DIW BERLIN

2036

Discussion Papers

Deutsches Institut für Wirtschaftsforschung 2023

Have the Effects of Shocks to Oil Price Expectations Changed?

Evidence from Heteroskedastic Proxy Vector Autoregressions

Martin Bruns and Helmut Lütkepohl

IMPRESSUM

DIW Berlin, 2023

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

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

ISSN electronic edition 1619-4535

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

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

Have the Effects of Shocks to Oil Price Expectations Changed?

Evidence from Heteroskedastic Proxy Vector Autoregressions

Martin Bruns
University of East Anglia, School of Economics,
Norwich Research Park, NR4 7TJ, Norwich, United Kingdom
email: martin.j.bruns@gmail.com

Helmut Lütkepohl
DIW Berlin and Freie Universität Berlin, Mohrenstr. 58, 10117 Berlin, Germany
email: hluetkepohl@diw.de

This version: April 21, 2023[1]

Abstract. Studies of the crude oil market based on structural vector autoregressive (VAR) models typically assume a time-invariant model and transmission of shocks or they consider a time-varying model and shock transmission. We assume a heteroskedastic reduced-form VAR model with time-invariant slope coefficients and test for time-varying impulse responses in a model for the global crude oil market that includes key macroeconomic variables. We find evidence for changes in the transmission of shocks to oil price expectations during the last decades which can be attributed to heteroskedasticity.

Key Words: Structural vector autoregression, heteroskedastic VAR, proxy VAR, crude oil market

JEL classification: C32

[1]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 changes in the paper. Thanks are also due to Riccardo Degasperi for providing detailed information on his proxy construction. An earlier version of this paper was circulated under the title *Time-Varying Transmission of Oil Supply News.*

1 Introduction

In this study we explore the impact of shocks to oil price expectations on key macroeconomic variables. We assume that the data are generated by a time-invariant reduced-form vector autoregressive (VAR) process with heteroskedastic residuals. This heteroskedasticity may lead to changes in both the structural shock variances and the transmission of these structural shocks to the underlying economy. In other words, despite the time-invariance of the reduced-form VAR process, the transmission of structural shocks can change due to the heteroskedasticity.

Previous studies of the global oil market present mixed evidence on the stability of the transmission of oil market shocks. For example, Kilian (2009), Kilian and Murphy (2014), Lütkepohl and Netšunajev (2014), Baumeister and Hamilton (2019), and Känzig (2021) consider time-invariant impulse responses for periods starting in the 1970s and ending in the new millennium. In contrast, Baumeister and Peersman (2013), Blanchard and Gali (2007), Blanchard and Riggi (2013) use models that allow for time-varying parameters and find evidence of changes in the transmission of oil market shocks in their sample periods starting in the 1960s or 1970s and ending in the new millennium. These studies allow for time-varying shock transmission due to time-varying VAR slope coefficients. In contrast to these studies, we assume that the VAR slope coefficients are time-invariant and investigate whether volatility changes may have driven time-variation in the transmission of oil market shocks.

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 Degasperi (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. It is investigated whether the changes in the residual volatility induce a time-varying transmission of the structural shocks. Specifically, we consider the case of oil price expectation shocks proposed by Känzig (2021) and Degasperi (2021). We explore the issue of a time-invariant transmission of these shocks using statistical tests as proposed by Lütkepohl and Schlaak (2022) and Bruns and Lütkepohl (2022b). The tests explore the time-invariance of the impact effects of the shocks and find evidence for time-varying impulse responses at 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 conclusions follow in Section 4. An Appendix collects further details on the data and additional supporting results.

2 Model Setup and Methodology

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

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

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

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

The M volatility regimes $\mathcal{T}_m = \{T_{m-1}+1, \ldots, T_m\}$ $(m = 1, \ldots, M)$ are assumed to be associated with consecutive time periods, with volatility changes occurring at time periods T_m, for $m = 1, \ldots, M-1$, 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, $w_t = (w_{1t}, \ldots, w_{Kt})'$, is related to the reduced-form errors, u_t, by a linear transformation. Formally, the 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 $B(1) = \cdots = B(M)$. Clearly, the structural parameters $B(m)$ and $\Sigma_w(m)$ are not identified without additional assumptions. In the empirical analysis we will identify the structural parameters of interest by instruments or proxy variables.

Note that the VAR slope coefficients, $A_1, \ldots, A_p$, are assumed to be time-invariant and so are the reduced-form impulse responses which can be computed recursively from the slope coefficients using the recursions $\Phi_j = \sum_{i=1}^{j} \Phi_{j-i} A_i$ for $j = 1, 2, \ldots$, with $\Phi_0 = I_K$ and $A_i = 0$ for $i > p$. Hence, time-varying structural impulse responses $\Theta_j(m) = \Phi_j B(m)$, $j = 0, 1, \ldots$, 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

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

based on the asymptotic normal distribution of $\sqrt{T}(\hat{\beta}(m) - \beta(m))$. Here $\hat{\beta}(m)$ denotes a suitable estimator of $\beta(m)$. The Wald statistic,

$$T\left(\hat{\beta}(m) - \hat{\beta}(k)\right)' \hat{\Omega}^{-1} \left(\hat{\beta}(m) - \hat{\beta}(k)\right), \tag{3}$$

with $\hat{\Omega}$ being a suitable estimator of the covariance matrix of the asymptotic distribution of $\sqrt{T}(\hat{\beta}(m) - \hat{\beta}(k))$, has an asymptotic $\chi^2(K-1)$-distribution under $\mathbb{H}_0$ and related t- or χ^2-statistics for testing individual elements of $\beta(m)$ can be constructed analogously.

