2146

Discussion Papers

Deutsches Institut für Wirtschaftsforschung 2025

Revisiting Oil Supply News Shocks: Proxy vs. Non-Gaussian Structural Vector Autoregressions

Helmut Lütkepohl and Till Strohsal

IMPRESSUM

DIW Berlin
German Institute for Economic Research
Anton-Wilhelm-Amo-Str. 58
10117 Berlin

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

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 Econstor:
http://ideas.repec.org/s/diw/diwwpp.html
https://www.econstor.eu/handle/10419/10

Revisiting Oil Supply News Shocks: Proxy vs. Non-Gaussian Structural Vector Autoregressions

Helmut Lütkepohl
DIW Berlin and Freie Universität Berlin, Boltzmannstr. 20, 14195 Berlin, Germany
email: luetkepohlh@gmail.com

Till Strohsal
HWR Berlin and Freie Universität Berlin, Badensche Str. 52, 10825 Berlin, Germany
email: till.strohsal@hwr-berlin.de

November 13, 2025[1]

Summary
We replicate a study by Känzig (*American Economic Review*, 111 (2021), 1092-1125), who employs structural vector autoregressive techniques to examine the impact of changes in oil supply expectations on the price of oil and other macroeconomic aggregates. Känzig identifies an oil supply news shock by constructing a proxy from OPEC announcements about their production plans. As this proxy is a controversial instrument for oil supply news, we use the non-Gaussianity of the data to identify independent structural shocks and find that one of them corresponds closely to Känzig's oil supply news shock, implying that the proxy is not necessarily needed to obtain a shock with the same characteristics.

Key Words: Structural vector autoregression, non-Gaussian shocks, proxy SVARs, instruments, shock labeling.

JEL classification: C32

[1]The authors thank Diego Känzig and Lutz Kilian for valuable comments on an earlier version of the paper.

1 Introduction

Känzig (2021) uses a structural vector autoregressive (VAR) analysis to study the effects of changes in oil supply expectations on the price of oil and key macroeconomic variables. He considers a six-dimensional benchmark model 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}). The sample frequency is monthly and the sample period is 1984M1-2017M12. The reduced-form model for the data generating process (DGP) is a VAR model,

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

with lag order $p = 12$. Here $y_t = (rp_t, prod_t, inv_t, ip_t^{World}, ip_t^{US}, cpi_t^{US})'$ and u_t is a zero-mean white noise process with covariance matrix Σ_u. The vector ν is a constant term, as in Känzig's study.

In structural VAR analysis, a range of tools is used to identify the structural form of the model and, hence, the shocks (see, e.g., Kilian and Lütkepohl (2017)). Känzig (2021) identifies his oil supply news shock by a proxy constructed by measuring the changes in oil price futures on the day of OPEC announcements. He implicitly assumes that changes in oil futures prices on the day of the announcement are driven by revisions in supply expectations. However, Degasperi (2025) argues that OPEC announcements also contain information on aggregate demand that has to be taken into account in constructing an oil supply news proxy. Moreover, Kilian (2024) criticizes the construction of Känzig's proxy on different grounds. He points out, for example, that the availability of oil futures prices and the market liquidity has changed during the time period used for construction of Känzig's proxy such that a heterogeneous set of futures price series was used. Kilian also questions the way the changes in futures prices on announcement days are aggregated to monthly frequency for the proxy series and shows that accounting for such issues may substantially change the impulse responses to a shock identified by the proxy. Thus, the Känzig proxy is a controversial measure of oil supply news.

Therefore, in this study, we will consider another way of identifying the structural shocks. As the data display non-normal features we will consider identifying shocks by non-Gaussianity and assess whether one of the shocks identified in that way is similar to Känzig's shock based on his proxy such that his controversial proxy is not needed for its construction. Note that, in our setup, in a six-variate model there can be at most six independent shocks that drive the DGP. Identifying these six shocks requires that at most one of them is Gaussian and all the others are non-Gaussian. In that case we can identify six independent shocks. As these shocks are identified purely on the basis of their distributional, statistical properties, it is not clear whether one of them qualifies as an oil supply news shock. We will show, however, that one of the shocks is very similar to Känzig's shock, thereby reinforcing his results.

