



Have the effects of shocks to oil price expectations changed? Evidence from heteroskedastic proxy vector autoregressions

Martin Bruns^{a,*}, Helmut Lütkepohl^b

^a University of East Anglia, School of Economics, Norwich Research Park, NR4 7TJ, Norwich, United Kingdom

^b DIW Berlin and Freie Universität Berlin, Mohrenstr. 58, 10117 Berlin, Germany

ARTICLE INFO

JEL classification:
C32

Keywords:
Structural vector autoregression
Heteroskedastic VAR
Proxy VAR
Crude oil market

ABSTRACT

Studies of the crude oil market based on structural vector autoregressive (VAR) models typically assume a time-invariant model and transmission of shocks and possibly allow for heteroskedasticity by using robust inference procedures. We assume a heteroskedastic reduced-form VAR model with time-invariant slope coefficients and explicitly consider the possibility of time-varying shock transmission due to heteroskedasticity. We study a model for the global crude oil market that includes key world and U.S. macroeconomic variables and find evidence for changes in the transmission of shocks to oil price expectations during the last decades which can be attributed to heteroskedasticity.

1. Introduction

A typical assumption in structural vector autoregressive (VAR) analysis is the time-invariance of the reduced-form VAR slope coefficients and the implied structural impulse responses across the sample period considered. This assumption is also used in a number of structural VAR studies on the global market for crude oil and its impact on macroeconomic variables (e.g., Kilian (2009), Kilian and Murphy (2014), Lütkepohl and Netšunajev (2014), Baumeister and Hamilton (2019), Känzig (2021), and Degasperis (2021)). At the same time, many authors allow for heteroskedasticity in the model residuals and account for it by using heteroskedasticity-robust inference (e.g., Känzig (2021) and Degasperis (2021)) or modelling time-varying volatility explicitly (e.g., Lütkepohl and Netšunajev (2014)). In this study we point out that heteroskedasticity in the VAR residuals may be a sign of changes in the structural shock variances, the transmission of these structural shocks to the underlying economy or both. In other words, despite the time-invariance of the reduced-form VAR process, the transmission of structural shocks can change if there is heteroskedasticity.

We follow the recent literature that uses event studies to identify the causal effects of revisions to oil price expectations on the macroeconomy. Känzig (2021) has proposed a proxy based on oil supply surprises measured by changes in oil price futures around OPEC announcements, and Degasperis (2021) has shown how this surprise series can be used to identify shocks to both oil supply and oil demand expectations. We use these shocks, consider a sample period from 1984 to 2019, and assume that the data are generated by a time-invariant VAR process

with heteroskedastic residuals. We explore whether the changes in the residual volatility come with a change in the transmission of the structural shocks using statistical tests as proposed by Lütkepohl and Schlaak (2022) and Bruns and Lütkepohl (2022). We find evidence that the impulse responses are different before and after the time of the 1990/91 gulf war.

The model and methodology used in the present study are briefly laid out in the following section. The empirical analysis is presented in Section 3 and more detailed results and information on specific issues are provided in a Supplement.

2. Model setup and methodology

We consider a K -dimensional heteroskedastic VAR model of order p ,

$$y_t = v + A_1 y_{t-1} + \dots + A_p y_{t-p} + u_t,$$

where u_t is a zero-mean white noise process with covariances

$$\mathbb{E}(u_t u_t') = \Sigma_t = \Sigma_u(m) \quad \text{for } t \in \mathcal{T}_m, \quad m = 1, \dots, M. \quad (1)$$

The M volatility regimes $\mathcal{T}_m = \{T_{m-1} + 1, \dots, T_m\}$ ($m = 1, \dots, M$) are assumed to be associated with consecutive time periods, with volatility changes occurring at time periods T_m , for $m = 1, \dots, M - 1$, with $T_0 = 0$ and T_M is the overall sample size, i.e., $T_M = T$. This reduced-form model setup is also used by Lütkepohl and Schlaak (2022).

The vector of structural shocks, $w_t = (w_{1t}, \dots, w_{Kt})'$, is related to the reduced-form errors, u_t , by a linear transformation. Formally, the

* Corresponding author.

