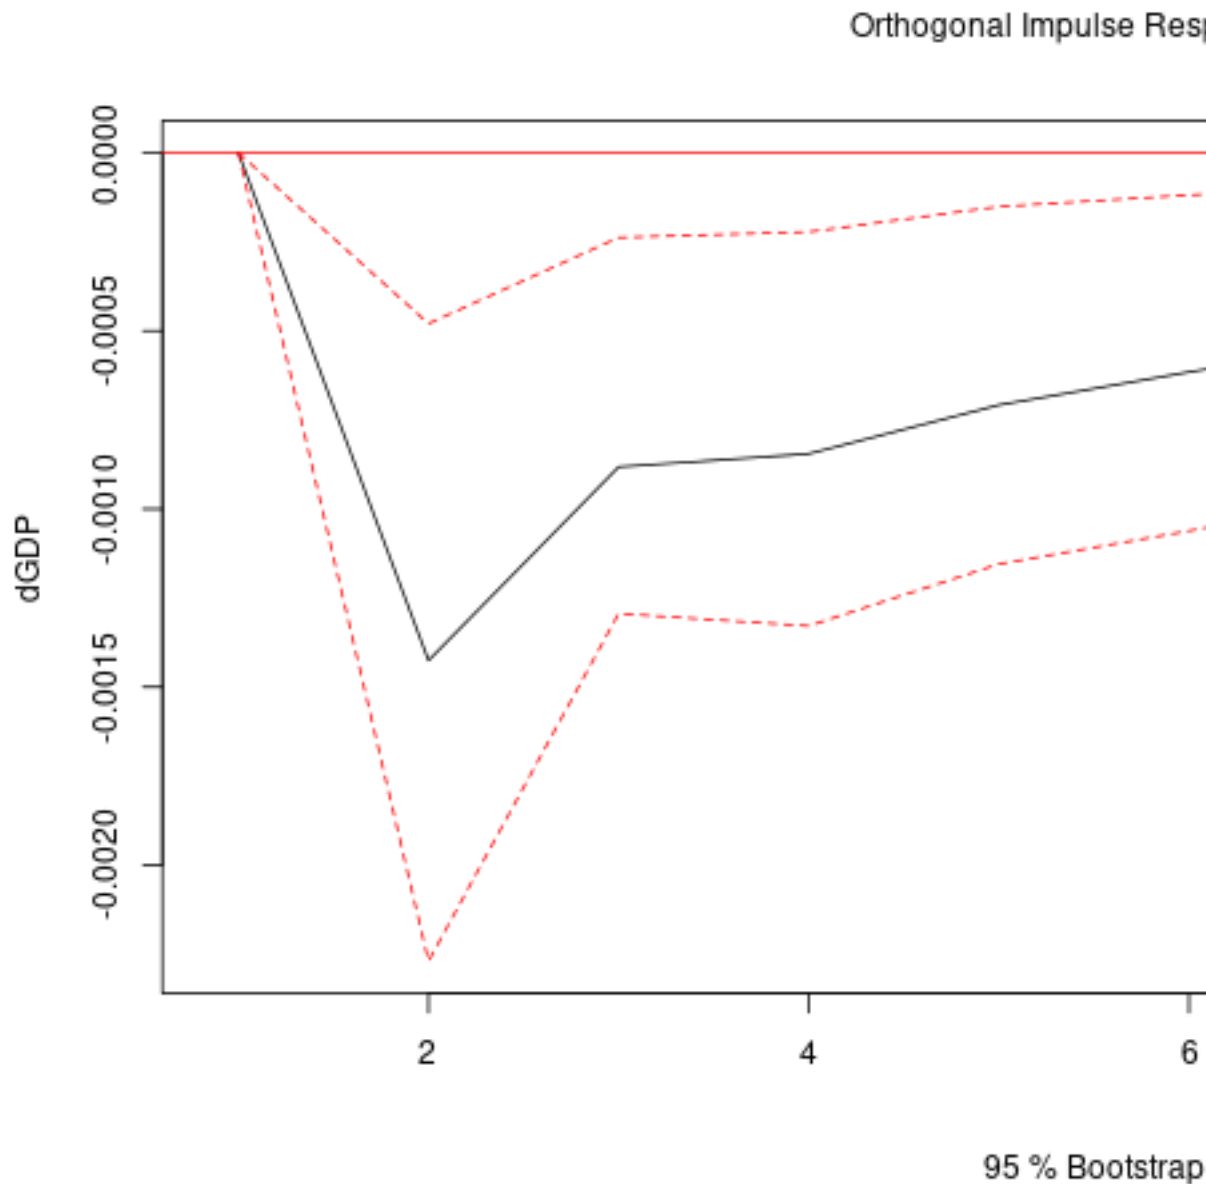
which confirms that RelSpread Granger-causes dGDP.
This is also shown in the impulse-response

```
> reg.irf <- irf(reg,response="dGDP",impulse="RelSpread")
> plot(reg.irf)
```