In the next section we will briefly present the methodology for identifying shocks by non-Gaussianity and in Section 3 we present our results and compare them to Känzig's findings. Section 4 concludes and some additional results are presented in the Appendix.

2 Identifying Shocks in Non-Gaussian VARs

Identifying structural shocks through non-Gaussianity in VAR analysis has gained popularity in recent years as documented in a growing number of studies that use this identification tool (see, e.g., Lanne and Lütkepohl (2010), Herwartz and Plödt (2016), Herwartz (2018, 2019), Hafner, Herwartz and Wang (2022), Hafner and Herwartz (2023), Gouriéroux, Monfort and Renne (2017), Gourieroux and Monfort (2014), Moneta, Entner, Hoyer and Coad (2013), Lanne, Meitz and Saikkonen (2017), Maxand (2020)). Assuming that the shocks, say w_t, are linear transformations of the reduced-form errors, u_t, and vice versa, that is, $u_t = Bw_t$, the matrix B is unique up to sign changes and changes in column ordering if at most one of the components of w_t has a Gaussian distribution while all the other components are non-Gaussian. In this setup, B is the matrix of impact effects of the shocks on the variables. It contains the central structural parameters of interest. Clearly, if B is known, the shocks can be obtained from the reduced-form residuals. It is no problem that the column signs and column order are not determined by the data. They will be specified by the analyst anyway.

A necessary condition for using this approach for shock identification is that the reduced-form residuals are non-Gaussian. Thus, in a first step of the analysis one needs to find evidence for non-Gaussian u_t. Even if that is confirmed, it does not follow that all shocks, except perhaps one, are non-Gaussian since a single non-Gaussian component of w_t would imply a non-Gaussian distribution of u_t. Therefore, while non-Gaussian u_t is a necessary condition for considering this identification strategy, it does not guarantee that all shocks can be identified.

The next step is to find a matrix B that ensures independent components of $w_t = B^{-1}u_t$. Different approaches have been proposed and used in empirical work. For example, Lanne et al. (2017) assume a non-Gaussian distribution of the structural errors and set up the likelihood function accordingly. They propose a maximum likelihood approach for estimating B, while Gouriéroux et al. (2017) and Hafner et al. (2022) consider the case where the shock distribution is unknown and propose pseudo maximum likelihood (PML) and nonparametric methods.

As we do not know the true distribution of Känzig's data, we will use a PML approach in our empirical study in Section 3. It proceeds by setting up the pseudo likelihood based on t-distributions for the shocks, and estimating the parameters by maximizing that pseudo likelihood function. Gouriéroux et al. (2017) show that this approach provides consistent and asymptotically normal estimators under general conditions. In the empirical study in Section 3, we employ a three-step maximization algorithm for the pseudo log likelihood as proposed by Lanne et al. (2017) and use the corresponding R code provided by Lange, Dalheimer, Herwartz and Maxand (2021). The algorithm estimates the model parameters including the degrees-of-freedom (df) parameters of the assumed t-distributions of the shocks.

3 Empirical Study

As mentioned earlier, we use the benchmark model and data from Känzig (2021). The reduced-form estimates are therefore the same as in Känzig's article. Using a multivariate Jaque-Bera test for non-Gaussianity of the reduced-form residuals, we obtain a p-value which is basically zero. Moreover, the residuals show excess kurtosis. Thus, the results clearly support non-Gaussian reduced-form residuals and, hence, we proceed with estimating the structural parameters B by the PML approach described in Section 2. The estimated B matrix, $\hat{B}_{nG}$, is presented in the Appendix.