E-mail addresses: martin.bruns@uea.ac.uk (M. Bruns), luekepohlh@gmail.com (H. Lütkepohl).

structural shocks in volatility regime \mathcal{T}_m are obtained from the reduced-form errors, u_t , by a linear transformation, $w_t = B(m)^{-1}u_t$, such that the components are instantaneously uncorrelated with diagonal covariance matrix $\Sigma_w(m)$. Thus, we allow both the variances of the structural shocks, $\Sigma_w(m)$, as well as the transformation matrices, $B(m)$, which represent the impact effects of the structural shocks, to depend on the volatility regime m . If the impact effects are time-invariant and do not change such that all variation in the residual covariances $\Sigma_u(m)$ is captured by changes in the $\Sigma_w(m)$, then $B(1) = \dots = B(M)$. As the VAR slope coefficients, A_1, \dots, A_p , are assumed to be time-invariant, time-varying structural impulse responses can only be due to changes in the impact effects of the shocks because they are functions of the impact effects of the shocks and the reduced-form impulse responses.

In our empirical analysis we are primarily interested in an oil market shock which we place first in the vector of structural shocks, w_t , that is, w_{1t} is the oil market shock. It is identified by a proxy, z_t , which is correlated with w_{1t} and uncorrelated with all other shocks. Given that interest focusses on w_{1t} , we would like to test the impact effects associated with the first shock. In other words, we are interested in testing for time-varying elements in the first column, say $b(m) = (b_1(m), \dots, b_K(m))'$, of the $B(m)$ matrices. Lütkepohl and Schlaak (2022) normalize the impact effect of the first shock on the first variable to be 1 in all volatility regimes and consider the $(K - 1)$ -dimensional vectors $\beta(m) = (b_2(m), \dots, b_K(m))'$. For $m, k \in \{1, \dots, M\}$, $m \neq k$, they propose a test of the pair of hypotheses

$$\mathbb{H}_0 : \beta(m) = \beta(k) \quad \text{versus} \quad \mathbb{H}_1 : \beta(m) \neq \beta(k) \quad (2)$$

based on the asymptotic normal distribution of $\sqrt{T}(\hat{\beta}(m) - \beta(m))$. We use the χ^2 -statistics of Lütkepohl and Schlaak (2022) for testing the null hypothesis in (2) and related t -statistics for testing individual elements of $\beta(m)$ (see the Supplement).

An extension of the test for the case of identifying more than one shock by a set of proxies was proposed by Bruns and Lütkepohl (2022). It does not require identification of the individual shocks identified by the set of proxies and will be used in the empirical analysis when more than one shock is of interest.

3. Empirical analysis

3.1. Data and model setup

Given the importance of oil for industrialized economies, the oil price is an important macroeconomic variable. Kilian and Murphy (2014) and Känzig (2021) argue that oil prices are forward looking and, hence, driven by expectations. We use a VAR(12) model with a constant term as Känzig (2021) and Degasperri (2021) for the real price of oil (rp_t), world oil production ($prod_t$), world oil inventories (inv_t), world industrial production (ip_t^{World}), U.S. industrial production (ip_t^{US}), and the U.S. consumer price index (cpi_t^{US}) such that

$$y_t = (rp_t, prod_t, inv_t, ip_t^{World}, ip_t^{US}, cpi_t^{US})'$$

The oil market variables (rp_t , $prod_t$, inv_t) are standard variables in VAR models for the global crude oil market as proposed by Kilian and Murphy (2014) (see also Zhou (2020)). All variables are in logs. Känzig uses monthly data from January 1974 to December 2017. We follow Degasperri (2021) and use a sample period 1984M1-2019M12, i.e., we start the sample at the beginning of the Great Moderation period (see, e.g., Stock and Watson (2003)) and terminate the sample at the start of the Covid-19 pandemic. Our gross sample size is 432. It includes 12 presample values for estimating the VAR(12) model such that our $T = 420$. Following Kilian and Murphy (2014), we include seasonal dummy variables in the model to account for seasonal variation in oil inventories. In addition, we employ a time dummy variable for 1985M12 to account for the sharp decline in real oil prices that foreshadows the collapse of OPEC in 1986, which might otherwise distort our test results. Precise variable specifications and data sources

can be found in the Supplement. Note that we are using updated data also for the period overlapping with Känzig's sample period.

