



Economics Working Paper Series

2015/004

Changes in the Global Oil Market

Erdenebat Bataa, Marwan Izzeldin and Denise R.Osborn

The Department of Economics
Lancaster University Management School
Lancaster LA1 4YX
UK

© Authors

All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission, provided that full acknowledgement is given.

LUMS home page: <http://www.lancaster.ac.uk/lums/>

Changes in the Global Oil Market

Erdenebat Bataa

National University of Mongolia

Marwan Izzeldin

University of Lancaster

Denise R.Osborn*

University of Manchester & University of Tasmania

February 2015

Abstract

Using a new iterative algorithm that tests for possible breaks in the coefficients and residual variances of recursively identified structural equations, we examine changes in the parameters of the oil market model of Kilian (2009). Our analysis reveals breaks in the coefficients of the oil production and price equations, together with volatility shifts in all equations. In particular, the medium term response of production to aggregate demand shocks increases after 1980 and the price response to supply shocks is more persistent from the mid-1990s. All variables evidence changes in the relative contributions of individual shocks to their forecast error variances.

JEL classification: E42, Q43.

Keywords: Oil price shocks, multiple breaks, breaks in SVAR

*Corresponding author. Denise Osborn, Centre for Growth and Business Cycles Research, Economics, School of Social Sciences, University of Manchester, Oxford Road, Manchester, UK, M13 9PL. Tel: +44(0)161-275-4861, denise.osborn@manchester.ac.uk. Erdenebat Bataa, School of Economic Studies, National University of Mongolia, Baga Toiruu 4, Ulaanbaatar, Mongolia. Tel: +976-95343365, tsors79@yahoo.com. Marwan Izzeldin, Department of Economics, Lancaster University Management School, Lancaster, UK, LA1 4YW, Tel: +44-1524-594674, m.izzeldin@lancaster.ac.uk

Acknowledgements: The first author gratefully acknowledge financial assistance from the the Asian Development Bank's Capacity Building for Planning Project in Mongolia and the Ministry of Economic Development of Mongolia. However, neither the Bank nor the Ministry bears responsibility for the contents of this paper. We would like to thank to Lutz Kilian and Pierre Perron for making their datasets and computer programs available on the Internet, the modifications of which were used in this paper. We are also grateful to comments received from seminar and conference participants at Lancaster University, University of Manchester, National University of Mongolia, Korea University, Open Society Institute AFP Meeting, OxMetrics conference, Yonsei Global Research Team Workshop, and Mark Wohar and Andrew Vivian for their constructive comments.

1 Introduction

A growing literature suggests that apparent asymmetries and/or changes in the response of the global economy to oil prices may be due to an inadequate distinction between demand and supply shocks, rather than to the nature of the underlying responses. In particular, Kilian (2009) argues that traditional vector autoregressive (VAR) modeling methodologies do not separate oil demand and supply. If the price of oil rises following an expansion in the level of global economic activity, such failure can wrongly attribute the consequent increase as an oil price shock. Therefore, Kilian (2009) proposes ordering restrictions to define a structural VAR (SVAR) model that captures the essential features of the world oil market and, consequently, disentangles demand and supply shocks.

However, the nature of the oil market has undergone several major changes over the four decades following the oil price shocks of the mid-1970s. The transition from the ‘official price’ regime of that period, when the oil price was set by long-term contracts, to the current market-based system of direct trading in spot and futures markets, has seen the power balance shift away from the Organization of the Petroleum Exporting Countries (OPEC). Moreover, OPEC itself experienced internal structural instability, leading to its near-collapse in 1986, and the demand side has seen emerging economies play an increasing role over the last ten years (Hamilton, 2014). Against this background, the present paper uses recent developments in the literature on structural break testing to examine changes in the nature of demand and supply shocks in the oil market through the model proposed by Kilian (2009). Specifically, we extend and apply the methodology of Bataa, Osborn, Sensier and van Dijk (2013), that allows separate estimation of coefficient, volatility and correlation breaks in a VAR framework, to an SVAR context. Since the procedure is quite flexible in permitting different break dates to apply in different parts of the model, we allow the data to inform us whether and how parameters have changed over time. Although our focus is the world oil market, the structural break testing methodology developed here can be used within any recursively ordered SVAR model.

Our paper is related to Baumeister and Peersman (2013b), who detect changes in the volatilities of oil market shocks around the mid-1980s and early 1990s affecting both supply and demand, with short-run price elasticities also changing. Their approach employs a time-varying parameter SVAR specification which assumes parameters change each period. The role of volatility change is also emphasized by Lütkepohl and Netšunajev (2014), who apply a Markov switching model for the volatility of oil market shocks in a specification with other parameters constant, with their results suggesting an apparently permanent shift in the mid-1980s. The results of these studies, therefore, imply that the oil market may have been stable since the early 1990s (or perhaps the mid-1980s), and a structural breaks analysis is well designed to investigate this question. The present study also goes further than these previous analyses of change since we have available data subsequent to the onset of the Global Financial Crisis in 2008. It is clearly of interest to examine how this major world event may have affected the oil market.

An important implication of the analysis of Baumeister and Peersman (2013b) is that the

SVAR coefficients and the volatilities of shocks in this market may change at different dates, which we allow by extending the testing approach of Bataa, Osborn, Sensier and van Dijk (2013) to the SVAR context. Our extension is very flexible, since (in a recursive specification) each equation can be considered separately. Consequently, our methodology not only permits coefficients and volatilities to change at different dates, but dates of change may also differ over equations. A further methodological innovation of the current study is that we allow for the long lags required to capture oil market responses (Hamilton and Herrera, 2004; Kilian, 2008, 2009) in a parsimonious specification through the use of restrictions analogous to those of the heterogeneous autoregression of Corsi (2009).

Our results confirm changes in the parameters of the global oil market model of Kilian (2009) over the period 1975 to 2013. More specifically, although the broad conclusions of Kilian (2009) are generally supported, we find that oil supply reacts positively and significantly to aggregate demand shocks from 1981 onwards, in contrast to the conclusion drawn from a constant parameter analysis. Oil price responses also change, particularly seen in a reduced persistence of its responses to an oil-specific demand shock. Volatility changes over the period lead to important changes in the forecast error variance decompositions, with oil production shocks contributing less and oil specific demand shocks more for all variables over the second half of our sample period.

The structure of the paper is as follows. Section 2 provides some background discussion of the oil market and argues that structural breaks are important in this context. The data series and their properties are discussed in Section 3. Section 4 presents the methodology employed for testing for structural breaks in a SVAR model. Substantive results are given in Section 5, with conclusions in a final section.

2 Background

In an SVAR framework, Kilian (2009) proposes contemporaneous ordering restrictions to identify oil supply shocks, demand shocks that drive all industrial commodities, and oil market-specific demand shocks. His approach leads to very clear conclusions: the main drivers of oil price changes since the 1970s have come from demand, not oil supply, and hence oil prices must be treated as endogenous in modeling the world or US economy. Nevertheless, his analysis assumes constant parameters over a period of marked change in the nature of the world oil market.

Mabro (2006) gives a detailed account of the changes in pricing regimes witnessed by the oil market. In 1973/4 OPEC inherited a regime in which the price of oil was effectively set by fiat¹. When the OPEC countries nationalized their upstream hydrocarbon assets, the administered price effectively was the price at which oil was sold and bought in arms-length transactions from the exporting countries. As can be seen from the real oil price panel of Figure 1, price

¹The ‘Seven Sisters’, namely the seven oil companies, which formed the ‘Consortium for Iran’, dominated the global petroleum industry from the mid-1940s to the 1970s. This group set the ‘posted price’ for a barrel of crude oil or a tonne of petroleum products.

variation in the 1970s was dominated by step-type increases in the middle and at the end of the decade. However, increasing non-OPEC supplies at a time of stagnant world demand resulted in considerable surplus capacity within the OPEC region, which led to competitive price cutting during the early 1980s. The administered (or fixed) price system collapsed in 1985. This was followed by a near-collapse of OPEC in 1986, with its members unable to coordinate their production levels until 1998. Ten members of OPEC, excluding Iraq, however pledged in March 1999 to cut production, which appears to have arrested the price fall (see Figure 1).

Important events involving major oil producers also threw the market into disarray, including the 1973-74 embargo by OPEC Countries, the Iranian Revolution in 1978, the Iran-Iraq War in 1980, and the First Persian Gulf War in 1990. Further, because oil is an imperishable and exhaustible resource, not extracting it is a form of investment. This is especially true when producing nations have no urgent budgetary needs, oil prices are increasing or good financial investment opportunities at home or abroad are small. It is plausible that producers will have responded to higher prices with higher production in the 1970s and 1980s, but this is not necessarily the case now. One reason is that the value of sovereign wealth funds, mostly of oil producing nations, stands at US\$4 trillion or 27% of US GDP in 2011 (PwC, 2011), even after a 20% reduction between 2007 and early 2009 due to the Global Financial Crisis (GFC) (Balin, 2010). Further, Hamilton (2014) points to geopolitical disturbances and geological limitations holding back world oil production since 2005.

Figure 2 illustrates movements in the spot price of light sweet crude oil on the New York Mercantile Exchange (NYMEX) from 1983 to 2010, together with the expected price as indicated by futures contracts of various maturities, as in De Gregorio, Landerretche, Neilson, Broda and Rigobon (2007). The figure indicates that futures prices were initially strongly correlated with subsequent movements in oil price, which is consistent with the Efficient Market Hypothesis. For example, immediately after the Iraq invasion of Kuwait in 1990 the futures market provided an accurate forecast of the subsequent decline in the oil price. However, from early 2000 that pattern was broken, as the futures market continued to indicate a price reduction or no-change, only to find price rising in the subsequent period. When futures contracts were first traded in 1983, the volume was very low compared to recent years, with growth (especially from 2004) due to internet trading and wider participation. Using data over 1990 to 2011, Hamilton and Wu (2014) document changes in oil futures risk premia since 2005, suggesting that the growth in the volume of futures trading might have reversed the direction in which the risk premium operates from the long side to the short side of the contract.

Other changes in the world oil market include the geographical locations for the final consumption of oil products. The share of the developed world (as measured by countries in the OECD) in total demand has shrunk from 75% in 1971 to 51% in 2012, with a corresponding increase in developing (non-OECD) countries². Indeed, Hamilton (2014) points out that China alone has accounted for more than half of the increase in global oil consumption since 2005.

²See BP Statistical Review of World Energy 2013 available at <http://www.bp.com/en/global/corporate/about-bp/energy-economics/statistical-review-of-world-energy-2013.html>

These changes reflect developments in the global economy over the last four decades. The US share of world GDP reached its peak in 1985³ and, alongside this, the manufacturing share in US GDP declined from 24.3% in 1970 to 12.8% in 2010, with a corresponding worldwide decline from 26.6% to 16.2% (Perry, 2012). A quarter of the global electricity was also generated by oil in 1973 compared to about 2% in 2010 (International Monetary Fund, 2011). Helbling (2013) points out that consumption of heavy residual fuel oil by utilities fell by two-thirds between 1970 and 1983, accounting for 15 percent of the global decline in oil consumption and petroleum fuel consumption in the transportation sector fell by 15 percent between 1979 and 1983 in the US. Arguably, therefore, supply shocks had the potential for more dramatic impacts on global economic activity in the earlier period than more recently, as would be predicted from the theoretical models of Lilien (1982), Bernanke (1983), and Davis and Haltiwanger (2001).

Despite the changes just discussed, the vast majority of oil market literature is based on time invariant models, with any analysis of time variation conducted through sample splitting or the use of rolling windows (Baumeister and Peersman, 2013a). Peersman and Van Robays (2009, 2012), Baumeister, Peersman and Van Robays (2010) and Baumeister and Peersman (2013a, 2013b) use a time-varying parameter VAR (TVP-VAR) framework to model both the transmission mechanism and the error covariance matrix to check the robustness of Kilian's (2009) results against possible time variation.

There are two main approaches to modeling changes in parameters over time: one can estimate a TVP-VAR model where the parameters change with each new observation, usually according to a random walk, implying structural breaks occur with each new sample observation⁴. However, Koop, Leon-Gonzalez and Strachan (2009) warn that this may allow an undue amount of variation in highly parametrized VAR models. Alternatively, one can test for breaks in individual parts of the model and allow these breaks only if they are statistically significant, as in Inoue and Rossi (2011). Our model follows this second approach. The parameter estimates within this framework are less affected by noise and outliers, since time variation is not imposed in each period, and are easier to interpret. Further, we separate coefficient from volatility breaks, as we believe the nature of these may differ, with impulse responses presented in relation to shocks of fixed magnitudes. In contrast, those of Baumeister and Peersman (2013b, Figure 2), for example, conflate volatility and response effects by providing impulse responses to a time-varying one standard deviation shock.

An important strand of recent literature on the oil market, including Peersman and Van Robays (2009, 2012), Baumeister, Peersman and Van Robays (2010), Baumeister and Peersman (2013a, 2013b), aims to identify analogous shocks to those of Kilian (2009) using sign restrictions. However, Kilian and Murphy (2012) respond that these sign restrictions are insufficient to identify the separate oil demand and supply shocks. Sign restrictions are further developed in Kilian

³The US then accounted for 32.74% of global nominal GDP; see http://en.wikipedia.org/wiki/Economy_of_the_United_States

⁴Another potential issue with the TVP-VAR model concerns the prior distributions of the parameters. In particular, estimating these using data prior to 1973 conflates different pricing policies; see Mabro (2006).

and Murphy (2014) by additional restrictions on short-run elasticity and impact responses and in Lütkepohl and Netšunajev (2014) by introducing a Markov switching volatility mechanism. However, these restrictions yield a range of possible responses to a given shock, leading to an interpretation problem (see Fry and Pagan, 2011). Sign restriction methodologies are also less amenable to the application of formal structural break inference than the exclusion restrictions imposed by Kilian (2009).

3 Data and Preliminary Analysis

Following Kilian (2009), our three oil market variables consist of monthly global crude oil production, the real price of oil and global economic activity. Oil prices and production are analysed in logarithmic form, while the global activity measure is expressed as percentage deviations from trend⁵; the variables are plotted in Figure 1 over our sample period, from January 1973 to February 2014⁶. As in Kilian (2009) and Baumeister and Peersman (2013a, 2013b), our oil price variable is refiner oil acquisition cost. Although the data are available only from 1974, we follow and Kilian (2009) and backdate this series using the US producer price index for oil. The resulting series is deflated by US seasonally adjusted CPI to derive the real price of oil.