Given that the VAR model is six-dimensional, we estimate six shocks $\hat{w}_t = \hat{B}_{nG}^{-1}\hat{u}_t$ which are obtained via the assumed non-Gaussianity of the data. To strengthen the evidence for proper identification, we have tested the non-Gaussianity of the estimated shocks and show the results in Table 1. The univariate Jarque-Bera test applied to each shock series individually produces very large test statistics with corresponding p-values very close to zero. Thus, it provides strong evidence for non-Gaussianity of all six shocks. In Table 1 we also present the estimated degrees-of-freedom parameters of the t-distributions of the shocks from our PML estimation procedure. The point estimates are all below 10 and the standard deviations are reasonably small. This also suggests non-Gaussianity of the shocks. Overall, there is substantial evidence for non-Gaussianity of all six shocks and, hence, of full identification of the shocks via non-Gaussianity.

A priori it is not clear that any one of these shocks is close to Känzig's

Table 1: Assessment of Non-Gaussianity of Structural Shocks

Shock	JB test	p-value	$\widehat{\text{df}}$	$se(\widehat{\text{df}})$
w_1	52.431	0.000	7.414	2.389
w_2	964.337	0.000	3.274	0.534
w_3	28.974	0.000	6.781	2.403
w_4	54.316	0.000	7.576	2.759
w_5	160.517	0.000	6.589	1.850
w_6	83.241	0.000	6.769	0.410

Note: The JB test is the Jarque-Bera test of $\mathbb{H}_0$: Gaussian vs. $\mathbb{H}_1$: non-Gaussian distribution based on the test statistic $JB = \frac{T}{6}sk^2 + \frac{T}{24}(\kappa - 3)^2$, where $sk = T^{-1}\sum_{t=1}^{T}\hat{w}_{kt}^3$ measures the skewness and $\kappa = T^{-1}\sum_{t=1}^{T}\hat{w}_{kt}^4$ the kurtosis of the shocks. Under $\mathbb{H}_0$, JB is asymptotically $\chi^2(2)$ distributed. $\widehat{\text{df}}$ is the PML estimate of the degrees-of-freedom parameter of the t-distributions of the shocks and $se(\widehat{\text{df}})$ is the corresponding estimated standard error.

shock because the non-Gaussian shocks are identified purely by statistical means. To compare them to Känzig's shock, we computed the correlations of our statistically identified shocks with Känzig's proxy VAR shock and present them in Table 2.[2] The first non-Gaussian shock, denoted as w_t^{nG}, has a correlation of 0.911 with Känzig's shock, henceforth w_t^K. Thus, the two shocks estimated by very different approaches are highly correlated. The correlation of Känzig's shock with any of the other shocks is much smaller (less than 0.25).

Table 2: Correlations of Non-Gaussian Structural Shocks with Känzig's Shock

Shock	w_1^{nG}	w_2	w_3	w_4	w_5	w_6
Correlation	0.911	0.171	0.190	0.126	-0.234	0.238

To make sure that w_t^{nG} can be interpreted in the same way as Känzig's shock, we first compare the corresponding impulse responses to those of w_t^K in Figure 1.[3] They are rather similar and suggest the same conclusions regarding the effects of the two alternative shocks on the model variables. Specifically, a negative shock leads to an immediate rise in the oil price.

[2]The shock series is computed by first estimating its impact effects, say b_K, and then using the relation $w_t^K = b_K'\Sigma_u^{-1}u_t/b_K'\Sigma_u^{-1}b_K$ from Stock and Watson (2018).

[3]Following Känzig, we computed 90% and 68% confidence bands using a moving block bootstrap with block size 24 and 10,000 replications.

Global oil production declines in a persistent manner and inventories increase. It may be worth noting, however, that, unlike the positive response of inventories to Känzig's shock, the impact response of inventories to the non-Gaussian shock is almost zero. One would expect an immediate positive response of inventories if the shock actually represents an oil supply news shock. As the confidence intervals also include positive values, the actual impact response may well be positive, however. Overall, as in Känzig (2021), the sharp upward reaction of the oil price, combined with the gradual reduction in production and the positive response of inventories, aligns with the interpretation of a news shock with respect to future oil supply.