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 (2022b). 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

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 consider a model proposed by Känzig (2021) and also used by Degasperi (2021). It contains the oil market variables from an oil market model proposed by Kilian and Murphy (2014) as well as world industrial production (IP), U.S. IP and a U.S. consumer price index as additional macroeconomic variables. In the following we present the reduced-form model setup first and discuss the proxies used for structural estimation. Then we test for time-varying shock transmission and consider impulse responses.

3.1 Empirical Model Setup

Känzig (2021) and Degasperi (2021) consider a six-dimensional baseline VAR(12) model with a constant term 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 Degasperi (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. Some other authors have presented evidence that the impact of oil market shocks on U.S. macroeconomic variables has changed around the middle of the 1980s (see, e.g., Blanchard and Gali (2007), Blanchard and Riggi (2013)) which supports our decision to start the sample in 1984. 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 Appendix A.2. 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 are in fact quite likely and have been discussed previously. 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. Thus, it makes sense to allow for heteroskedasticity in our VAR model.

To determine more precise time periods of potential volatility changes, we have first of all plotted the residuals of the reduced-form VAR(12) model in Figure 1. The plots display obvious volatility changes in at least some of them during the sample period. In particular, there are volatility changes around the time of the gulf war of 1990/91 in the rp_t, $prod_t$, and inv_t residuals which are quite plausible, given that the war had a major impact on the oil market (see Barsky and Kilian (2004)).

To assign volatility change points to specific sample periods we have used a statistical procedure for evaluating potential volatility changes linked to the Gaussian likelihood based criterion function

$$\sum_{m=1}^{M}(T_m - T_{m-1})^{-1} \log\det\widehat{\Sigma}_m, \tag{4}$$

where $\widehat{\Sigma}_m = (T_m - T_{m-1})^{-1}\sum_{t=T_{m-1}+1}^{T_m} \hat{u}_t\hat{u}_t'$, $m = 1,\dots,M$, are computed from ordinary least squares (OLS) residuals $\hat{u}_t$. We have allowed for up to $M = 3$ volatility regimes (two change points). The objective function (4) is minimized sequentially to search for the volatility change points T_m. First the function is minimized for T_1, assuming $M = 2$. The resulting change point is 1990M9. Then $M = 3$ is assumed and $T_1 = 1990\text{M}9$ is fixed while T_2 runs over the other sample points. Thereby we obtain an additional potential change point $T_2 = 2005\text{M}4$. These potential volatility change points are plausible given discussions in the previous literature.[2]

We have used an LM test for heteroskedasticity as described in Lütkepohl (2005, pp. 600-601) to investigate whether the residual covariances in all three potential volatility regimes are in fact distinct. The tests yield very small p-values below 0.1% and thereby support the notion of different variances in the three volatility regimes (see Table A.2 in Appendix A.3).

3.2 The Proxies

We use the proxies for oil market shocks constructed by Känzig (2021) and Degasperi (2021) from OPEC announcements about their production plans. OPEC announcements are arguably a relevant indicator of future oil supply as a substantial share of the world's oil production has taken place within OPEC during our sample period.

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.

[2]We have also searched for a fourth volatility regime and found $T_3 = 2008\text{M}12$. However, our tests for time-varying impact effects for the additional volatility regime did not find significant changes in 2008M12. Therefore we focus the following analysis on up to $M = 3$ volatility regimes only and use inference procedures that account for further volatility changes within the volatility regimes considered.

However, Degasperi (2021) argues that OPEC announcements may not only reveal news on oil supply but may also be determined by demand conditions in the oil market. Those arguments are also in line with earlier discussions of oil supply and demand shocks by Kilian and Murphy (2014). Kilian and Zhou (2023) suggest that the Känzig shocks are better viewed as shocks to oil price expectations. 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 to avoid confusion with other proxies and shocks.

As a consequence of his critique of Känzig's oil supply news shock, Degasperi (2021) constructs a supply and a demand proxy by separating the surprises in the Känzig proxy in oil supply and demand surprises. He does so 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 signify the two proxies obtained in this way as Degasperi-supply and Degasperi-demand proxy (or shock), respectively (see Appendix A.1 for details on the construction of all three proxies). In the following we will assess the time-invariance of the responses to the shocks identified by the Känzig and Degasperi proxies.

3.3 Testing for Time-Varying Shock Transmission

Before we apply the tests for time-varying impact effects in our framework which assumes time-invariant VAR slope coefficients, it may be worth considering the suitability of the tests proposed by Lütkepohl and Schlaak (2022) and Bruns and Lütkepohl (2022b). An obvious issue is whether our volatility model in equation (1), which assumes abrupt changes in volatility at given time points, is suitable. Looking at the residual series in Figure 1, the model may indeed be too simple to capture all the dynamics in the second moments. The good news is, however, that the tests are to some extent robust to misspecifying the volatility change points. Moreover, they are applicable even if the volatility change happens gradually rather than abruptly. Also, if some of the volatility regimes contain further heteroskedasticity this does not invalidate the tests. Fortunately, Lütkepohl and Schlaak (2022) and Bruns and Lütkepohl (2022b) have shown that the tests may well detect changes in the impact effects of a shock even if such deviations from the model assumptions are present, as long as the volatility in the different regimes is clearly distinct, although not perfectly well described by the stylised model. They point out, however, that such deviations from the ideal conditions may lead to reduced power of the tests.

Moreover, Lütkepohl and Schlaak (2022) and Bruns and Lütkepohl (2022b) show that, even if the ideal model conditions hold, the power of the tests for time-varying impact effects depends on the strength of the proxy as an instrument for estimating the impact effects, the size of the model and the length of the volatility regimes. A weaker proxy, a larger model and shorter volatility regimes reduce the power of the tests. Clearly, our six-dimensional VAR model of lag order 12 is rather large and some of the volatility regimes are relatively short. Thus, we are working in a low-power environment, which is worth keeping in mind when interpreting the test results.