Our sample extends over four decades and hence covers a variety of periods, including the oil price rises of the 1970s, the partial breakdown of OPEC and price collapse in 1986, the period of the Great Moderation and the GFC; see Section 2. We are particularly interested in whether recent events, including the GFC, changed the oil market and the nature of the shocks. In this sense, our analysis differs from recent studies which exclude the latest period on the assumption that standard monetary functions do not apply (e.g., Kilian and Lewis, 2011, whose sample ends in June 2008) or that it would unduly inflate the role of oil shocks when output is decreasing (e.g. Blanchard and Riggi, 2013, whose sample ends in December 2007). Since we test for structural breaks, the data is allowed to determine whether the relationships over recent years do, indeed, differ from earlier periods. The changing features over our extended sample appear particularly evident in the graphs in the right-hand panel of Figure 1, which shows month to month changes (after logarithmic transformations for oil production and real oil prices).

As a preliminary to our SVAR analysis, we examined the unit root properties of the series using the procedure of Kejriwal and Perron (2010) that is robust to breaks in the trend. Although results are not shown to conserve space⁷, oil production was clearly judged to be $I(1)$, with most results pointing to a similar conclusion for the real oil price. The latter is compatible with Blanchard and Riggi (2013), who note that the real price of oil shows a near random walk response. However, conclusions relating to a unit root in the activity measure are fairly marginal.

⁵The last is constructed by Kilian (2009) and based on detrended real bulk dry cargo freight rates. More specifically, the original series in US dollars is deflated by the US CPI and detrended. The underlying cargo rates are not readily available, and we employ the data provided by Kilian after these transformations.

⁶Our data sources are the US Department of Energy for the oil variables and Kilian's website for global activity. In addition, CPI is obtained from the FRED database of the Federal Reserve Bank of St. Louis.

⁷These are available from the authors on request.

Nevertheless, this variable is persistent and structural break tests do not perform well for such data (see Diebold and Chen, 1996, and Prodan, 2008, among others). Therefore, we prefer to analyze the structural stability of the SVAR estimated after differencing all three variables. This is in line with a number of recent studies, including Baumeister and Peersman (2013a, 2013b), Kilian and Lewis (2011), and Jo (2014).

3.1 Univariate analysis

A univariate analysis of the oil market data series aids understanding of their characteristics. To this end, Table 1 provides the results of a structural breaks analysis using the methodology of Bataa, Osborn, Sensier and van Dijk (2014), where observed changes in an oil market variable ΔY_t are decomposed into components capturing the level (L_t), deterministic seasonality (S_t), outliers (O_t) and dynamics (y_t) using

$$\Delta Y_t = L_t + S_t + O_t + y_t \quad (1)$$

where structural breaks are permitted in all components, except O_t . Since our series are expressed as differences, L_t consists of a mean (which is allowed to change over regimes), while S_t is defined so that seasonality sums to zero over the calendar year. Dynamics are captured through an AR model without an intercept, with breaks permitted in the AR coefficients and disturbance variance for y_t . All breaks (level, seasonality, dynamics and variance) can occur at distinct points of time, and hence regimes are specific to these individual components, while outliers are detected allowing for any mean and seasonality shifts. The procedure is based on significance tests, which are conducted at a 5% significance level, with 15% trimming employed for L_t and dynamics in y_t , but 20% for S_t since month-specific seasonal effects are observed only once per year.

Seasonality is included in (1) since there is evidence that oil market variables are seasonal, due to both demand and supply factors; see, for example, International Energy Agency (1996). In particular, demand for heating oil surges during winter months and petroleum consumption rises during northern hemisphere holiday periods (Moosa, 1995). Seasonality is weaker on the supply side, but there are periods during which annual maintenance is undertaken for refineries and climatic influences (such as the hurricane season in the Gulf of Mexico) can affect supply on a seasonal basis. Perhaps as a result of these factors, world oil prices tend to be strongest in the autumn and weakest in the spring⁸. Further, the global economic activity measure of Kilian (2009) is based on shipping freight rates, which (according to Stopford, 2009) follow seasonal demand patterns. Indeed, Kilian and Murphy (2014) recognize the importance of seasonal variation by including seasonal dummies in a VAR model similar to the one we employ. Nevertheless, unit root tests (not reported) rule out the presence of nonstationary stochastic seasonality⁹.

⁸See Oil Market Basics, an online publication of US Energy Information Administration (EIA), at http://www.eia.gov/energyexplained/index.cfm?page=oil_prices

⁹Results of seasonal unit root tests are available from the authors on request. It is also noteworthy that Gallo, Mason, Shapiro and Fabritius (2010) also find little support for seasonal unit roots in oil market data, including

After differencing, Table 1 indicates that the individual oil market variables have largely time-invariant properties, except for volatility. Detected seasonality in the real price of oil changed at January 1986 and November 1998, with the former possibly associated with the fixed pricing regime giving way to a market determined one and the latter is very close to the OPEC’s synchronized cut in production to resist a price fall. Further, differencing largely removes persistence, with low AR orders selected by the Hannan-Quinn criterion and no dynamics at all for oil production. Nevertheless, oil price inflation is moderately persistent.

Baumeister and Peersman (2013b) emphasize changes in the volatility of oil market variables, and our results confirm the presence of two or more volatility changes in all three variables. Production growth is most volatile in the turbulent oil market period to 1990, with the lowest volatility in the most recent period, from late 2004. On the other hand, changes in activity show greatest volatility from the beginning of the GFC, with subdued volatility over the Great Moderation period, here dated to begin in 1980 and end in October 2008; see also Figure 1. Real oil price inflation, on the other hand, shows highest volatility between 1986 and 2009. Although the estimated volatility break dates differ across series, they all point to changes in the first decade of the twenty first century, providing an indication of changes in the nature of shocks to the global oil market over this period. Confidence intervals for volatility break dates are fairly tight, although these should be taken as only indicative in the context of our iterative methodology.

Our subsequent analysis uses data with outliers, seasonality and (constant) means removed. The only outliers detected relate to oil production in the latter part of the 1970s, and these outlier observations are replaced by the median of the neighbouring six observations. Note also that we apply a relatively loose criterion for defining outliers at seven times the interquartile range, in order to avoid losing observations containing valuable economic information, including that relating to the recent global recession. The removal of seasonality and means reduces the burden of coefficient estimation in the subsequent VAR analysis.

4 Structural VAR Methodology

This section first describes the structural VAR (SVAR) methodology proposed by Kilian (2009) to identify oil supply, global demand and oil-specific demand shocks. Subsection 4.2 then outlines an iterative method based on Bataa, Osborn, Sensier and van Dijk (2013) that we use to disentangle structural breaks in the SVAR.

4.1 The oil market model

Based on the SVAR of Kilian (2009) our model is estimated using monthly data for $\Delta \mathbf{z}_t = (\Delta prod_t, \Delta rea_t, \Delta rpo_t)$, where $\Delta prod_t$ is the change in log global crude oil production, Δrea_t oil price and supply.

refers to changes in the global real economic activity measure discussed in Section 3 (to proxy aggregate demand for industrial commodities), and Δrpo_t denotes the one month change in the log real price of oil. The SVAR can be written as

$$\begin{aligned} \mathbf{A}_0 \Delta \mathbf{z}_t &= \boldsymbol{\alpha} + \sum_{i=1}^{24} \mathbf{A}_i \Delta \mathbf{z}_{t-i} + \boldsymbol{\varepsilon}_t \\ E(\boldsymbol{\varepsilon}_t \boldsymbol{\varepsilon}_t') &= \boldsymbol{\Sigma} = \begin{bmatrix} \sigma_{oils}^2 & 0 & 0 \\ 0 & \sigma_{aggd}^2 & 0 \\ 0 & 0 & \sigma_{oild}^2 \end{bmatrix} \end{aligned} \quad (2)$$

where $\boldsymbol{\varepsilon}_t = (\varepsilon_{oils,t}, \varepsilon_{aggd,t}, \varepsilon_{oild,t})'$ denotes a vector of structural shocks with variances of oil supply, aggregate demand and oil demand shocks σ_{oils}^2 , σ_{aggd}^2 , σ_{oild}^2 , respectively; $\boldsymbol{\varepsilon}_t$ is both serially and mutually uncorrelated. Our specification differs from Kilian (2009) only in differencing real oil price and global activity and in the removal of outliers and seasonality.

As in Kilian (2009), a lag length of $p = 24$ is employed in (2), which allows for relatively long lagged responses to shocks. Hamilton and Herrera (2004) show that sufficiently long lags are required to capture responses to oil market changes, while Kilian (2008) also provides an industry level explanation for these. It is nevertheless noteworthy that conventional order selection criteria do not point to the use of a high order in (2); for example, Lütkepohl and Netšunajev (2014) use $p = 3$ for monthly data, as selected by AIC over their sample period.

However, the use of an unrestricted lag order of 24 creates difficulties for structural break inference, since (omitting intercepts) each equation will contain 72 coefficients, with 216 coefficients in the whole system. Testing for structural changes in such a large number of coefficients is infeasible, especially when economic arguments suggest several potential changes early in the sample. In order to counter the curse of dimensionality while retaining lags to a maximum of two years, we adopt the heterogeneous autoregression (HAR) approach of Corsi (2009). This convenient coefficient reduction technique is widely used in the finance literature to capture long memory; see Wang, Bauwens and Hsiao (2013) and Bandi, Russell and Yang (2013), among others. With the frequencies for lagged dynamics often selected as the previous day, previous week and previous month (and perhaps previous quarter), the HAR is found to capture well the dynamics of daily volatility data. Although employed here for lower frequency macroeconomic data, the importance of long memory is implicit in the lags required to capture oil market dynamics.

Our specific demeaned structural heterogeneous VAR (SHVAR) is given by

$$\mathbf{A}_0 \Delta \mathbf{z}_t = \boldsymbol{\Psi}_1 \Delta \mathbf{z}_{t-1} + \boldsymbol{\Psi}_2 \Delta \mathbf{z}_{t-3} + \boldsymbol{\Psi}_3 \Delta \mathbf{z}_{t-6} + \boldsymbol{\Psi}_4 \Delta \mathbf{z}_{t-12} + \boldsymbol{\Psi}_5 \Delta \mathbf{z}_{t-24} + \boldsymbol{\varepsilon}_t \quad (3)$$

where the k -lag difference operator is $\Delta_k = (1 - L^k)$ and L is the conventional lag operator. Since $\Delta_k \mathbf{z}_t = \mathbf{z}_t - \mathbf{z}_{t-k} = \sum_{j=0}^{k-1} \Delta \mathbf{z}_{t-j}$, the SHVAR specification (3) allows the current monthly change $\Delta \mathbf{z}_t$ to depend on short term, medium term and longer term dynamics. The location of

important lagged dependency nodes are specified to coincide with important macroeconomic data frequencies, namely the previous month, previous quarter, previous half year, previous year and previous two years. Note that our SHVAR is a restricted SVAR model with the responses over 24 lags for each of the three variables captured through 15 coefficients. The specific coefficient restrictions in (3) on the SVAR of (2) are:

$$\begin{aligned} \mathbf{A}_1 &= \sum_{j=1}^5 \Psi_j, \mathbf{A}_2 = \mathbf{A}_3 = \sum_{j=2}^5 \Psi_j, \mathbf{A}_4 = \mathbf{A}_5 = \mathbf{A}_6 = \sum_{j=3}^5 \Psi_j, \\ \mathbf{A}_7 &= \mathbf{A}_8 = \dots = \mathbf{A}_{12} = \sum_{j=4}^5 \Psi_j, \mathbf{A}_{13} = \mathbf{A}_{14} = \dots = \mathbf{A}_{24} = \Psi_5. \end{aligned} \quad (4)$$

There are 57 coefficient restrictions in each SHVAR equation and we test these against the general form SVAR, as explained in the following subsection. After imposing these restrictions (when valid), estimates from the SHVAR can be used to conduct conventional SVAR analyses, such as the examination of impulse responses.

Kilian (2009) identifies separate shocks for oil supply, aggregate demand and oil-specific demand shocks through contemporaneous ordering restrictions in \mathbf{A}_0 of (2). More specifically, he assumes this matrix is lower triangular, so that production is unaffected by any within-month demand shocks (the short-run supply curve for crude oil is vertical), while aggregate global demand is contemporaneously unaffected by oil-specific demand shocks. Since the last of these is captured through the real oil price equation, oil price shocks are therefore assumed to influence aggregate demand only with a lag of at least one month, which Kilian (2009) justifies by noting the sluggish nature of aggregate demand responses to historical oil price increases. In effect, oil-specific demand shocks then capture all contemporaneous influences on oil prices that are not reflected in the oil supply and aggregate demand shocks, and include effects arising from fluctuations in the precautionary demand for oil.

4.2 Testing for multiple structural breaks

Based on the procedure developed by Bataa, Osborn, Sensier and van Dijk (2013, 2014), we employ an iterative methodology to examine structural breaks in the coefficients and shock variances of the global oil market model. Using the recursive ordering in the structural model of

(2), the individual equations can be written as:

$$\Delta prod_t = \sum_{i=1}^{24} \mathbf{a}_{i,(1)} \Delta \mathbf{z}_{t-i} + \varepsilon_{oils,t} \quad (5)$$

$$\Delta rea_t = a_{21} \Delta prod_t + \sum_{i=1}^{24} \mathbf{a}_{i,(2)} \Delta \mathbf{z}_{t-i} + \varepsilon_{aggd,t} \quad (6)$$

$$\Delta rpo_t = a_{31} \Delta prod_t + a_{32} \Delta rea_t + \sum_{i=1}^{24} \mathbf{a}_{i,(3)} \Delta \mathbf{z}_{t-i} + \varepsilon_{oid,t} \quad (7)$$

where $\mathbf{a}_{i,(j)}$, $j = 1, 2, 3$ are the j^{th} rows of \mathbf{A}_i , and a_{ij} are elements of \mathbf{A}_0 in (2). The restrictions of (4) are imposed when these are compatible with the data, so that equations (5) to (7) then constitute a SHVAR model.

Following the structural break testing strategy of Bataa Osborn, Sensier and van Dijk (2013), our procedure employs the multiple structural break tests of Qu and Perron (2007), but iterates between coefficient and variance breaks. Although heteroskedasticity consistent inference is used in the initialization of the procedure to detect possible coefficient breaks, subsequent iterations employ appropriate GLS-transformed data, based on residual variance estimates that incorporate breaks. The procedure allows coefficient and variance breaks to be estimated and dated separately, and the simulations of Bataa, Osborn, Sensier and van Dijk (2013) show it generally works well.

As discussed in Hansen (2000), structural changes in the marginal distribution of regressors can lead to size distortions when inference is applied to the stability of regression coefficients. Therefore, rejections indicated by the use of asymptotic critical values based on stable marginal distributions are subject to further test through a bootstrap procedure that takes account of breaks and the recursive nature of the three equations, as discussed below. For computational feasibility, this bootstrap procedure treats break dates estimated through the asymptotic procedure as if they are known dates of potential change.