In his Figure 5, Känzig (2021) presents the contribution of his shock to the historical decomposition of the price of oil. We have also computed the corresponding contribution of w_t^{nG} and compare it with the contribution of Känzig's shock in our Figure 2. Notably, during the first part of the sample, up to 1991 the contributions of the two alternative shocks are very similar. In other words, the purple line (contribution of w_t^{nG}) is very close to the black line (contribution of w_t^K). In fact, the former is almost for the full sample included in the confidence bands of the contribution of w_t^K. Moreover, from 1990 until about 2010, the contribution of w_t^{nG} is very close to the red dotted line which represents the percent deviations from the mean of the real price of oil. Hence, the development of the real price of oil is to a large extent determined by w_t^{nG} shocks during that period.

Känzig (2021) augments his benchmark model by a range of variables which he adds to the model one-by-one and repeats his analysis with a number of seven-dimensional models. We have also applied our statistical identification approach to those additional models that are based on monthly data. In Table A.1 in the Appendix we present the correlations of the most highly correlated non-Gausian shock with Känzig's shock from the corresponding seven-dimensional model. In most cases it turns out that the correlation exceeds 0.8. Thus, even if an additional variable is added to the benchmark model, one of the non-Gaussian shocks is highly correlated to the corresponding Känzig shock, indicating that the results from the benchmark model are quite robust to model extensions when identification through non-Gaussianity is used. There are some exceptions, however. If the financial uncertainty index VXO is used as an additional variable, the correlation between the Känzig shock and the non-Gaussian shock with the highest correlation is only 0.573. Thus, the two shocks are apparently a bit different. The reason may be that none of the non-Gaussian shocks is a suitable candidate for an oil supply news shock in a model augmented by the VXO. Similarly, when the SVAR is augmented with an effective exchange rate (broad or narrow index), the correlation between our candidate non-Gaussian and