We have also investigated the strength of the proxies for the different volatility regimes by considering the usual F-statistics and heteroskedasticity-robust F-statistics resulting from a regression of the OLS residuals of the real price of oil on the proxy in the regimes of interest. It turns out that most F-values are below 10 for the full sample and our subsample periods (see Table A.3 in Appendix A.3). As the standard threshold for a strong proxy used in the related literature is 10 (see Stock, Wright and Yogo (2002)), our proxies are relatively weak which again undermines the power of the tests for time-varying impact effects of the shock of interest. We also confirm Degasperi (2021)'s finding that the oil supply proxy is weak, while an F-test value of more than 10 is found for the oil demand proxy over the whole sample. However, overall the F-tests in Table A.3 as well as the other features of the model suggest that we are working in a low power environment.

For testing for time-varying impact effects, we consider the proxies one-by-one as well as the two Degasperi proxies jointly. The reason for considering a joint test of the impact effects of the two Degasperi proxies 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. Test results for $M = 2$ and $M = 3$ volatility regimes based on individual proxies are presented in Table 1 and results obtained by considering the two Degasperi proxies jointly are given in Table 2.

Results for Känzig's Shock

Looking first at the p-values of the joint tests of the null hypothesis in expression (2) based on the test statistic (3) in Table 1, the tests find evidence for a change in the impact effects of a Känzig shock in September 1990 but no evidence for a change in April 2005. Given the low-power features discussed earlier, we interpret p-values below 10% as indication of evidence againt $\mathbb{H}_0$. If only one volatility change is assumed in 1990M9, the joint test has a p-value of 6.6%. If two volatility changes are assumed, the joint test does not detect changes in the impact effects of the Känzig shocks, which is likely the result of the low power environment. Therefore, for further analyzing time-varying impulse responses, we focus on $M = 2$ volatility regimes with a change point in 1990M9. Looking now more closely at the impact effects on the individual variables in the subperiods associated with $M = 2$, the p-value of 0.054 in Table 1 is evidence that in particular the response of world oil inventories has changed during the sample period. This result is supported for the case of $M = 3$ volatility change points with a p-value of 0.083 when testing the impact effect of oil inventories in the first against the second regime.

Results for Degasperi's Shocks

We first consider the two Degasperi shocks together and test all their impact effects jointly using the identification-robust test proposed by Bruns and Lütkepohl (2022b). The advantage of the test is that it does not require separately identifying the two shocks related to the two proxies. The p-values are presented in Table 2 and offer clear evidence in favor of time-varying impact effects of the shocks. For two volatility regimes, a p-value of 0.005 is obtained. Moreover, for $M = 3$ volatility regimes a test of $\mathbb{H}_0 : \beta(1) = \beta(3)$ results in a p-value of 0.061 and, hence, rejects at a 10% level. Thus, there is clear evidence in favor of a change during our sample. As there is no evidence for a change during the last two volatility regimes from 1990M10-2019M12, the tests overall support a change in 1990M9.

Of course, the joint test of all impact effects of both shocks jointly does not allow us to assess which ones of the impact effects are time-varying. To further investigate that issue, we consider the two shocks individually. The corresponding results are presented in Table 1. Interestingly, none of the p-values related to tests for the Degasperi-supply shock is smaller than 10% such that the tests do not support time-varying impact effects of the supply shock. Thus, the evidence for time-varying impact effects found with the joint test is likely to be due only to varying impact effects of the Degasperi-demand shocks. Indeed the evidence for a change in its impact effects in 1990M9 is quite clear in Table 1. For $M = 2$ volatility regimes the test rejects at a 10% level and for $M = 3$ the null hypothesis $\mathbb{H}_0 : \beta(1) = \beta(3)$ is even rejected with a p-value of 1.1%. Considering the p-values of the tests for the individual variables, there is some evidence for a change in the response of $prod_t$, inv_t, and cpi_t^{US} in 1990M9.

Overall our tests support time-varying impact effects of oil market shocks during our sample period. It is interesting, however, that no evidence is found for changes in the transmission of the Degasperi-supply shock. 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 consider the impulse responses in the following.

3.4 Impulse Responses

To assess the implications of our test results for the impulse responses more generally, we have computed the responses to the Känzig and Degasperi shocks and depict them separately for the period before and after 1990M9 in Figure 2. The confidence intervals in the figures are generated by a moving block bootstrap (MBB) to allow for the possibility that there are further changes in volatility before and after the gulf war. It was shown by Brüggemann, Jentsch and Trenkler (2016) and Jentsch and Lunsford (2019) that the MBB is asymptotically valid even in (conditionally) heteroskedastic structural VAR models. However, in small samples, the MBB tends to produce wide confidence intervals and may not be very reliable (see Brüggemann et al. (2016), Lütkepohl and Schlaak (2019), and Bruns and Lütkepohl (2022a)).[3] Unfortunately, the first volatility regime from 1984M1-1990M9 is rather short so that the confidence intervals may not be very reliable.

The confidence intervals for the impulse responses for the pre- and post-1990M9 subperiods in Figure 2 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 shocks which are worth noting.

In line with the test results, the Känzig shocks, depicted in the left column of Figure 2, lead 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 responses to the Degasperi-supply shock depicted in the middle column of Figure 2 are quite different. Only for oil production the impulse responses do not overlap on impact. This result is not surprising given that the tests for the Degasperi-supply shock produce relatively large p-values and bearing in mind that this test has low power given the weakness of the proxy. It is obvious that the confidence intervals for the very short first volatility regime (1984M1-1990M9) are rather wide. Overall there is some evidence of time-varying impulse responses of the Degasperi-supply shock for our sample period. In particular, we confirm Degasperi (2021)'s finding of a decline in world and U.S. industrial production in response to an oil supply shock only for the period before 1990M9.

