Oil and the macroeconomy: A structural VAR analysis with sign restrictions*

Francesco Lippi
University of Sassari
EIEF and CEPR

Andrea Nobili
Bank of Italy

May 2008

Abstract

We consider an economy where the oil price, industrial production, and other macroeconomic variables fluctuate in response to a variety of fundamental shocks. We estimate the effects of different structural shocks using robust sign restrictions suggested by theory using US data for the 1973-2007 period. The estimates show that identifying the shock underlying the oil price change is important to predict the sign and the magnitude of its correlates with the US production. The results offer a natural explanation for the smaller correlation between oil prices and US production in the recent years compared to the seventies.

JEL classification: C32; E3; F4.

Keywords: Business cycle; Oil prices; Structural VAR.

*We thank Fernando Alvarez, Fabio Canova, Luca Dedola, Stefano Neri and Harald Uhlig for helpful discussions and suggestions. We also benefited from the comments by participants to seminars at the Bank of Italy, the Center for European Integration Studies (Bonn), the Study Center Gerzensee. The views in this paper should not be attributed to the Bank of Italy. All remaining errors are ours.

1 Introduction

Large fluctuations in oil prices are a recurrent feature of the macroeconomic environment. Despite of oil's relatively small share as a proportion of total production costs, such dynamics raise the specter of the seventies, worrying consumers, producers and policy makers. This view is supported by some empirical evidence that large and persistent oil price upswings lead to economic recessions and higher inflation rates. Hamilton (1983) pointed out that nine out of ten of the U.S. recessions since World War II were preceded by a spike in oil prices. The reduced-form relation between oil prices and US macroeconomic variables, however, has not been stable over time. The evidence available since the mid eighties detects smaller real effects of oil prices on the US economy.[1]

Two routes have been followed to tackle this instability. One is to invoke a misspecification of the true relationship. Several authors thus propose alternative asymmetric or non-linear transformations of the original oil price series in the attempt to recover a stable relationship. The second one, discussed e.g. by Barsky and Kilian (2004), calls for a structural interpretation of the reduced form correlations mentioned above. This is the route followed by this paper. We take the view that oil prices (and quantities) respond to a variety of shocks. This has two consequences. First, similar oil price fluctuations are associated with different macroeconomic consequences depending on the nature of the underlying fundamental economic shock. Second, oil prices are not exogenous to the US business cycle, at least in principle. For instance, while the disruption of oil supplies in the middle east raises oil prices and reduces US production −through a higher cost of the energy input− a similar increase in the price of oil can be observed following a demand boom of the US economy that increases the demand for oil imports. The identification of the funda-

[1]See Hooker (1996) and Hamilton (2008). Similar conclusions are reached for European countries by Mork, Olsen and Mysen (1994) and Cunado and Perez de Gracia (2003).

mental shocks underlying oil price movements is thus key to interpret the reduced form correlation between oil prices and the US economy. Blanchard and Gali (2007) also conjecture the possibility that the time varying effect of "oil shocks" on US production may reflect shocks of different nature. Given their partial identification strategy, however, they do not explore this hypothesis quantitatively. This paper fills that gap.

The model offers a natural explanation of the unstable (unconditional) correlation between oil price fluctuations and the US GDP. As the sign of this correlation depends on the type of shock that hits the economy, a negative correlation emerges in periods when oil-supply shocks dominate, while a positive correlation emerges in periods of strong US demand shocks or supply shocks in RoW. The unconditional correlation detected by OLS estimates between oil prices and US production is tenuous because it blends conditional correlations with different signs.

The paper is structured as follows. Section 2 uses a simple model economy, adapted from Backus and Crucini (2000), to outline our view of the interactions, and the nature of fundamental shocks, that characterize the oil market and the US macroeconomy. This model lays the foundation of our empirical analysis. Section 3 discusses the empirical approach, which is based on a structural vector autoregression (VAR) in which fundamental shocks are identified by means of sign restrictions taken from the theory, as in Canova and Nicolo (2002) and Uhlig (2005). This approach allows us to identify the structural shocks using some robust properties of the model economy presented in Section 2, without imposing the whole structure of our theoretical model on the data, i.e. allowing for −and quantifying− a degree of "model uncertainty". Section 4 presents the findings of our estimation. Section 5 discusses the sensitivity of our findings with respect to alternative hypotheses put forth by Blanchard and Gali (2007), Kilian (2007) and Leduc and Sill (2004). Section 6 concludes.

2 Theoretical frame

We present a three-country theoretical economy that is useful to organize ideas about the US macroeconomy and its interaction with the oil market. The model is taken from Backus and Crucini (2000) who extend the two-good two-country economy of Backus, Kehoe and Kydland (1994) and incorporate a country that produces oil. The model features supply shocks z_j in each country j. We enrich the model by introducing stochastic preference shocks. Below we present the essential ingredients of the theoretical economy and discuss the implications that will be used in the empirical analysis. More information on the model solution and calibration is given in Appendix A.

