Multiple Time Series Analysis

Alexander Benkwitz

This chapter demonstrates how to use XploRe for specifying, estimating, and interpreting **vector autoregressive** (VAR) **models**. VAR modeling belongs to **multiple time series analysis**. It is one approach in econometrics to describe a system of more than one equation and was introduced by Sims (1980). The classical multi equation modeling (e.g. Dhrymes 1978) is not considered here.

Using the XploRe quantlib `multi` we will build a model for aggregate money demand. The first section shows how to prepare data and how to start the interactive menu part of the quantlib. The second section explains what could be done in a preliminary analysis. The specification, estimation, and validation of a full VAR model is carried out in the third section. Finally, model interpretation is left for the last section.

For a detailed monograph on multiple time series analysis see Lütkepohl (1993). The quantlib `multi` is based on a software of Haase, Lütkepohl, Claessen, Moryson, and Schneider (1992).

1 Getting Started

The quantlib `multi` is an interactive tool for specifying and analyzing multiple time series models. In the next sections we will try to build a small German money demand model. Our approach is to specify a model in full VAR form. This can be justified as a first step since the full VAR is a quite general model. It is also chosen here for exposition purpose.

Such a model should obviously include the money demand itself and some of its determinants. Money demand is set equal to the money circulating in the economy (henceforth referred as M). Two determinants are included: the transaction volume of the economy which is given by the gross national product

(GNP, henceforth referred as Y) and the costs of holding money which is given by an interest rate (henceforth referred as I). From economic theory we know that Y has a positive whereas I has a negative effect on M. We might want to find out whether these theoretical considerations are supported by some observed data and whether a full VAR model is a satisfactory description of the relation between the variables involved.

But before starting the analysis we need to prepare the data and start the quantlet `domulti`.

1.1 Data Preparation

The time series need to be provided as a numeric $T \times K$ data matrix `x`, where T is the length of the time series, and K its dimension. The first K-dimensional vector of observations is in the first row of `x`, the second vector of observations is in the second row of `x`, and so forth. The data matrix must not contain verbal information. Additional data information (variable names, observation period) are given directly to `domulti`.

This chapter will use the dataset `mts` which contains three time series:

- German real money—M1 (M),
- German real gross national product (Y), and
- German interest rate (I).

The sample begins in 1960, 1-st quarter and ends 1989, 4-th quarter. It contains 120 quarterly observations. M and Y are seasonally adjusted and in prices of 1985. It is stated in Billion of German Marks. All data is taken from the OECD.

1.2 Starting `multi`

The quantlet `domulti` takes up to five inputs: The data matrix x, the main period of the first observation (1960), the sub period of the first observation (1), the periodicity (4), and the variable names (M, Y, and I).

Call `domulti` by

```
library("multi")
x=read("mts.dat")
domulti(x,1960,1,4,"M"|"Y"|"I")
```

mts01.xpl

The main menu

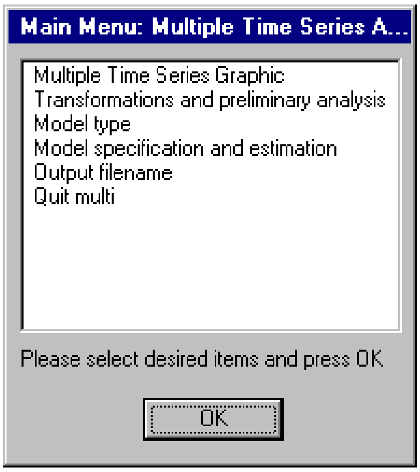

will appear which means that we have entered the menu driven part of the `multi` quantlib. We can start the analysis now.

2 Preliminary Analysis

This section should generally answer the question what model type corresponds best to the given data. Since the task is to fit a full VAR model the consideration is restricted to the question whether the given data set fits well in a full VAR framework. For this we note that the inference we want to make in Sections 3 and 4 requires data generated by a stable process. Stability implies mean and variance stationarity of the data. These features will be of interest in the following preliminary analysis.

2.1 Plotting the Data

It is good practice to start time series investigation by just visual inspection of the data graphs. We can view all time series in one chart or separate charts. Since we deal with multiple time series analysis we choose option one. This gives the following picture:

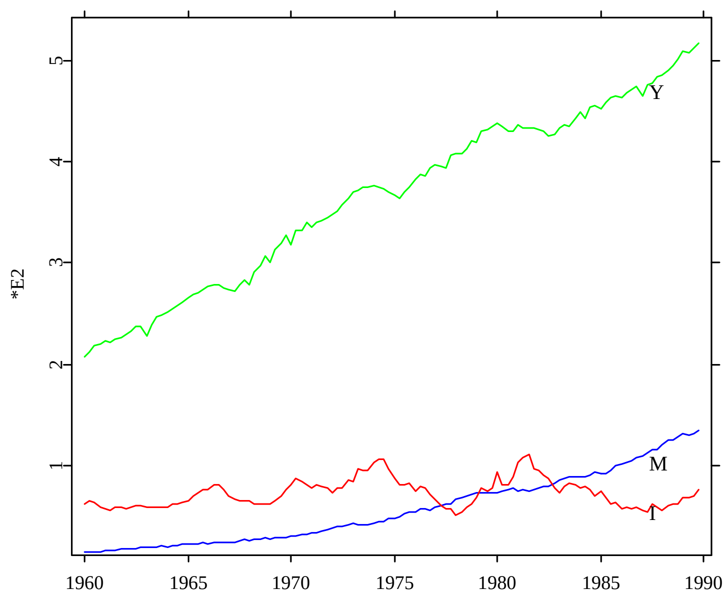

Sample period 1960 (1) -- 1989 (4), 120 observations

In order to display the interest series (I) together with the other two series in one panel we changed its scaling by factor 10. It seems that M and Y may be subject to linear trend and/or exponential growth. It is clear from this graph that M and Y have no stationary mean.