Due to the contemporaneous causality assumption embodied in (5)-(7), together with the diagonal covariance matrix of a SVAR (or SHVAR) model, each equation can be validly estimated by ordinary least squares (OLS). Therefore, we directly test for structural breaks in these individual equations, thus reducing the burden of testing for multiple breaks compared with a system approach while adding flexibility in allowing different break dates across equations. Such equation-wise testing strategy is relatively old and used in, for example, Bernanke and Mihov (1998) and Bagliano and Favero (1998).

The algorithm for the first equation, namely (5), is:

1. (a) Test the overall null hypothesis of no coefficient breaks using the heteroskedasticity robust ‘double maximum’ $WDmax$ test statistic against the possibility of $m_c \leq M_c$ breaks, where m_c is unknown and the maximum number of breaks M_c is pre-specified. The statistic allows both the number of breaks and their dates to be unknown, with

asymptotic critical values applied. If the $WDmax$ test rejects the null hypothesis of no coefficient breaks, sequential F -type tests (with their asymptotic critical values) are used to estimate the number of breaks and their locations, as recommended by Bai and Perron (1998). If l breaks are detected at dates $\hat{T}_1^{(c)}, \dots, \hat{T}_l^{(c)}$ regime-specific observations $t = \hat{T}_{k-1}^{(c)} + 1, \dots, \hat{T}_k^{(c)}$ are used to obtain $\hat{\mathbf{a}}_{i,(1)}^{(k)}$ ($i = 1, \dots, 24$) and the corresponding residuals $\hat{\varepsilon}_{oils,t}$ for each regime $k = 1, \dots, l + 1$.

- (b) Verify the significance of each detected coefficient break $k = 1, \dots, l$ through a heteroskedasticity robust *Wald*-statistic test of the null hypothesis $\mathbf{a}_{i,(1)}^{(k)} = \mathbf{a}_{i,(1)}^{(k+1)}$, with inference conducted conditional on all other $l - 1$ estimated breaks. The computed statistic is compared to the corresponding empirical distribution obtained from a bootstrap data generating process (DGP) that employs estimated coefficients but restricted through $\mathbf{a}_{i,(1)}^{(k)} = \mathbf{a}_{i,(1)}^{(k+1)}$, and a wild bootstrap process for $\varepsilon_{oils,t}$ in equation (5)¹⁰.
 - (c) If not all coefficient breaks $k = 1, \dots, l + 1$ are individually significant, reduce the number of coefficient breaks to $l - 1$ and estimate new break dates, together with corresponding coefficients and residuals, and return to step 1(b). Repeat until all coefficient breaks are individually significant.
2. (a) Based on the residuals $\hat{\varepsilon}_{oils,t}$ obtained in step 1 (or 3 below), apply the asymptotic ‘double maximum’ likelihood ratio-type test statistic to test the null $H_0 : \sigma_{oils,1}^2 = \dots = \sigma_{oils,m_v+1}^2$ for an unknown number of volatility regimes with $m_v \leq M_v$. If the null hypothesis is rejected, the exact number of breaks is obtained using a sequential test procedure, now based on likelihood ratio-type tests and asymptotic critical values (see Bataa Osborn, Sensier and van Dijk, 2013). If m volatility breaks are detected at dates $\hat{T}_1^{(v)}, \dots, \hat{T}_m^{(v)}$ obtain $\hat{\sigma}_{oils,j}^2$ for each regime $j = 1, \dots, m + 1$, estimated from observations $t = \hat{T}_{j-1}^{(v)} + 1, \dots, \hat{T}_j^{(v)}$.
 - (b) For each break $j = 1, \dots, m$ identified in 2(a), and conditioning on all other $m - 1$ breaks, compute the usual quasi-likelihood ratio test statistic, LR , for the null hypothesis $\sigma_{oils,j}^2 = \sigma_{oils,j+1}^2$. For inference on break j , the shock vector $\hat{\varepsilon}_{oils,t}$ for $t = \hat{T}_{j-1}^{(v)} + 1, \dots, \hat{T}_j^{(v)}, \dots, \hat{T}_{j+1}^{(v)}$ is randomly *i.i.d.* re-sampled, with a wild bootstrap employed in other regimes to create the bootstrapped shocks $\hat{\varepsilon}_{oils,t}^*$ ¹¹. Then $\hat{\varepsilon}_{oils,t}^*$, together with the $(l + 1)$ sets of coefficient estimates found in step 1 (or 3), form the bootstrap DGP that is used to obtain the empirical null distribution, and hence the empirical p -value, for LR ¹².

¹⁰Based on the Monte Carlo studies of the wild bootstrap (Godfrey and Orme, 2004, Davidson and Flachaire, 2008), we set $\varepsilon_{oils,t}^* = \omega_t \hat{\varepsilon}_{oils,t}$, $t = 1, \dots, T$, in which the scalar random variable ω_t has the Rademacher distribution, taking the two possible values $+1, -1$ with probabilities of 0.5.

¹¹The *i.i.d.* bootstrap within the regimes under test enforces the null hypothesis of unchanged variance, whereas the wild bootstrap for the remaining observations allows variances to differ across the other regimes.

¹²The coefficients and shocks are re-estimated for the bootstrap DGP, but the coefficient break dates are assumed known at $\hat{T}_1^{(c)}, \dots, \hat{T}_l^{(c)}$.

- (c) If not all variance breaks are individually significant, reset m to the previous value minus one; then estimate new variance break dates and regime-specific variances. Return to step 2(b) until all m variance breaks are individually significant.
3. Re-estimate the number and dates of coefficient breaks using a feasible generalized least squares (GLS) approach, which is achieved by pre-multiplying all observations entering (5) for $t = \hat{T}_{j-1}^{(v)} + 1, \dots, \hat{T}_j^{(v)}$ by $1/\hat{\sigma}_{oils,j}$, $j = 1, \dots, m + 1$. Re-apply the coefficient test procedure as in step 1, but now apply the homoscedastic version of the multiple breaks test procedure to the coefficient vector $\mathbf{a}_{i,(1)}^{(k)}$ in the GLS-transformed equation.
 4. Iterate between steps 1 and 3 until the numbers and dates of coefficient and variance breaks do not change. Note that step 2 is always applied to the residuals calculated using the original observed (not GLS-transformed) values.

Analogous procedures are applied to detect breaks in the economic activity and price equations, (6) and (7), with the relevant contemporaneous coefficients included in the coefficient break tests. There is, however, an important difference in the bootstrap break testing procedure applied for these latter equations. For the first equation of the system, bootstrap production data are generated recursively from (5), with observed values of the other two series used in the bootstrap replications¹³ together with wild bootstrap shocks. For the activity equation (6), however, values for production are bootstrapped employing the breaks detected in the coefficients of (5) and the wild bootstrap for its shocks; these generated values are used when bootstrapping activity. Similarly, when the procedure is applied to the price equation, the bootstrapped values of oil production and economic activity are generated, conditional on the breaks in each of their respective equations and their wild bootstrapped shocks, and these values are employed in (7) when generating bootstrapped price data.

In addition to break date estimates, the methodology of Qu and Perron (2007) is used to obtain associated confidence intervals. Nevertheless, due to the iterative procedure employed, these should be taken as merely indicative.

In implementing this procedure, the maximum number of iterations is set to 40. It is, however, possible that the procedure may converge to a cycle of two or more sets of break dates, rather than a unique set. In such a case, we select among the sets in the cycle using the modification of the BIC criterion proposed by Hall, Osborn and Sakkas (2013) for structural break inference. In particular, we minimize the information criterion

$$BIC = \sum_{i=1}^{m+1} \ln(RSS/T) + K(l, m)\ln(T)/T \quad (8)$$

where RSS is the residual sum of squares of an individual SVAR equation obtained using T sample observations with m variance breaks and l breaks in the coefficients. The penalty term

¹³Hansen (2000) shows that a fixed bootstrap procedure works well for structural break testing in a single equation from a VAR system.

involves the total number of parameters (variance and coefficients) and the number of breaks, i.e. $K(l, m) = (m + 1) + p(l + 1) + 3(l + m)$, where p is the number of coefficients estimated in each coefficient regime. In particular, based on the arguments of Hall, Osborn and Sakkas (2013), the estimated number of breaks is weighted by 3 in this penalty function.

Although the structural breaks analysis is based on the SHVAR model which imposes the coefficient restrictions of (4) on the unrestricted SVAR (2), these restrictions are subject to test. More explicitly, to take account of any detected volatility breaks in the shocks of (5)-(7), the restrictions are tested equation by equation using GLS transformed data, with tests applied both over the whole sample and each of the potential regimes defined by the coefficient breaks. The test is a conventional F -test of the coefficient restrictions compared to (2), with a finite sample *i.i.d.* bootstrap employed for this purpose¹⁴. The bootstrap in all cases uses 10,000 replications.

5 Results

Subsection 5.1 examines possible time variation in the structural VAR parameters through the application of the procedure outlined in subsection 4.2, while subsections 5.2 and 5.3 examine the implications of the detected breaks through impulse responses and forecast error variance decompositions, respectively. As discussed in section 3, the analysis is based on the three oil market variables after (first) differencing, with means and deterministic seasonality removed, together with three outliers in the production series. To facilitate comparisons, all results are expressed in percentage terms.

5.1 Breaks in structural parameters

Owing to the large number of coefficients (namely 15 to 17 in each equation of the SHVAR), a maximum of five breaks is permitted in the coefficients of each equation ($M_c = 5$), with a minimum of 15% of the sample required in each regime. With this specification, the earliest and the latest break can happen at November 1980 and April 2008, respectively, and there has to be at least 5 years and 11 months between two breaks. With only one variance parameter, the maximum number of breaks is set at eight ($M_v = 8$), with a minimum of 10% of the total sample in each volatility regime. The results of the structural break tests are reported in Table 2, together with asymptotic critical values for a 5% level of significance in brackets¹⁵. The estimation of the VAR with long lags implies that two years of data are ‘lost’, so that the sample period for estimation starts in January 1975.

It may be noted that the iterative procedure converges quickly, after 3 iterations for the

¹⁴The asymptotic 5% critical value under homoskedasticity from $F(60, \infty)$ is 1.32.

¹⁵Qu and Perron’s (2007) test statistics have the same limit distributions as those in Bai and Perron (1998), who tabulate critical values up to 10 parameters. Although Bai and Perron (2003) provide response surfaces for estimating critical values, Hall and Sakkas (2013) show these surfaces can lead to misleading inferences with a large number of parameters. Therefore we simulate critical values suitable for our cases; details are available from the authors on request.

production equation, 2 for the economic activity and 4 for the real price, as can be seen from panel III of Table 2. Also note that 2 iterations here imply that convergence takes place immediately, the initialization and one subsequent iteration, as the procedure finds no coefficient break.

Panel I (A) shows that the $WDmax$ statistic for the oil production equation, at 255.42, easily rejects the null hypothesis of constant coefficients in relation to the asymptotic 5% critical value (38.82). The asymptotic sequential test does not reject the null of one break (against two), with the bootstrap procedure confirming the existence of this single break (p -value 0.1%). This break is dated at December 1980, with a tight 90% confidence interval, of October 1980 to February 1981. This break occurs early in the period of surplus capacity and competitive price cuts by OPEC members (see section 2), during which (as evident in Figure 1) production declined; the implications of this coefficient break form part of the impulse response functions and forecast error variance decompositions examined in the following subsections. The corresponding procedure for volatility (panel II (A)) finds two breaks, in 1990 and 2004, which are essentially those identified in the univariate analysis of subsection 3.1. These both mark declines in the volatility of oil production shocks; this pattern is in line with the findings of Baumeister and Peersman (2013b), whose sample ends in 2010.

The role of our finite sample bootstrap procedure to confirm asymptotically detected breaks is seen in the results for economic activity in panel I (A) of Table 2. Although the asymptotic $WDmax$ statistic (206.82 compared with a critical value of 40.86) strongly rejects the null hypothesis of no break in the coefficients, and the asymptotic sequential test indicates one break, the bootstrap test delivers a p -value of nearly 25% at the estimated date, and hence does not confirm the existence of this break¹⁶. Based on the Monte Carlo experiments of Bataa, Osborn, Sensier and van Dijk (2013) that finds bootstrap inference to be more reliable in finite samples, we conclude there is no break in the coefficients of (6). The contemporaneous coefficient a_{21} of Panel I (B) indicates that activity has a negative contemporaneous estimated point response to increases in oil production throughout the period, although no evidence on the significance of this coefficient is provided. The volatility breaks asymptotically detected for activity in 1979 and 2008, shown in panel II (A), are robust to the bootstrap test. The first of these marks the end of the turbulence of the 1970s, resulting in a volatility reduction, and the latter the onset of the GFC. Although Lütkepohl and Netšunajev (2014) also find that economic activity experiences a general volatility reduction over their sample period to 2007, the Markov switching volatility model they employ applies to the system and the switch is estimated to occur later (in 1987). The flexibility of our procedure in allowing volatility changes to occur at different dates over equations is evident here.

In contrast to the stability of the coefficients of the oil production and economic activity equations (the former from 1980), oil price inflation responses exhibit breaks in May 1988 and October 1994. Notice from panel I (A) that the use of asymptotic critical values indicates four breaks, but the bootstrap procedure finds one of these to have a p -value of 85%, and deletion

¹⁶The deleted coefficient break is dated at April 2008. If this break is included, the first volatility break remains unchanged, while the second is dated four months earlier.

of this finds a further break that is not significant (p -value of 83%). In addition, three breaks are uncovered in the volatility of shocks¹⁷. In line with Baumeister and Peersman (2013b), the coefficient and volatility breaks detected in (7) point to important changes taking place in the determination of oil prices over the last two decades of the twentieth century. However, our results imply that breaks are not confined to volatility shifts, as assumed in the model of Lütkepohl and Netšunajev (2014). Competition amid declining prices, not only among OPEC member countries such as Iraq and Iran but also non-OPEC countries, led to precautionary demand shocks of relatively small magnitude until the near-collapse of OPEC in 1986, with volatility being particularly small in the first half of that decade in relation to subsequent periods. Indeed, we find both coefficient and volatility breaks occur between 1986 and 1988, and also in the mid- to late-1990s, with the 1990s breaks possibly associated with the development of the oil futures market.