[3] Unlike other proxy VAR studies, we use Hall intervals instead of the usual percentile intervals because the former have a build-in bias correction, see Kilian and Lütkepohl (2017, Chapter 12).

Finally, the impulse responses of the Degasperi-demand shock in the right column of Figure 2 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}. For cpi_t^{US}, but to a lesser degree also for $prod_t$, that is well in line with our test results. Especially the responses of $prod_t$ to supply and demand shocks indicates that major changes must have occurred in the way oil producers respond to oil production in other areas of the world. Perhaps the shale oil boom in the U.S. in the post-1990M9 period has contributed to that change in the oil market. Interestingly, a negative Degasperi-demand shock seems to lead to a more severe reduction in U.S. IP in the more recent post-1990M9 period than a Degasperi-supply shock. Both the Degasperi-supply shock and the Degasperi-demand shock lead to an increase in U.S. prices, confirming Degasperi (2021)'s conclusion that these types of shocks present challenges to monetary authorities.

Overall our results present substantial evidence 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. Thus, it may be worth taking that possibility into account in future studies of the global crude oil market and its impact on the U.S. and the world economy.

4 Conclusions

We have pointed out that, in heteroskedastic structural VAR models, changes in the transmission of the structural shocks can occur even if the reduced-form VAR slope coefficients are time-invariant. Using this insight we have investigated the transmission of oil market shocks that affect oil price expectations based on monthly data for the period 1984M1-2019M12. Our model includes six variables. Three of them are related to the global market for crude oil (the real price of oil, world oil production, world oil inventories) and three are key macro economic variables (world industrial production, U.S. industrial production, U.S. consumer prices). We consider three oil market shocks constructed from proxies that were proposed in the related literature. One of the shocks is constructed to capture news on oil supply, another one is capturing demand news, and a third one potentially reflects both types of components. All three shocks affect oil price expectations. Although we assume time-invariant slope coefficients of the VAR model for our sample period, there is substantial evidence for a change in the transmission of all three oil market shocks. In particular, our evidence suggests that the responses of oil production, oil inventories, IP, and U.S. consumer prices are different before and after the 1990/91 gulf war. The specific change in the responses of the variables depends on the shock considered.

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.

References

Barsky, R. B. and Kilian, L. (2004). Oil and the macroeconomy since the 1970s, *Journal of Economic Perspectives* **18**: 115–134.

Baumeister, C. and Hamilton, J. D. (2019). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks, *American Economic Review* **109**(5): 1873–1910.

Baumeister, C. and Peersman, G. (2013). The role of time-varying price elasticities in accounting for volatility changes in the crude oil market, *Journal of Applied Econometrics* **28**: 1087–1109.

Blanchard, O. J. and Gali, J. (2007). The macroeconomic effects of oil shocks: Why are the 2000s so different from the 1970s?, *Working Paper 13368*, National Bureau of Economic Research, NBER.

Blanchard, O. J. and Riggi, M. (2013). Why are the 2000s so different from the 1970s? A structural interpretation of the changes in the macroeconomic effects of oil prices, *Journal of the European Economic Association* **11**: 1032–1052.

Brüggemann, R., Jentsch, C. and Trenkler, C. (2016). Inference in VARs with conditional heteroskedasticity of unknown form, *Journal of Econometrics* **191**: 69–85.

Bruns, M. and Lütkepohl, H. (2022a). An alternative bootstrap for proxy vector autoregressions, *Computational Economics* (forthcoming).

Bruns, M. and Lütkepohl, H. (2022b). Heteroskedastic proxy vector autoregressions: An identification-robust test for time-varying impulse responses in the presence of multiple proxies, *Technical report*, DIW Berlin.

Degasperi, R. (2021). Identification of expectational shocks in the oil market using OPEC announcements, *Technical report*, University of Warwick.

Jentsch, C. and Lunsford, K. G. (2019). The dynamic effects of personal and corporate income tax changes in the United States: Comment, *American Economic Review* **109**: 2655–2678.

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

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

Kilian, L. and Murphy, D. (2014). The role of inventories and speculative trading in the global market for crude oil, *Journal of Applied Econometrics* **29**: 454–478.

Kilian, L. and Zhou, X. (2023). The econometrics of oil market VAR models, *Advances in Econometrics* **45B**: 65–95.

Lütkepohl, H. (2005). *New Introduction to Multiple Time Series Analysis*, Springer-Verlag, Berlin.

Lütkepohl, H. and Netšunajev, A. (2014). Disentangling demand and supply shocks in the crude oil market: How to check sign restrictions in structural VARs, *Journal of Applied Econometrics* **29**: 479–496.

Lütkepohl, H. and Schlaak, T. (2019). Bootstrapping impulse responses of structural vector autoregressive models identified through GARCH, *Journal of Economic Dynamics and Control* **101**: 41–61.

Lütkepohl, H. and Schlaak, T. (2022). Heteroscedastic proxy vector autoregressions, *Journal of Business and Economic Statistics* **40**: 1268–1281.

Stock, J. H. and 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.

Stock, J. H., Wright, J. H. and Yogo, M. (2002). A survey of weak instruments and weak identification in generalized method of moments, *Journal of Business & Economic Statistics* **20**(4): 518–529.

Zhou, X. (2020). Refining the workhorse oil market model, *Journal of Applied Econometrics* **35**: 130 –140.

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
Degasperi-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
Degasperi-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 (2022b) 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.