2.2 Data Transformation

It is common to handle linear trend by differencing the data. Exponential growth can be transformed by applying the natural logarithm. Exactly these two transformations are supported. If both transformations are chosen the logarithmic transformation is automatically performed first.

Further transformations may be performed with XploRe before the data matrix x is given to `domulti`.

Here we choose both transformations for the series M and Y. Since we deal with seasonally adjusted data we use the default differencing lag of 1.

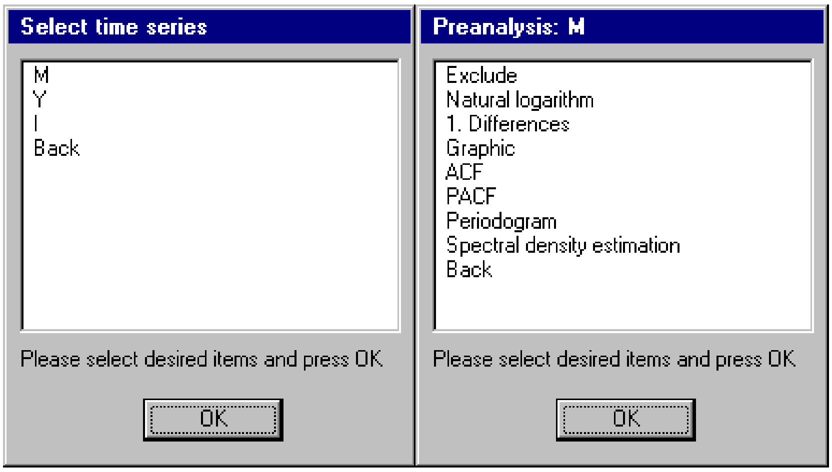

After performing the transformations with the two menus above we plot the transformed time series which gives the following graphics:

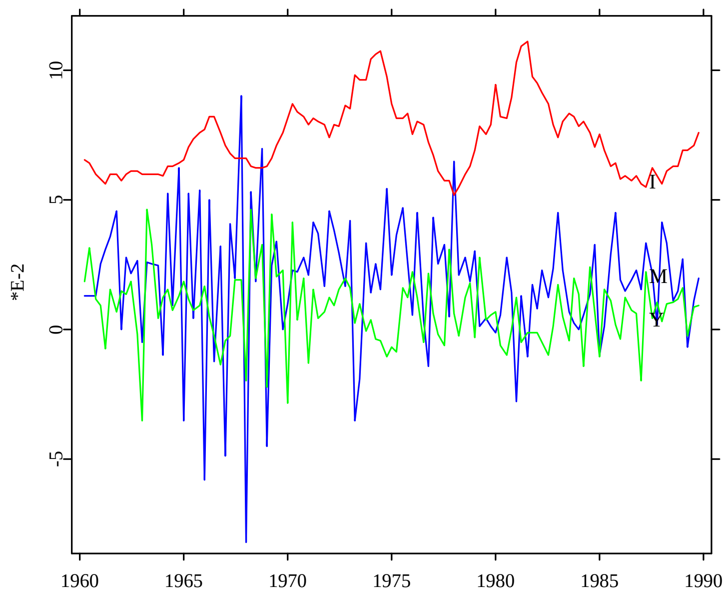

Sample period 1960 (2) -- 1989 (4), 119 observations

Note that our sample size reduced to 119 observations after differencing once. In the last picture the interest rate series is scaled by factor 10^{-2}.

Now it is reasonable to assume mean and variance stationarity of the M and Y series. However, at the beginning of both series we still observe a period of high fluctuations compared with the end. We might keep this feature in mind for later steps of the analysis.

3 Specifying a VAR Model

The starting point in the Sims methodology (Sims 1980) is the formulation of an unrestricted VAR model. We will specify a VAR model of order p (VAR(p)) in the following general form:

$$y_t = \nu + A_1 y_{t-1} + \ldots + A_p y_{t-p} + u_t \ , \qquad t = 1, \ldots, T \ , \tag{1}$$

where

- $y_t = (y_{1t}, \ldots, y_{Kt})^T$ is the K-dimensional vector of the time series at time t,
- $\nu = (\nu_1, \ldots, \nu_K)^T$ is a vector of ones (intercept),
- $u_t = (u_{1t}, \ldots, u_{Kt})^T$ is a K-dimensional disturbance vector with covariance matrix Σ_u,
- A_i is the parameter matrix of y_{t-i}, $i = 1, \ldots, p$.

To analyze a model in the form (1) we need to set the model type to Full VAR after selecting Model type in the Main Menu:

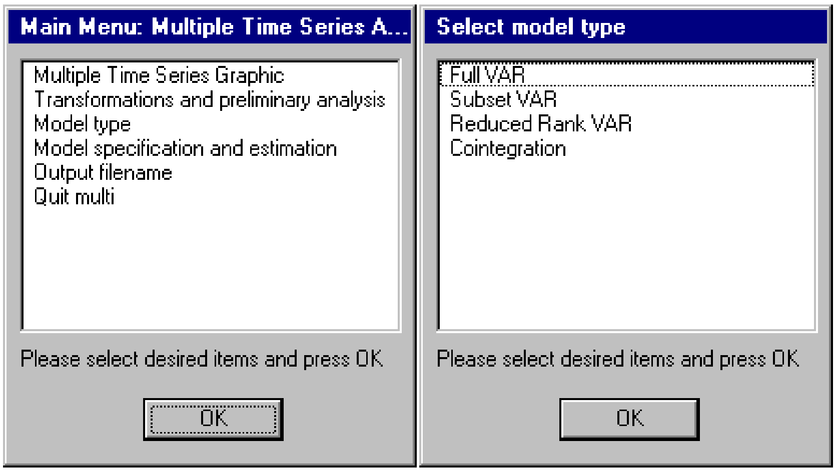

In our model y_t is specified with $y_t = (\Delta(\ln M), \Delta(\ln Y), I)_t^T$. The variable order in the vector y_t is given by the data matrix x. The next steps include finding a suitable model order p and estimating the model parameter matrices $\nu, A_1, \ldots, A_p$. Having done this we should conduct a residual analysis to check the whiteness assumption. If the residual analysis is satisfactory we can use the model for interpretation and forecasting.

3.1 Process Order

We can use economic theory or information contained in the data for specifying the model order p. Since we have no a priori knowledge from theory we use statistical tools for choosing an appropriate p. The quantlib `multi` provides the FPE, AIC, HQ and SC criteria (see **?**, Chapter 4). They all compare different VAR(p) models with $p = 0, \ldots, p_{\max}$ with respect to some objective function. The order p which optimizes the function is the recommended order.

Before we apply the order selection criteria we must set the highest possible order $p_{\max}$. This can be difficult: In order to avoid an optimum at the edge and to restrict the parameter space not too much $p_{\max}$ should be reasonable large. On the other hand $p_{\max}$ must not be too large since we need at least $p_{\max} + 1$ presample values which reduces the sample size T and results in unprecise estimates or worst in a model that cannot be estimated. Since we deal with quarterly data we should consider at least the periodicity as a possible process order. Moving a bit "away" from the periodicity we set $p_{\max} = 7$.

In the Main Menu we select Model specification and estimation and prepare the subsequent call to the model selection criteria:

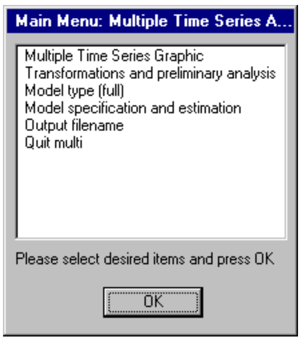

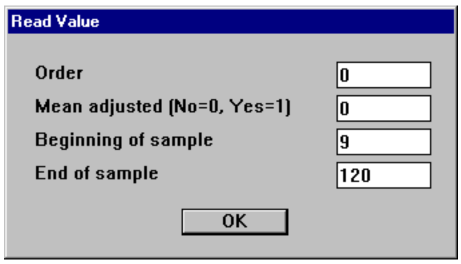

For that we must divide the data set in a presample and sample period. We can ignore the input fields Order and Mean adjusted (set to 0). The presample period must contain at least $p_{\max}$ observations. Since we have differenced the data once one more observation is 'lost'. Therefore the Beginning of sample is set to $p_{\max} + 2 = 9$. If the sample is not split appropriately an error message appears in the output window which indicates the problem.

Press OK to enter the menu of VAR estimation results (main results menu) and select VAR order criteria. Here we are asked to input $p_{\max}$:

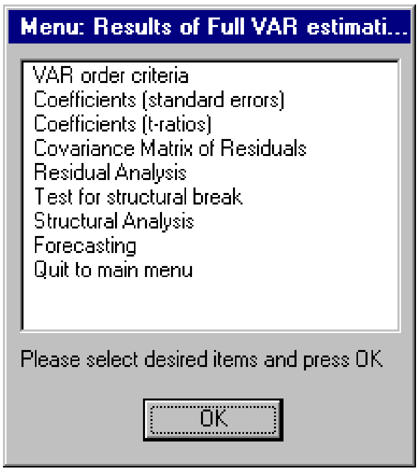

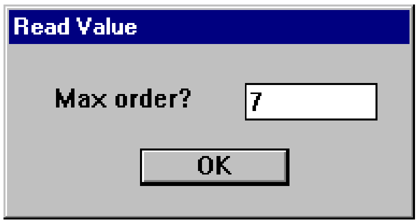

The results of the order selection criteria will be presented in a separate window. The optimum for every criterion is found at the minimum value:

Order	ln(FPE)	AIC	HQ	SC
0	-15.3010	-15.3546	-15.3546	-15.3546
1	-17.5302	-17.5839	-17.4952	-17.3654
2	-17.6266	-17.6806	-17.5033	-17.2437
3	-17.6536	-17.7086	-17.4427	-17.0532
4	-17.9254	-17.9821	-17.6276	-17.1083
5	-17.8242	-17.8836	-17.4405	-16.7914
6	-17.7468	-17.8103	-17.2785	-16.4996
7	-17.6602	-17.7293	-17.1089	-16.2002

The optimum values are FPE(4) $= -17.9254$, AIC(4) $= -17.9821$, HQ(4) $= -17.6276$, and SC(1) $= -17.3654$. The recommendation of the SC-criterion is quite different from the others. Such a result is not uncommon. For a detailed discussion about the properties of the criteria see Lütkepohl (1993, Chapter 4).

We start our analysis with $p = 4$ but should keep in mind the other possible process order. Thus we start with a VAR(4) which is the most general model supported by the data. This also includes a VAR(1) by setting $A_2 = A_3 = A_4 = 0$.

