

1259

Discussion
Papers

Identifying Structural Vector Autoregressions via Changes in Volatility

Opinions expressed in this paper are those of the author(s) and do not necessarily reflect views of the institute.

IMPRESSUM

© DIW Berlin, 2012

DIW Berlin
German Institute for Economic Research
Mohrenstr. 58
10117 Berlin

Tel. +49 (30) 897 89-0
Fax +49 (30) 897 89-200
<http://www.diw.de>

ISSN print edition 1433-0210
ISSN electronic edition 1619-4535

Papers can be downloaded free of charge from the DIW Berlin website:
<http://www.diw.de/discussionpapers>

Discussion Papers of DIW Berlin are indexed in RePEc and SSRN:
<http://ideas.repec.org/s/diw/diwwpp.html>
<http://www.ssrn.com/link/DIW-Berlin-German-Inst-Econ-Res.html>

December 4, 2012

Identifying Structural Vector Autoregressions via Changes in Volatility¹

Helmut Lütkepohl

Department of Economics, Freie Universität Berlin and DIW Berlin
Mohrenstr. 58, D-10117 Berlin, Germany
email: hluetkepohl@diw.de

Abstract. Identification of shocks of interest is a central problem in structural vector autoregressive (SVAR) modelling. Identification is often achieved by imposing restrictions on the impact or long-run effects of shocks or by considering sign restrictions for the impulse responses. In a number of articles changes in the volatility of the shocks have also been used for identification. The present study focusses on the latter device. Some possible setups for identification via heteroskedasticity are reviewed and their potential and limitations are discussed. Two detailed examples are considered to illustrate the approach.

Key Words: Markov switching model, vector autoregression, heteroskedasticity, vector GARCH, conditional heteroskedasticity

JEL classification: C32

¹This paper was written while the author was a Bundesbank Professor. Comments by Lutz Kilian on an earlier version of the paper are gratefully acknowledged.

1 Introduction

In structural vector autoregressive (SVAR) analysis, typically a critical issue is the identification of economically meaningful shocks. Restrictions on the impact effects, the long-run responses of the variables, or the signs of the impulse responses are usually used to identify or characterize the shocks of interest (e.g., Sims (1980), Amisano and Giannini (1997), Blanchard and Quah (1989), King, Plosser, Stock and Watson (1991), Faust (1998), Canova and De Nicoló (2002), Uhlig (2005)). Such restrictions are sometimes controversial. Taking that into account, another strand of the literature uses statistical data properties for identification. In particular, changes in the volatility of the residuals or the variables is used as additional identification information. For instance, Rigobon (2003), Rigobon and Sack (2003), Lanne and Lütkepohl (2008) use heteroskedasticity while Normandin and Phaneuf (2004), Bouakez and Normandin (2010) and Lanne, Lütkepohl and Maciejowska (2010) consider conditional heteroskedasticity for the identification of shocks.

In the present study the general principle is discussed and specific modelling strategies are considered. First the general modelling strategy is presented and its advantages and limitations are discussed in Section 2. The general idea is that the structural shocks are extracted by transforming the reduced form residuals. The basic assumption that structural shocks should be instantaneously uncorrelated is insufficient to uniquely determine or identify the shocks. Therefore further restrictions have to be imposed for uniquely identifying the shocks or at least characterizing them sufficiently to make them distinguishable. In a conventional analysis this is typically done through exclusion or sign restrictions. In the present study possibilities are considered how changes in volatility can be used for this purpose. It is assumed that the impulse responses remain invariant even if the volatility of the shocks is changing, which provides additional restrictions for the shocks that can even uniquely determine the shocks if the changes in volatility are sufficiently heterogeneous.

Some specific models for changes in volatility that can be used in this context are considered in Section 3. In particular, a standard heteroskedasticity model with exogenously generated shifts in variance is discussed, an approach based on a multivariate generalized autoregressive heteroskedastic (MGARCH) model for the movements in volatility is presented and a Markov regime switching mechanism is considered for capturing changes in the residual volatility. Detailed examples are provided in Sections 4 and 5. The first example compares different identification schemes for monetary policy shocks that have been considered in the literature. The second example considers

models for the crude oil market. Finally, Section 5 concludes.

2 The Model Setup

2.1 The basic model

The point of departure is a K -dimensional reduced form VAR(p),

$$y_t = \nu + A_1 y_{t-1} + \cdots + A_p y_{t-p} + u_t. \quad (1)$$

Here ν is a $(K \times 1)$ constant term, A_j ($j = 1, \dots, p$) are $(K \times K)$ VAR coefficient matrices and u_t is a zero-mean white noise error term. The model may contain cointegrated variables and may in fact be set up as a vector error correction model. Also there may be more general deterministic terms such as seasonal dummies or a polynomial trend.

The standard SVAR analysis assumes that the structural shocks, say ε_t , are obtained from the reduced form residuals by a linear transformation, say $\varepsilon_t = B^{-1}u_t$ or, equivalently, $B\varepsilon_t = u_t$. In a conventional SVAR analysis B is chosen such that the structural shocks, that is, the components of ε_t are instantaneously uncorrelated. In other words, ε_t has a diagonal covariance matrix, that is, $E(\varepsilon_t \varepsilon_t') = \Sigma_\varepsilon$ is a diagonal matrix. In practice, B is often chosen such that Σ_ε is an identity matrix so that $\varepsilon_t \sim (0, I_K)$ and, hence, $E(u_t u_t') = \Sigma_u = BB'$ if u_t is homoskedastic and, thus, has a time-invariant covariance matrix. Notice that by substituting $B\varepsilon_t$ for u_t in (1), the matrix B is easily recognized as the matrix of impact effects of the structural shocks.

The matrix B is not uniquely determined by the relation $\Sigma_u = BB'$ because Σ_u is symmetric and, hence, has at most $K(K+1)/2$ different elements while B has K^2 elements. Consequently, at least $K(K-1)/2$ further relations or restrictions are needed to uniquely specify the transformation matrix B and thereby characterize the shocks. As mentioned earlier, restrictions on the impact effects of the shocks, that is, on the elements of B directly, or indirect constraints on the long-run effects are often used in a conventional SVAR analysis to specify unique structural shocks. Alternatively, sign restrictions on the impulse responses may be imposed. They will not fix unique structural shocks in general but will specify a range of structural shocks. Such an approach may be sufficient for answering at least some questions of interest.

Another possibility for specifying unique shocks opens up if the variances of the reduced form and, hence, also the structural shocks change during the sample period. To see this suppose that $E(u_t u_t') = \Sigma_1$ if t is from the first part of the sample, say $t = 1, \dots, T_1$, and $E(u_t u_t') = \Sigma_2$ in the second part of the sample ($t > T_1$) and $\Sigma_1 \neq \Sigma_2$. Moreover, suppose that all

other VAR parameters in (1) remain time-invariant. Then there is a matrix decomposition result that ensures the existence of a matrix B and a diagonal matrix $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_K)$ such that

$$\Sigma_1 = BB' \quad \text{and} \quad \Sigma_2 = B\Lambda B'. \quad (2)$$

Using this B matrix to obtain structural shocks from the reduced form errors, that is, defining $\varepsilon_t = B^{-1}u_t$, we get

$$E(\varepsilon_t \varepsilon_t') = \begin{cases} I_K, & t = 1, \dots, T_1, \\ \Lambda, & t > T_1. \end{cases}$$

Since Λ is diagonal, the structural errors obtained in this way satisfy the basic requirement of being instantaneously uncorrelated. Moreover, Lanne et al. (2010) show that the matrix B is unique apart from changes in the signs and permutations of the columns if the diagonal elements of Λ are all distinct. The latter condition can be checked easily, of course. If it is satisfied, we have unique shocks by just imposing the basic requirement that the structural shocks are instantaneously uncorrelated. Of course, using the same transformation matrix B for the whole sample period implies that the impact effects of the shocks are time-invariant as well and only the variances change. All other requirements for uniqueness of the joint decomposition of the two covariance matrices do not affect the shocks substantively. The possible changes in sign just mean that we may consider negative instead of positive shocks and vice versa. Also the fact that the columns of B can be permuted reflects just that the ordering of the shocks can be chosen freely. It is worth looking at an example in some detail.