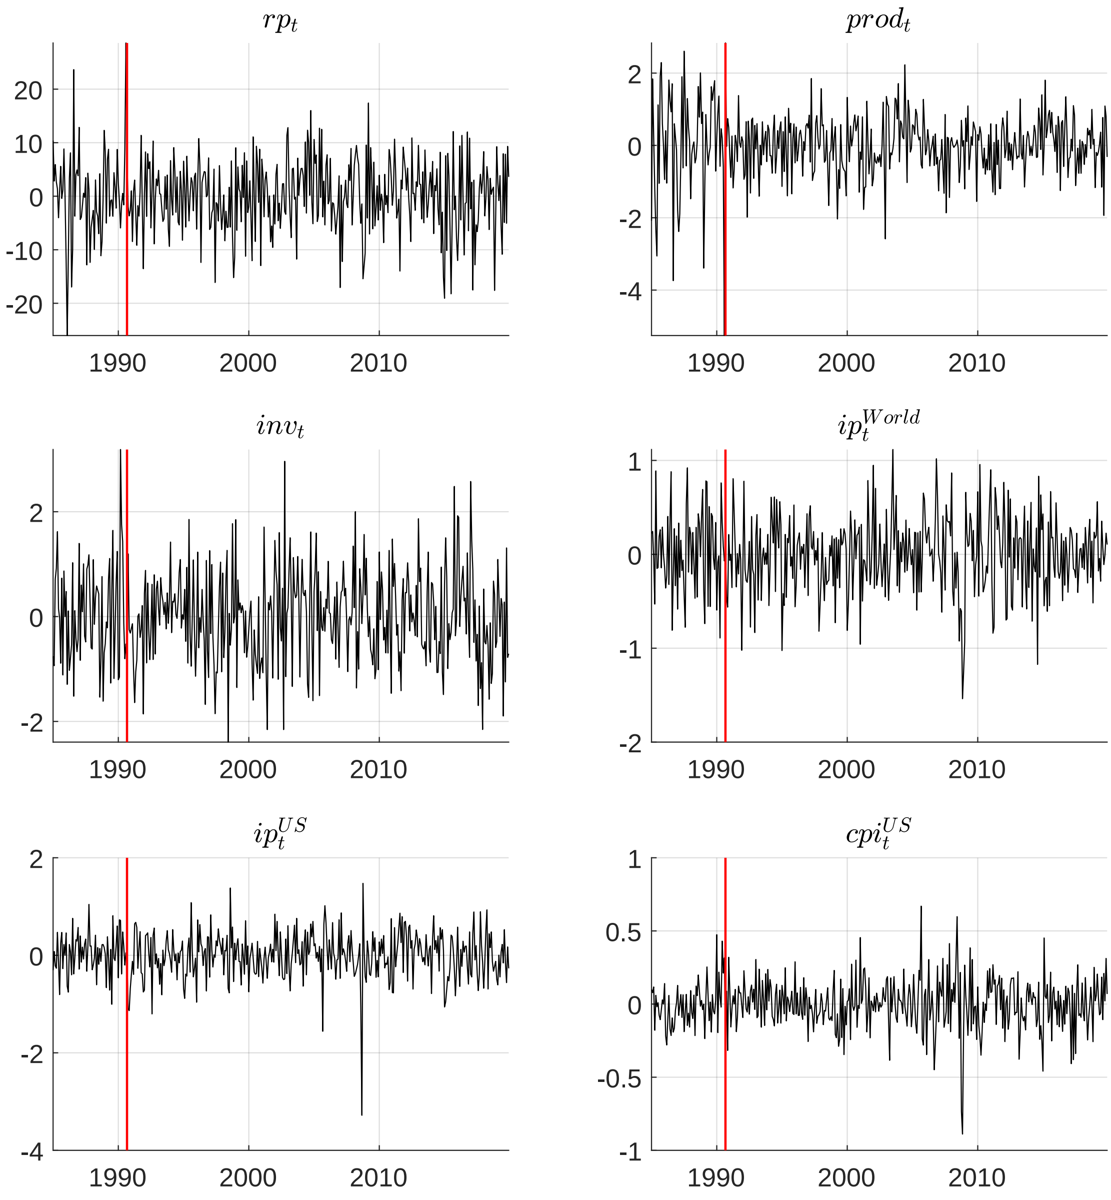

Figure 1: OLS residuals of reduced-form VAR(12) model. The red line marks September 1990.