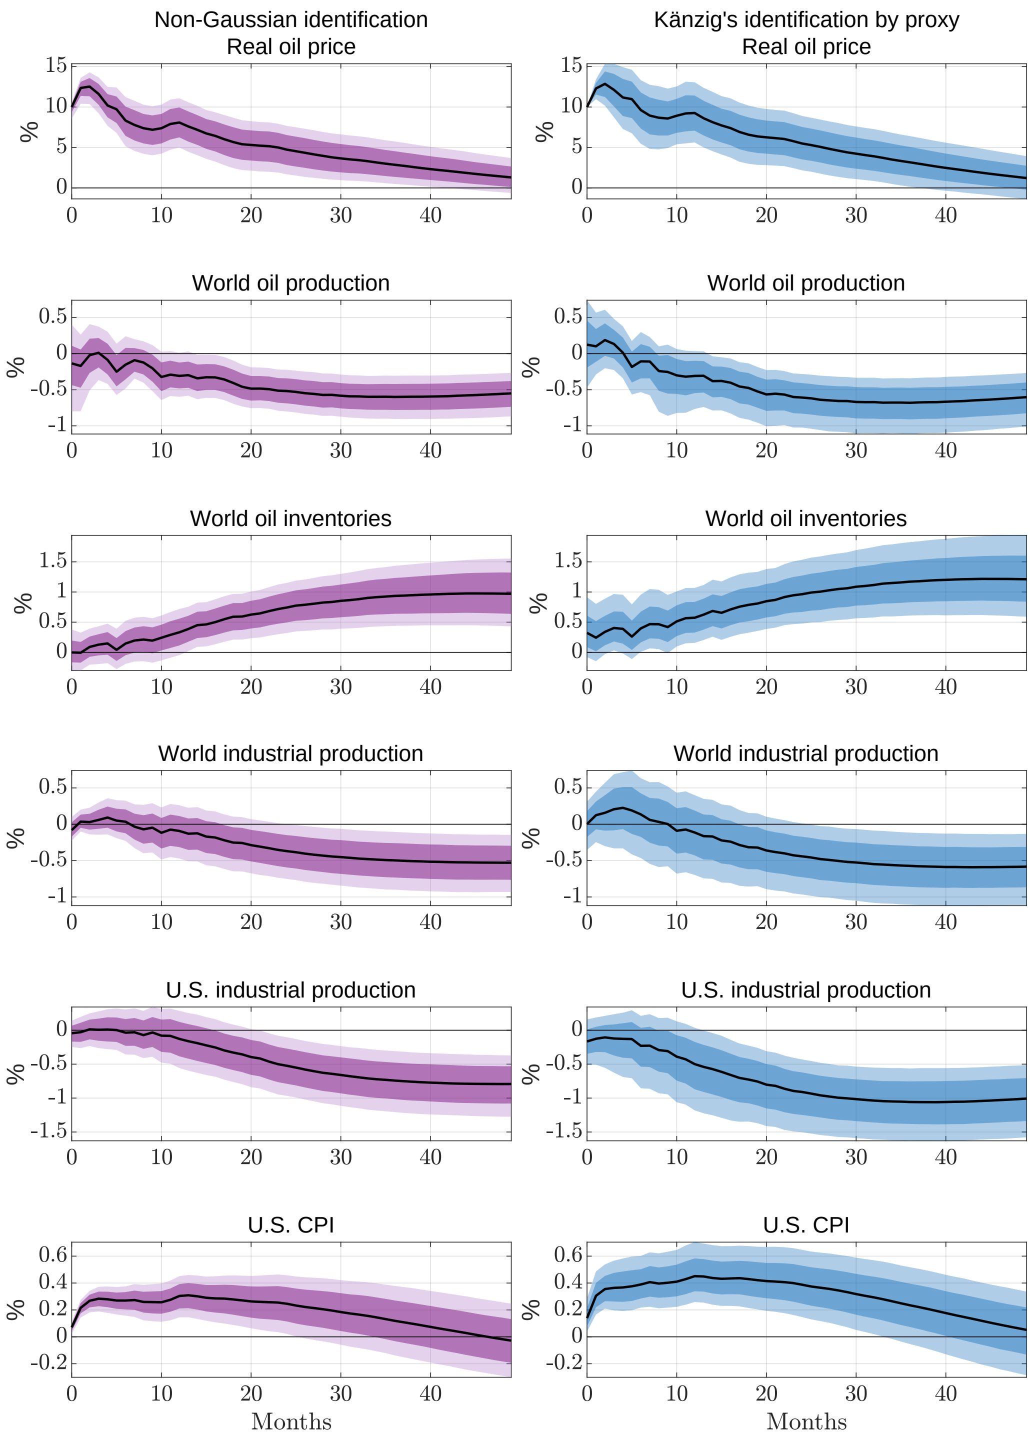

Figure 1: Comparison of impulse response functions. Figures with purple bands show IRFs for the non-Gaussian shock with 68% and 90% pointwise intervals. Figures with blue bands show the IRFs from Känzig's (2021) shock with 68% and 90% pointwise intervals.