As anticipated for a model capturing demand and supply, positive oil production shocks immediately lead to lower prices (a_{31}) while prices increase in response to an aggregate demand shock (a_{32}); see the contemporaneous coefficients of panel I (B). However, with prices early in the period set in advance (see section 2), effectively no contemporaneous price response to a production shock is found until the latter part of the 1980s. The extent of the price response to a production change is tempered from the mid-1990s, with the estimated contemporaneous response declining from -1.81 to -0.42. On the other hand, the contemporaneous oil price response to an aggregate demand shock is reduced over the middle coefficient regime (the end of the 1980s to the mid-1990s), but is then essentially restored to its value as in the early part of the sample period. Therefore, at least as seen in the contemporaneous coefficients, oil prices were largely unresponsive to demand changes in the early 1990s, with the restoration of the link possibly associated with the rise of China and other large emerging economies. These are, however, only point estimates of the contemporaneous responses, with a fuller discussion of dynamic effects in the following subsection.

Finally, panel IV of Table 2 provides evidence that the structural coefficient restrictions of (4) in the SHVAR form, examined equation by equation, are compatible with the data. This is the case whether the tests are applied to the whole sample of data, assuming no breaks in coefficients, or within the identified coefficient regimes. Note, however, that the test statistic cannot be calculated in some regimes, due to the large number of coefficients in the unrestricted equation compared with the number of observations in the relevant sub-sample.

5.2 Impulse responses

Impulse response functions computed across sub-samples, as in Inoue and Rossi (2011), provide a quantitative comparison of time-variation in the coefficients of the model. In our case, four regimes are given by the breaks identified in panel I of Table 2, namely January 1975 to December

¹⁷Including the four asymptotic coefficient breaks would lead to omission of the 1981 volatility break, illustrating the role of iteration between coefficient and volatility breaks.

1980, January 1981 to May 1988, June 1988 to October 1994, and November 1994 to February 2014. Although the first break arises from the production equation and the remaining two from the oil price equation, all affect the dynamic responses of the system. Figures 3 to 5 show the estimated impulse responses associated with each of these regimes, together with (in the first column) the impulse responses implied by the SHVAR model estimated over the full sample with constant coefficients. Clearly, the transmission of shocks is obtained from the SVAR coefficients, and hence estimated responses to a given shock are unaffected by volatility changes. Indeed, the breaks identified in Table 2 imply that responses have been unchanged since the mid-1990s and hence unaffected by the GFC. As in Kilian (2009), each shock is normalized such that an increase in the oil price is anticipated, and the magnitude of the shock applied across all regimes is equal to one standard deviation of the shock as estimated using the whole sample with no breaks. Further, although our SHVAR is estimated in differences, the cumulated impulse responses are shown; the responses we present are therefore comparable to those of Kilian (2009, Figure 3).

Indeed, the patterns of responses seen in the first columns of Figures 3 to 5 are qualitatively very similar to those reported by Kilian (2009), so that our differencing of the economic activity measure and log real oil prices has little effect on these responses. Perhaps the most marked difference in our results compared with Kilian (2009) is that real activity responds positively to an oil-specific demand shock for only six months after the shock and is significant (compared to the one standard deviation band) for only three months in our extended sample.

To facilitate comparison, not only is a shock of constant magnitude applied across regimes, but (with a few exceptions) a common vertical scale is employed. As in Kilian (2009) one and two-standard deviation confidence bands (approximate 66% and 95% confidence intervals) are obtained using a recursive-design wild bootstrap procedure, which employs a bootstrap data generation process with parameters as estimated in the SHVAR with breaks. Using the wild bootstrap to take account of volatility changes, the SHVAR parameters are re-estimated within each replication over the sub-periods given by the (equation-specific) coefficient break dates, which are treated as known. This leads to the distribution of impulse responses summarized through the confidence bands in the figures.

Interesting differences apply across the identified subperiods in Figure 3, which tracks the impact of a supply shock. In particular, the responses in the first regime (to 1980) indicate that the initial production loss is almost wiped out after 2 years, implying that the shock is not permanent and production eventually returns to its trend level. After 1980, the shock leads to a permanent production loss, albeit estimated to be less severe than the magnitude of the initial shock. Oil supply shocks have relatively little effect on real activity, with these never significant according to the two standard error confidence bands.

Negative supply shocks¹⁸ lead to increases in real oil prices, often after a delay, but these effects are short-lived in most sub-periods. The pattern is particularly notable in the regime extending from June 1988 to October 1994, the period after the near-collapse of OPEC in 1986.

¹⁸Note that a different scale is used when plotting the price responses for the last two sub-samples as compared to other periods.

Here the price response is fast, with prices significantly increasing in the month of the shock. However, even in this regime, the real price of oil subsequently falls and is depressed after a delay of around six months. Although the shock represents bad news, excess capacity was then at a record high level. Thus, a production disruption causes prices to increase immediately, but a relatively quick recovery in production is also possible. The current regime, from 1994, shows a distinctive price response to the supply shock compared with earlier sub-periods. In particular, the effect of the shock is persistent (although not always statistically significant), with all estimated responses positive at all horizons examined. This pattern is similar to that obtained from the whole sample estimates, but not typical of the entire period when breaks are taken into account.

The transmission mechanism for the effects of aggregate demand shocks has also changed; see Figure 4. In particular, aggregate demand shocks have more immediate and longer lasting effects on the real price of oil after 1994 than previously, with these being statistically significant and leading to persistent effects in the final regime. As in Figure 3 (and also Figure 5), the period between 1988 and 1994 contrasts with other sub-periods for the effects of shocks on oil prices. In particular, after a delay of a few months, positive demand shocks are found to lead to significant (and perverse) real price declines. The negative response of oil production to this shock in the earliest sub-period also appears perverse, perhaps indicating that production was largely set in the light of political rather than market considerations during the 1970s¹⁹. This pattern carries over at longer lags to the full sample estimates without breaks, illustrating the danger of not recognizing changes in the market over these four decades. The results of Kilian (2009) also imply that production does not respond significantly to aggregate demand shocks. Allowing for breaks, however, Figure 4 shows that aggregate demand shocks result in positive and significant production responses, as anticipated, from 1981 onwards.

Finally, Figure 5 shows the responses to an oil-specific demand shock for each variable in the system. After a short-lived positive effect, the oil-specific demand shock depresses oil production before 1980; otherwise, the production responses are generally not significant according to the two standard error bands, in line with Kilian (2009). The response of aggregate activity to an oil-specific demand shock response varies relatively little over time (except for the 1988-1994 subperiod), with a short-run positive response followed by decline. The final regime, from 1994, however, sees a different response of the price of oil to this shock than earlier periods. Whereas the price effect is persistent in the sub-periods to 1988, and to a lesser extent between 1988 and 1994, the post-1994 regime sees the effect of the shock effectively disappear after two years.

¹⁹Although it falls outside our estimation period, this is illustrated by the Organisation of Arab Petroleum Exporting Countries imposing an embargo on exports to the US and other countries in response to their supplying Israel with arms during the Yom Kippur War.

5.3 Forecast error variance decompositions

To shed further light on the nature of changes in the oil market over the four decades of our sample, Figure 6 plots forecast error variance decompositions (FEVDs) for each variable assuming time-invariant SHVAR parameters (coefficients and volatilities), while Figure 7 takes account of the breaks identified in Table 2. Note that volatility shifts contribute to the patterns seen in Figure 7. In particular, the SHVAR coefficients are constant from late 1994, and hence changes after that date in Figure 7 are due entirely to changes in the volatilities of oil market shocks. The decompositions apply to the differenced variables as employed in the SVAR, and hence examine the effects of the different oil market shocks on forecast errors for future changes in the respective variables. The graphs employ a combination of line type and colour coding. The four colours indicate six months intervals, with green being up for a forecast horizon to 6 months, followed by red (7 to 12 months), blue (13 to 18 months) and cyan (19 to 24 months). Within each interval, six line types indicate shorter to longer forecast horizons.

If parameter constancy and homoskedastic structural shocks are assumed (Figure 6), demand shocks (aggregate together with oil-specific) clearly dominate forecast error variances for future changes in oil production. Supply shocks explain about 20% at short horizons, but their contribution quickly reduces to zero as the horizon increases. Aggregate demand shocks contribute most to the short-run error variance, but their role switches with oil specific demand shocks as the forecast horizon increases. While aggregate demand shocks also dominate the picture for real activity, the share of oil specific demand shocks increases to 30% at a two year horizon. However, the forecast error variance of the real price of oil is almost entirely explained by its own shocks, with aggregate demand and supply shocks playing little role for changes in the oil price in the constant parameter specification at any horizon considered.

The implications of the SVAR model with breaks in Figure 7 are quite different²⁰. While demand shocks continue to explain much of the forecast error variance of oil production at all horizons, nevertheless supply shocks play a substantial role until around 1990. As late as the aftermath of OPEC's near-collapse in 1986, these contribute almost 40% of the short-run forecast error variance for oil production. Kilian (2008) estimates that oil producers had capacity utilization rates near 90% (effectively operating at full capacity) until 1979, then precipitously falling to record low levels, reaching below 65% in OPEC and below 85% worldwide in 1985. In times of abundant oil supply with relatively low volatility and competitive pricing from the 1990s, oil-specific demand shocks exert the strongest influence, with aggregate demand also important during the volatile GFC period.

Indeed, the model implies that oil-specific demand shocks explain most of the forecast error variances for all variables in the two decades marking the run-up to the GFC. In the current slack economic and high volatility environment, however, aggregate demand shocks explain almost 40% of the forecast variance error of production, more than 40% for real activity and more than

²⁰Note that the 'spikes' apparent in this graph can result when there are small differences between the coefficient and volatility break dates as estimated in Table 2.

10% for the real oil price, all of which are higher than during any sub-period after the OPEC near-collapse of 1986.

For real activity and the real price of oil, the forecast horizon has a more important effect in Figure 7 compared with the stable model of Figure 6. In the Great Moderation period from the mid-1980s to essentially the GFC, aggregate demand shocks explain at most 60% of the variation in forecasts errors of real activity within 6 months, decreasing to 20% at a horizon of more one year. However increases in the volatility of oil-specific demand shocks in the otherwise tranquil period leading up to the GFC (Table 2) causes their contribution to real activity forecast error variances to rise to around 90% at long horizons. This is consistent with Hamilton (2009), who argues that developments in this market may have triggered the 2008 US recession. Nevertheless, as pointed out by Baumeister and Peersman (2013b), in this model with no other forward-looking variables, such shocks are not necessarily associated with the oil market alone, but can be driven by revisions to expectations about aggregate activity.

Lütkepohl and Netšunajev (2014) provide a FEVD analysis of the real price of oil obtained from their Markov-switching volatility model, for which their state 1 is essentially the period from around 1986 to the end of their sample in 2007 while state 2 is the predominant regime over 1975 to 1986. The implication from Figure 7 that oil specific shocks explain almost all of the forecast error variance for the real price of oil from the mid-1980s to the GFC, with aggregate demand shocks playing a substantial role in the earlier period, agrees with their findings. However, our results effectively decompose the earlier period into more regimes (associated with coefficient as well as volatility breaks) and show that the role of aggregate activity is highest in the 1970s and declines in the first half of the 1980s, prior to its very small role for oil price changes over the two decades of the Great Moderation. Our extended sample also shows that the role of real activity is re-asserted when its volatility rises as a consequence of the economic disruption associated with the GFC.

6 Conclusions

This paper extends the structural break methodology of Bataa, Osborn, Sensier and van Dijk (2013), which iteratively tests for coefficient and covariance breaks that may occur at different time points, to examine changes in an SVAR model of the global oil market. The extension relies on single equation methods and can be employed in any structural model identified through contemporaneous causal ordering restrictions.

Changes in the oil market over the four decades from the mid-1970s are examined through the model of Kilian (2009), which disentangles demand and supply shocks using a three equation SVAR representing oil supply, aggregate demand and oil-specific demand. We find strong statistical evidence that both the transmission mechanism and the volatilities of shocks have changed. Baumeister and Peersman (2013a, 2013b) employ a continuous time-varying SVAR model in order to take account of changes in both structural coefficients and volatilities in this

market, while Lütkepohl and Netšunajev (2013) allow for volatility (but not coefficient) changes through a Markov switching specification. Although our results are broadly in line with the implications of these studies, our use of formal structural break tests allows us to focus more explicitly on whether change has occurred and the nature of any such change. In particular, while the coefficients of the oil supply and oil-specific demand equations both exhibit changes over the two decades to 1994, no subsequent coefficient breaks are uncovered.

In line with both Baumeister and Peersman (2013b) and Lütkepohl and Netšunajev (2013), we find breaks in the volatilities of shocks to be a feature of this market. Whereas volatility changes for aggregate demand shocks can be associated with the Great Moderation and the Global Financial Crisis, those in oil-specific demand shocks appear to be distinct from these features. Further, the volatility of oil production shocks has been at a historic low since 2004, making supply more predictable.

Recognition of breaks leads to time variation in the impulse responses and forecast error variance decompositions for the three variables of the oil market model. A key finding relates to the responses of oil production to an aggregate demand shock. In particular, a constant parameter specification implies that oil production is unresponsive to an aggregate demand shock for up to a year, followed by a medium run decline. However, taking account of breaks (particularly in the coefficients of the oil supply equation) shows responses to be positive and significant from 1981 onwards. Partly because production has become more predictable due to declines in its volatility, oil supply shocks contribute little to the forecast error variance for any variable in the system from 1995, but these do play a role prior to that date.

The real price of oil responds to aggregate demand shocks more strongly (and with greater statistical significance) from the mid-1990s onwards, at the same time becoming less persistent in response to oil specific demand shocks. In terms of its forecast error variance, however, aggregate demand shocks play very little role over the low volatility period of the Great Moderation although they are important until 1985 and re-emerge during the post-GFC period. Nevertheless, it is notable that no changes in SVAR coefficients, and hence no changes in impulse responses for the oil market model, are uncovered that can be attributed to the GFC.