2.2 An illustrative example

Consider a bivariate system so that

$$u_t = \begin{bmatrix} u_{1t} \\ u_{2t} \end{bmatrix} = \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} \begin{bmatrix} \varepsilon_{1t} \\ \varepsilon_{2t} \end{bmatrix}.$$

For this case the relations (2) are

$$\begin{bmatrix} \sigma_{1,1}^2 & \sigma_{12,1} \\ \sigma_{12,1} & \sigma_{2,1}^2 \end{bmatrix} = \begin{bmatrix} b_{11}^2 + b_{12}^2 & b_{11}b_{21} + b_{12}b_{22} \\ b_{11}b_{21} + b_{12}b_{22} & b_{21}^2 + b_{22}^2 \end{bmatrix}$$

and

$$\begin{bmatrix} \sigma_{1,2}^2 & \sigma_{12,2} \\ \sigma_{12,2} & \sigma_{2,2}^2 \end{bmatrix} = \begin{bmatrix} \lambda_1 b_{11}^2 + \lambda_2 b_{12}^2 & \lambda_1 b_{11} b_{21} + \lambda_2 b_{12} b_{22} \\ \lambda_1 b_{11} b_{21} + \lambda_2 b_{12} b_{22} & \lambda_1 b_{21}^2 + \lambda_2 b_{22}^2 \end{bmatrix}.$$

Thus, we get six relations,

$$\begin{aligned}
\sigma_{1,1}^2 &= b_{11}^2 + b_{12}^2, \\
\sigma_{12,1} &= b_{11}b_{21} + b_{12}b_{22}, \\
\sigma_{2,1}^2 &= b_{21}^2 + b_{22}^2, \\
\sigma_{1,2}^2 &= \lambda_1 b_{11}^2 + \lambda_2 b_{12}^2, \\
\sigma_{12,2} &= \lambda_1 b_{11}b_{21} + \lambda_2 b_{12}b_{22}, \\
\sigma_{2,2}^2 &= \lambda_1 b_{21}^2 + \lambda_2 b_{22}^2,
\end{aligned}$$

from which we can solve for the six structural parameters $b_{11}, b_{12}, b_{21}, b_{22}, \lambda_1, \lambda_2$. Here the latter two parameters are the diagonal elements of Λ . The solution is unique (up to sign) if λ_1 and λ_2 are distinct and are ordered in some specific way. For instance, $\lambda_1 < \lambda_2$. Notice that the variances of the structural shocks are normalized to one in the first part of the sample. Hence, the λ_i indicate the change in variance from the first to the second variance regime. In other words, they can be interpreted as relative variances in the second part of the sample. The uniqueness requirement just means that the change in variance is not homogeneous across both variables. In fact, only the variance of one of the variables has to change. This already suffices for a unique solution.

It should be clear, however, that this is a purely technical way of getting unique shocks and does not necessarily result in economically meaningful shocks. Just like using a Choleski decomposition of the residual covariance matrix to obtain a unique set of orthogonal shocks does not lead automatically to economically meaningful shocks but requires further arguments to make them meaningful. For example, Rigobon (2003) discusses the case of a bivariate system with demand and supply shocks. Although in this case it is indeed difficult to disentangle the shocks in the absence of economic restrictions and, hence, the statistical identification approach comes in handy, it must be kept in mind that the method just normalizes the shocks in a mathematical way and the shocks identified in this way need not correspond to supply and demand shocks. Such an interpretation requires further economic insights and needs to be viewed with some caution. What we can say if the identification conditions are satisfied is that there are only two specific shocks that satisfy the basic requirements of structural shocks and induce time-invariant responses throughout the full sample. The latter assumption is, of course, implicit in a conventional analysis that ignores the changes in volatility or does not use them for identification purposes. On the other hand, one may question the assumption that the instantaneous effects of the shocks are not affected by the volatility changes. Indeed, also the VAR coefficients are assumed to be time-invariant and that assumption may be unrealistic in some cases. It is required, however, to make the approach

work. If also the VAR coefficients change over time, then the assumption of time-invariant impulse responses does not make sense any more and without that assumption enforcing time-invariant impact effects is meaningless.

2.3 The general model

In general there may be more than two volatility regimes. Suppose there are M such regimes and the corresponding covariance matrices are $\Sigma_1, \dots, \Sigma_M$. In that case a decomposition

$$\Sigma_1 = BB', \quad \Sigma_i = B\Lambda_m B', \quad m = 2, \dots, M, \quad (3)$$

with diagonal matrices $\Lambda_m = \text{diag}(\lambda_{m1}, \dots, \lambda_{mK})$ ($m = 2, \dots, M$) may not exist if arbitrary covariance matrices Σ_m ($m = 1, \dots, M$) are allowed for. In fact, the decomposition (3) imposes testable restrictions on the covariance matrices. Thereby we can check whether the data are compatible with the decomposition and, hence, we can use B to transform the reduced form residuals into structural errors with time-invariant impact effects. Lanne et al. (2010, Proposition 1) show that the structural shocks obtained in this way are unique (apart from ordering and sign) if for any two subscripts $k, l \in \{1, \dots, K\}$, $k \neq l$, there is a $j \in \{2, \dots, M\}$ such that $\lambda_{jk} \neq \lambda_{jl}$. Although this condition for exact (local) identification is more complicated than in the 2-state case, it shows that not much is needed for identification of structural shocks via heteroskedasticity. Moreover, the identification condition can be checked with statistical tests because, if there are M distinct volatility regimes, then the diagonal elements of the Λ_m matrices are identified and, hence, can be estimated consistently with a proper asymptotic distribution under common assumptions. We will discuss the issue further when we consider specific models for the volatility regimes in Section 3 and we will present examples in Sections 4 and 5.

If unique structural shocks are obtained via heteroskedasticity, then any further restrictions become over-identifying. In particular, restrictions on the impact or the long-run effects imposed in conventional SVAR analyses will be over-identifying and, hence, can be tested against the data. This implication of the present framework is very convenient if there are controversial identifying restrictions in a conventional analysis. In the present framework the data can speak up against such restrictions or confirm that they are compatible with the data. Of course, it can happen that even this framework does not allow for discrimination between competing economic theories. In other words, different models may both be in line with the data in the present framework. Moreover, rejecting a particular model in the present framework

could be the result of the underlying assumptions of the statistical model not being compatible with the data. For example, the assumption of regime invariant impact effects may be problematic and then a particular set of restrictions may be rejected although it is actually the model framework that is too narrow for the test.

So far we have considered what is known as a B -model in the SVAR literature (see Lütkepohl (2005, Chapter 9)). In practice A -models are sometimes preferred because they allow to impose restrictions on the instantaneous interactions between the observed variables more easily. An A -model has the form

$$A_0^* y_t = \nu^* + A_1^* y_{t-1} + \cdots + A_p^* y_{t-p} + \varepsilon_t. \quad (4)$$

In this model the A_i^* ($i = 0, 1, \dots, p$) are structural form VAR coefficient matrices. The matrix A_0^* represents the instantaneous relations between the observed elements of y_t . The B -model can be obtained from the A -model by pre-multiplying (4) with $B = A_0^{*-1}$. Thus, the two models are closely related. However, in the A -model the matrix A_0^* is typically assumed to have unit diagonal elements. Thereby each of the K variables can appear as dependent variable on the left-hand side of one of the equations. In that case, assuming unit variances of the structural errors is not reasonable. Thus, the error term ε_t in the A -model is usually assumed to have diagonal covariance matrix.

In that situation it may also be preferable not to normalize the variances of the structural shocks in the first regime if there are different volatility regimes. Hence, one may prefer to choose a decomposition

$$\Sigma_m = A_0^{*-1} \Lambda_m (A_0^{*-1})', \quad m = 1, \dots, M, \quad (5)$$

where A_0^* has a unit main diagonal. Although this modification of our basic model is in principle easy to deal with as far as estimation is concerned, the identification conditions are more difficult to state in this model. Hence, it may be better to leave the diagonal elements of A_0^* unrestricted and normalize the variances of the structural shocks in the first regime. Such a modification leaves zero restrictions on A_0^* unaffected. Thus, they may be imposed as in a conventional analysis. Alternatively, one may set up a so-called AB -model of the form

$$A_0^* y_t = \nu^* + A_1^* y_{t-1} + \cdots + A_p^* y_{t-p} + B \varepsilon_t \quad (6)$$

and specify B to be a diagonal matrix and require the structural shocks ε_t to have unit variances throughout the sample in a conventional analysis or in the first volatility regime in a heteroskedastic model.

There are two other generalizations of the present setup that have been used in practice. First, one may include additional unmodelled right-hand side variables. As long as time-invariance of all parameters apart from those in the residual covariance structure can be justified, the extension is straightforward. The same holds if there are cointegrated variables and a vector error correction model is specified. In that case one may estimate the cointegration relations in a first step without accounting for heteroskedasticity and fix the cointegration parameters at those estimates in the subsequent analysis.

3 Specific Models for Changes in Residual Volatility

In this section different models for the changes in residual volatility are presented. The first one assumes exogenous changes in given time periods. The second model supposes that the volatility changes are generated by a vector generalized autoregressive conditional heteroskedasticity (MGARCH) process and the third approach models the volatility changes with the help of a Markov process. All three models have been used in applications in the literature. Which one of them is most suitable in a particular situation is best judged by the analyst and subject matter considerations may play a role for the choice. For instance, for higher frequency financial market data conditional heteroskedasticity is often a plausible assumption whereas in low frequency macro data models assuming that some exogenous event has caused a change in the variance at a specific point in time may be plausible. Also the standard tools for model checking can be used to confirm whether a particular model captures the volatility features of the data. In practice it may be difficult, however, to discriminate between different models purely on the basis of statistical tools because they capture similar data features.