Two industrialized countries, the US and RoW (rest of the industrial world), produce imperfectly substitutable consumption goods, a and b, using capital (k), labor (n) and oil (o). The US produce good a using the technology

$$y_t = z_t n_t^{\alpha} [\eta k_t^{1-\nu} (1-\eta) o_t^{1-\nu}]^{(1-\alpha)/(1-\nu)} \tag{1}$$

where z denotes a stochastic (unit-mean) productivity variable. An analogous technology is used for the production of b by RoW. The oil supply, y^o, is determined according to $y_t^o = z_t^o + (n_t^o)^\alpha$ where z_t^o is an exogenous stochastic oil supply component and $(n_t^o)^\alpha$ an endogenous supply by the third country, which one can think of as the union of OPEC and other (non US) oil producing countries. Goods a and b are aggregated into final consumption (c) and investment (i) goods using the CES aggregator

$$x(a,b) = [\psi a^{1-\mu} + (1-\psi) b^{1-\mu}]^{1/(1-\mu)} \tag{2}$$

The consumption bundle is subject to stochastic AR(1) preference shocks, $s_t = (1-\rho_s) + \rho_s s_{t-1} + \tilde{s}_t$ such that the weight in the consumption good aggregate is given by $\tilde{\psi} \equiv s_t \cdot \psi$.[2] Capital obeys the accumulation equation $k_{t+1} = (1-\delta)k_t + k_t\, \phi(i_t/k_t)$, where $\phi(\cdot)$ is a concave function positing adjustment costs in capital formation as in Baxter and Crucini (1993). Consumers in each country maximize the expected value of lifetime utility

$$\max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t\ U\left(c_t, 1-n_t\right) \tag{3}$$

where $U(c_t, 1-n_t)$ is a CES function defined in Appendix A and $\beta < 1$ is the intertemporal discount.

The model allows us to examine the effects of supply side (productivity) and demand (preference) shocks in each economy. The solid lines in Figure 1 report impulse responses produced using the benchmark parametrization based on Backus and Crucini (2000), see Appendix A. Each column of the figure reports (from top to bottom) the impulse response functions of oil prices, oil production, the price of output and the output level to the structural shock shown below the column. All prices are expressed relative to the US consumption deflator, i.e. the US consumption price index is chosen as the numeraire, as done later in the empirical analysis.

The first column describes the effect of a positive oil supply shock ($z^o > 0$). The shock causes a negative correlation between the response of oil price and quantity. As the price of the energy input falls, production prices in the US rise moderately (i.e. the consumption good depreciates relative to domestic production) and production increases. This represents the prototype textbook case of an exogenous oil supply shock. Innovations to productivity ($z^* > 0$) or preferences ($s^* > 0$) in the RoW also affect the market for oil and that for US goods. As shown in the second and third columns of Figure 1, the sign of the response of the oil price and of oil quantity to each of these shocks is the same, i.e. these shocks appear to an observer of the oil market as typical "demand shocks" causing oil prices and quantities to move in the same direction (solid lines). The bottom panels of columns two and three show

[2]Similar effects are obtained by considering shocks to the intertemporal discount factor.

Figure 1: Structural shocks in the model economy

(% deviation from steady state)