7 References

- Bagliano, F.C. and Favero, C.A. 1998. Measuring monetary policy with VAR models: An Evaluation. *European Economic Review* 42(6), 1069-1112.
- Bai, J. and Perron, P. 1998. Estimating and testing linear models with multiple structural changes. *Econometrica* 66, 47-78.
- Bai, J. and Perron, P. 2003. Critical values for multiple structural change tests. *Econometrics Journal* 6, 72-78. (unpublished table Appendix available at <http://people.bu.edu/perron/>).
- Balin, B.J. 2010. The impact of the global economic crisis on sovereign wealth funds. *Asian-Pacific Economic Literature* 24(1), 1-8.
- Bandi, F.M., Russell, J.R. and Yang, C. 2013. Realized volatility forecasting in the presence of time-varying noise, *Journal of Business and Economic Statistics* 31(3), 331-345.
- Bataa, E., Osborn, D.R., Sensier, M., and van Dijk, D. 2013. Structural breaks in the international dynamics of inflation. *Review of Economics and Statistics* 95(2), 646-659.
- Bataa, E., Osborn, D.R., Sensier, M., and van Dijk, D. 2014. Identifying changes in mean, seasonality, persistence and volatility for G7 and Euro Area inflation. *Oxford Bulletin of Economics and Statistics* 76(3), 360-388.
- Baumeister, C. and Peersman, G. 2013a. Time-varying effects of oil supply shocks on the U.S. economy. *American Economic Journal: Macroeconomics* 5(4), 1-28.
- Baumeister, C. and Peersman, G. 2013b. The role of time-varying price elasticities in accounting for volatility changes in the crude oil market. *Journal of Applied Econometrics* 28(7), 1087-1109.
- Baumeister C., Peersman G. and Van Robays, I. 2010. The economic consequences of oil shocks: differences across countries and time. In *Inflation in an Era of Relative Price Shocks*, Fry, J. and Kent (eds), Reserve Bank of Australia, 91-128.
- Bernanke, B.S. 1983. Irreversibility, uncertainty, and cyclical investment. *Quarterly Journal of Economics* 98(1), 85-106 .
- Bernanke, B.S. and Mihov, I. 1998. Measuring Monetary Policy. *Quarterly Journal of Economics* 113(3), 869-902.
- Blanchard, O.J. and Riggi, M. 2013. Why are the 2000s so different from the 1970s? A structural interpretation of changes in the macroeconomic effects of oil prices in the US. *Journal of the European Economic Association* 11(5), 1032-1052.
- Corsi, F. 2009. A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics* 7(2), 174-196.
- Davis, S.J. and Haltiwanger, J. 2001. Sectoral job creation and destruction responses to oil price changes. *Journal of Monetary Economics* 48, 465-512.
- Davidson R. and Flachaire E. 2008. The Wild bootstrap tamed at last. *Journal of Econometrics* 146(1), 162-169.
- De Gregorio, J., Landerretche, O., Neilson, C., Broda, C., and Rigobon, R. 2007. Another pass-through bites the dust? oil prices and inflation. *Economia* 7(2), 155-208.
- Diebold, F. X. and Chen, C. 1996. Testing structural stability with endogenous breakpoint: A size comparison of analytic and bootstrap procedures. *Journal of Econometrics* 70(1), 221-241.
- Fry, R. and Pagan, A. 2011. Sign restrictions in structural vector autoregressions: a critical review. *Journal of Economic Literature* 49, 938-960.

- Gallo, A., Mason, P., Shapiro, S. and Fabritius, M. 2010. What is behind the increase in oil prices? Analyzing oil consumption and supply relationship with oil price. *Energy* 35, 4126-4141.
- Godfrey L.G, Orme C.D. 2004. Controlling the finite sample significance levels of heteroskedasticity-robust tests of several linear restrictions on regression coefficients, *Economics Letters* 82, 281-287.
- Hall, A. and Sakkas, N. 2013. Approximate p-values of certain tests involving hypotheses about multiple breaks. *Journal of Econometric Methods* 2(1), 53-67.
- Hall, A., Osborn, D.R., and Sakkas, N. 2013. Inference on structural breaks using information criteria. *Manchester School* 81(S3), 54-81.
- Hamilton, J.D. 2009. Causes and consequences of the oil shock of 2007-08. *Brookings Papers on Economic Activity* 40(1), 215-283.
- Hamilton, J.D. 2014. The changing face of world oil markets. Mimeo, University of California San Diego.
- Hamilton, J.D. and Herrera, A.M. 2004. Oil shocks and aggregate macroeconomic behavior: the role of monetary policy. *Journal of Money, Credit and Banking* 36(2), 265-286.
- Hamilton, J.D. and Wu, C. 2014. Risk premia in crude oil futures prices. *Journal of International Money and Finance* 42(1), 9-37.
- Hansen, B.E. 2000. Testing for structural breaks in conditional models. *Journal of Econometrics* 97, 93-115.
- Helbling 2013. On the edge. *Finance and Development* 50(3), 16-18.
- Inoue, A. and Rossi, B. 2011. Identifying the sources of instabilities in macroeconomic fluctuations. *Review of Economics and Statistics* 93(4), 1186-1204.
- International Energy Agency. 1996. Global supply, demand seasonality and the role of stocks. Monthly Oil Market Report. October.
- International Monetary Fund. 2011. Oil scarcity, growth and global imbalances. Chapter 3 in World Economic Outlook, April.
- Jo, S. 2014. The effect of oil price uncertainty on global real economic activity. *Journal of Money, Credit and Banking*, 46(6), 1113-1135.
- Kejriwal, M. and Perron, P. 2010. A sequential procedure to determine the number of breaks in trend with an integrated or stationary noise component. *Journal of Time Series Analysis* 31, 305-328.
- Kilian, L. 2008. Exogenous oil supply shocks: How big are they and how much do they matter for the U.S. economy? *Review of Economics and Statistics* 90(2), 216-240.
- Kilian, L. 2009. Not all oil price shocks are alike: disentangling demand and supply shocks in the crude oil market. *American Economic Review* 99(3), 1053-1069.
- Kilian, L. and Lewis, L.T. 2011. Does the FED respond to oil price shocks? *Economic Journal* 121, 1047-1072.
- Kilian, L. and Murphy, D.P. 2012. Why agnostic sign restrictions are not enough: Understanding the dynamics of oil market VAR model. *Journal of the European Economic Association* 10(5), 1166-1188.
- Kilian, L. and Murphy, D.P. 2014. The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics* 29(3), 454-478.
- Koop G., Leon-Gonzalez, R. and Strachan R. W. 2009. On the evolution of the monetary policy transmission mechanism. *Journal of Economic Dynamics and Control* 33, 997-1017.
- Lilien, D.M. 1982. Sectoral shifts and cyclical unemployment *Journal of Political Economy* 90(4),

777-793.

- Lutkepohl, H. and Netsunajev, 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(3), 479-496.
- Mabro, R. 2006. *Oil in the Twenty-First Century: Issues, Challenges, and Opportunities*. Oxford University Press.
- Moosa, I.A. 1995. Deterministic and stochastic seasonality in the oil imports of the Group of Seven countries. *OPEC Energy Review* 19(3), 181-201.
- Peersman, G. and Van Robays, I. 2009. Oil and the Euro are economy. *Economic Policy*, October, 603-651.
- Peersman, G. and Van Robays, I. 2012. Cross country differences in the effects of oil shocks. *Energy Economics* 34, 1532-1547.
- Perry, M.J. 2012. Manufacturing's declining share of GDP is a global phenomenon, and it's something to celebrate. US Chamber of Commerce blog. <http://emerging.uschamber.com/blog/2012/03/manufacturing%E2%80%99s-declining-share-gdp>
- Prodan, R. 2008, Potential pitfalls in determining multiple structural changes with an application to purchasing power parity. *Journal of Business and Economic Statistics* 26, 50-65.
- PwC. 2011. The impact of Sovereign Wealth Funds on economic success. www.pwc.com/cy/en/Issues/Assets/sovereign-funds.pdf.
- Qu, Z.J. and Perron, P. 2007. Estimating and testing structural changes in multivariate regressions. *Econometrica* 75, 459-502.
- Stopford, M.C. 2009. *Maritime Economics*, 3rd edition. Routledge,
- Wang, C.S.-H., Bauwens, L. and Hsiao, C. 2013. Forecasting a long memory process subject to structural breaks. *Journal of Econometrics* 177, 171-184.

Table 1. Univariate iterative decomposition

	Oil Production	Economic Activity	Real Oil Price
Level component			
Number of breaks	0	0	0
Seasonality			
Number of breaks	0	0	2
Break dates			1986.01 (1983.08-1988.06) 1998.11 (1997.07-2000.03)
Outliers			
Dates	1975.12, 1977.03, 1978.03	None	None
Dynamics			
AR order	0	2	1
Number of breaks	0	0	0
Persistence	0.05	0.17	0.46
Volatility			
Number of breaks	2	2	2
Break dates	1990.11 (1988.05-1991.02) 1980.06 (1974.08-1981.01) 1977.04 (1974.03-1977.12) 2004.10 (2000.12-2005.12) 2008.10 (2008.09-2011.06) 2009.12 (2005.02-2011.07)		
Standard deviations			
Across regimes	1.87, 0.90, 0.60	6.33, 3.94, 12.12	4.33, 2.17, 7.07, 4.85
[Ignoring breaks]	[1.36]	[6.10]	[5.85]
Number of iterations for convergence			
Main loop	2	2	5
Dynamics sub-loop	2	2	2

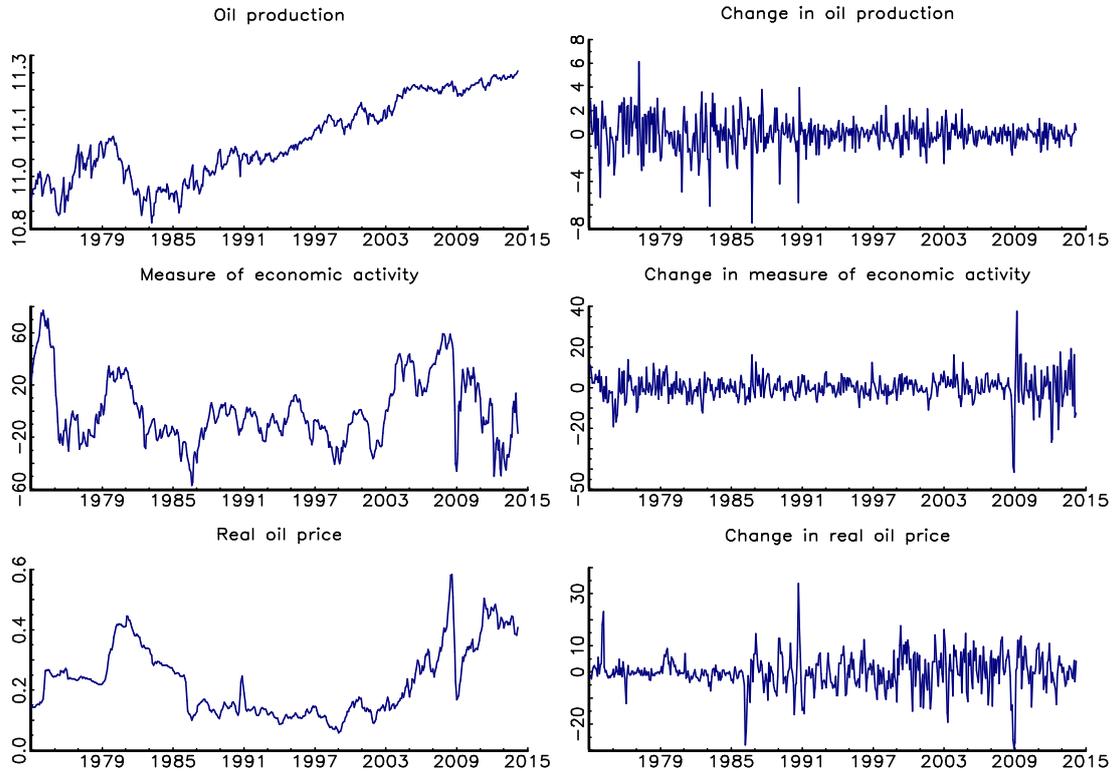
Notes: Decomposition as in (1) using the iterative method of Bataa *et al.* (2014), with breaks detected using the Qu and Perron (2007) test using a 5% significance level (trimming 15% and maximum of 5 breaks, except for the seasonal component with 20% trimming and maximum of 2 breaks). Outliers are defined as being 7 times the interquartile range from the median after removal of level and seasonality breaks. The AR order of the dynamic component is selected according to the Hannan-Quinn (HQ) information criterion, and used at entry to the dynamic/volatility sub-loop; an AR(1) is employed if 0 lags are selected by HQ. Persistence is computed as the sum of the estimated AR coefficients within the relevant dynamic regime, with the standard deviation of disturbances also shown within volatility regimes. When breaks are detected, the estimated breaks dates are shown in **bold** together with their asymptotic 90% confidence intervals in parentheses. Finally, the numbers required to achieve convergence of the loops of the algorithm are shown.

Table 2. Coefficient and variance break test results for the oil market SHVAR

	Oil Production	Economic Activity	Real Oil Price
<i>I (A). Coefficients breaks</i>			
Overall test	255.42* [38.82]	206.82* [40.86]	243.78* [42.61]
Seq(2/1)	33.26 [37.63]	36.71 [39.37]	51.29* [41.22]
Seq(3/2)			57.51* [42.59]
Seq(4/3)			57.51* [43.89]
bootstrap <i>p</i> -value for each break	(0.1)	(24.47)	(85.42) (83.34) (1.08) (1.42)
Break 1	1980.12 [80.10-81.02]		1988.05 [88.02-88.08]
Break 2			1994.10 [94.07-95.01]
<i>I (B). Contemporaneous coefficients</i>			
	n.a.	a_{21}	a_{31} a_{32}
Across regimes [Ignoring breaks]	-	-0.17 [-0.17]	0.06, -1.81, -0.42 0.17, 0.03, 0.20 [-0.28] [0.16]
<i>II (A). Variance breaks</i>			
Overall test	137.75* [10.67]	171.81* [10.67]	126.69* [10.67]
Seq(2/1)	17.11* [10.97]	24.06* [10.97]	21.79* [10.97]
Seq(3/2)	9.89 [11.88]	5.34 [11.88]	21.63* [11.88]
Seq(4/3)			12.14 [12.49]
bootstrap <i>p</i> -value for each break	(0.0) (0.31)	(0.04) (0.0)	(0.40) (0.0) (0.15)
Break 1	1990.10 [88.02-91.01]	1979.09 [73.04-80.02]	1981.03 [79.06-81.11]
Break 2	2004.09 [99.08-06.04]	2008.09 [08.08-11.06]	1986.02 [86.01-86.12]
Break 3			1998.09 [98.01-09.11]
<i>II (B). Shock variances</i>			
	σ_{oils}^2	σ_{aggd}^2	σ_{oild}^2
Across regimes [Ignoring breaks]	2.79, 0.74, 0.36 [1.71]	38.56, 15.56, 141.10 [35.76]	7.05, 2.53, 21.50, 45.17 [31.43]
<i>III. Number of iterations required to converge</i>			
	3	2	4
<i>IV. F statistic (bootstrap p-value) for SHVAR restrictions</i>			
Regime I	n.c.	1.28 (14.5)	1.08 (33.8)
Regime II	0.92 (72.6)	-	n.c.
Regime III	-	-	1.56 (19.5)
All sample	0.94 (57.2)	1.28 (14.5)	1.41 (11.15)