3.1 Exogenous changes in volatility

3.1.1 The general setup

In the original article on identifying structural shocks via heteroskedasticity in a SVAR analysis, Rigobon (2003) considers changes in the variances of the shocks at specific time points. Although he also discusses the implications of misspecifying the change points, the baseline model is one where the change in volatility occurs exogenously. Thus, it is assumed that the reduced form

residual covariance matrix in time period t is

$$E(u_t u_t') = \Sigma_t = \begin{cases} \Sigma_1 & \text{for } t \in \mathcal{T}_1, \\ \vdots & \\ \Sigma_M & \text{for } t \in \mathcal{T}_M, \end{cases} \quad (7)$$

where $\mathcal{T}_m = \{T_{m-1} + 1, \dots, T_m\}$ ($m = 1, \dots, M$) are M given volatility regimes that typically consist of consecutive time periods. Here $T_0 = 0$ and $T_M = T$ is assumed. The T_m , for $m > 0$, represent the points of volatility changes.

3.1.2 Inference

Assuming that the error term in (1) is Normal, that is, $u_t \sim \mathcal{N}(0, \Sigma_t)$, with covariance matrix structure as in (7), estimation of the model can be done by maximum likelihood (ML). The log-likelihood function is

$$\log l(\beta, \boldsymbol{\sigma}) = -\frac{KT}{2} \log 2\pi - \frac{1}{2} \sum_{t=1}^T \log |\Sigma_t| - \frac{1}{2} \sum_{t=1}^T u_t' \Sigma_t^{-1} u_t, \quad (8)$$

where $\boldsymbol{\sigma}$ contains all unknown covariance parameters and the time-invariant VAR parameters including deterministic terms are collected in β (see Lütkepohl (2005, Chapter 17) for further details). One may also use a feasible GLS procedure by estimating (1) with equationwise OLS first and then using the residuals obtained, say \hat{u}_t , for estimating the covariance matrices as

$$\hat{\Sigma}_m = \frac{1}{T_m - T_{m-1}} \sum_{t \in \mathcal{T}_m} \hat{u}_t \hat{u}_t'.$$

In the next step these estimates can then be used in a GLS estimator

$$\hat{\beta} = \left(\sum_{t=1}^T Z_{t-1} Z_{t-1}' \otimes \hat{\Sigma}_t^{-1} \right)^{-1} \left(\sum_{t=1}^T (Z_{t-1} \otimes \hat{\Sigma}_t^{-1}) y_t \right), \quad (9)$$

where $\beta = \text{vec}[\nu, A_1, \dots, A_p]$, $Z_{t-1} = (1, y'_{t-1}, \dots, y'_{t-p})'$ and $\hat{\Sigma}_t = \hat{\Sigma}_m$ for $t \in \mathcal{T}_m$. If the VAR process is stable, these estimators have standard asymptotic properties and can be used accordingly.

The structural parameters can be estimated by substituting BB' for Σ_1 , $BA_m B'$ for Σ_m ($m = 2, \dots, M$) and $y_t - (Z_t' \otimes I_K)\beta$ for u_t in (8) and using a full maximization of the log-likelihood. Alternatively, the β parameters

can be estimated in a first step and in the second step the concentrated log-likelihood is optimized only with respect to the structural parameters B and Λ_m ($m = 2, \dots, M$).

Again the estimators have standard asymptotic properties under common assumptions. These estimators can thus be used for checking the crucial assumptions underlying the present analysis. For example, since the estimators of the diagonal elements of the Λ_m are asymptotically Normal, they can be used to test for distinct diagonal elements and, hence, for identification of the shocks. Moreover, a likelihood ratio (LR) test for checking the null hypothesis that the covariance matrices can be decomposed as in (3) can be set up by comparing the log-likelihood maxima with and without the restriction. The LR statistic has an asymptotic χ^2 distribution with

$$\frac{1}{2}MK(K+1) - K^2 - (M-1)K$$

degrees of freedom under the null hypothesis (see Lanne et al. (2010)).

Notice, however, that the asymptotic properties of the estimators and test statistics may differ substantially from the small sample properties because we are in a dense parameter environment and macro time series are typically relatively short. Recall also that a certain volatility regime may be present only for a relatively short time span. Hence, the interpretation of the statistical results should be done cautiously. Remember also that the volatility regimes are assumed known which is not true in practice. Instead there will be uncertainty regarding the change points and perhaps some pretesting may have occurred to determine the change points (see, e.g., Ehrmann, Fratzscher and Rigobon (2011)). All these problems do not add to the reliability of the inference in this context. Still, the fact that the data have a chance to speak up regarding some of the crucial issues related to the identification of the shocks is an advantage over having no information on these issues.

Although this type of model requires nonlinear estimation techniques, doing a bootstrap for inference on impulse responses should not result in insurmountable problems.

3.1.3 Applications

Rigobon (2003) applies this approach to investigate the relationship between sovereign bonds in three Latin American countries, Argentina, Brazil and Mexico, using daily bond yields for the period January 1994 to December 2001. There is no natural instantaneous causal direction of the relation between the bond yields in these three countries. In fact, shocks hitting the system of the three series of interest are expected to have an instantaneous

effect on all of them. Also, the long-term impact of shocks does not provide easily acceptable identifying restrictions. Hence, using the changes in volatility makes sense.

Rigobon and Sack (2003) point out the instantaneous nature of the relation between short-term interest rates and stock prices and use identification via heteroskedasticity to disentangle the relation between these variables. In particular, they use differences in volatility shifts in the stock market and in monetary policy for identification. More precisely, they assume that the volatility of monetary shocks is constant whereas the volatility in equity prices varies. They distinguish four different volatility regimes that are determined by 30-day rolling variance estimates. Their analysis is based on daily US data from March 1985 to December 1999. They find that monetary policy reacts to unexpected movements in the stock market.

Rigobon and Sack (2004) address related problems as in the previous application. In other words, they also study the links between monetary policy and the stock market. Using identification through heteroskedasticity they find that earlier results from event studies may be biased.

Lanne and Lütkepohl (2008) use identification through heteroskedasticity to compare different models for the US money market that have been considered in the earlier literature. Their analysis will be discussed in more detail as an example in Section 4.

Ehrmann et al. (2011) use the approach to investigate the linkages between the US and euro area money, bond and equity markets by analyzing a system of short-term interest rates, long-term bond rates, stock returns from both regions and an exchange rate. They use two-day returns for a period from 1989-2008. The change in volatility is determined by a rolling window procedure that results in various volatility regimes for their dataset. Again identification through heteroskedasticity is used because there are not enough identifying exclusion restrictions for the impact effects of the shocks.

3.2 Vector GARCH residuals

Rather than assuming heteroskedastic residuals one may consider conditional heteroskedasticity. Normandin and Phaneuf (2004) propose to use multivariate GARCH models to capture changes in volatility. Their setup is presented in this section.

3.2.1 The general setup

The reduced form errors u_t are assumed to be generated by a multivariate GARCH (MGARCH) process of the form

$$\Sigma_{u,t|t-1} = E(u_t u_t' | u_{t-1}, \dots) = B \Sigma_{\varepsilon,t|t-1} B', \quad (10)$$

where $\Sigma_{\varepsilon,t|t-1} = \text{diag}(\sigma_{1,t|t-1}^2, \dots, \sigma_{K,t|t-1}^2)$ is a diagonal matrix with

$$\sigma_{k,t|t-1}^2 = \gamma_{k0} + \sum_{j=1}^q \gamma_{kj} \varepsilon_{t-j}^2 + \sum_{j=1}^s g_{kj} \sigma_{k,t-j|t-j-1}^2, \quad k = 1, \dots, K. \quad (11)$$

In other words, the structural shocks ε_t are assumed to be instantaneously uncorrelated and have a diagonal MGARCH(q, s) structure. This is basically the generalized orthogonal GARCH (GO-GARCH) model of van der Weide (2002). A closely related model was also proposed by Vrontos, Dellaportas and Politis (2003). These authors consider a model of the form (10) with a triangular matrix B . This choice simplifies the identification issue. Normandin and Phaneuf (2004) assume an MGARCH(1,1) model which is a standard model in practice for variables with GARCH effects. By sticking to the simple first order class they avoid a more difficult GARCH order selection step.