3.2 Model Estimation

In order to estimate a VAR(4) we need to go back to the Main Menu and select Model specification and estimation again. However, this time we set the Order to 4 for estimating a VAR(4) model:

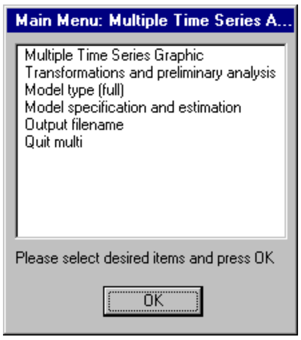

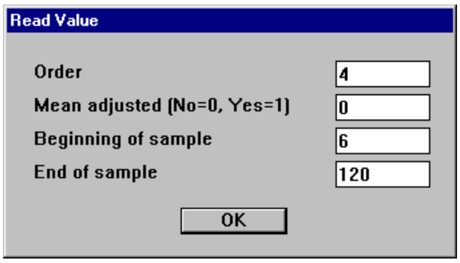

Press OK to enter the results main menu. The VAR(4) is estimated by multivariate least squares. Next we view the estimates $\hat{\nu}, \hat{A}_1, \ldots, \hat{A}_4$ and their t-values:

```
Intercept  (t-ratios in parentheses)

 0.03 ( 1.86)
 0.02 ( 2.24)
 0.36 ( 0.93)
```

```
 A1

-0.28 (-3.05)   0.15 ( 1.00)  -0.02 (-4.04)
-0.06 (-1.00)  -0.28 (-2.98)   0.00 ( 0.50)
 1.91 ( 0.82)   0.93 ( 0.24)   1.09 (11.04)
```

```
 A2

 0.03 ( 0.27)  -0.35 (-2.32)   0.01 ( 1.83)
 0.02 ( 0.41)  -0.09 (-0.94)   0.00 (-0.90)
-0.05 (-0.02)   4.11 ( 1.06)  -0.11 (-0.72)
```

```
 A3

-0.14 (-1.60)   0.03 ( 0.18)   0.00 (-0.30)
 0.14 ( 2.52)  -0.06 (-0.61)   0.00 (-0.85)
 1.88 ( 0.85)   0.87 ( 0.23)   0.07 ( 0.49)
```

```
 A4

 0.44 ( 5.09)  -0.02 (-0.16)   0.01 ( 1.31)
 0.19 ( 3.41)   0.22 ( 2.47)   0.00 ( 1.05)
 2.23 ( 1.00)   1.32 ( 0.37)  -0.13 (-1.17)
```

It can be seen that not all elements of the parameter matrices are significant different from zero. Especially in $\hat{A}_2$ and $\hat{A}_3$ we observe only one significant value. This could be the starting point for choosing a subset VAR where single elements of A_i are restricted to zero.

Selecting Covariance matrix of residuals from the main results menu displays the estimated residual covariance and correlation matrices.

```
Covariance Matrix of Residuals
Determinant =   1.16E-08

 3.56e-04   2.00e-05  -8.31e-04
 2.00e-05   1.42e-04   5.57e-04
-8.31e-04   5.57e-04   2.37e-01

Correlation Matrix of Residuals

 1.000   0.089  -0.090
 0.089   1.000   0.096
-0.090   0.096   1.000
```

The correlation matrix tells us something about the contemporaneous correlation structure in the residual vector u_t. We note here that there is no correlation in u_t. At a later point we will come back to the implications of this feature.

3.3 Model Validation

In Subsection 3.2 we have estimated a VAR(4)–model. Since we did not know the "correct" order we used statistical tools to find a reasonable one. Some estimation results were presented. Partly they are based on properties of the estimator (limiting normal distribution) which assume certain conditions. Whether these conditions hold is checked in this subsection. One can think of a residual analysis, tests for nonnormality and tests for structural change. Here we will consider the residual analysis and a test for nonnormality in more detail.

Checking the whiteness of the residuals is a prerequisite for drawing valid conclusions from the t-values presented above. If we want to compute reliable forecast intervals we need to check the normality of the residuals in addition.

From the main results menu we select Residual Analysis which enables us to go through the three steps of residual analysis in `multi`:

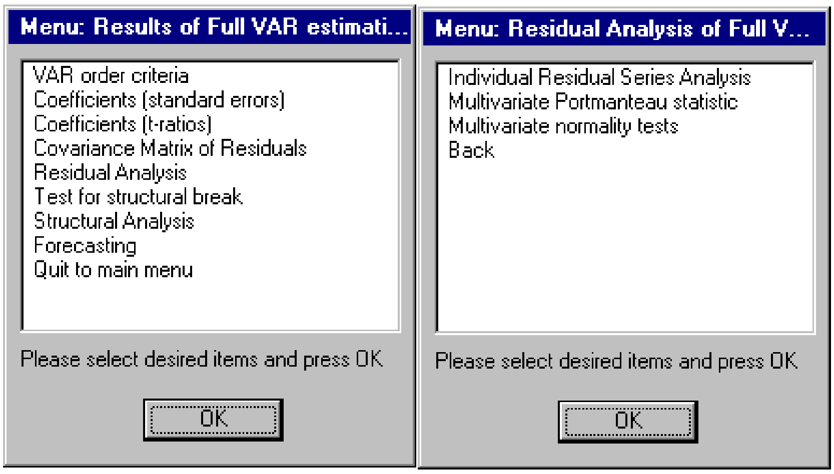

Individual residual analysis

First we have to select one equation. Then we have the chance to do some transformations to the estimated residuals $\hat{u}_{i,t}$. We selected here Residuals which leaves $\hat{u}_{i,t}$ untransformed.