Orthogonal Impulse Res[

95 % Bootstrap

1.4 Predicting GDP Norway with spread – Out of Sample Analysis

We look at an example from Norway. Predict (changes in) GDP for mainlaind Norway, using the relative spread as the explanatory variable.

To confirm predictive power, one usually need to look into the out-of-sample performance of a given predictive variable to really conclude that the forecasting power is there.

We will therefore do some forecasting exercises.

Let us first show the R code for doing this

The reading of the data is similar to before

```
> library(forecast)
> outdir <- "../../results/2016_02_out_of_sample_forecasts/"
```

```
> source ("/home/bernt/data/2016/ssb/read_quarterly_gdp.R")
> head(qGDP)
1978(1) 1978(2) 1978(3) 1978(4) 1979(1) 1979(2)
 247094  249853  250733  255300  257066  257132
>
> source("../../data/read_quarterly_norwegian_liquidity_data.R")
> head(RelSpread)
  1980(1)    1980(2)    1980(3)    1980(4)    1981(1)    1981(2)
0.0391982 0.0396895 0.0386271 0.0338552 0.0496760 0.0494737
> names(RelSpread) <- "RelSpread"
>
> dGDP <- diff(log(qGDP))
> names(dGDP) <- "dGDP"
>                                            # now do out - of - sample analysis
> lRelSpread <- lag(RelSpread,-1)
> names(lRelSpread) <- "lRelSpread"
>                                            # use two lags of growth since the current lag is
> ldGDP <- lag(dGDP,-2)
> data <- merge(dGDP,ldGDP,RelSpread,lRelSpread,all=FALSE)
```

The above has gotten the data in the right form for analysis. Here I have aligned the predictive lagged variables to match the dependent variables, you may want to check that. Is the current dGDP and Spread aligned with the prevous quarter Spread, and two quarters ago dGDP?

```
> head(data)
                 dGDP          ldGDP RelSpread lRelSpread
1980(2) -0.011668339  0.013237630 0.0396895  0.0391982
1980(3)  0.002693454  0.013019691 0.0386271  0.0396895
1980(4)  0.004024743 -0.011668339 0.0338552  0.0386271
1981(1)  0.017495227  0.002693454 0.0496760  0.0338552
1981(2) -0.002339087  0.004024743 0.0494737  0.0496760
1981(3) -0.004511960  0.017495227 0.0458140  0.0494737
> dGDP <- data$dGDP
> ldGDP <- data$ldGDP
> RelSpread <- data$RelSpread
> lRelSpread <- data$lRelSpread
```

Now, ready to do the forecasting exercise. Do it by using the first 10 years (40 observations) to start estimation. Then a for loop over the remaining days, each time doing the prediction.

```
> n <- length(dGDP)
[1] 142
> actual <-NULL
> predicted0 <- NULL
> predicted1 <- NULL
>
> START <- 40
>
> for (t in START:n){
```

```
+    g <- dGDP[1:(t-1)]
+    lg <-   ldGDP[1:(t-1)]
+    lsp <- lRelSpread[1:(t-1)]
+
+    pred0 <- lm( g ~ lg )
+
+    pred1 <- lm( g ~ lg + lsp)
+
+    nd0 <- data.frame(lg <- ldGDP[t])
+    fc0 <- forecast.lm(pred0, nd0 ,h=1)
+    predicted0 <- c(predicted0, fc0$mean)
+
+    nd1 <- data.frame(lsp <-lRelSpread[t], lg <- ldGDP[t])
+    fc1 <- forecast.lm(pred1, nd1, h=1)
+    predicted1 <- c(predicted1, fc1$mean)
+
+    actual <- c(actual,dGDP[t])
+ }
```