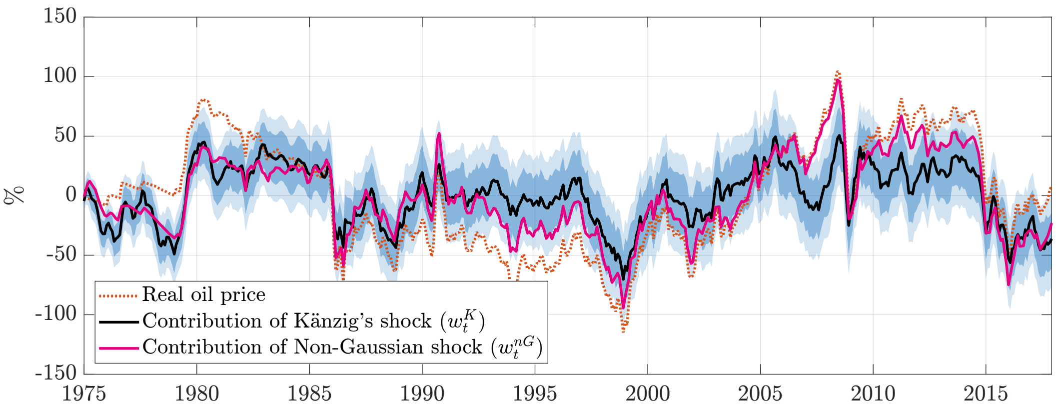

Figure 2: Comparison of Historical Decompositions.

Känzig's shock is relatively low (0.632 and 0.650) meaning that none of the non-Gaussian shocks may qualify as an oil supply news shock.

Given that in most models a highly correlated non-Gaussian shock is found, our results overall reinforce the robustness of Känzig's identification. Across different methods and model extensions, the same structural content is recovered.

4 Conclusions

The overall conclusion from our analysis is that, in the benchmark model, using shock identification through non-Gaussianity provides one shock that has the characteristics of Känzig's oil supply news shock in different respects. In particular, it is highly correlated with Känzig's shock, produces very similar impulse responses and leads to a very similar contribution of the corresponding historical decomposition of the price of oil. Given that there is only one shock in the set of shocks identified by non-Gaussianity that has these properties, our results support Känzig's interpretation of his shock even if one disagrees with the construction of his corresponding proxy.

Having said this, it may be worth noting that no empirical study is perfect and that alternative results can be obtained by using different methods and looking at the data from different angles. This has been done with the Känzig data, for example, by Degasperi (2025) who constructs a proxy for oil supply expectations and another one for an information shock capturing revisions of the expectations about aggregate demand. He finds that the shock to oil supply expectations affects some variables differently than Känzig's shock. Moreover, Kilian (2024) points out some potential problems with the construction of Känzig's proxy. He notes that the availability of the oil futures prices has changed during the period for which Känzig constructs the proxy and shows that using only the futures prices for a period where consistent series are available makes a difference for the impulse responses. He also notes that the way the daily changes in the oil futures prices on announcement days are aggregated to monthly proxy frequency is problematic and gets different results by an alternative aggregation method. He questions that the proxy can be viewed as a measure for oil supply news. Bruns and Lütkepohl (2023) also consider the Känzig study and allow for the possibility of a change in the shock transmission during the sample period. They find evidence that the transmission may have changed at the time of the 1990/91 gulf war. Forni, Franconi, Gambetti and Sala (2025) allow for nonlinearity of the model and find evidence for asymmetric responses to shocks that induce increases and decreases of oil prices. Although such results throw additional light on Känzig's findings, the results of the present study clearly support his main conclusions when looking at the data through the lens of a non-Gaussian identification technique for the shocks.

Appendix. Additional Results

We obtain the following point estimates of the impact effects matrix B:

$$\hat{B}_{nG} = \begin{pmatrix} 6.438 & 0.017 & -0.427 & 1.703 & 0.022 & 0.452 \\ -0.087 & 1.321 & -0.250 & 0.090 & -0.037 & -0.138 \\ -0.001 & 0.184 & 0.841 & -0.215 & -0.014 & 0.135 \\ -0.052 & -0.008 & 0.108 & 0.433 & 0.085 & -0.018 \\ -0.028 & 0.076 & 0.047 & 0.163 & 0.506 & 0.013 \\ 0.045 & 0.016 & -0.028 & 0.028 & -0.006 & 0.174 \end{pmatrix}$$

In contrast, Känzig (2021) estimates the impact effects of his shock as:

$$\hat{b}_K = (6.596, 0.082, 0.213, 0.002, -0.109, 0.092)'$$

Obviously, the first column of $\hat{B}_{nG}$ is similar to this vector, thus indicating that the first non-Gaussian shock may correspond to Känzig's shock.