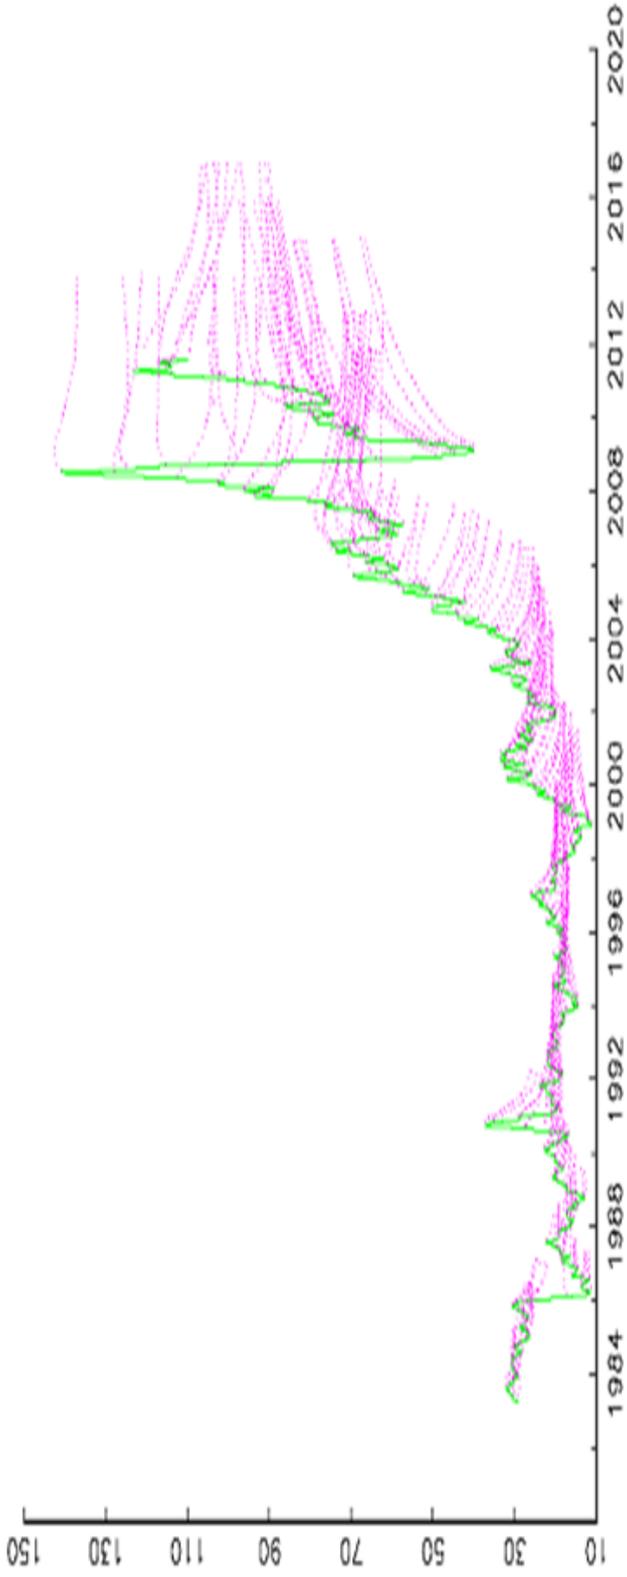
Notes: Values reported in panels I (A) and II (A) are at convergence of the iterative procedure of Bataa, Osborn, Seniser and van Dijk (2013). The overall test examines the null hypothesis of no break against an unknown number of breaks, to a maximum of 5 breaks for each SHVAR equation and 8 for the variance. If the overall statistic is significant at 5%, sequential tests are applied starting with the null hypothesis of one break and continuing until the relevant statistic is not significant. Asymptotic critical values for the 5% significance level are reported next to respective test statistics [in square brackets]. * indicates the statistic is significant at 5%. The estimated break dates and their 90% confidence intervals [in square brackets] are also reported. Bootstrap *p*-values corresponding to the null hypothesis that an asymptotically detected break does not exist are also reported in brackets. Panels I (B) and II (B) show estimated contemporaneous coefficients and shock variances respectively across sub-samples defined by the break dates. Also shown in square brackets are those quantities that ignore the breaks. Panel III reports the number of iterations required to converge in coefficient and volatility break dates. The last panel reports on the validity of the SHVAR restrictions by *F* tests over the sub-samples defined by the mean break dates and over the whole sample. The asymptotic 5% critical value from $F(60, \infty)$ is 1.32. n.c. indicates it was impossible to compute the *F* statistic over the sub-sample due to non-invertability issue (too many parameters with not long enough sub-sample).

Figure 1. Data



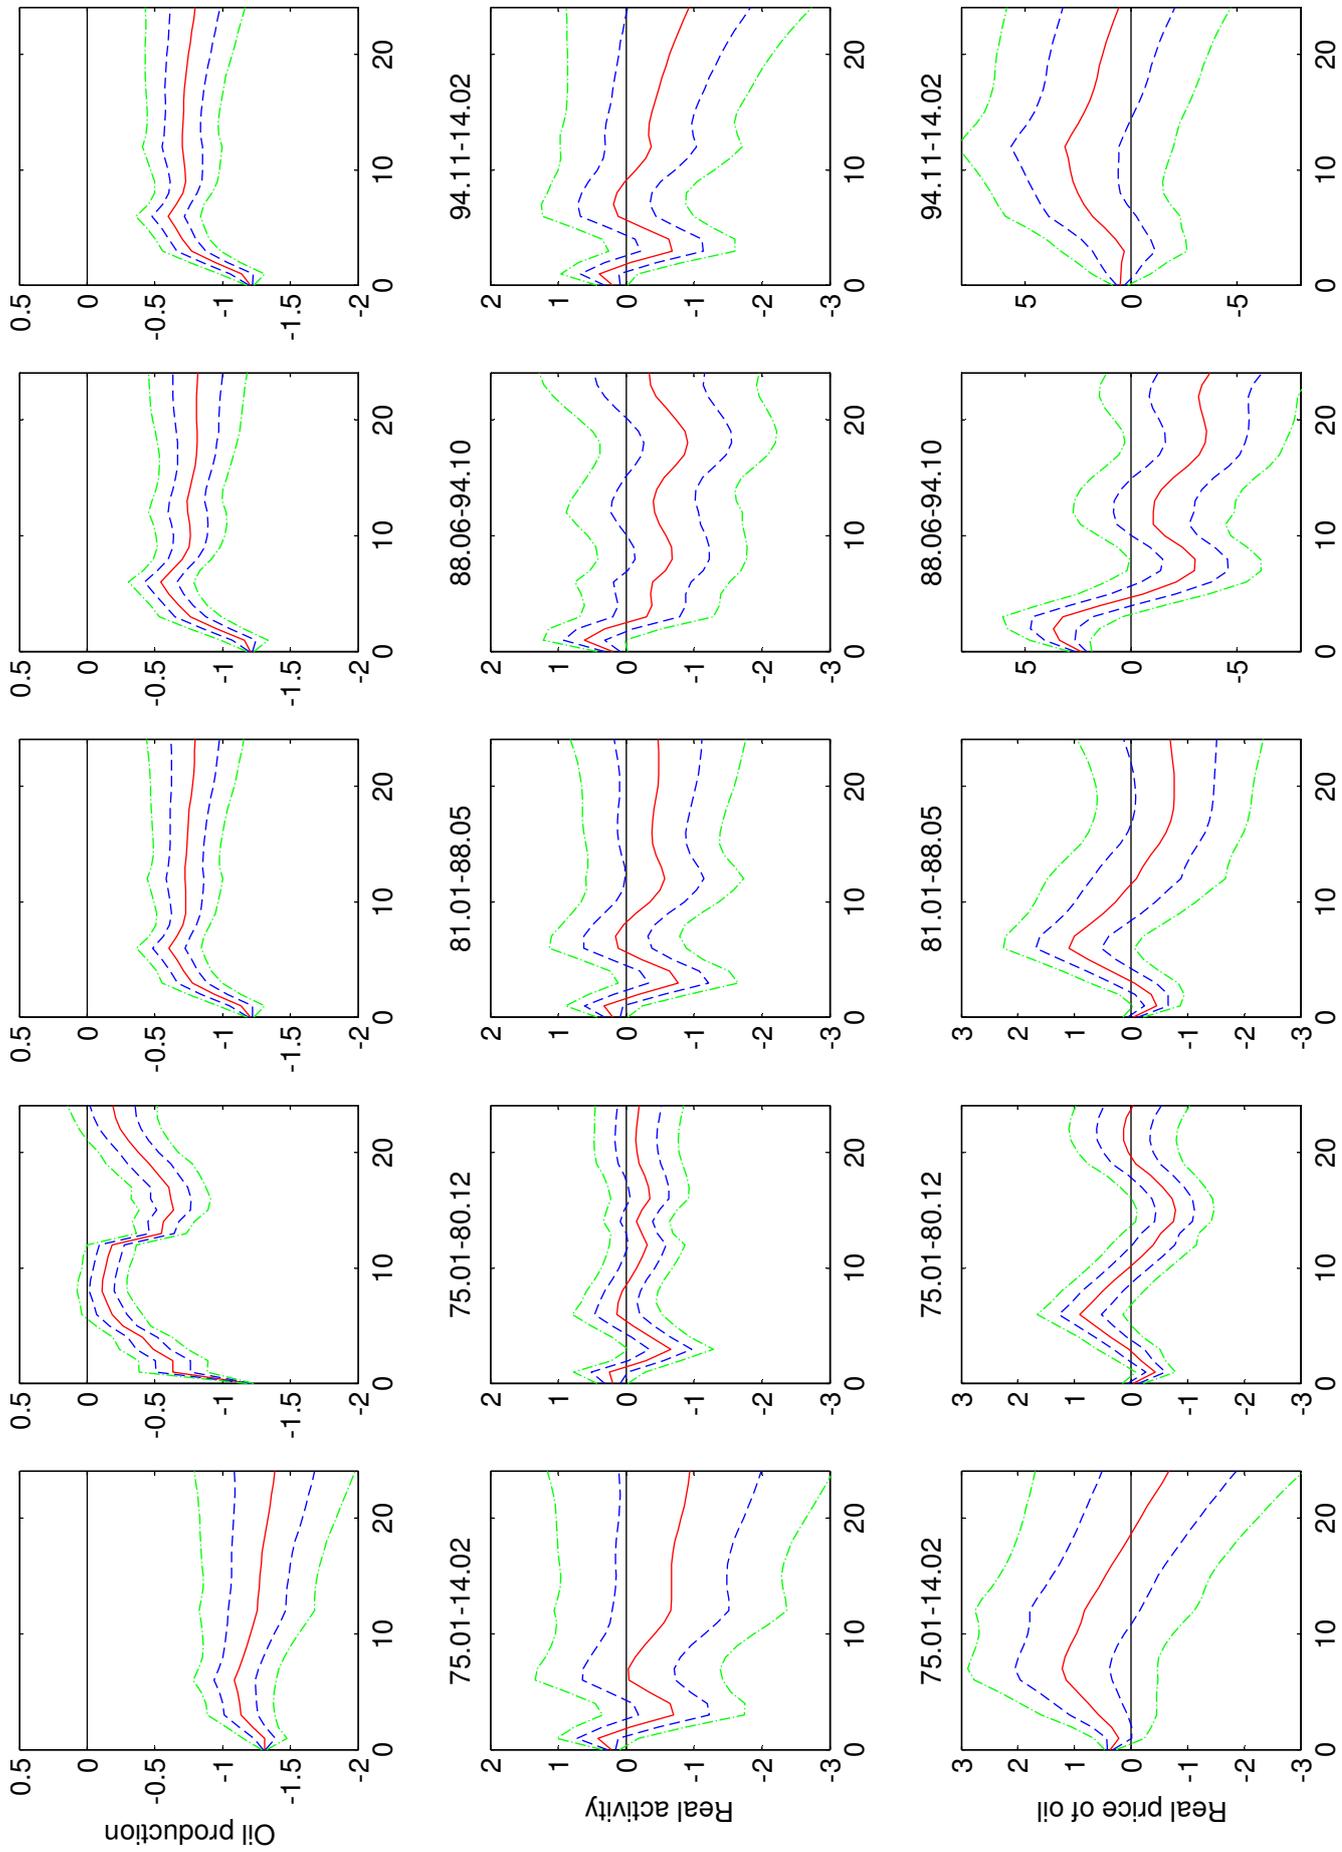
Notes: In the first column, oil production is the logarithm of the number of barrels produced per day, the measure of economic activity is the detrended series of Kilian (2009), downloaded from <http://www-personal.umich.edu/~lkilian/>, the real oil price is the logarithm of the price measured in US dollars and deflated by the US CPI. The data in panel b are differenced (oil production and real oil price are also multiplied by 100), with deterministic seasonality and outliers removed; see Section 3 of the text.