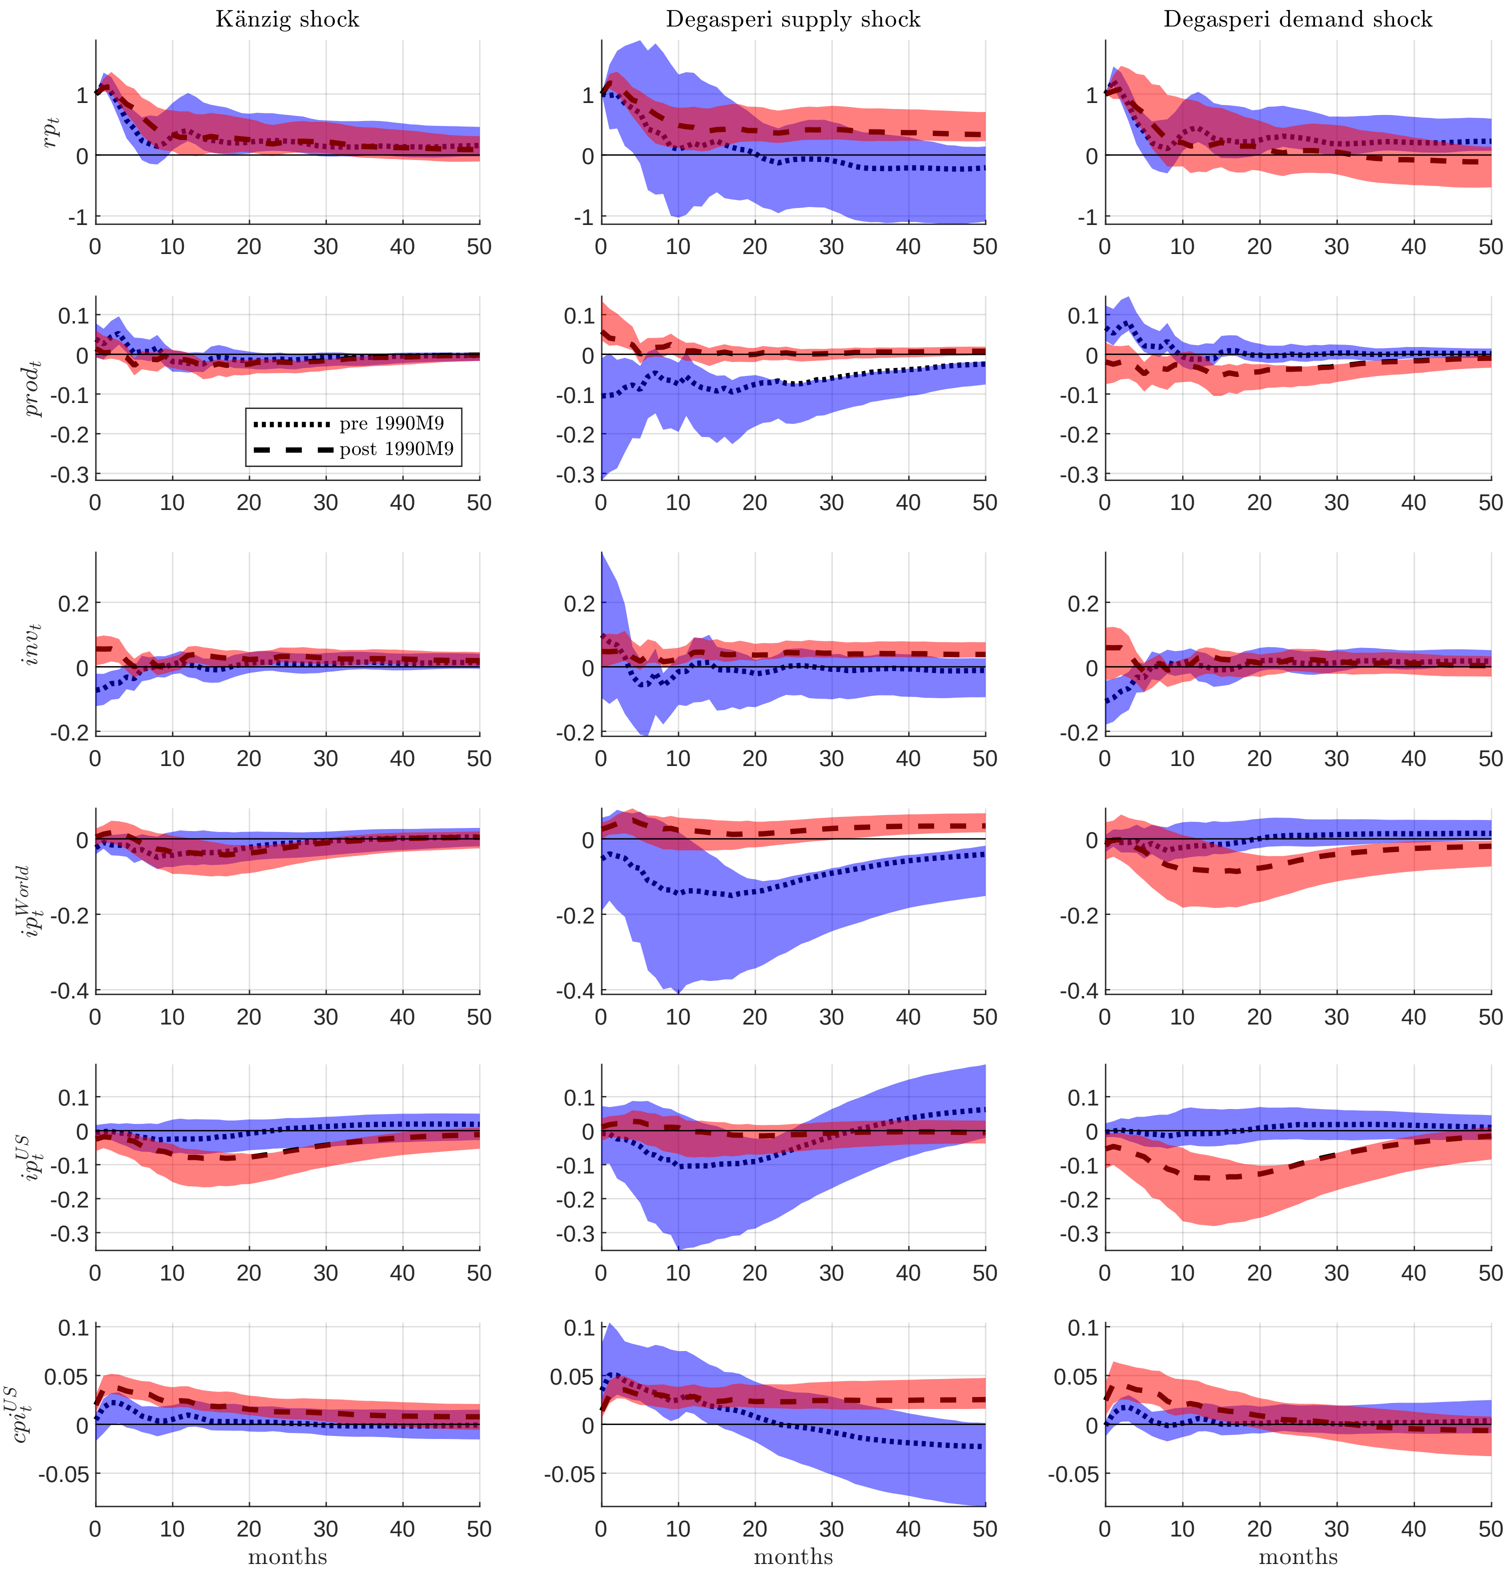

Figure 2: Responses to Känzig's shock (left column), Degasperi-supply shock (middle column) and Degasperi-demand shock (right column) with pointwise 68% confidence bands. Blue areas indicate confidence intervals for the period from 1984M1-1990M9 and red areas represent confidence intervals for the period from 1990M10-2019M12 (Hall intervals based on MBB with 5000 bootstrap replications).