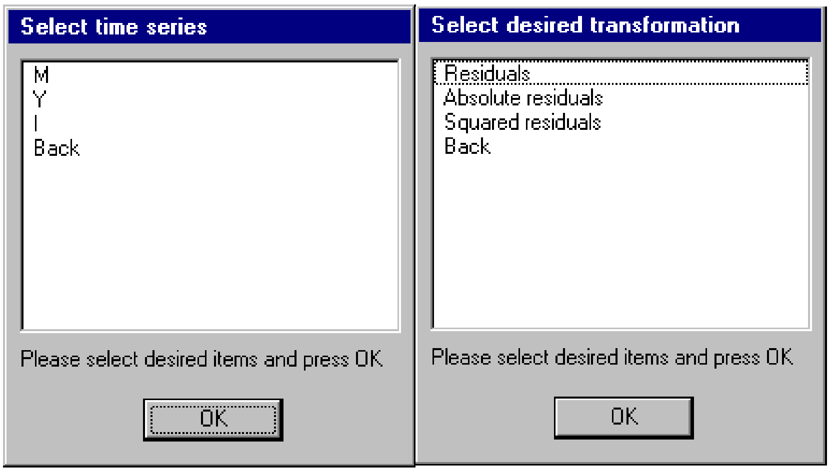

From the following menu we present here the Plot of residuals. We do this for all equations.

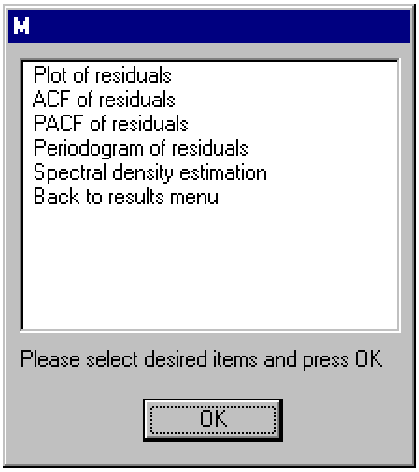

In the residual plots the unit of measurement is one standard deviation. In other words, the residuals are normalized to have unit variance. Thus, if many residuals exceed 2 in absolute value this may be evidence for nonnormality or nonlinear features that are not adequately captured by the model. Furthermore we might look for distinct patterns in the residual plots that rule out whiteness.

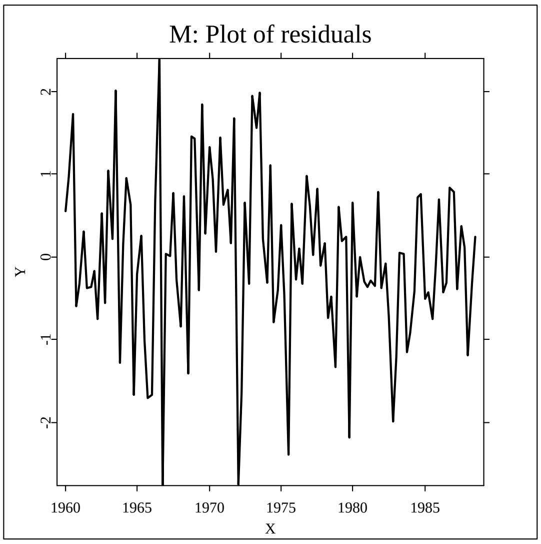

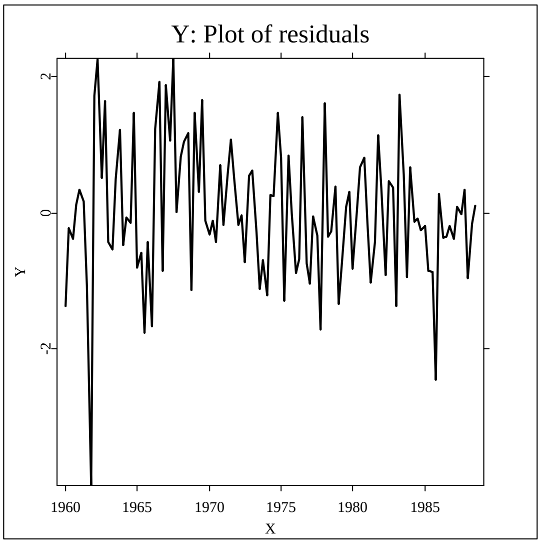

Y: Plot of residuals
2
0
-2
Y
1960
1965
1970
1975
1980
1985
X

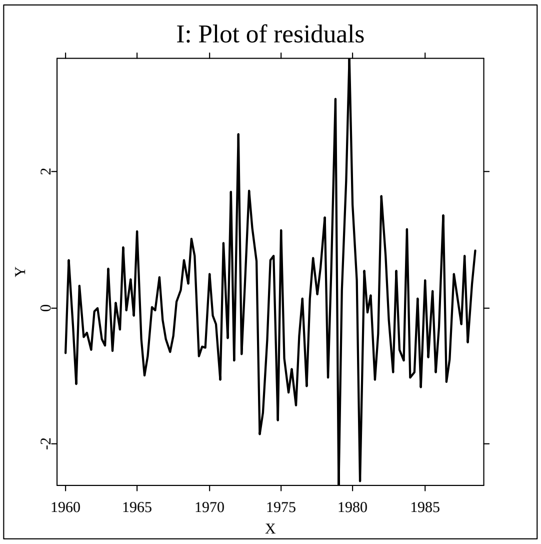

I: Plot of residuals
2
0
-2
Y
1960
1965
1970
1975
1980
1985
X

Multivariate portmanteau statistic

Checking the white noise assumption for the residuals is a central issue. Many inferential procedures rely on this assumption.