Table A.1: Correlations of Non-Gaussian Structural Shocks with Känzig's Shocks from Extended Models

Extension	Correlation
U.S. Civilian unemployment rate	0.847
U.S. Personal consumption expenditures	0.894
Oil price expectations	0.905
Michigan survey inflation expectations	0.902
Financial uncertainty index VXO	0.573
Geopolitical risk	0.795
Federal funds rate	0.907
Excess bond premium	0.931
S&P 500	0.810
Effective exchange rate (broad index)	0.632
Effective exchange rate (narrow index)	0.650
U.S. terms of trade	0.873

Note: See Känzig (2021, Online Appendix, Section B.2) for precise definitions of the variables and data sources.

References

Bruns, M. and Lütkepohl, H. (2023). Have the effects of shocks to oil price expectations changed? Evidence from heteroskedastic proxy vector autoregressions, *Economics Letters* **233**: 111416.

Degasperi, R. (2025). Identification of expectational shocks in the oil market using OPEC announcements, *Technical report*, Bank of Italy.

Forni, M., Franconi, A., Gambetti, L. and Sala, L. (2025). Asymmetric transmission of oil supply news, *Quantitative Economics* **16**: 947–979.

Gourieroux, C. and Monfort, A. (2014). Revisiting identification in structural VAR models, *Technical report*, CREST and University of Toronto.

Gouriéroux, C., Monfort, A. and Renne, J.-P. (2017). Statistical inference for independent component analysis: Application to structural VAR analsysis, *Journal of Econometrics* **196**: 111–126.

Hafner, C. M. and Herwartz, H. (2023). Dynamic score-driven independent component analysis, *Journal of Business & Economic Statistics* **41**: 298–308.

Hafner, C. M., Herwartz, H. and Wang, S. (2022). Causal inference with (partially) independent shocks and structural signals on the global crude oil market, *Technical report*, University of Göttingen.

Herwartz, H. (2018). Hodges-Lehmann detection of structural shocks – an analysis of macroeconomic dynamics in the euro area, *Oxford Bulletin of Economics and Statistics* **80**: 736–754.

Herwartz, H. (2019). Long-run neutrality of demand shocks: Revisiting Blanchard and Quah (1989) with independent structural shocks, *Journal of Applied Econometrics* **34**(5): 811–819.

Herwartz, H. and Plödt, M. (2016). The macroeconomic effects of oil price shocks: Evidence from a statistical identification approach, *Journal of International Money and Finance* **61**: 30–44.

Känzig, D. R. (2021). The macroeconomic effects of oil supply news: Evidence from OPEC announcements, *American Economic Review* **111**(4): 1092–1125.

Kilian, L. (2024). How to construct monthly VAR proxies based on daily surprises in futures markets, *Journal of Economic Dynamics and Control* **168**: 104966.

Kilian, L. and Lütkepohl, H. (2017). *Structural Vector Autoregressive Analysis*, Cambridge University Press, Cambridge.

Lange, A., Dalheimer, B., Herwartz, H. and Maxand, S. (2021). svars: An R package for data-driven identification in multivariate time series analysis, *Journal of Statistical Software* **97**(5): 1–34.

Lanne, M. and Lütkepohl, H. (2010). Structural vector autoregressions with nonnormal residuals, *Journal of Business & Economic Statistics* **28**: 159–168.

Lanne, M., Meitz, M. and Saikkonen, P. (2017). Identification and estimation of non-Gaussian structural vector autoregressions, *Journal of Econometrics* **196**: 288–304.

Maxand, S. (2020). Identification of independent structural shocks in the presence of multiple Gaussian components, *Econometrics and Statistics* **16**: 55–68.

Moneta, A., Entner, D., Hoyer, P. O. and Coad, A. (2013). Causal inference by independent component analysis: Theory and applications, *Oxford Bulletin of Economics and Statistics* **75**: 705–730.

Stock, J. H. and Watson, M. W. (2018). Identification and estimation of dynamic causal effects in macroeconomics using external instruments, *Economic Journal* **128**: 917–948.