A Appendix

A.1 Construction of Oil Price Expectation Proxies

Känzig (2021) constructs his proxy in the following steps: First, he collects changes, on the day of OPEC announcements, in oil price futures of maturities between 1 and 12 months. Second, he takes the first principal component of these changes. Third, he aggregates the series to a monthly frequency by summing up the changes within the same month and by setting the proxy to 0 for months without an announcement. We use the proxy provided on Känzig's website at a monthly frequency.

To replicate Degasperi (2021)'s proxy construction, we proceed in the following steps: First, we download Känzig (2021)'s composite oil surprise series at OPEC meeting frequency from his website. Second, we construct a series of day-on-day S&P500 returns at OECD meeting frequency (based on the S&P500 composite price index; S&PCOMP provided via Datastream). Third, we match these two series. Following Degasperi (2021), when the stock market is closed on a day of an OPEC meeting, we employ the growth rate of the day when the stock market reopens. Fourth, we construct an oil supply surprise series as oil surprises on days when oil surprises and the S&P500 return have the same sign and an oil demand surprise series as oil surprises on days when the two have opposite signs. Lastly, we construct the Degasperi oil supply and the Degasperi oil demand proxy by summing up these two series separately to a monthly frequency, setting months without a surprise to zero.

Our construction of oil supply and oil demand proxies is based on the composite oil surprise measure at OECD meeting frequency provided by Känzig (2021) and differs from Degasperi (2021) in the following ways: First, Degasperi (2021) employs 6 more dates than Känzig (2021) since Känzig (2021) takes only the last day of OPEC conferences, while Degasperi (2021) takes also days that are not the last for which there is a press release available on the OPEC website. Second, Degasperi (2021) does not include the series for the 12 month ahead contract when estimating the composite surprise series via principal components. Känzig (2021) includes this series. Third, Degasperi (2021)'s proxies start from 1984M1 because the prices for the 11 months ahead future contracts are available only from that date. Känzig (2021) replaces the missing prices with zeros in the estimation of the principal components. Fourth, there are three announcements (29/09/1993, 24/11/1993, and 28/09/2016) for which Känzig (2021) computes the surprise on the day after the OPEC announcement (presumably because markets were closed by the time the announcement was made), while Degasperi (2021) computes it on the same day.

A.2 Variable Definitions and Data Sources

Table A.1: Data Description, Sources, and Sample Periods

Variable	Description	Source	Sample period	Transformation
Instrument				
Känzig proxy	WTI crude oil futures composite measure spanning the first year of the term structure (settlement price); see A.1 for details	Känzig's website	1984M1 - 2019M12	none
Degasperi-supply	Degasperi oil supply series; see A.1 for details	own calculations	1984M1 - 2019M12	none
Degasperi-demand	Degasperi oil demand series; see A.1 for details	own calculations	1984M1 - 2019M12	none
Baseline				
rp_t	WTI spot crude oil price (WTISPLC) deflated by seasonally adjusted U.S. CPI (CPIAUCSL)	FRED	1984M1 - 2019M12	100*log
$prod_t$	World oil production	Datastream	1984M1 - 2019M12	100*log
inv_t	OECD crude oil inventories, calculated based on OECD petroleum stocks (EIA1976) and U.S. crude oil and petroleum stocks (EIA1533, EIA1541), as in Kilian and Murphy (2014)	Datastream/ own calculations	1984M1 - 2019M12	100*log
ip_t^{World}	Industrial production of OECD + 6 (Brazil, China, India, Indonesia, Russia and South Africa) from Baumesiter and Hamilton (2019)	Baumeister's website	1984M1 - 2019M12	100*log
ip_t^{US}	U.S. industrial production index (INDPRO, seasonally adjusted)	FRED	1984M1 - 2019M12	100*log
cpi_t^{US}	U.S. CPI for all urban consumers: all items (CPIAUCSL, seasonally adjusted)	FRED	1984M1 - 2019M12	100*log