Now, let us compare the forecasts
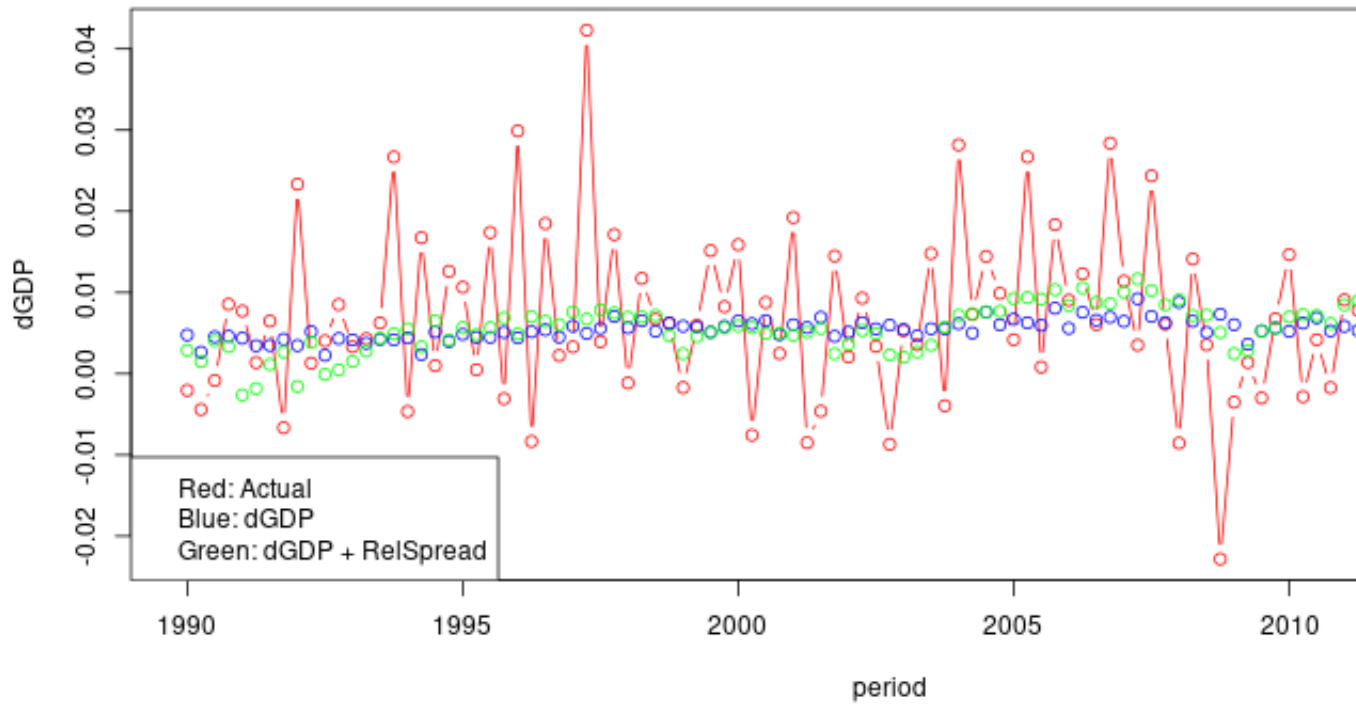The following is the R code for creating this picture.

```
> x <- index(dGDP)[START:n]
>
> filename <- paste0(outdir, "comparing_forecasts_with_without_spread.eps")
> postscript(filename,height=5,width=10,horizontal=FALSE)
> plot(c(predicted0,predicted1,actual)~c(x,x,x),
+       type="n",xlab="period",ylab="dGDP",
+       main="Comparing Forecasts with / without Spread")
> lines(actual~x,type='b',col="red")
> points(predicted0~x,,col="blue")
> points(predicted1~x,,col="green")
>
> text(1995,-0.013,"Red: Actual")
> text(1995,-0.016,"Blue: ldGDP")
> text(1995,-0.019,"Green: ldGDP + lSpread")
> dev.off()
```