Figure 2. Spot and futures oil prices



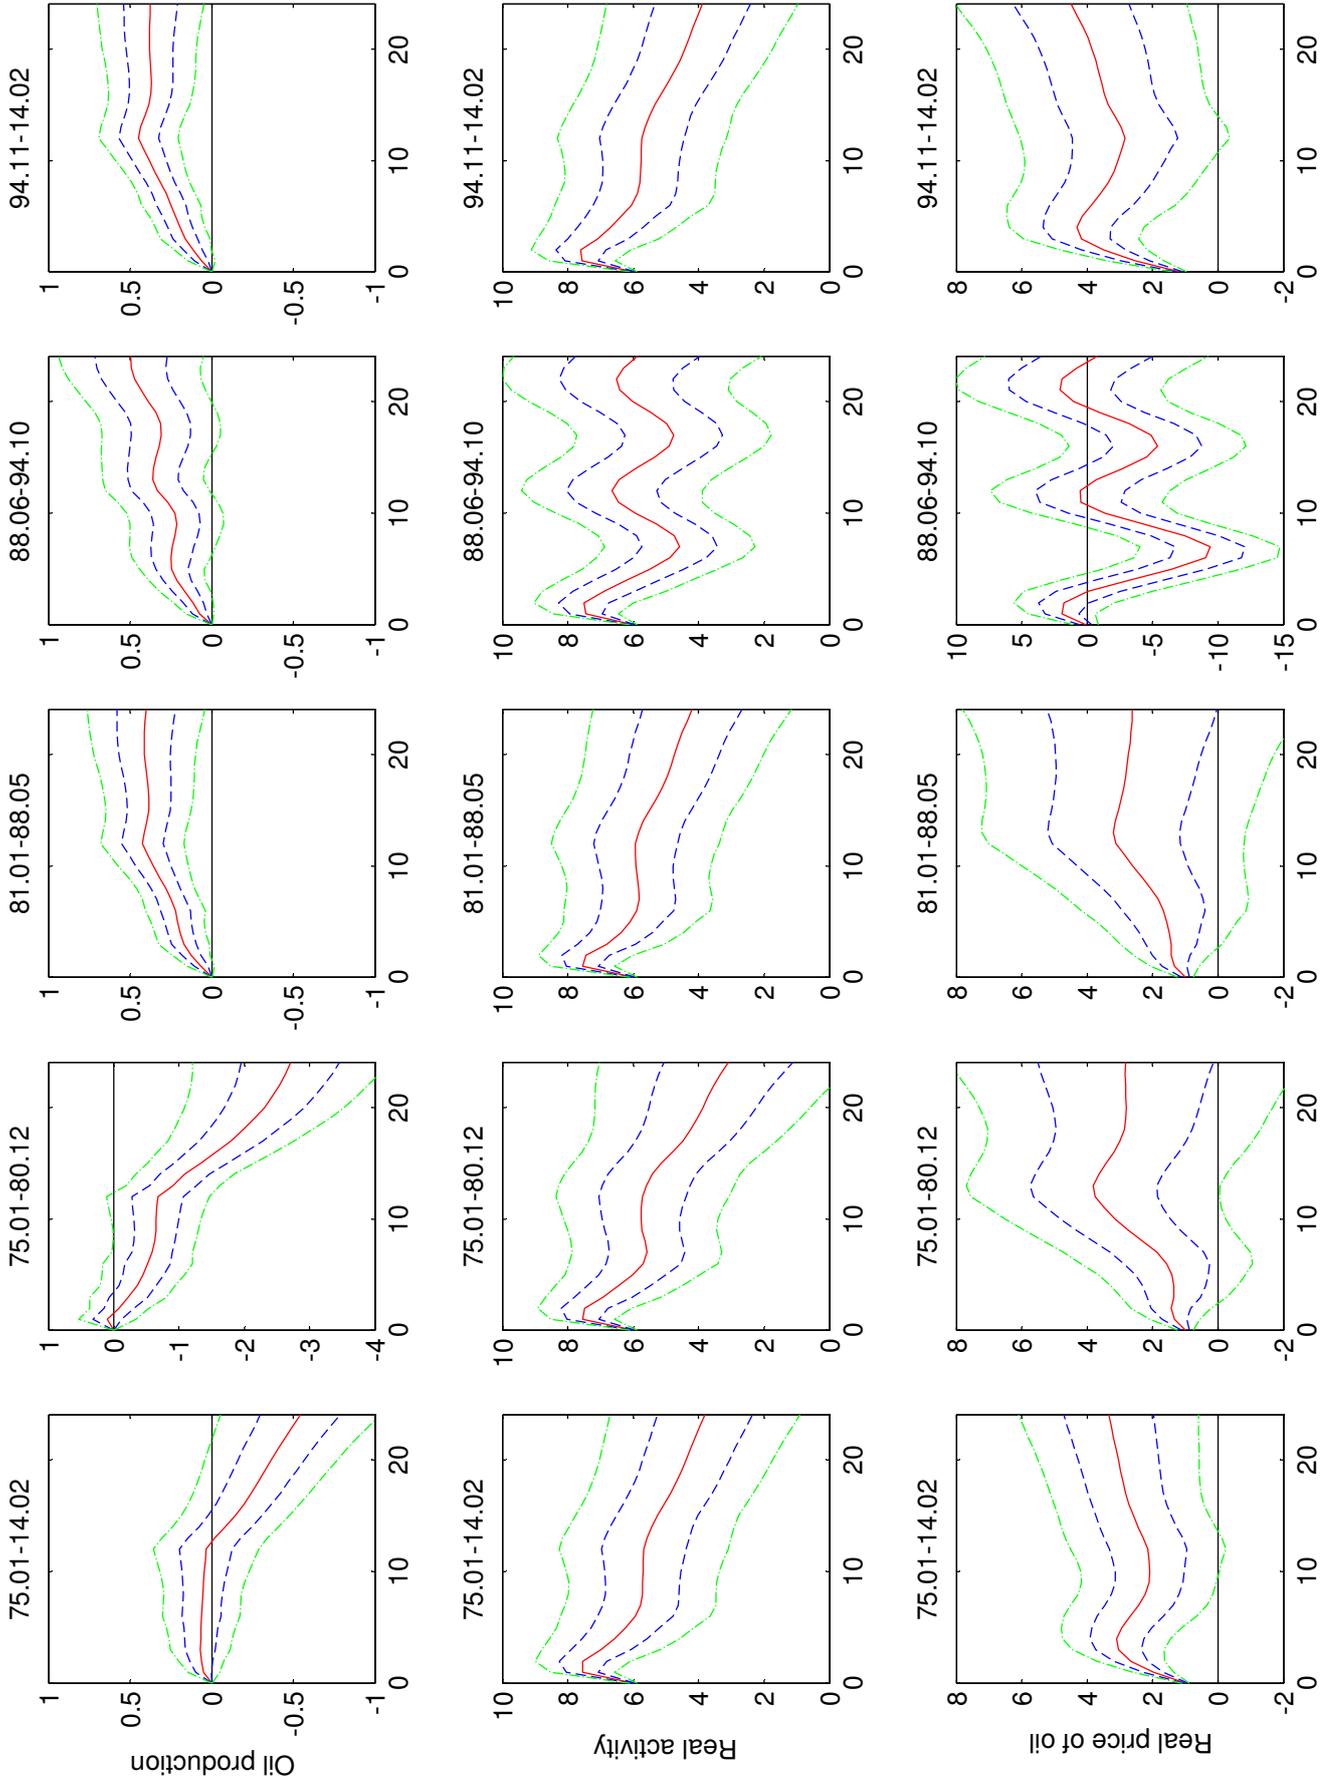
Notes: Spot (solid-green line) and futures prices (dotted-pink lines) of light sweet crude oil in NYMEX. US\$/barrel.

Figure 3. Response to a supply shock



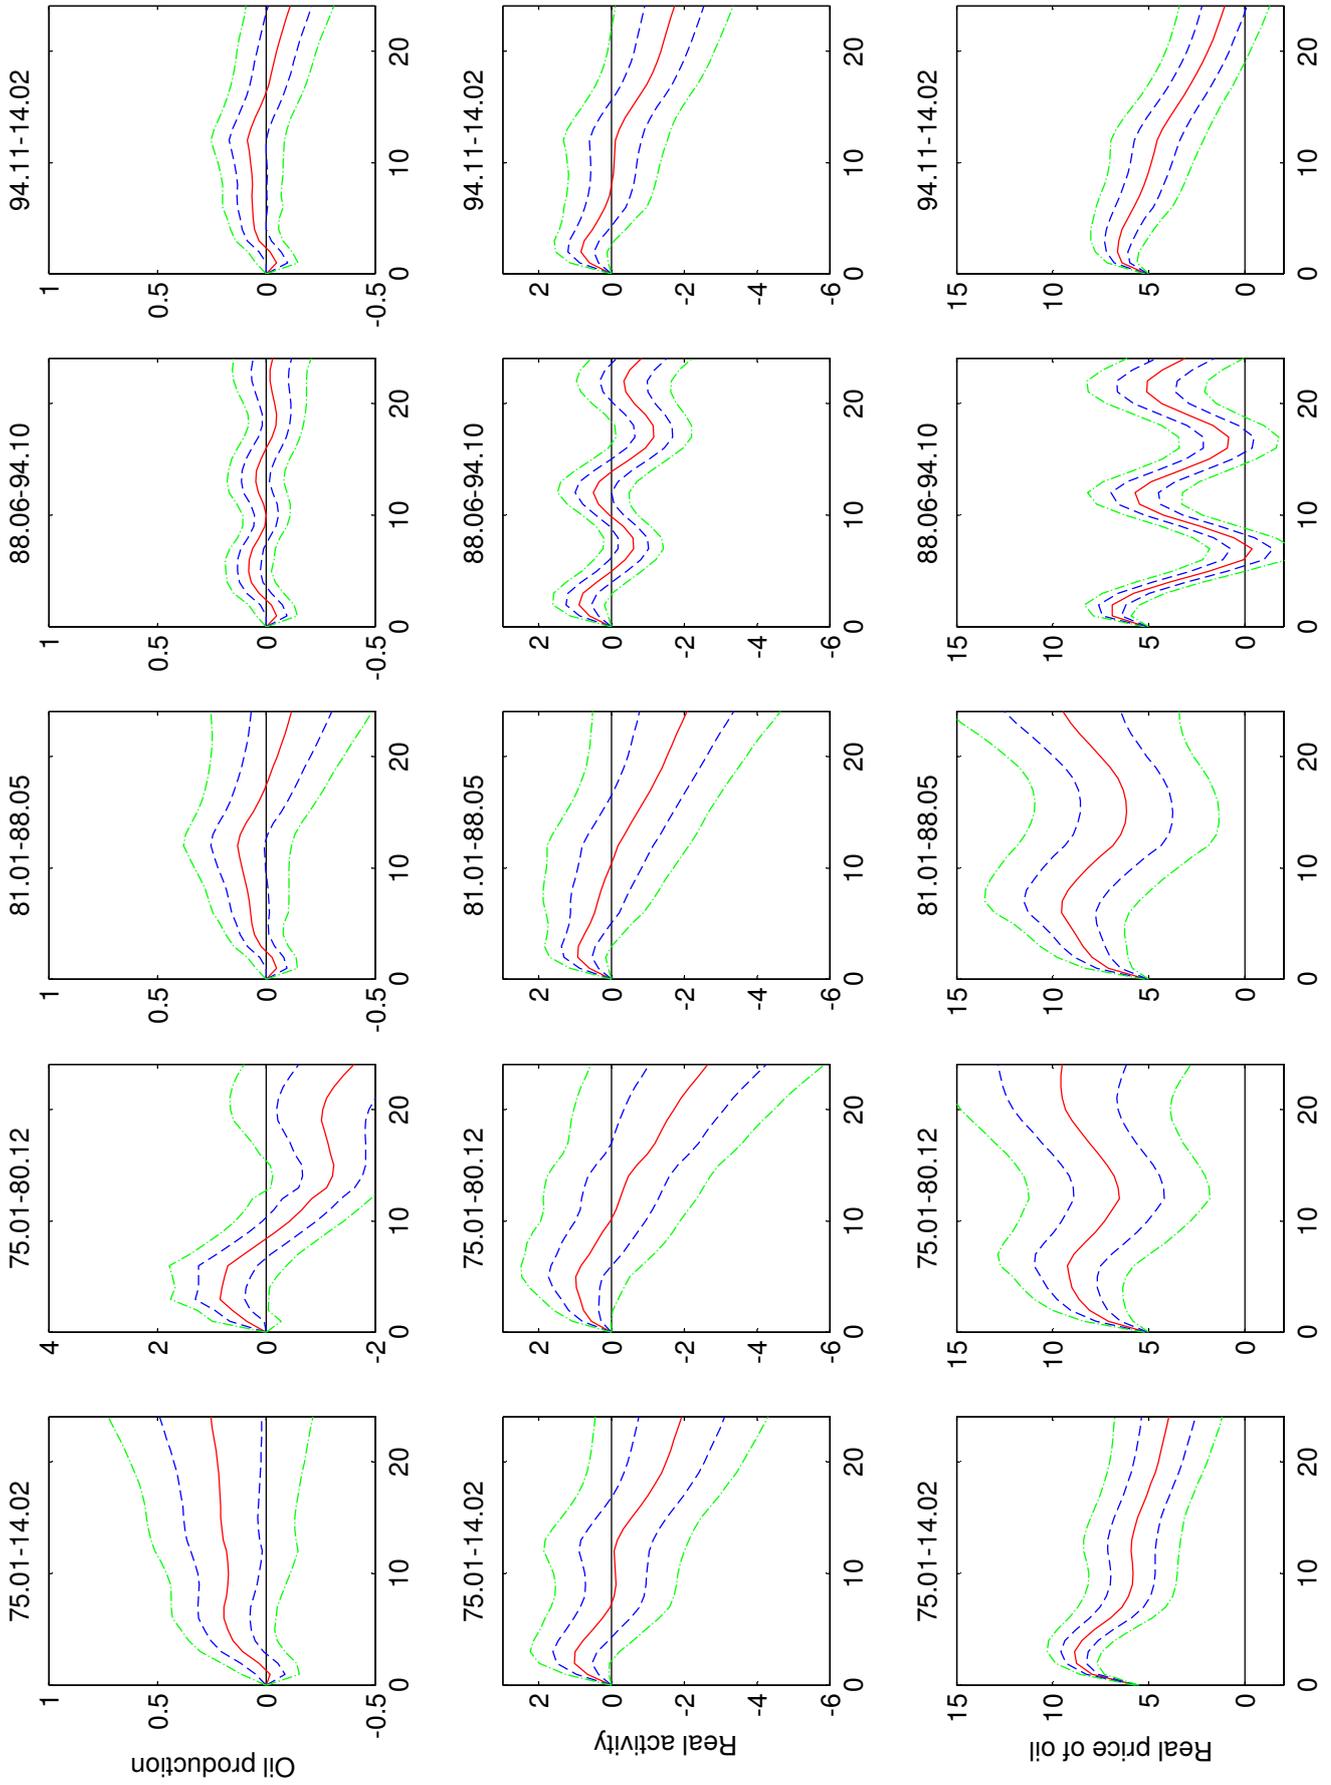
Notes: Each graph shows the cumulated impulse responses to a shock of a common magnitude, equal to one standard deviation estimated over the whole sample with no breaks. There are five columns representing the whole sample (assuming no breaks) and the sub-samples defined by the coefficient break dates in Table 2. In each plot there are one (blue dashed) and two (green dotted) standard deviation confidence intervals are obtained using a recursive wild bootstrap procedure, as in Kilian (2009).