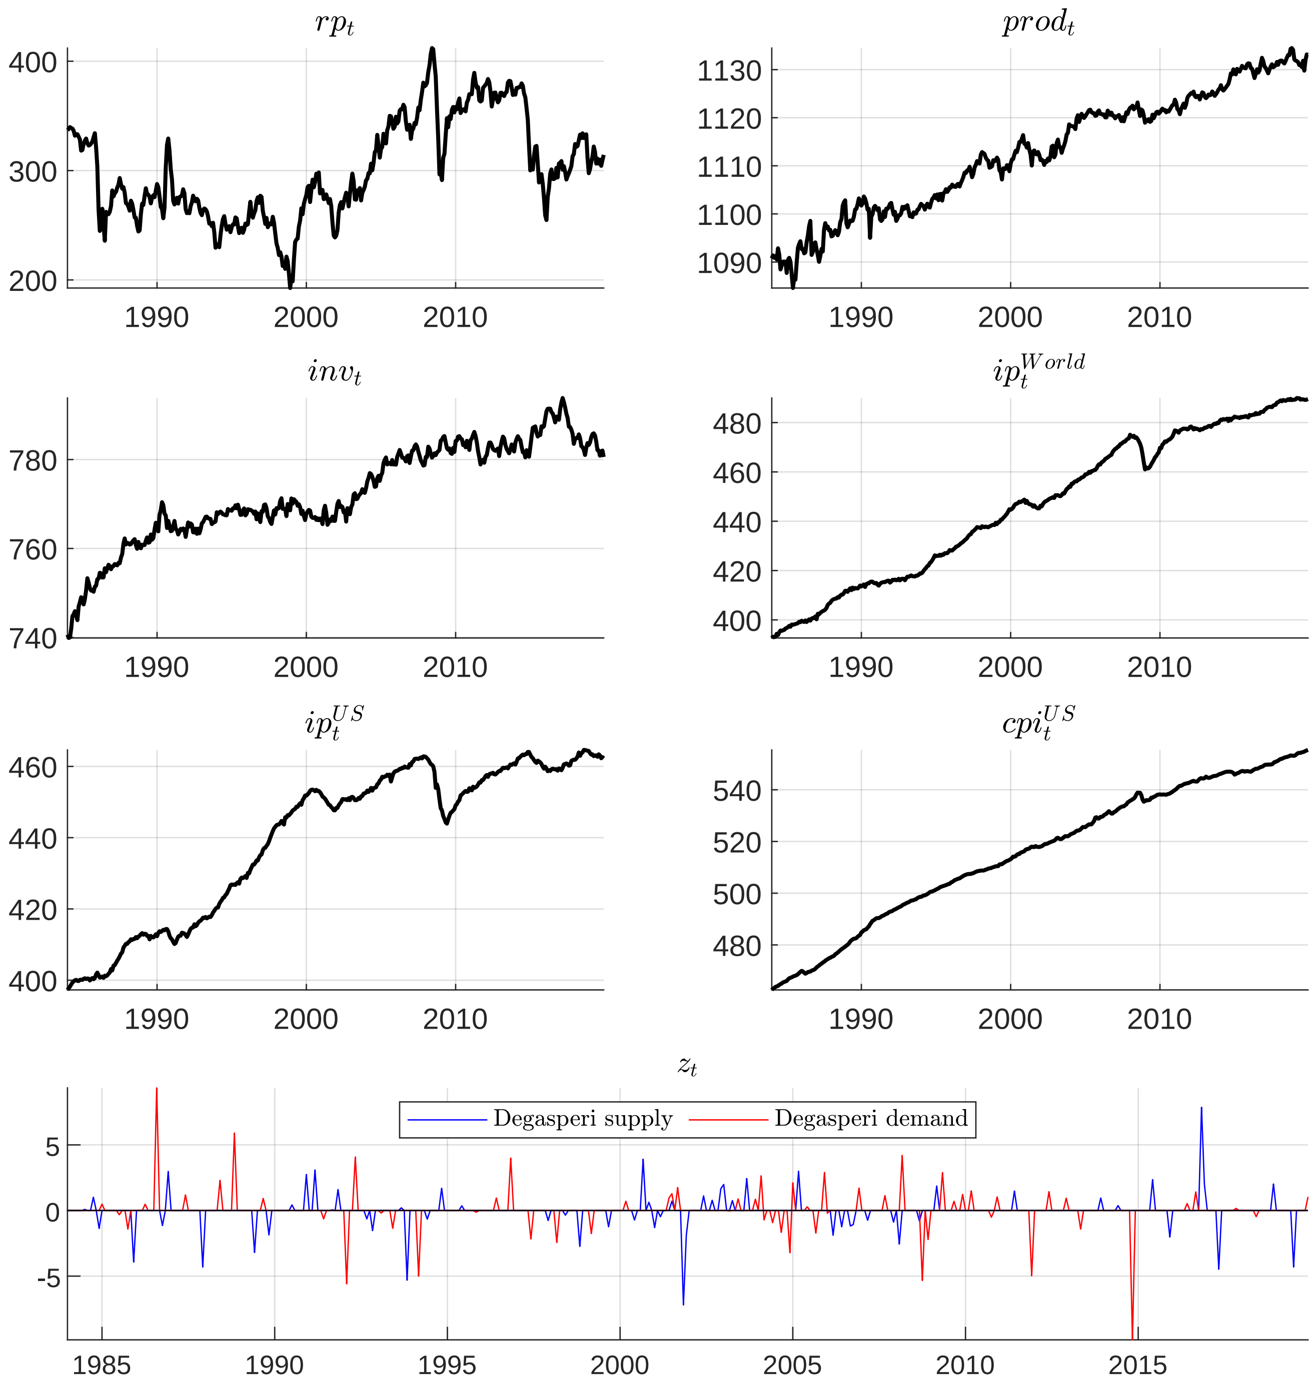

Figure A.1: Data series used in the VAR (top) and proxies (bottom). See Section A.1 for details on the construction of the Degasperi-supply proxy (blue), the Degasperi-demand proxy (red) and the Känzig proxy (blue and red).

A.3 More Detailed Results

Table A.2: Tests for Heteroskedasticity for Reduced-Form VAR Model

Period 1	versus	Period 2	Test statistic	p-value
1984M1-1990M9	versus	1990M10-2019M12	179.46	0
1984M1-1990M9	versus	1990M10-2005M4	89.62	1.88E-10
1984M1-1990M9	versus	post-2005M4	141.83	0
1990M10-2005M4	versus	post-2005M4	56.80	3.84E-5

Note: The LM test for heteroskedasticity described in Lütkepohl (2005, pp. 600-601) was used to test whether the residual covariance in period 1 is different from the residual covariance in period 2.

Table A.3: Tests for Strengths of Proxies

	Känzig proxy		Degasperi-supply proxy		Degasperi-demand proxy	
Sample period	F-test	robust F-test	F-test	robust F-test	F-test	robust F-test
1984M1 - 2019M12	14.16	8.54	3.91	3.08	10.54	5.77
1984M1 - 1990M9	5.21	4.71	0.49	4.40	5.41	4.73
1990M10 - 2019M12	8.47	6.00	3.59	2.20	4.81	4.03