Comparing Forecasts without and with spread

```
> predicted2 <- NULL
>                                                # now doing forecast rolling
> for (t in START:n){
+    g <- dGDP[(t-START+1):(t-1)]
+    lg <-  ldGDP[(t-START+1):(t-1)]
+    lsp <- lRelSpread[(t-START+1):(t-1)]
+    pred2 <- lm( g ~ lg + lsp)
+    nd2 <- data.frame(lsp <-lRelSpread[t], lg <- ldGDP[t])
+    fc2 <- forecast.lm(pred2, nd2 ,h=1)
+    prediction <- fc2$mean
+    predicted2 <- c(predicted2, prediction)
+ }
>
```
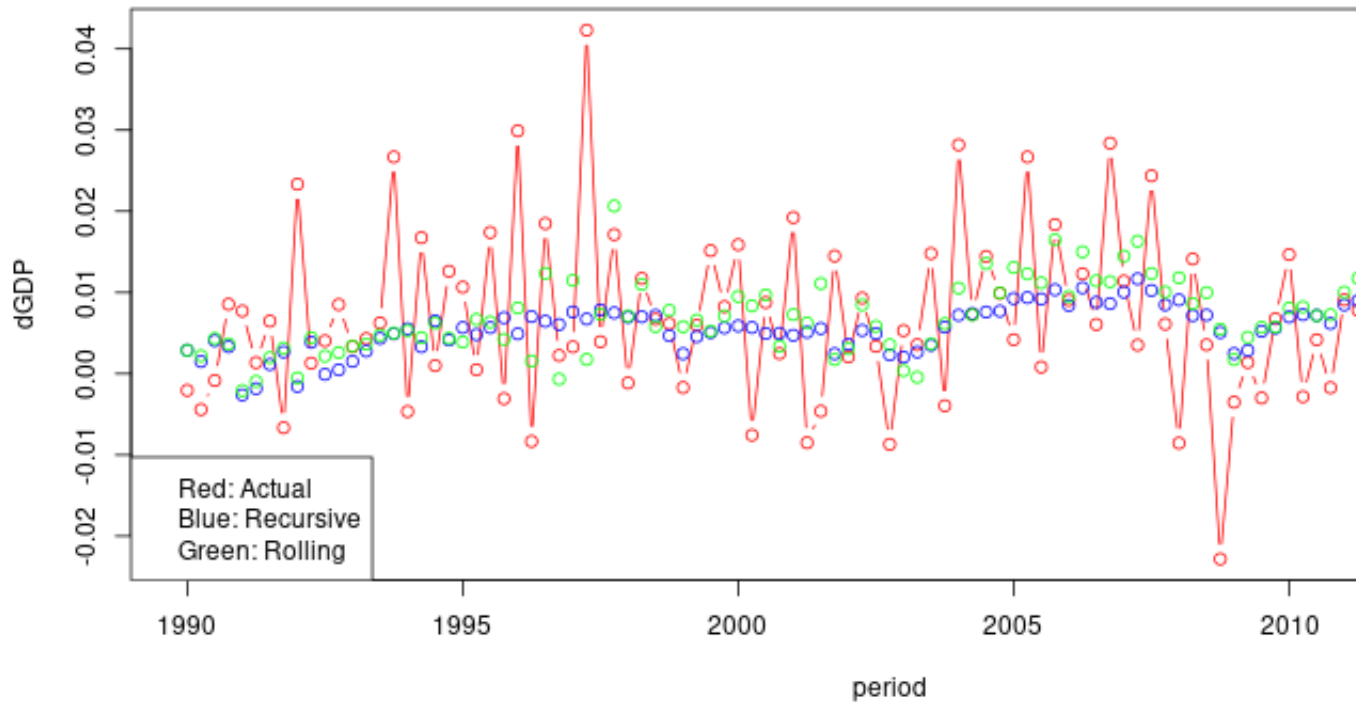
Compare the two methods graphically