Sentana and Fiorentini (2001) provide general results that imply identification of the structural shocks (apart from permutation and sign changes) if all but one of the components are actually conditionally heteroskedastic and $\Gamma' \Gamma$ is invertible, where $(\sigma_{k,1|0}^2, \dots, \sigma_{k,T|T-1}^2)$ is the k th row of Γ' . In other words, in the GARCH(1,1) case we can have for at most one $k \in \{1, \dots, K\}$ that $\gamma_{k1} = g_{k1} = 0$. Invertibility of $\Gamma' \Gamma$ ensures that the conditional heteroskedasticity is not driven by a small number of processes like in a factor model where a single process may determine the development of a larger set of variables. Of course, this requirement corresponds to the condition that there must be sufficient heterogeneity in the volatility changes when there are only finitely many volatility regimes (see Section 2.3). Provided the identification conditions are satisfied, the setup in (10) ensures that, if we choose $\varepsilon_t = B^{-1} u_t$, the structural shocks will be unique and the impulse responses including the impact effects, are time-invariant.

3.2.2 Inference

For parameter estimation Normandin and Phaneuf (2004) propose a two-step procedure that fits a reduced form VAR(p) process in the first step and then estimates the GARCH and structural parameters by maximizing the

corresponding Gaussian log-likelihood. In other words, B and the GARCH parameters are estimated by maximizing

$$\log l^c(B, \gamma) = -\frac{1}{2} \sum_{t=1}^T [\log \det(B\Sigma_{\varepsilon,t|t-1}B') + \hat{u}'_t(B\Sigma_{\varepsilon,t|t-1}B')^{-1}\hat{u}_t], \quad (12)$$

where \hat{u}_t , $t = 1, \dots, T$, are the reduced form residuals from the first estimation stage and γ is the vector of all GARCH parameters. For example, for a GARCH(1,1),

$$\gamma = (\gamma_{01}, \dots, \gamma_{0K}, \gamma_{11}, \dots, \gamma_{1K}, g_{11}, \dots, g_{1K})'.$$

Obviously, the estimation problem is a highly nonlinear optimization task and Bouakez and Normandin (2010) mention that a full optimization of the function in (12) may be difficult for high dimensional systems. In that case further simplifications may be necessary. For example, one may split up the system and introduce further estimation stages, in particular, if suitable constraints on B are justifiable.

In principle, the identification conditions can be checked with statistical tools. In particular, it can be tested whether the GARCH parameters in (11) are significantly different from zero. Moreover, the nonsingularity of the matrix $\Gamma'\Gamma$ can be tested. Normandin and Phaneuf (2004) basically check only the nonsingularity of the empirical counterpart of this matrix. It would be desirable, of course, to develop a proper statistical test of the null hypothesis that the population matrix $\Gamma'\Gamma$ is singular.

Once structural shocks are identified, an impulse response analysis can be performed. Usual methods for generating confidence intervals for them with a bootstrap may be computationally demanding, given the difficulties in maximizing the likelihood function. Normandin and Phaneuf (2004) mention that they used a Bayesian method from Sims and Zha (1999) for that purpose without giving details.

3.2.3 Applications

Normandin and Phaneuf (2004) use their MGARCH framework to compare identification schemes for monetary policy shocks similar to the example discussed in Section 4 based on US monthly data for the 1982M11 - 1998M12 period. They reject some of the assumptions for identifying monetary policy shocks in a conventional framework. In Section 4 a slightly different model will be considered in a related analysis that illustrates how identification through changes in volatility can be used in this context.

Bouakez and Normandin (2010) study the effects of US monetary policy shocks on exchange rates between the US and the G7 countries. The monetary policy shocks are identified via GARCH in the residuals. They use monthly data for 1982M11 to 1998M12 for one set of countries and from 1982M11 to 2004M10 for the remaining countries. They find that monetary policy shocks have a substantial impact on exchange rate fluctuations.

3.3 Markov switching in residual volatility

The first proposal for modelling volatility changes assumes a finite number of states that are exogenously generated while the previous proposal of an MGARCH process generating the volatility changes amounts to assuming an endogenously generated continuum of possible volatility regimes. In the present subsection an intermediate proposal is presented where only a finite number of volatility regimes is assumed that are, however, generated endogenously via a Markov process. In the context of structural shock identification through changes in volatility this approach was first suggested by Lanne et al. (2010). A more complete methodology for using it in the SVAR framework is presented by Herwartz and Lütkepohl (2011).

3.3.1 The general setup

In this approach the distribution of the reduced form error term u_t is assumed to depend on a discrete Markov process s_t ($t = 0, \pm 1, \pm 2, \dots$) with M states. In other words, $s_t \in \{1, \dots, M\}$. The transition probabilities between the states are

$$p_{ij} = \Pr(s_t = j | s_{t-1} = i), \quad i, j = 1, \dots, M.$$

Lanne et al. (2010) assume that the conditional distribution of u_t , given the state s_t , is Normal,

$$u_t | s_t \sim \mathcal{N}(0, \Sigma_{s_t}). \tag{13}$$

This assumption allows them to do ML estimation but is otherwise not critical. Notice that the VAR coefficients do not depend on the Markov process s_t and are assumed time-invariant. Hence, our model is more restrictive than the Markov switching (MS) models considered by Rubio-Ramirez, Waggoner and Zha (2005), Sims and Zha (2006) and Sims, Waggoner and Zha (2008), for example.

Although there is only a finite number of volatility states, the model can mix these states by assigning probabilities strictly between zero and one to

them in any particular period t . Thus, it may capture gradual transitions from one state to another and one may argue that it can also generate a continuum of states. They are centred at a finite number of states, however, which may simplify the interpretation. The state covariance matrices $\Sigma_1, \dots, \Sigma_M$ are used for the identification of shocks, as before. In other words, the state covariance matrices are decomposed as in (3) and uniqueness of the structural shocks obtained as $\varepsilon_t = B^{-1}u_t$ is ensured if the diagonal elements of the Λ_m matrices have a fixed ordering and are sufficiently heterogeneous, as specified in Section 2.3.

3.3.2 Inference

The statistical analysis has to address the following issues in addition to the usual VAR inference: The number of volatility states has to be determined, the parameters of the overall model have to be estimated, the diagonal elements of the Λ_m matrices have to be compared and possibly identifying restrictions from a conventional analysis have to be tested within the present framework. Under the assumption of a conditional Normal distribution given the states, the likelihood function can be set up easily,

$$\log l(\beta, B, \boldsymbol{\lambda}, P|\mathbf{y}) = \sum_{t=1}^T \log \left(\sum_{m=1}^M \Pr(s_t = m|Y_{t-1}) f(y_t|s_t = m, Y_{t-1}) \right), \quad (14)$$

where $\boldsymbol{\lambda}$ is the vector of all diagonal elements of $\Lambda_2, \dots, \Lambda_M$, P is the matrix of transition probabilities, \mathbf{y} is the full sample, $Y_{t-1} = (y'_{t-1}, \dots, y'_{t-p})'$ and

$$f(y_t|s_t = m, Y_{t-1}) = (2\pi)^{-K/2} \det(\Sigma_m)^{-1/2} \exp \left\{ -\frac{1}{2} u'_t \Sigma_m^{-1} u_t \right\}.$$

Optimizing this log-likelihood function is a highly nonlinear optimization problem that poses a number of computational challenges. A suitable EM algorithm is provided by Herwartz and Lütkepohl (2011) who adopt an algorithm from Krolzig (1997) for the present model. Unfortunately, this algorithm seems to work reliably only for reasonably small systems with a small number of variables K , a small number of volatility states M and a moderate number of lags p . A detailed discussion of the problems associated with the optimization of the log-likelihood function is given by Herwartz and Lütkepohl (2011). In practice, if the conditional Normality of the residuals given State m is not satisfied, the procedure is a quasi ML procedure. The Normal distribution assumption is made for convenience. It is not essential for the asymptotic properties of the estimates to hold, provided some general conditions are satisfied.

Inference for the λ parameters can then be done with standard procedures. For example, equality of some of them can be tested with standard Wald or LR tests. The latter tests have the drawback in the present situation that they involve two difficult log-likelihood maximizations. Hence, Wald tests may be preferable from a computational point of view. Unfortunately, they are known to have poor small sample properties in similar situations and, hence, may also not be fully satisfactory. In any case, in principle these tests can be used to check the identification conditions. Moreover, if the shocks can be identified via the volatility structure, other restrictions can be tested as well with conventional tests.

Standard model selection criteria can be used for choosing between models with different numbers of volatility states. Psaradakis and Spagnolo (2003, 2006) find that they work reasonably well. This is important because testing models with different numbers of states against the data is not easy. The usual tests have nonstandard properties because of unidentified parameters under the null hypothesis (e.g., Hansen (1992), Garcia (1998)). Of course, model selection criteria can also be used to compare models with other sets of restrictions.