For our sample period, 1984M1-2019M12, volatility changes in the data have been discussed previously in the related literature. For example, Lütkepohl and Netšunajev (2014) find that some of the special events in the oil market discussed by Barsky and Kilian (2004) resulted in high volatility episodes notably at the time of the 1990/91 gulf war. Based on a detailed heteroskedasticity analysis considering special events in the oil market and statistical tools, we find volatility changes in $T_1 = 1990M9$ and $T_2 = 2005M4$. These two volatility change points will be used in the empirical analysis. Details of our heteroskedasticity analysis are given in the Supplement.

We use the proxies for oil market shocks constructed by Känzig (2021) and Degasperri (2021) from OPEC announcements about their production plans. Känzig (2021) calls the identified shock an 'oil supply news shock'. He assumes that changes in oil futures on the day of the announcement are driven exclusively by revisions in the expectations of market participants due to the announcements, rather than other factors such as oil demand or geopolitical shocks, given the tight window of one day around the announcement. Thus, they can be assumed to be exogenous to the global economic outlook. However, Degasperri (2021) argues that OPEC announcements may not only reveal news about future oil supply but may also be determined by news about future demand conditions in the oil market. He separates the surprises in the Känzig proxy in oil supply and demand surprises by classifying a surprise as a supply surprise if the day-on-day growth rate of the S&P500 index declines and a surprise is a demand surprise if stock returns go up. In the following we will refer to the proxy proposed by Känzig (2021) as the Känzig proxy and the related shock as the Känzig shock. The two Degasperri proxies are signified as Degasperri-supply and Degasperri-demand proxy (or shock), respectively (see the Supplement for details on the construction of all three proxies).

3.2. Analysis of time-invariant shock transmission

As mentioned earlier, we use tests proposed by Lütkepohl and Schlaak (2022) and Bruns and Lütkepohl (2022) to investigate the time-invariance of the impact effects of the shocks. These authors show that their tests are robust to misspecifying the volatility change points or the number of volatility regimes. They are even applicable if some of the volatility regimes contain further heteroskedasticity or the change in volatility is gradual rather than abrupt. They point out, however, that such deviations from the ideal conditions may lead to reduced power of the tests. This point is important as these model deficiencies are likely to be present in our data. Further features that may reduce the power of the tests are small sample sizes within the individual volatility regimes, many VAR lags and large number of variables as well as proxies that are not strongly correlated with the shocks of interest. In the Supplement we show that such features are present in our model and data. Hence, we are working in a low power environment. Therefore we interpret p -values below 10% as indication of evidence against \mathbb{H}_0 . Test results for $M = 2$ and $M = 3$ volatility regimes based on individual proxies are presented in Table 1.

Some low p -values in Table 1 suggest that there has been a change in the impact effects of the shocks in September 1990. For the Känzig shock p -values below 10% are obtained for the joint test of $\mathbb{H}_0 : \beta(1) = \beta(2)$ for $M = 2$ volatility states as well as for the individual impact effect of inv_t . The latter result is reinforced by the small p -value for $M = 3$ when testing a possible change in the impact effect in 1990M9.

Interestingly, none of the p -values related to tests for the Degasperri-supply shock is smaller than 10% such that the tests do not support time-varying impact effects of this shock. On the other hand, there are small p -values for the Degasperri-demand shock for $M = 2$ (see the joint test and the individual tests for inv_t and cpi_t). Moreover, for $M = 3$ there is some evidence for a change in the response of $prod_t$, inv_t , and cpi_t^{US} in 1990M9.

Table 1
Tests for time-varying impact effects (p -values).