The menu point Multivariate Portmanteau statistic provides two tools. Here we look at the residual vector u_t at time points t and $t-i$. For these we compute the i-th autocorrelation matrix $R_u(i)$. White noise means zero autocorrelation for all $i \geq 1$. Before checking the autocorrelation functions and carrying out an overall test we are asked to input a maximum lag h we want to check autocorrelation for:

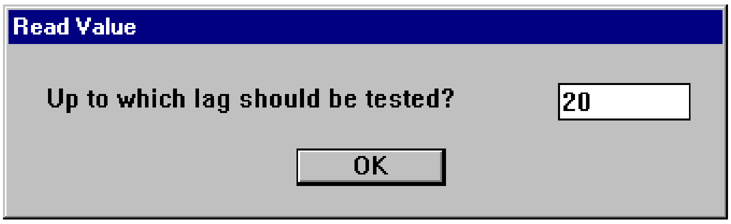

For the overall test to work the maximum lag of the autocorrelations to be included must exceed the order of the process as otherwise negative degrees of freedom of the approximating χ^2 distribution will result. If a lag h less than or equal to the VAR order is specified a warning is given and the statistic is computed for the smallest feasible lag. Generally, the χ^2 approximation to the true distribution of the modified portmanteau statistic may be inappropriate for small lags h. At the same time we must make sure that $h < T$ for obvious reasons.

Here we have chosen $h = 20$. The resulting plots of the autocorrelation functions appear. The autocorrelation plots come along with approximate $\pm 2/\sqrt(T)$ confidence bounds. These plots do not exhibit significant autocorrelations. Especially the $\hat{\rho}_i, t$ with $t \leq p$ are much smaller than the approximate confidence bound which is a good result since the exact confidence bound for smaller autocorrelation lag can be much smaller than the approximate.

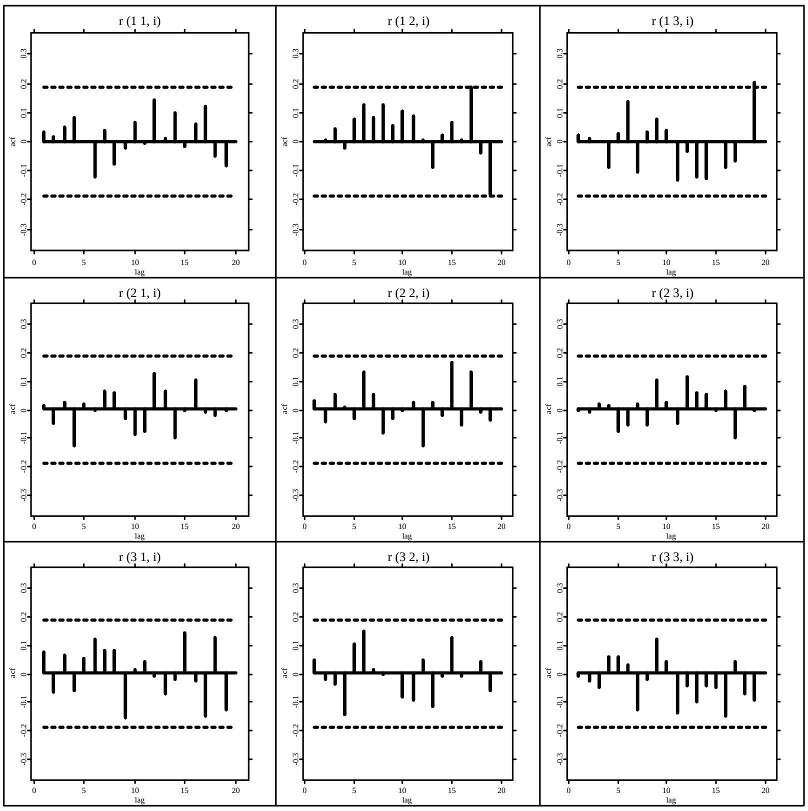

A test for the overall significance of the residual autocorrelations up to lag h appears in a second display. It is the result of the test

$$H_0 : \boldsymbol{R}_h = (R_u(1), \ldots, R_u(h)) = 0 \qquad \text{vs.} H_1 : \boldsymbol{R}_h \neq 0 \tag{2}$$

The value of the modified portmanteau statistic $\bar{P}_h$ (see Lütkepohl 1993, Chapter 4) is shown:

```
 Multivariate Portmanteau statistic

Lag order h                        20
Model order p                       4
Degrees of freedom  dim^2(h-p)   144
statistic P                    145.74
chi2(P,df)                     0.5562

Reject e.g. if chi2(P,df) > 0.95
```

As expected we cannot reject H_0 at a 95% significance level. This result is in line with the shown residual autocorrelation functions above. The residuals of our model do not exhibit autocorrelation.

Multivariate normality test

Multivariate Normality test displays the χ^2-statistics associated with the skewness and kurtosis of the residuals which may be used for tests of nonnormality.

	x	df	Chi2(df,x)
Skewness:	4.69	3	0.80
Kurtosis:	1.43	3	0.30
Skewness and Kurtosis:	6.12	6	0.59

```
  x: Value of Statistic

  df: degrees of freedom

  Chi2(df,x): Chi-square cdf

  Reject normality e.g. if Chi2(df,x) > 0.95 !
```

The result shows that normality of the residuals is not rejected on grounds of the test of kurtosis and the joint test of skewness and kurtosis. The implication is that the confidence intervals computed for the forecasts are reliable.

4 Structural Analysis

The interpretation of VAR models based on parameter matrices $A_1, \ldots, A_p$ is clearly restricted. Therefore concepts and tools were developed to interpret VAR models easily. The most important are the causality concepts, forecast error variance decomposition and the impulse response analysis. In this section we will deal with the impulse response analysis.