When a structural model with identified shocks is specified, it can be used for structural analysis. Impulse responses are a standard tool for that purpose. In a classical ML analysis confidence bands for impulse responses are typically generated with bootstrap methods. Clearly, such methods are problematic when a single estimation is difficult. To mitigate the computational problems, Herwartz and Lütkepohl (2011) propose a fixed design wild bootstrap procedure. Such a procedure makes sense in the present framework because it can take care of conditional heteroskedasticity, as shown by Goncalves and Kilian (2004). The bootstrap samples are constructed conditionally on the ML estimates using

$$y_t^* = \hat{\nu} + \hat{A}_1 y_{t-1} + \dots + \hat{A}_p y_{t-p} + u_t^*. \quad (15)$$

Here the right-hand side variables for a given t are always the original lagged variables and the errors $u_t^* = \eta_t \hat{u}_t$, where η_t is a binary random variable with values 1 and -1 that have equal probability. The procedure preserves heteroskedasticity and the pattern of contemporaneous dependence of the data. Herwartz and Lütkepohl (2011) propose to bootstrap parameter estimates θ^* of $\theta = \text{vec}[\nu, A_1, \dots, A_p]$ and B^* of B , conditionally on the initially estimated transition probabilities and λ to limit the computational burden. Computing the bootstrap impulse responses requires nonlinear optimization of the log-likelihood even with these simplifications and is computationally demanding.

3.3.3 Applications

Lanne et al. (2010) illustrate their approach by considering a small US macro model from Primiceri (2005) consisting of inflation, unemployment and an interest rate. They find that previously used identification restrictions are problematic and are not supported by the data when changes in volatility are taken into account. As a second example they consider a US model from Sims et al. (2008) for log GDP, inflation and a short-term interest rate. Based on quarterly data from 1959Q1-2005Q4 they find support for conventional identification restrictions, as used by Sims et al. (2008).

In a study that advances the methodology needed for a structural MS-VAR analysis, Herwartz and Lütkepohl (2011) consider an example from Peersman (2005) that aims at investigating the causes of the early millennium economic slow down. They use a four-variable system consisting of the price of oil, output, a consumer price index and a short-term interest rate. Peersman uses zero restrictions on the impact effects and the long-run effects of the shocks for identification. Using quarterly US data from 1980Q1 to 2002Q2 and models with two and three volatility states they find that some but not all of the identifying restrictions used by Peersman are in line with the data when changes in volatility are accounted for. They question that the model based on conventional identifying restrictions properly reflects the causes of the economic downturn at the beginning of the millennium.

Lütkepohl and Netšunajev (2012) illustrate how the MS-SVAR approach can be used for checking models identified through sign restrictions. They consider models for the crude oil market from Kilian (2009) and Kilian and Murphy (2011). That analysis is discussed as a more detailed example in Section 5 and is therefore not elaborated upon here.

Netšunajev (2012) applies the MS-SVAR methodology to reconsider alternative approaches for identifying technology shocks and sheds new light on the conflicting evidence on the impact on hours worked in the bivariate system. He does not find clear evidence of an increase in hours worked in response to a positive technology shock.

4 Comparison of Money Market Models

This section is based on Lanne and Lütkepohl (2008) (henceforth LL). Alternative identification schemes for monetary shocks are reviewed in the next subsection and then an empirical comparison using identification of shocks via heteroskedasticity is performed.

4.1 Alternative identification schemes

The models for the US money market considered in the following use the variables:

- TR_t – total reserves;
- NBR_t – nonborrowed reserves;
- BR_t – borrowed reserves; and
- FF_t – federal funds rate.

Because $TR = BR + NBR$, only three of these variables are really modelled. Hence, we define $y_t = (TR_t, NBR_t, FF_t)'$.

Bernanke and Mihov (1998) and Christiano, Eichenbaum and Evans (1999) consider the following general structural model for the market for bank reserves: The demand for total reserves is described as

$$TR_t = -\alpha FF_t + f_{TR}(\text{policy information}) + \sigma_d \varepsilon_t^d,$$

where ε_t^d is a demand disturbance with variance 1 and α and σ_d are parameters. The demand for borrowed reserves is given by

$$BR_t = \beta FF_t - \gamma NBR_t + f_{BR}(\text{policy information}) + \sigma_b \varepsilon_t^b,$$

where ε_t^b is the borrowing function disturbance with variance 1 and β , γ and σ_b are parameters. Finally, the Fed policy rule for setting nonborrowed reserves is

$$NBR_t = f_{NBR}(\text{policy information}) + \phi^d \sigma_d \varepsilon_t^d + \phi^b \sigma_b \varepsilon_t^b + \sigma_s \varepsilon_t^s.$$

The quantity ε_t^s represents the exogenous monetary policy shock, again with variance 1, and ϕ^d , ϕ^b and σ_s are parameters. It is assumed that the Fed knows the disturbances to the demand for total reserves and borrowed reserves when it makes policy decisions. The policy information set consists of instantaneous and lagged values of the variables

- gdp_t – log of Gross Domestic Product (GDP);
- p_t – log of the implicit GDP deflator; and
- $pcom_t$ – index of commodity prices.

The functions $f_*(\cdot)$ are all linear functions. Thus, we have a three-equation VAR model with unmodelled variables $x_t = (gdp_t, p_t, pcom_t)'$. For an impulse response analysis these variables have to be modelled as well. However, for the present purposes where we focus on comparing alternative identification schemes for monetary policy shocks, there is no need to model them. The reduced form model is

$$y_t = \nu + A_1 y_{t-1} + \dots + A_p y_{t-p} + C_0 x_t + C_1 x_{t-1} + \dots + C_p x_{t-p} + u_t \quad (16)$$

and the reduced form error term u_t is related to the structural shocks $\varepsilon_t = (\varepsilon_t^d, \varepsilon_t^b, \varepsilon_t^s)'$ as $u_t = B\varepsilon_t$, where

$$B = \begin{bmatrix} \sigma_d \frac{\beta - \phi^d \alpha \gamma + \phi^d \alpha}{\beta + \alpha} & -\alpha \sigma_s \frac{\gamma - 1}{\beta + \alpha} & -\alpha \sigma_b \frac{-1 + \phi^b \gamma - \phi^b}{\beta + \alpha} \\ \sigma_d \phi^d & \sigma_s & \sigma_b \phi^b \\ \sigma_d \frac{\phi^d \gamma - \phi^d + 1}{\beta + \alpha} & \sigma_s \frac{\gamma - 1}{\beta + \alpha} & \sigma_b \frac{-1 + \phi^b \gamma - \phi^b}{\beta + \alpha} \end{bmatrix}. \quad (17)$$

The nine elements of B are determined by the eight free parameters $\psi = (\alpha, \beta, \gamma, \phi^d, \phi^b, \sigma_d^2, \sigma_b^2, \sigma_s^2)'$. Since six restrictions are obtained from the relations $\Sigma_u = BB'$, we need two more restrictions to uniquely determine the eight parameters and, hence, the structural shocks in this general setup. Christiano et al. (1999) list a number of studies that report specific identifying assumptions. In particular, we consider the following four identification schemes from the given references.

- *FF policy shock*: $\phi^d = 1/(1 - \gamma)$ and $\phi^b = -\phi^d$ (Bernanke and Blinder (1992)). These restrictions imply that the Fed targets the federal funds rate, hence the abbreviation *FF* policy shock.
- *NBR policy shock*: $\phi^d = \phi^b = 0$ (Christiano and Eichenbaum (1992)). With these restrictions the policy shocks can be associated with the errors of the equation for nonborrowed reserves.
- *NBR/TR policy shock*: $\alpha = \phi^b = 0$ (Bernanke and Mihov (1998), Strongin (1995)). These restrictions assume that shocks to total reserves are demand shocks that are accommodated by the Fed.
- *BR policy shock*: $\phi^d = 1$, $\phi^b = \alpha/\beta$ and $\gamma = 0$ (Cosimano and Sheehan (1994)). These restrictions assume that the Fed targets borrowed reserves.

The first three sets of restrictions are just-identifying and, hence, they cannot be checked against the data in a conventional setting. Therefore LL use the heteroskedasticity in the residuals as additional identifying information. Thereby the foregoing restrictions on B become over-identifying and can be checked against the data.

4.2 Empirical analysis

LL use monthly US data from Bernanke and Mihov (1998) for the period 1965M1-1996M12. The monthly values for GDP and its deflator are constructed from quarterly data by using a set of monthly interpolator variables.

A crucial question in the present framework is, of course, whether there are changes in volatility during the sample period. Indeed changes are found by other authors as well. In fact, it is argued that the monetary policy during the Volcker era differs from that of other periods. Of course, one would like the data to confirm corresponding changes in the volatility of the shocks. For that purpose, models with heteroskedastic residuals must be estimated. LL do this with a multi-step iterative procedure. The model (16) is first estimated with equation-wise OLS. The different covariance matrices are estimated from these residuals and then a feasible GLS estimator is used.