Känzig proxy								
$M = 2$ volatility regimes								
T_1	\mathbb{H}_0	joint	$prod_t$	inv_t	ip_t^{World}	ip_t^{US}	cpi_t^{US}	
1990M9	$\beta(1) = \beta(2)$	0.066	0.661	0.054	0.295	0.538	0.223	
$M = 3$ volatility regimes								
T_1	T_2	\mathbb{H}_0	joint	$prod_t$	inv_t	ip_t^{World}	ip_t^{US}	cpi_t^{US}
1990M9	2005M4	$\beta(1) = \beta(2)$	0.632	0.909	0.083	0.651	0.836	0.775
		$\beta(1) = \beta(3)$	0.150	0.468	0.176	0.290	0.535	0.121
		$\beta(2) = \beta(3)$	0.805	0.635	0.794	0.499	0.617	0.176
Degasper-supply proxy								
$M = 2$ volatility regimes								
T_1	\mathbb{H}_0	joint	$prod_t$	inv_t	ip_t^{World}	ip_t^{US}	cpi_t^{US}	
1990M9	$\beta(1) = \beta(2)$	0.527	0.193	0.705	0.140	0.740	0.253	
$M = 3$ volatility regimes								
T_1	T_2	\mathbb{H}_0	joint	$prod_t$	inv_t	ip_t^{World}	ip_t^{US}	cpi_t^{US}
1990M9	2005M4	$\beta(1) = \beta(2)$	0.372	0.196	0.740	0.150	0.740	0.149
		$\beta(1) = \beta(3)$	0.963	0.652	0.730	0.431	0.860	0.985
		$\beta(2) = \beta(3)$	0.992	0.849	0.844	0.857	0.957	0.667
Degasper-demand proxy								
$M = 2$ volatility regimes								
T_1	\mathbb{H}_0	joint	$prod_t$	inv_t	ip_t^{World}	ip_t^{US}	cpi_t^{US}	
1990M9	$\beta(1) = \beta(2)$	0.082	0.103	0.065	0.903	0.262	0.064	
$M = 3$ volatility regimes								
T_1	T_2	\mathbb{H}_0	joint	$prod_t$	inv_t	ip_t^{World}	ip_t^{US}	cpi_t^{US}
1990M9	2005M4	$\beta(1) = \beta(2)$	0.521	0.705	0.301	0.591	0.525	0.704
		$\beta(1) = \beta(3)$	0.011	0.021	0.090	0.479	0.333	0.041
		$\beta(2) = \beta(3)$	0.871	0.683	0.623	0.409	0.950	0.243

Table 2
Joint tests for time-varying impact effects (p -values) of Degasperi Shocks.

$M = 2$ volatility regimes				
	T_1	\mathbb{H}_0	joint	
	1990M9	$B_{1:2}(1) = B_{1:2}(2)$	0.005	
$M = 3$ volatility regimes				
	T_1	T_2	\mathbb{H}_0	joint
	1990M9	2005M4	$B_{1:2}(1) = B_{1:2}(2)$	0.931
			$B_{1:2}(1) = B_{1:2}(3)$	0.061
			$B_{1:2}(2) = B_{1:2}(3)$	0.996

Note: p -values based on the identification-robust test of Bruns and Lütkepohl (2022) for time-varying impact effects of both shocks jointly. $B_{1:2}(m)$ signifies the elements of the impact effects matrix considered in volatility regime m .

In Table 2 the test results for considering the two Degasperi shocks jointly present even stronger evidence for a change in 1990M9. The p -value for $M = 2$ volatility regimes is below 1% and, for $M = 3$, a p -value of 0.061 is obtained for $\mathbb{H}_0 : B_{1:2}(1) = B_{1:2}(3)$. The reason for considering the impact effects of the two Degasperi proxies jointly is that it is not clear whether the two proxies actually identify the shocks separately. Note that the movement in the stock index may be driven by demand and supply effects jointly and an increasing or declining stock index may just reflect which of these effects is dominant but still may incorporate both effects.

Overall our tests support time-varying impact effects of oil market shocks during our sample period. Of course, the fact that changes in the transmission have occurred does not necessarily mean that such changes are substantial. To investigate that issue in more detail, we have computed the impulse responses. The responses to the Känzig and Degasperi shocks are depicted separately for the period before and after 1990M9 in Fig. 1. The confidence intervals in the figures are generated

by a moving block bootstrap (MBB) as explained and justified in the Supplement.