Figure 4. Response to an aggregate demand shock



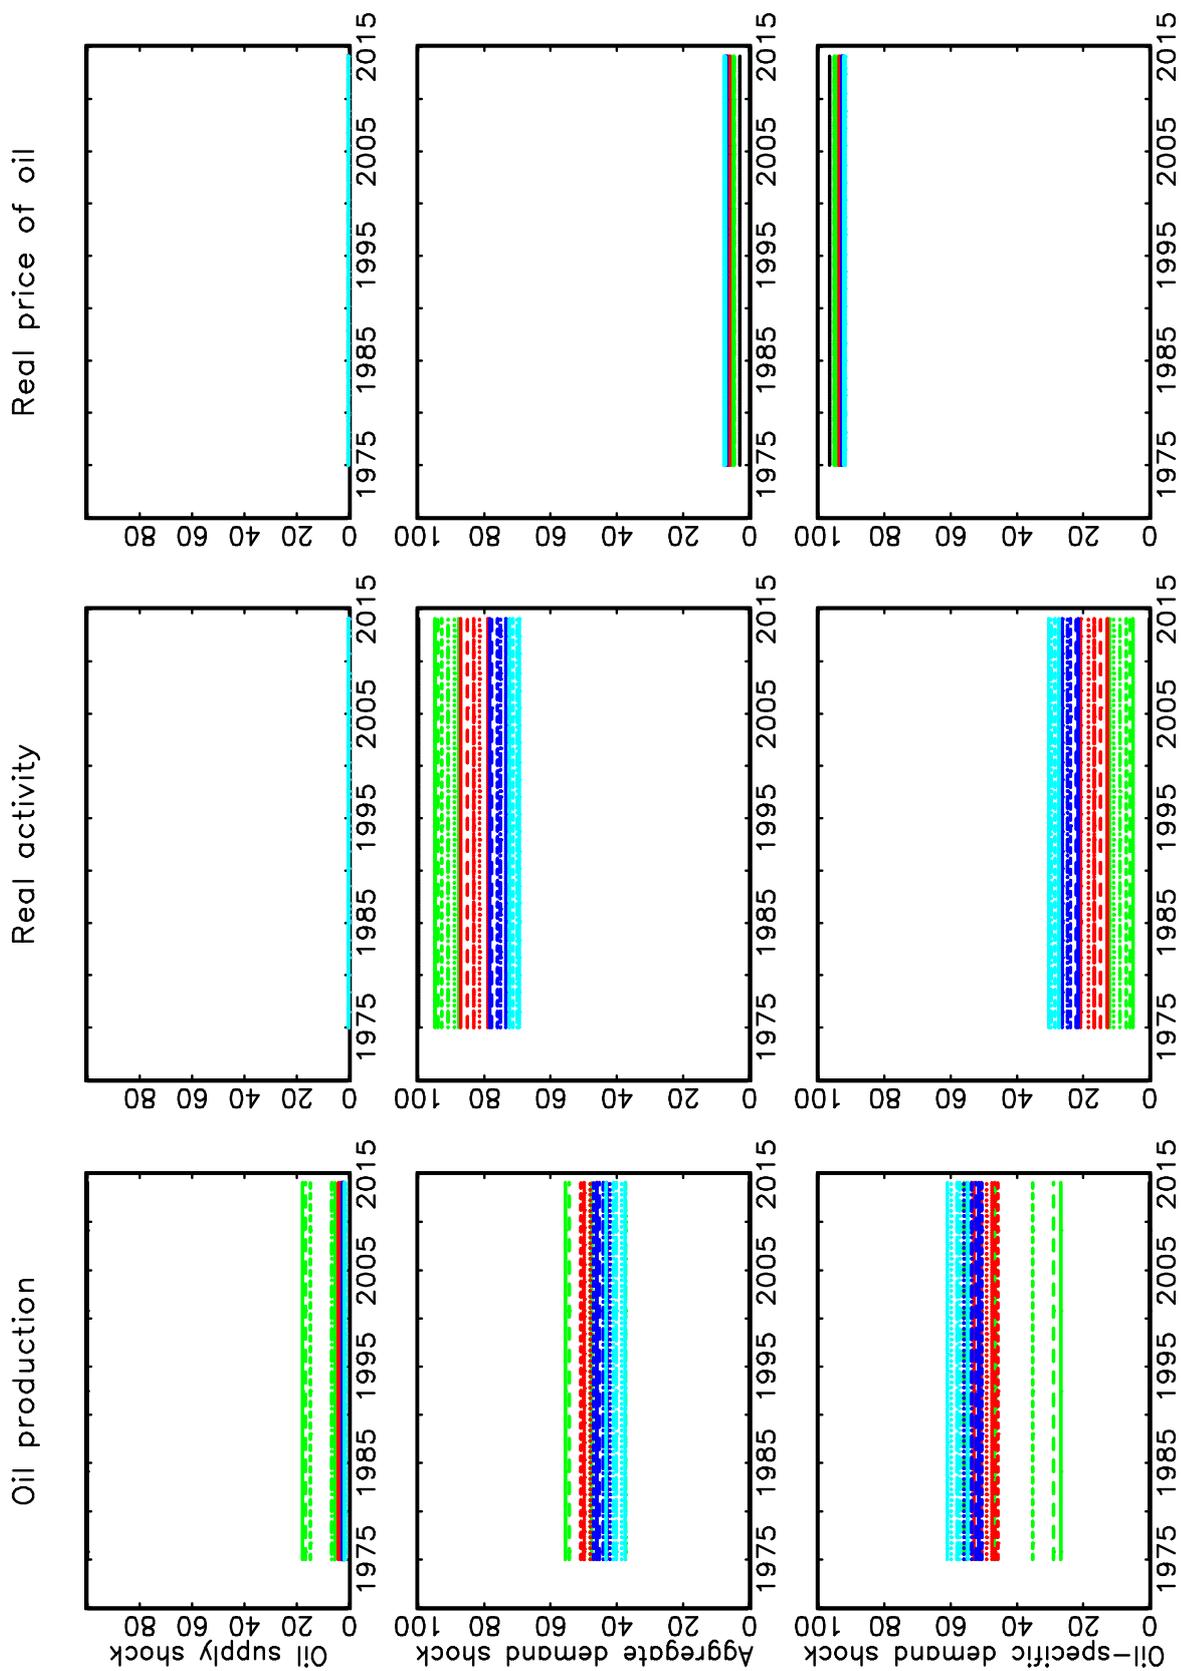
Notes: See Figure 3.

Figure 5. Response to an oil specific demand shock



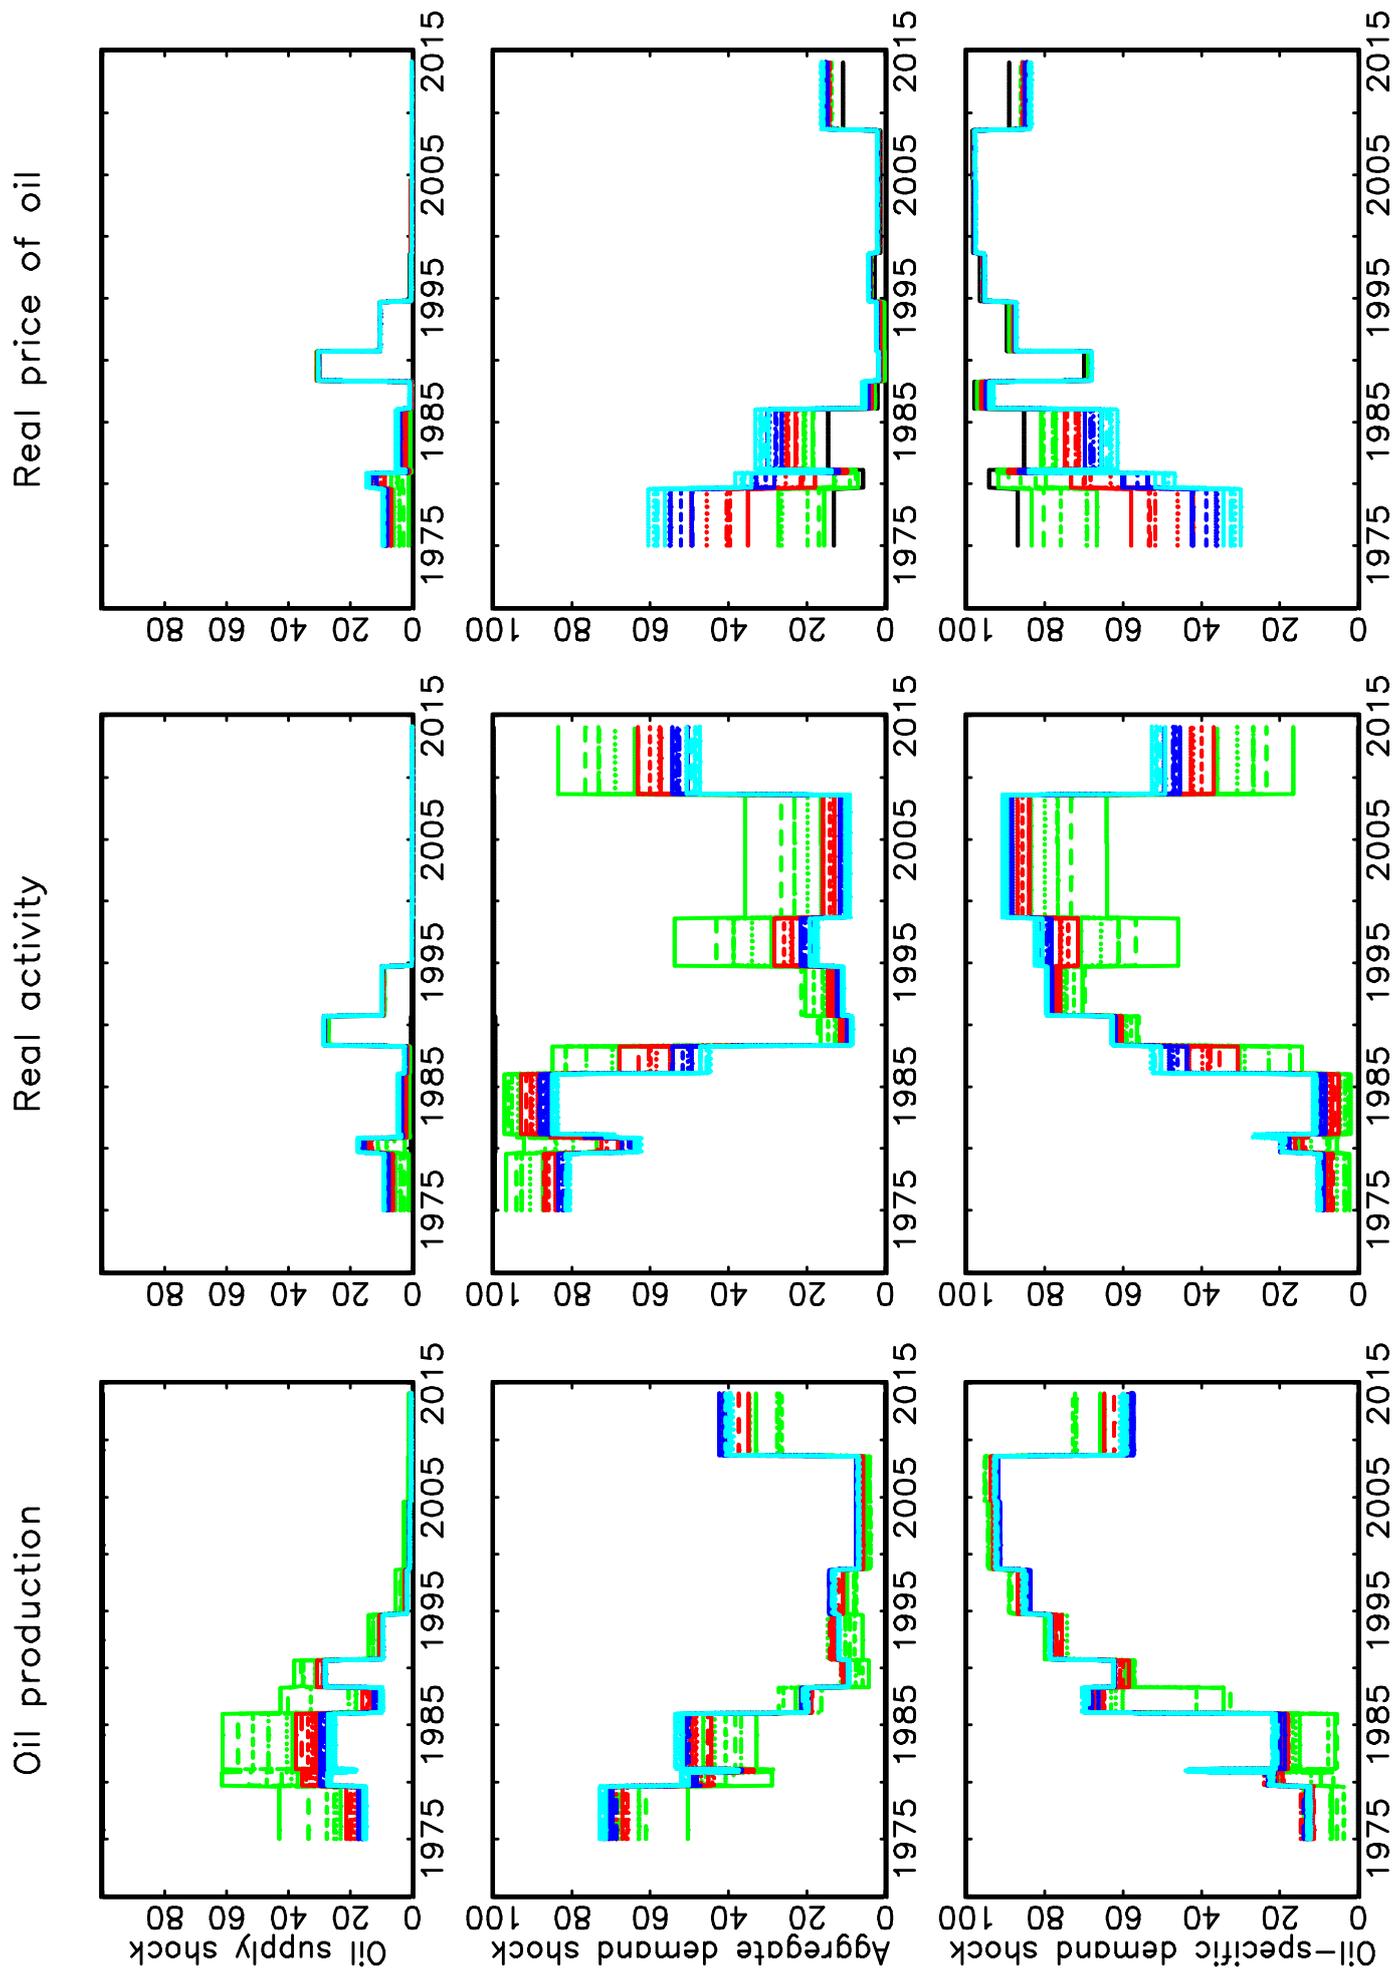
Notes: See Figure 3.

Figure 6. FEVD with no break



Notes: The decomposition shows the proportion of the forecast error variance of each of differenced variable that can be explained by oil market shocks up to 24 months into the future. Line type and colour coding are employed: green, red, blue and cyan lines indicate closer to the present to further into the future (in that order) and, within each colour, solid, dashed, short dashes, dots and dashes, dotted, and closely spaced dots are also arranged from closer to the present to further into the future. They assume no structural breaks throughout the sample.

Figure 7. FEVD with structural breaks



Notes: The color and line type coding are the same as in Figure 6.