Based on these GLS estimators heteroskedasticity is checked with LR-type tests. It is confirmed that there is strong evidence for a change in volatility in 1984M2 in line with other literature and there is also some evidence that another change has occurred in 1979M10. Because only one volatility change is needed for identification of the shocks, LL go on under the assumption of a break in 1984M2 and leave open the possibility of another volatility change in 1979M10, that is, they leave the covariance matrix for the period prior to 1979M10 unrestricted and also do not use it for identification. Estimation under this assumption gives the diagonal elements of Λ shown in Table 1.² Apparently they are quite different and one-standard error intervals around the estimates do not overlap. Although this is not a particularly strict criterion, it shows that there is some evidence for distinct λ_m values. Given that these are reduced form parameters, we could, in fact, perform formal tests of equality. In Section 5 we discuss how to do that and, for the moment, pretend that the structural shocks in the present model are identified via heteroskedasticity.

Thus, all the identification schemes presented in the previous subsection can be tested against the data. The results of the corresponding LR tests are presented in Table 2, where the degrees of freedom used for evaluating

²Note that we report the corrected values from LL.

Table 1: Estimates of Structural Parameters with Standard Errors in Parentheses from LL

Parameter	Estimate
λ_1	1.1463 (0.2583)
λ_2	0.0144 (0.0032)
λ_3	0.5520 (0.1246)

Table 2: LR Tests of Over-Identifying Restrictions from LL

Identification scheme	H_0	LR statistic	p -value
<i>FF</i>	$\phi^d = 1/(1 - \gamma)$ and $\phi^b = -\phi^d$	8.021	0.018
<i>NBR</i>	$\phi^d = \phi^b = 0$	55.880	7.342e-13
<i>NBR/TR</i>	$\alpha = \phi^b = 0$	66.119	2.890e-14
<i>BR</i>	$\phi^d = 1, \phi^b = \alpha/\beta$, and $\gamma = 0$	55.947	7.101e-13

the p -values are equal to the number of restrictions specified under H_0 . They are the asymptotically correct ones if the shocks are uniquely identified via heteroskedasticity, that is, under the assumption of distinct λ_m . If not all but only some of the λ_m are distinct, a smaller number of degrees of freedom may apply. However, even with the larger number assumed here, the p -values of all the tests are very small and at a 5% significance level all null hypotheses and, hence, all identification schemes are rejected. It is worth recalling, however, that LR tests in the presence of many parameters may be biased and in particular may reject too often. Thus, in the present case, one may want to work with very small significance levels. Even if a 1% level were chosen, only the *FF* scheme would survive. Thus, in the present framework only the Bernanke and Blinder (1992) model, based on the assumption that the Fed targets the federal funds rate, is not clearly rejected by the data. Our results agree to a large extent with those of Normandin and Phaneuf (2004) who also find that the data are not supportive of the standard identification procedures for monetary policy shocks used in much of the previous literature. Note, however, that these authors use a slightly different sample period, 1982M11 - 1998M12, and a different model for changes in volatility.

What the example shows is that it may pay to use the additional information coming from the heteroskedasticity of the data to discriminate between competing models that cannot be tested against the data in a conventional setting. Of course, rejecting all but one model is a very special case. Our framework could have ended up not discriminating between the alternative identification schemes. A further example is presented in the next section.

5 A Model for the Market of Crude Oil

In this subsection a model from Kilian (2009) for the crude oil market is reconsidered. It consists of three variables,

- $\Delta prod_t$ – change in oil production;
- q_t – real activity; and
- p_t – real price of oil.

Thus, $y_t = (\Delta prod_t, q_t, p_t)'$. Monthly data for the period 1973M2 - 2006M12 from Kilian (2009) are used. The change in oil production is measured as log differences of world crude oil production in millions of barrels pumped per day (averaged by month), as obtained from the US Department of Energy. For q_t Kilian constructs a measure based on dry cargo single voyage ocean freight rates. It is meant to capture worldwide real economic activity. Finally, the real price of oil, p_t , is based on a series from the US Department of Energy for refiner acquisition cost of imported crude oil. Thus, all three variables refer to the world market.

Kilian (2009) considers three shocks:

- $\varepsilon_t^{\text{oil-s}}$ – oil supply shocks;
- $\varepsilon_t^{\text{aggr-d}}$ – aggregate demand shocks; and
- $\varepsilon_t^{\text{oil-d}}$ – oil market specific demand shocks that are innovations to the real price of oil that cannot be explained by oil supply shocks or aggregate demand shocks.

He identifies the shocks by imposing exclusion restrictions on the impact effects. More precisely, he assumes that oil-market specific demand shocks ($\varepsilon_t^{\text{oil-d}}$) do not have an instantaneous impact on oil production and real activity and aggregate demand shocks ($\varepsilon_t^{\text{aggr-d}}$) do not affect oil production instantaneously. Thereby B is a lower-triangular matrix and the reduced form and structural residuals are related as follows:

$$\begin{bmatrix} u_t^{\Delta prod} \\ u_t^q \\ u_t^p \end{bmatrix} = \begin{bmatrix} b_{11} & 0 & 0 \\ b_{21} & b_{22} & 0 \\ b_{31} & b_{32} & b_{33} \end{bmatrix} \begin{bmatrix} \varepsilon_t^{\text{oil-s}} \\ \varepsilon_t^{\text{aggr-d}} \\ \varepsilon_t^{\text{oil-d}} \end{bmatrix}. \quad (18)$$

Of course, one may question these identification restrictions and, in fact, doubts are expressed by Kilian and Murphy (2011) and Baumeister and Peersman (2010) who consider sign restrictions instead. We will return to

their studies later but will report results from Lütkepohl and Netšunajev (2012) (henceforth LN) first, who use heteroskedasticity in the residuals to check the identification assumptions. Thus, the following empirical study draws heavily on their results.

5.1 Empirical analysis

NL consider a VAR(3) model with intercept because in a standard order selection exercise not accounting for changes in volatility, AIC favours that model. In contrast, Kilian (2009) uses a VAR(24). Such a large model would be difficult to estimate by ML with an MS structure for the residual variance, however. In any case, a finite order VAR should only be viewed as an approximation to the DGP. Below, a VAR(3), which preserves features of special interest in the present context, is used for expositional purposes. In particular, the reduced form residuals display clear changes in volatility. They are plotted in Figure 1. The volatility of the residuals of the oil production equation are more volatile in the earlier part of the sample while the volatility in the price of oil residuals is larger in the later part of the sample. One may also spot some change in volatility in the equation for real activity although that change is not quite so obvious. Of course, visual inspection can usually not be a substitute for a formal statistical analysis that will be reported shortly. Before turning to that, it is worth pointing out that the kind of volatility changes observed by visual inspection are precisely what is needed for our analysis. Recall that the change in volatility has to be sufficiently heterogeneous in order to get identification through heteroskedasticity.

Table 3: Comparison of MS-VAR(3) Models for $y_t = (\Delta prod_t, q_t, p_t)'$ (results from LN, Table 2)

Model	AIC	SC
VAR(3) (without MS)	8615	8759
MS(2)-VAR(3)	8302	8478
MS(3)-VAR(3)	8206	8422
MS(4)-VAR(3)	8159	8423

Note: L_T – likelihood function, $AIC = -2\log L_T + 2 \times \text{no of free parameters}$, $SC = -2\log L_T + \log T \times \text{no of free parameters}$.

A formal statistical analysis of volatility changes is reported in Table 3 where AIC and SC values are presented for models with different numbers of volatility states. Notice that a VAR(p) model with m volatility states is




Figure 1: Residuals of VAR(3) model from NL.

denoted as $MS(m)$ -VAR(p). From the results in Table 3 it is clear that both AIC and SC favour models that allow for volatility changes over a standard VAR model without MS. They do not agree on the number of states, however. AIC favours the largest model with four different volatility states while SC prefers a model with only three states. Since an $MS(4)$ model is the largest one in our choice set, one may wonder whether AIC may even prefer more than four states. Admittedly, a maximum of four states was chosen for computational reasons. Even an $MS(4)$ -VAR(3) model is already difficult to estimate from the given dataset.


Since the model selection criteria are not very reliable in choosing the number of states, we plot the estimated smoothed probabilities associated with the states in Figure 2. The first two states of the $MS(3)$ model are quite similar to those of the $MS(2)$ model and the third state picks out some special events in the crude oil market (see, e.g., Barsky and Kilian (2004) for a discussion of such events). Thus, a three state model makes good sense from a subject matter point of view. In contrast, the $MS(4)$ model has one state (the third state) that comes up only in a very few periods. While it may well be plausible to single out one state that covers only very special periods, such a situation is problematic from a computational and estimation point of view. Clearly, such a model may require that some parameters are basically estimated from very few observations and, hence, the estimates may be unreliable. If the estimation algorithm actually associates a separate state

with a single observation, the associated covariance matrix estimate becomes singular and the algorithm runs into a problem because it cannot invert the matrix in evaluating the log-likelihood function. Given that a three state model produces quite plausible states, we favour overall this model, like NL, and continue the analysis with an MS(3)-VAR(3) model.