Comparing Recursive and Rolling Forecasts

Again, R code that produced this graph

```
> filename <- paste0(outdir, "comparing_forecasts_recursive_rolling.eps")
> postscript(filename,height=5,width=10,horizontal=FALSE)
> plot(c(predicted1,predicted2,actual)~c(x,x,x),
+       type="n",xlab="period",ylab="dGDP",
+       main="Comparing Recursive and Rolling Forecasts")
> lines(actual~x,type='b',col="red")
> points(predicted1~x,,col="blue")
> points(predicted2~x,,col="green")
>
> text(1995,-0.013,"Red: Actual")
> text(1995,-0.016,"Blue: Recursive")
> text(1995,-0.019,"Green: Rolling")
```

Now, finally, some numbers for doing the comparison of forecasting performance.
First, just the first few actual numbers (which is after year 10).

```
> head(cbind(actual,predicted0,predicted1,predicted2))
            actual   predicted0    predicted1     predicted2
[1,] -0.0020493868 0.004777306  0.002887876  0.0028878764
[2,] -0.0044126042 0.002648651  0.001499642  0.0021917203
[3,] -0.0008367517 0.004537217  0.004064054  0.0043157129
[4,]  0.0085492579 0.004616494  0.003324567  0.0036604472
[5,]  0.0077134356 0.004370217 -0.002618097 -0.0020811801
[6,]  0.0013654204 0.003442451 -0.001833177 -0.0009533538
```

9

The various summary stats.

```
> p <- length(actual)
> res0 <- actual-predicted0
> mse0  <- sum(res0^2)/p
> mse0
[1] 9.560838e-05
> mae0  <- sum(abs(res0))/p
> mae0
[1] 0.006948177
>
> res1 <- actual-predicted1
> mse1  <- sum(res1^2)/p
> mse1
[1] 9.499064e-05
> mae1  <- sum(abs(res1))/p
> mae1
[1] 0.007300568
>
> res2 <- actual-predicted2
> mse2  <- sum(res2^2)/p
> mse2
[1] 9.768683e-05
> mae2  <- sum(abs(res2))/p
> mae2
[1] 0.007389973
>
> resR <- actual-predicted0
> resU <- actual-predicted1
> p=length(resR)
> h=1
> MSEr <- sum(resR^2)/p
> MSEr
[1] 9.560838e-05
> MSEu <- sum(resU^2)/p
> MSEu
[1] 9.499064e-05
> ENCNEW <- (p-h+1)* (sum(resR^2-resU*resR)/p)/MSEu
> ENCNEW
[1] 3.319659
> MSEF <- (p-h+1)*(MSEr-MSEu)/MSEu
> MSEF
[1] 0.6698309
> MSEu/MSEr
[1] 0.9935388
>
>                                          # doing same comparison for the rolling forecasts
>
> resU <- actual-predicted2
> p=length(resR)
```

```
> h=1
> MSEr <- sum(resR^2)/p
> MSEr
[1] 9.560838e-05
> MSEu <- sum(resU^2)/p
> MSEu
[1] 9.768683e-05
> ENCNEW <- (p-h+1)* (sum(resR^2-resU*resR)/p)/MSEu
> ENCNEW
[1] 7.887039
> MSEF <- (p-h+1)*(MSEr-MSEu)/MSEu
> MSEF
[1] -2.191495
> MSEu/MSEr
[1] 1.021739
```

References

Randi Næs, Johannes A Skjeltorp, and Bernt Arne Ødegaard. Stock market liquidity and the Business Cycle. *Journal of Finance*, LXVI:139–176, February 2011.

Johannes Skjeltorp and Bernt Arne Ødegaard. The information content of market liquidity: An empirical analysis of liquidity at the Oslo Stock Exchange. Working Paper, Norges Bank (Central Bank of Norway), November 2009.