The confidence intervals for the impulse responses for the pre- and post-1990M9 subperiods in Fig. 1 in many cases overlap substantially implying that the changes in the transmission process between the two subperiods may not be dramatic. However, there are some more substantial changes in the responses of the variables to the different shock. For example, the Känzig shock, depicted in the left column of Fig. 1, leads to larger inventories in the post-1990M9 period. There is also some indication that the U.S. variables ip_t^{US} and cpi_t^{US} react more strongly to the shock post-1990M9. The impulse responses of the Degasperi-demand shock in the right column of Fig. 1 have non-overlapping confidence intervals for the pre- and post-1990M9 periods for $prod_t$, ip_t^{World} , ip_t^{US} , and cpi_t^{US} which indicates time-varying shock transmission and may reflect the low power of the previously applied tests. For cpi_t^{US} , but to a lesser degree also for $prod_t$, that is well in line with our test results. Thus, the impulse responses reinforce the

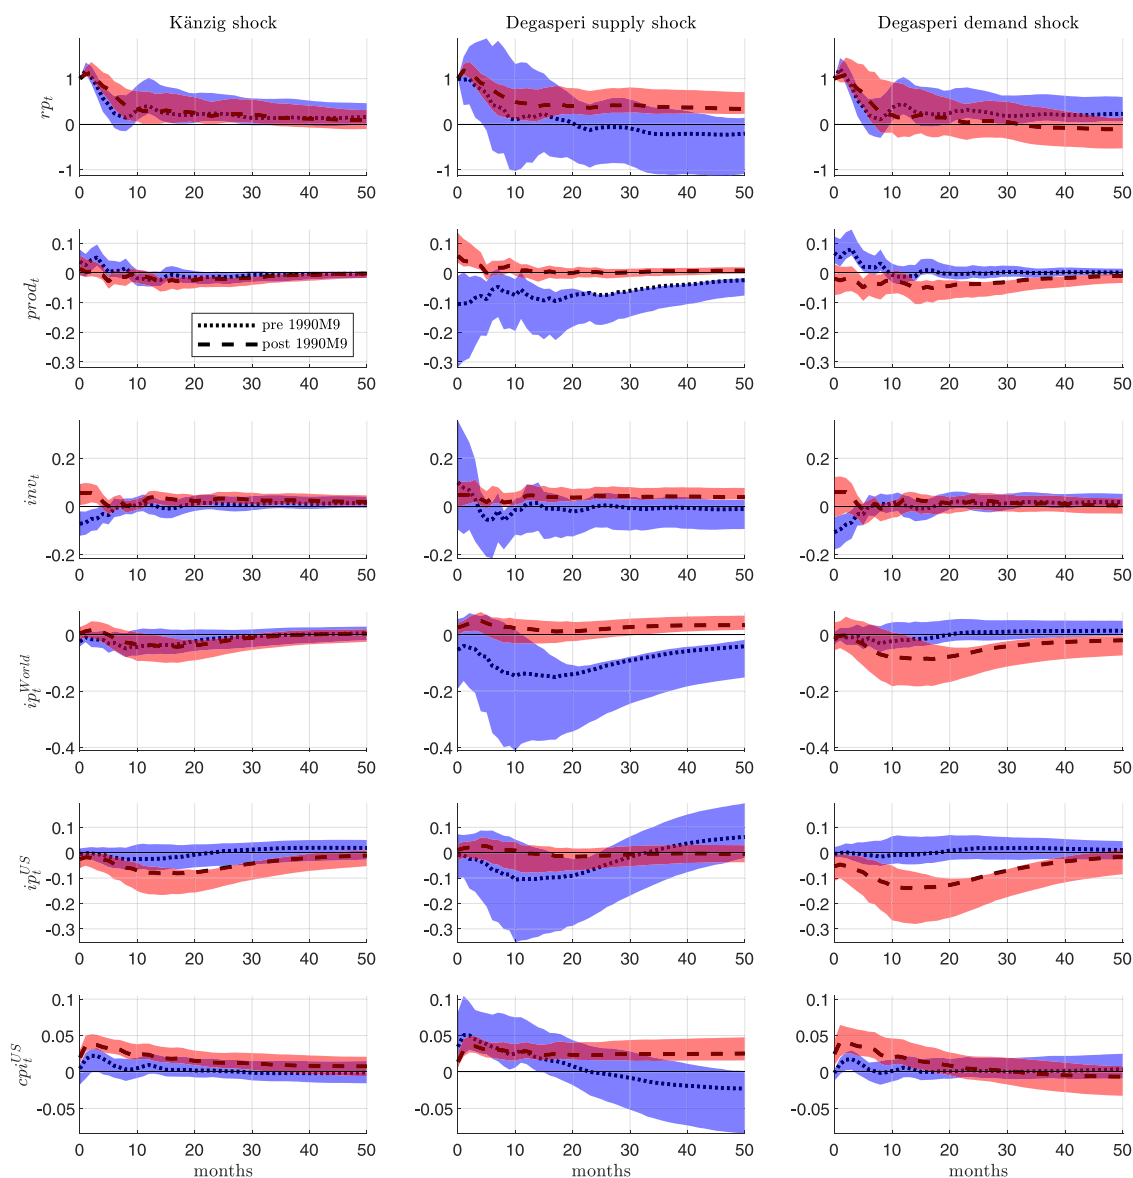


Fig. 1. Responses to Känzig’s shock (left column), Degasper supply shock (middle column) and Degasper demand shock (right column) with pointwise 68% confidence bands. Blue areas indicate confidence intervals for the period from 1984M1-1990M9 and red areas represent confidence intervals for the period from 1990M10-2019M12 (Hall intervals based on MBB with 5000 bootstrap replications). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

conclusions from our tests that even if we assume that the VAR slope coefficients are time-invariant across our sample period, there has been a change in the transmission of oil market shocks during our sample period due to the change in the volatility of the shocks.

Overall our findings suggest that it may be worth allowing for the possibility of time-varying shock transmission in studying the impact of the oil market on the macroeconomy if the model residuals are heteroskedastic instead of simply assuming time-invariance and using heteroskedasticity-robust inference.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgements

The authors thank Thore Schlaak for initiating this research by drawing their attention to the study by Känzig (2021) and for providing some evidence for time-varying shock transmission. They thank Diego Känzig for making his data available and commenting on an earlier draft of this paper. They also thank Lutz Kilian for insightful comments that resulted in substantial improvements in the paper. Thanks are also due to Riccardo Degasper for providing detailed information on his proxy construction and to Lukas Boer for comments. Martin Bruns thanks the British Academy for financial support, grant number TDA22 220046. An earlier version of this paper was circulated under the title *Time-Varying Transmission of Oil Supply News*.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.econlet.2023.111416>.

References

- Barsky, R.B., Kilian, L., 2004. Oil and the macroeconomy since the 1970s. *J. Econ. Perspect.* 18, 115–134.
- Baumeister, C., Hamilton, J.D., 2019. Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. *Amer. Econ. Rev.* 109 (5), 1873–1910.
- Bruns, M., Lütkepohl, H., 2022. Heteroskedastic Proxy Vector Autoregressions: An Identification-Robust Test for Time-Varying Impulse Responses in the Presence of Multiple Proxies. Technical report, DIW Discussion paper 2036, DIW Berlin.
- Degasperi, R., 2021. Identification of Expectational Shocks in the Oil Market Using OPEC Announcements. Technical report, University of Warwick.
- Känzig, D.R., 2021. The macroeconomic effects of oil supply news: Evidence from OPEC announcements. *Amer. Econ. Rev.* 111 (4), 1092–1125.
- Kilian, L., 2009. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *Amer. Econ. Rev.* 99, 1053–1069.
- Kilian, L., Murphy, D., 2014. The role of inventories and speculative trading in the global market for crude oil. *J. Appl. Econometrics* 29, 454–478.
- Lütkepohl, H., Netšunajev, A., 2014. Disentangling demand and supply shocks in the crude oil market: How to check sign restrictions in structural VARs. *J. Appl. Econometrics* 29, 479–496.
- Lütkepohl, H., Schlaak, T., 2022. Heteroscedastic proxy vector autoregressions. *J. Bus. Econom. Statist.* 40, 1268–1281.
- Stock, J.H., Watson, M.W., 2003. Has the Business Cycle Changed and Why? NBER Macroeconomics Annual 2002, Volume 17, NBER Chapters, National Bureau of Economic Research, Inc, pp. 159–230.
- Zhou, X., 2020. Refining the workhorse oil market model. *J. Appl. Econometrics* 35, 130–140.