4.1 Impulse Response Analysis

The impulse response analysis quantifies the reaction of every single variable in the model on an exogenous shock to the model. Two special cases of shocks can be identified: The single equation shock and the joint equation shock where the shock mirrors the residual covariance structure. In the first case we investigate forecast error impulse responses, in the latter orthogonalized impulse responses. The reaction is measured for every variable a certain time after shocking the system. The impulse response analysis is therefore a tool for inspecting the inter–relation of the model variables.

We enter the impulse response analysis directly when selecting the menu point Structural Analysis in the main results menu:

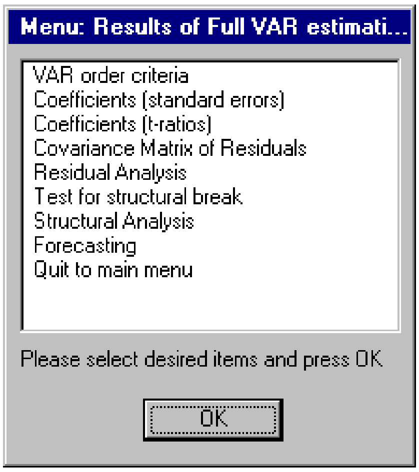

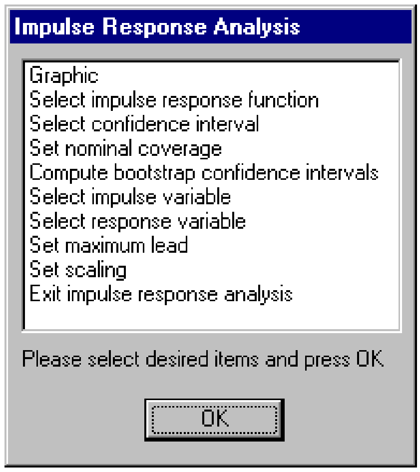

Now we have to decide about the time horizon. Here we have chosen 12 periods which is a time span of three years in our model. Next we select forecast error impulse response and orthogonalized impulse response functions in turn,

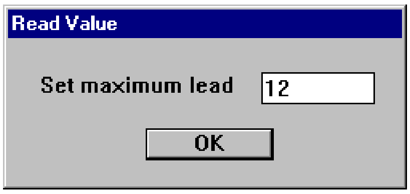

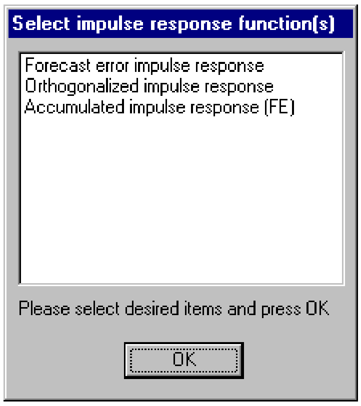

which gives the following pictures. On the left side we see the forecast error impulse responses, on the right side the orthogonalized impulse responses:

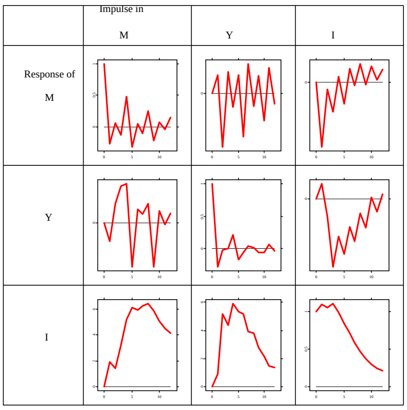

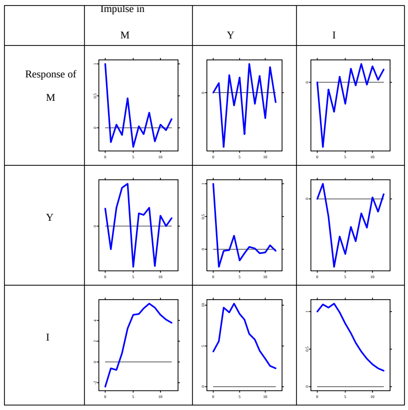

Not surprisingly there is no big difference between the two pictures. This was expected since looking at the residual correlation structure in Subsection 3.2 we do not see any strong correlation between the residuals of the equations.

When interpreting these charts we have to keep in mind that M and Y enter the model as $\Delta(\ln M)$ and $\Delta(\ln Y)$. In the first row of charts we see the response of money growth rate to a unit impulse in money growth rate, GNP growth rate, and interest, respectively. The second and third row of charts show the response pattern of GNP growth rate and interest.

The charts in the last row/last column are the reaction patterns we expected. They show the negative relation of interest and money/GNP. However, all money/GNP charts do not show a definite pattern. We could assume that after the initial impulse the true impulse responses are zero. A measure for checking the accuracy of the estimated impulse responses is desirable. It is provided in Subsection 4.2.

In our model it might be particularly interesting to analyze accumulated impulse responses. Accumulated impulse responses at time horizon h are obtained by summing up all impulse responses from 0 to h. Selecting this type of impulse responses function gives the following picture:

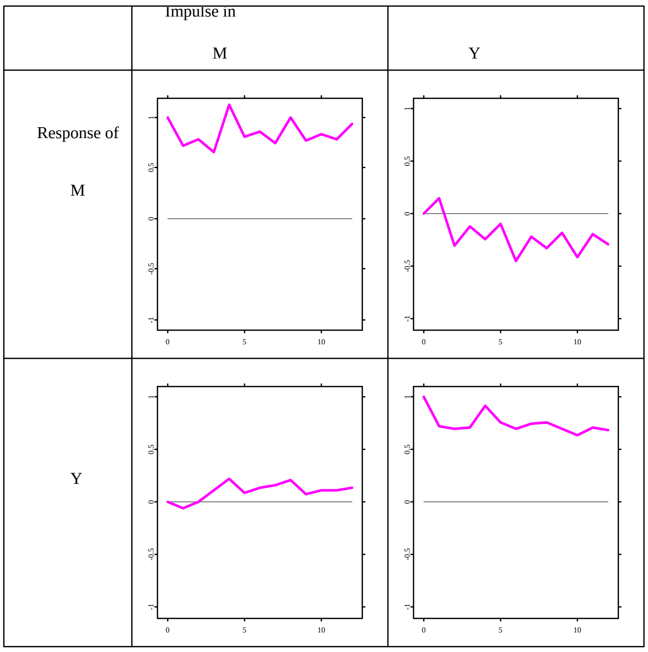

It shows that in the long run the total impulse response from money to money and from GNP to GNP is slightly below one and the cross total impulse responses are close to zero. The negative effect of GNP on money seems not very plausible. We therefore need some measure for checking the accuracy of the estimated impulse responses.

4.2 Confidence Intervals for Impulse Responses

The impulse response plots in Subsection 4.1 are based on the model estimates in Subsection 3.2 and therefore also estimates. In order to make inference on statistical grounds we need some measures for the reliability of these estimates. `multi` provides confidence intervals. There are the asymptotic normal distribution confidence interval (Lütkepohl 1993, Chapter 3), and two types of bootstrap confidence intervals (e.g. Benkwitz, Lütkepohl, and Wolters 2000):

Let ϕ be an arbitrary impulse response, $\hat{\phi}_T$ its estimate based on a sample size T, and $\hat{\phi}_T^*$ a bootstrapped impulse response. The bootstrap confidence intervals are based on the statistics $\hat{\phi}_T^*$ (Efron percentile) and $\hat{\phi}_T^* - \hat{\phi}_T$ (Hall Percentile). A precise description of these confidence intervals can be found in Efron and Tibshirani (1993) and Hall (1992). In order to compute bootstrap confidence intervals we have to set the number of bootstrap drawings. Here we have set this number to 1,000 drawings.

Confidence intervals based on the asymptotic normal distribution are known to fail even asymptotically in some cases (Lütkepohl 1993). Furthermore their small sample properties might be bad (Kilian 1998). The first problem cannot be solved by the bootstrap confidence intervals implemented in `multi` (Benkwitz, Lütkepohl, and Neumann 2000). However, the second might be tackled by the bootstrap.

Here we have decided to compute confidence intervals based on a nominal coverage of 95%.

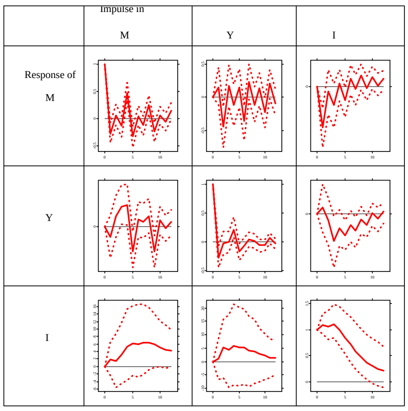

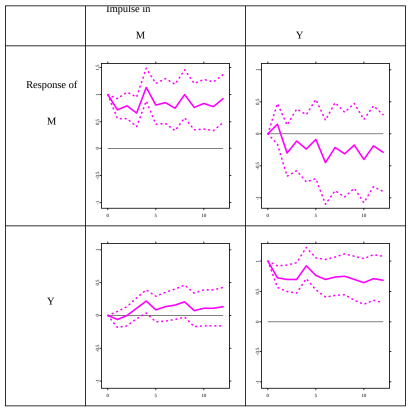

The left picture shows the 95% Hall percentile confidence intervals for the forecast error impulse responses. The right picture shows the 95% Hall percentile confidence intervals for the accumulated impulse responses. Note that the confidence intervals are computed in a pointwise manner. We are now able to determine which impulse response is significant. It turns out that the response of interest to an impulse in money growth rate and GNP growth rate (left picture) is insignificant. Furthermore, our hypothesis about the zero impulse responses seems to be validated.

To sum up we find many insignificant and very little significant impulse responses. This might be due to preliminary data transformation and/or model choice. One might attempt to build a subset or cointegration model. These models possibly better fit the data which results in better interpretation.

References

Benkwitz, A., Lütkepohl, H., and Neumann, M. (2000). Problems related to confidence intervals for impulse responses of autoregressive processes, *Econometric Reviews* **19**: 69–103.

Benkwitz, A., Lütkepohl, H., and Wolters, J. (2000). Comparison of bootstrap confidence intervals for impulse responses of german monetray systems, *Macroeconomic Dynamics*, to appear.

Dhrymes, P. (1978). *Introductory Econometrics*, Springer, New York.

Efron, B. and Tibshirani, R. (1993). *An Introduction to the Bootstrap*, Chapman & Hall, New York.

Haase, K., Lütkepohl, H., Claessen, H., Moryson, M., and Schneider, W. (1992). *MulTi: A Menu-Driven GAUSS Program for Multiple Time Series Analysis*, Institut für Statistik und Ökonometrie, Universität Kiel, Kiel.

Hall, P. (1992). *The Bootstrap and Edgeworth Expansion*, Springer, New York.

Kilian, L. (1998). Small-sample confidence intervals for impulse response functions, *The Review of Economics and Statistics* **80**: 186–201.

Lütkepohl, H. (1993). *Introduction to Multiple Time Series Analysis*, 2nd edn, Springer, Berlin.

Sims, C. (1980). Macroeconomics and reality, *Econometrica* **48**: 1–48.