Having decided on an MS(3)-VAR(3) model, the next important question in the context of the present analysis is whether the three covariance matrices can be decomposed as in (3). For answering that question we may use an LR test with $\frac{1}{2}MK(K+1) - K^2 - (M-1)K = 3$ degrees of freedom (see Section 3.1). Such a test results in a p -value of 0.63 and, hence, there is little basis for rejecting the more restricted model. In other words, the covariance matrix decomposition in (3) is in line with the data and we can work with a model with state-invariant matrix of initial effects of the shocks.

The next question then concerns the uniqueness of the B transformation matrix and, hence, the uniqueness of the shocks. That uniqueness depends on the diagonal elements of the Λ_m matrices. The estimated quantities for the MS(3)-VAR(3) model are given in Table 4 together with estimated standard errors. Taking into account the estimated λ_{mi} values and the estimation uncertainty reflected in the standard errors, there is clearly hope for having enough heterogeneity for identification. To check this formally, the corresponding tests are presented in Table 5. Recall from Section 3.3 that for each pair $i, j \in \{1, \dots, K\}$, $i \neq j$, there has to be an $m \in \{1, \dots, M\}$ such that $\lambda_{mi} \neq \lambda_{mj}$. Thus, identification is ensured if all three null hypotheses listed in Table 5 are rejected. Wald and LR tests are given in the table. The Wald tests are easier to compute because they only require estimates of the unrestricted model. The associated p -values are all below 10% and indicate that the identification conditions are satisfied. As Wald tests may not be very reliable in the present situation, we also report LR tests. They require two very nonlinear optimizations of the log-likelihood. Given that this function has many local optima and, hence, there is some danger of ending up in a local optimum and comparing inappropriate likelihood values in the LR tests, these tests are also problematic. Taking their values seriously, they do provide strong evidence for the identification conditions to be satisfied. Notice that the tests are performed for models where the λ_{2i} are ordered from largest to smallest. The ordering has to be fixed to ensure uniqueness of the decomposition of the covariance matrices.


Given the evidence in favour of the identification conditions, we may proceed with testing Kilian's identification conditions presented in (18). The corresponding LR test has a p -value of 0.84. Thus, we cannot reject the restrictions at conventional significance levels. In other words, in our framework the data do not object to Kilian's identification conditions. It is still



(a) Smoothed state probabilities for MS(2)-VAR(3)



(b) Smoothed state probabilities for MS(3)-VAR(3)



(c) Smoothed state probabilities for MS(4)-VAR(3)

Figure 2: Smoothed state probabilities of MS(m)-VAR(3) models for $m = 2, 3, 4$ from Fig. 2 of NL (state invariant B matrix imposed for MS(3) and MS(4)).

Table 4: Estimates of Structural Parameters of MS(3)-VAR(3) Models for $y_t = (\Delta prod_t, q_t, p_t)'$ with State-invariant B , adapted from LN

parameter	estimate	std.dev.
λ_{21}	5.384	0.933
λ_{22}	3.210	0.485
λ_{23}	0.056	0.009
λ_{31}	25.235	6.938
λ_{32}	3.387	1.242
λ_{33}	4.441	1.018

Note: Standard errors are obtained from the inverse of the outer product of numerical first order derivatives.

Table 5: Tests for Equality of λ_{mj} s for MS(3)-VAR(3) Model with State-invariant B (from LN)

H_0	Wald statistic	p -value	LR statistic	p -value
$\lambda_{21} = \lambda_{22}, \lambda_{31} = \lambda_{32}$	7.99	0.02	20.39	3.7×10^{-5}
$\lambda_{21} = \lambda_{23}, \lambda_{31} = \lambda_{33}$	7.87	0.02	21.04	2.7×10^{-5}
$\lambda_{22} = \lambda_{23}, \lambda_{32} = \lambda_{33}$	5.16	0.07	27.15	1.3×10^{-6}




Figure 3: Impulse responses with 68% confidence bounds of the MS(3)-VAR(3) model with state-invariant B (Fig. 4 from LN).

worthwhile considering the case where these restrictions are not imposed. The corresponding impulse responses will be discussed next.

5.2 Impulse response analysis

Since we have identified shocks by assuming the MS structure for the volatility, we can go on and use these shocks for an impulse response analysis. We present the resulting responses to structural impulses obtained as $\varepsilon_t = B^{-1}u_t$ and 68% bootstrap confidence intervals in Figure 3. These impulse responses are arbitrary in several respects. First, the scaling of the responses is determined by normalizing the shocks such that they have identity covariance matrix in the first state. Since the numbering of the states is arbitrary, this normalization is also arbitrary. The scaling does not affect the shape of the impulse responses, of course. Another element of arbitrariness results from the signs of the shocks. In Figure 3 they are chosen such that an interpretation becomes possible, as we will see shortly. Finally, the ordering of the shocks is arbitrary. As mentioned in Section 3.3, B is unique only up to column permutations. Because the columns of B correspond to the impact effects of the shocks, permuting the columns comes down to changing the order of the shocks.

Because the shocks are uniquely determined in our framework, they must correspond to the three shocks of interest in the system under investigation, namely an oil supply shock, an aggregate demand shock and an oil-market

specific demand shock, provided that these shocks are really the ones that drive the system. In our framework it is unclear *a priori* which of the three shocks deserves which label. However, by looking at the impulse responses we are able to assign labels and we have done so in Figure 3. The first shock is called an oil supply shock because a surprise reduction in oil production goes together with a reduction in output and a possible increase in the price of oil. Notice that the reaction of the oil price is not clear in our system. However, taking into account the uncertainty in the response of the oil price, a price increase due to a surprise reduction in oil production is quite possible.

The second shock is labelled aggregate demand shock because it increases oil production, output and the price of oil simultaneously, at least initially. Finally, the third shock is the only one that leads to a clear increase in the price of oil on impact. Therefore the shock is the only one that qualifies as an oil-market specific demand shock. The effects in Figure 3 are also in line with the sign restrictions imposed by Kilian and Murphy (2011) for a the same system of variables. They characterize the shocks by the signs of their impact effects and in a sense we are doing the same thing by labelling the shocks obtained in our framework accordingly. Notice, however, that in the present analysis the shocks are point identified and only the labelling is done by considering the effect or impulse responses whereas using sign restrictions results in set-identified shocks only that may be more difficult to interpret because their effects may be more diffuse.

Obviously, also some of the effects on our models are not all that clear, given the substantial sampling uncertainty reflected in the confidence intervals associated with the impulse responses. In general there is no guarantee that the shocks of interest are actually in line with the model under study. Even if in an underlying theoretical model the desired shocks are the most important driving forces, they may not be found in an associated empirical model. The empirical model usually differs from the theoretical one in a number of dimensions. For example, variables that are not of interest in the theoretical model may be quite important in the empirical model, that is, there may be an omitted variables problem. Moreover, the empirical variables may not be exactly the same as the corresponding quantities in the theoretical model. In other word, there may be measurement errors. For example, the theoretical model may not account for trends and seasonality that may be present in practice. Hence, there are a number of reasons why the empirical shocks may not be those of interest from a theoretical point of view. LN use this point as an argument for proposing their framework for checking the validity of sign restrictions.

Regarding the confidence intervals in Figure 3, it is remarkable that the original impulse responses in the upper left-hand and lower right-hand panels




Figure 4: Impulse responses with 68% confidence bounds of the MS(3)-VAR(3) model with state-invariant, lower-triangular B (Fig. 5 from LN).

fall largely outside the intervals. Such an outcome is not impossible, of course, but it is still a bit unusual and may indicate that there is room for improvement in constructing suitable confidence intervals for the impulse responses in our model. The intervals are standard percentile intervals and it may have been worth trying other intervals for impulse responses (e.g., Benkwitz, Lütkepohl and Wolters (2001)).

Given that Kilian's restrictions for the impact effects are not rejected in our analysis, it is also plausible to consider the corresponding impulse responses. They are presented in Figure 4. Because the triangularity of B is very much in line with the data, it is not surprising that the impulse responses in Figure 4 look quite similar to those in Figure 3, except, of course, that now some of the impact effects are restricted to zero. Thus, in contrast to the example in Section 4, in the present example using heteroskedasticity for the identification of the structural shocks supports the more stringent restrictions that were used in a previous analysis. The important point to note is that accounting for changes in volatility allows us to check the restrictions that had to be assumed in a conventional analysis.

6 Conclusions

In the present paper it is argued that taking into account changes in the volatility of the residuals can provide useful additional identifying infor-

mation for structural shocks. The basic assumptions are that changes in volatility occur during the sample period, that they just affect the volatility of the shocks but not the effects the shocks have on the variables and that the changes in volatility are sufficiently heterogeneous across variables. It is important to note that the requirements regarding the existence and heterogeneity of the volatility changes can be investigated by statistical procedures. Hence, the data are informative about some of the crucial conditions for identification. If full identification of the structural shocks is obtained in this framework, conventional identifying restrictions become over-identifying and, hence, can be checked with statistical tests. Thus, competing conventional identifying assumptions can be tested against the data, as illustrated in Section 4. Moreover, the framework allows us to check that the empirical model is in line with the underlying theoretical model more generally and we have argued that differences between the empirical and theoretical models may stem from omitted variables or measurement error problems, for example.

It must be emphasized, however, that the shocks obtained in the present framework are unique due to the specific statistical model setup. Hence, they are statistically identified or normalized. There is no guarantee that these shocks necessarily correspond to economically interesting shocks. Therefore using this framework without additional economic reasoning why the shocks allow us to understand the economic relations in a system of variables is problematic. It is a useful tool for checking that an empirical model is compatible with the data, however.

References

- Amisano, G. and Giannini, C. (1997). *Topics in Structural VAR Econometrics*, 2nd edn, Springer, Berlin.
- Barsky, R. B. and Kilian, L. (2004). Oil and the macroeconomy since the 1970s, *Journal of Economic Perspectives* **18**: 115–134.
- Baumeister, C. and Peersman, G. (2010). Sources of the volatility puzzle in the crude oil market, *Working Paper 2010/634*, Universiteit Gent.
- Benkwitz, A., Lütkepohl, H. and Wolters, J. (2001). Comparison of bootstrap confidence intervals for impulse responses of German monetary systems, *Macroeconomic Dynamics* **5**: 81–100.
- Bernanke, B. S. and Blinder, A. (1992). The federal funds rate and the channels of monetary transmission, *American Economic Review* **82**: 901–921.

- Bernanke, B. S. and Mihov, I. (1998). Measuring monetary policy, *Quarterly Journal of Economics* **113**: 869–902.
- Blanchard, O. and Quah, D. (1989). The dynamic effects of aggregate demand and supply disturbances, *American Economic Review* **79**: 655–673.
- Bouakez, H. and Normandin, M. (2010). Fluctuations in the foreign exchange market: How important are monetary policy shocks?, *Journal of International Economics* **81**: 139–153.
- Canova, F. and De Nicoló, G. (2002). Monetary disturbances matter for business fluctuations in the G-7, *Journal of Monetary Economics* **49**: 1131–1159.
- Christiano, L. J. and Eichenbaum, M. (1992). Identification and the liquidity effect of a monetary policy shock, in A. Cukierman, Z. Hercowitz and L. Leiderman (eds), *Political Economy, Growth, and Business Cycles*, MIT Press, Cambridge, MA, pp. 335–370.
- Christiano, L. J., Eichenbaum, M. and Evans, C. (1999). Monetary policy shocks: What have we learned and to what end?, in J. B. Taylor and M. Woodford (eds), *Handbook of Macroeconomics*, Vol. 1A, Elsevier, Amsterdam, pp. 65–148.
- Cosimano, T. and Sheehan, R. (1994). The federal reserve operating procedure, 1984-1990: An empirical analysis, *Journal of Macroeconomics* **16**: 573–588.
- Ehrmann, M., Fratzscher, M. and Rigobon, R. (2011). Stocks, bonds, money markets and exchange rates: Measuring international financial transmission, *Journal of Applied Econometrics* **26**: 948 – 974.
- Faust, J. (1998). The robustness of identified VAR conclusions about money, *Carnegie-Rochester Conference Series in Public Policy* **49**: 207–244.
- Garcia, R. (1998). Asymptotic null distribution of the likelihood ratio test in Markov switching models, *International Economic Review* **39**: 763–788.
- Goncalves, S. and Kilian, L. (2004). Bootstrapping autoregressions with conditional heteroskedasticity of unknown form, *Journal of Econometrics* **123**: 89–120.
- Hansen, B. E. (1992). The likelihood ratio test under nonstandard conditions: Testing the Markov switching model of GNP, *Journal of Applied Econometrics* **7**((erratum: 11, 195-198)): S61–82.
- Herwartz, H. and Lütkepohl, H. (2011). Structural vector autoregressions with markov switching: Combining conventional with statistical identification of shocks, *Working Paper series*, EUI, Florence.

- Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market, *American Economic Review* **99**: 1053–1069.
- Kilian, L. and Murphy, D. (2011). Why agnostic sign restrictions are not enough: Understanding the dynamics of oil market VAR models, *Journal of the European Economic Association* .
- King, R. G., Plosser, C. I., Stock, J. H. and Watson, M. W. (1991). Stochastic trends and economic fluctuations, *American Economic Review* **81**: 819–840.
- Krolzig, H.-M. (1997). *Markov-Switching Vector Autoregressions: Modelling, Statistical Inference, and Application to Business Cycle Analysis*, Springer-Verlag, Berlin.
- Lanne, M. and Lütkepohl, H. (2008). Identifying monetary policy shocks via changes in volatility, *Journal of Money, Credit and Banking* **40**: 1131–1149.
- Lanne, M., Lütkepohl, H. and Maciejowska, K. (2010). Structural vector autoregressions with Markov switching, *Journal of Economic Dynamics and Control* **34**: 121–131.
- Lütkepohl, H. (2005). *New Introduction to Multiple Time Series Analysis*, Springer-Verlag, Berlin.
- Lütkepohl, H. and Netšunajev, A. (2012). Disentangling demand and supply shocks in the crude oil market: How to check sign restrictions in structural vars, *Discussion Paper 1195*, Deutsches Institut für Wirtschaftsforschung.
- Netšunajev, A. (2012). Reaction to technology shocks in Markov-switching structural VARs: Identification via heteroskedasticity, *mimeo*, European University Institute.
- Normandin, M. and Phaneuf, L. (2004). Monetary policy shocks: Testing identification conditions under time-varying conditional volatility, *Journal of Monetary Economics* **51**: 1217–1243.
- Peersman, G. (2005). What caused the early millennium slowdown? Evidence based on vector autoregressions, *Journal of Applied Econometrics* **20**: 185–207.
- Primiceri, G. E. (2005). Time varying structural vector autoregressions and monetary policy, *Review of Economic Studies* **72**: 821–852.
- Psaradakis, Z. and Spagnolo, N. (2003). On the determination of the number of regimes in Markov-switching autoregressive models, *Journal of Time Series Analysis* **24**: 237–252.

- Psaradakis, Z. and Spagnolo, N. (2006). Joint determination of the state dimension and autoregressive order for models with Markov regime switching, *Journal of Time Series Analysis* **27**: 753–766.
- Rigobon, R. (2003). Identification through heteroskedasticity, *Review of Economics and Statistics* **85**: 777–792.
- Rigobon, R. and Sack, B. (2003). Measuring the reaction of monetary policy to the stock market, *Quarterly Journal of Economics* **118**: 639–669.
- Rigobon, R. and Sack, B. (2004). The impact of monetary policy on asset prices, *Journal of Monetary Economics* **51**: 1553–1575.
- Rubio-Ramirez, J. F., Waggoner, D. and Zha, T. (2005). Markov-switching structural vector autoregressions: Theory and applications, *Discussion Paper*, Federal Reserve Bank of Atlanta.
- Sentana, E. and Fiorentini, G. (2001). Identification, estimation and testing of conditionally heteroskedastic factor models, *Journal of Econometrics* **102**: 143–164.
- Sims, C. A. (1980). Macroeconomics and reality, *Econometrica* **48**: 1–48.
- Sims, C. A., Waggoner, D. F. and Zha, T. (2008). Methods for inference in large multiple-equation Markov-switching models, *Journal of Econometrics* **146**: 255–274.
- Sims, C. A. and Zha, T. (1999). Error bands for impulse responses, *Econometrica* **67**: 1113–1155.
- Sims, C. A. and Zha, T. (2006). Were there regime switches in U.S. monetary policy?, *American Economic Review* **96**: 54–81.
- Strongin, S. (1995). The identification of monetary policy disturbances: Explaining the liquidity puzzle, *Journal of Monetary Economics* **35**: 463–498.
- Uhlig, H. (2005). What are the effects of monetary policy on output? Results from an agnostic identification procedure, *Journal of Monetary Economics* **52**: 381–419.
- van der Weide, R. (2002). GO-GARCH: A multivariate generalized orthogonal GARCH model, *Journal of Applied Econometrics* **17**: 549–564.
- Vrontos, I. D., Dellaportas, P. and Politis, D. N. (2003). A full-factor multivariate GARCH model, *Econometrics Journal* **6**: 312–334.