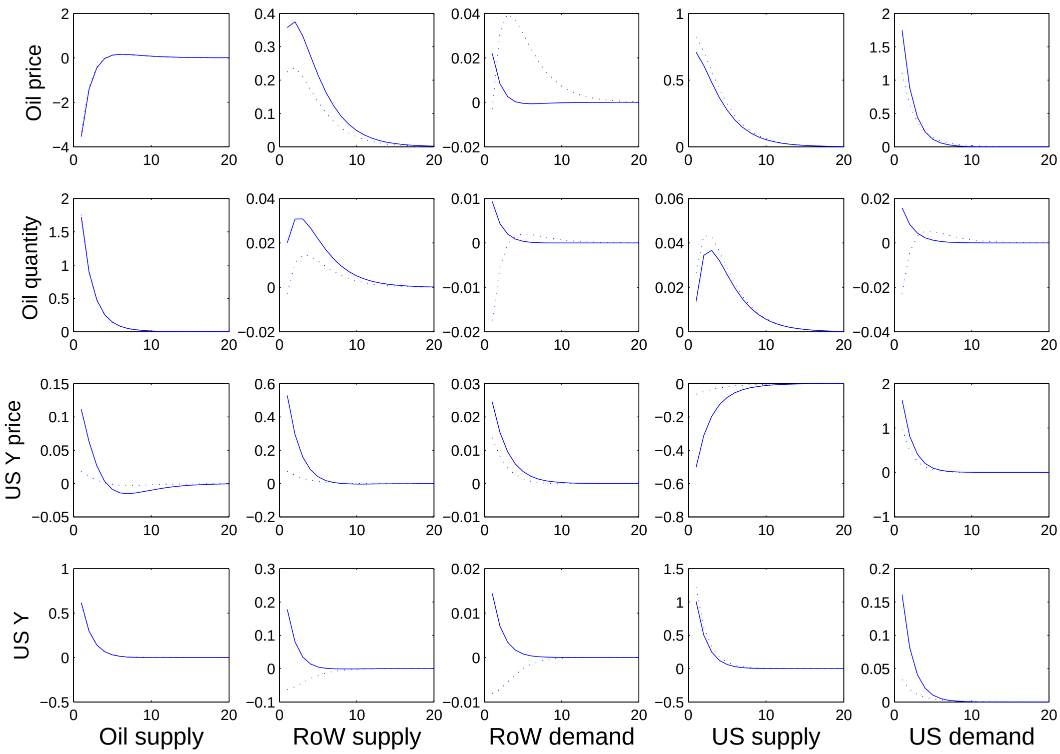

Note: Each column describes the effect of the structural shock indicated below the column. Each row corresponds to the variable indicated on the left vertical axis. All prices are expressed relative to the US consumption price index. The solid lines is the benchmark parametrization. The dashed lines show the responses under a parametrization with a high substitution elasticity between the goods of US and RoW ($\mu = 0.7$ instead of $\mu = 2$ in the benchmark case).

that US production increases following these shocks. The fourth and fifth columns describe the effects of shocks originated in the US economy. A positive productivity shock ($z > 0$) raises US production and reduces its price (relative to the CPI). The ensuing increase in oil demand ultimately leads to higher oil prices and output, though the latter is small in the benchmark parametrization. Finally, a US demand shock ($s > 0$) increases US production and its price. The increased demand spills over to the oil market, where prices and production increase (solid lines).

The quantitative features of the effects described above depend on the parameterization but some qualitative features of this model are robust across parametrizations. Four properties in particular will be used in the empirical analysis. First, the US production and its price respond with opposite signs to US productivity shocks and, second, they respond with the same sign to US demand shocks. Third, oil supply shocks always cause the oil price to move with a sign that is opposite to that of oil production. Fourth, supply shocks in the RoW appear on the oil market as demand shocks, moving oil prices and quantities in the same direction. Other predictions of the model are not robust across parametrizations, for instance the response of US production to shocks that occur in the RoW depends on the degree of substitutability between domestic and foreign goods. The benchmark case (solid lines) assumes a low substitutability $1/\mu < 1$, a case where the income effect dominates the substitution effect. Results for the case with high substitutability $1/\mu > 1$ are shown by the dashed lines. It appears that following a positive supply shock in the rest of the world (second column) the US supply increases when the substitutability between the goods is small, and falls otherwise. As the sign of some responses crucially hinge on parametrization chosen, the identification of the structural shocks presented below will restrict attention to the four robust features just discussed.

We conclude by noting that the model economy shows that the expected change of US production conditional on an oil price increase depends on the underlying fundamental shock. For instance, while the oil price hike caused by an oil supply shock is followed by a decrease of US production, the oil price hike caused by a US demand shock or a foreign supply shock is followed by an increase of US output. Therefore, it should not be surprising that over a long sample period the unconditional correlation between oil prices and US GDP appears tenuous, as it blurs conditional correlations with different signs. The empirical analysis will allow us to cast light on this conjecture.

3 The Empirical Analysis

The empirical analysis aims to identify a set of structural shocks affecting the oil market and the US business cycle and analyze their effects on the oil prices and US output. The identification method is based on sign restrictions, following the approach pioneered by Canova and Nicolo (2002) and Uhlig (2005). The idea is to identify structural shocks using some robust properties of the model, such as the sign of impulse responses discussed in the previous section, without imposing on the data the whole structure of the theoretical model i.e. allowing for some degree of "model uncertainty". This is convenient when, as in our case, the model economy is stylized and one is reluctant to assume that the model is the true data generating process and/or e.g. the magnitude of the substitution elasticity. We first briefly describe the identification method and then discuss the sign restrictions that are used, relating them to the theory of Section 2.

The analysis is based on the vector autoregression (VAR)

$$y_t = B(L)\ y_{t-1} + \epsilon_t \qquad \epsilon_t \sim N(0, \Sigma), \tag{4}$$

where $B(L)$ is a lag polynomial of order p and y_t contains four variables describing the US and the oil market. The first two are the (log of) US industrial production and the (log of) consumer price index. This set of variables is similar to those used in previous studies, such as Bernanke, Gertler and Watson (1997) and Hamilton and Herrera (2004).[3] Two additional variables describe the oil market: the (log of) average oil spot nominal price and the (log of) global oil production.[4] Estimation of the VAR is based on monthly data spanning the period January 1973 - December

[3]The VAR by Bernanke et al. (1997) also included the non-energy commodity price index the short-term and long-term interest rates. The effect of including these variables in our VAR model is discussed in Section 5.

[4]Production is in barrels per day, the spot price is from the International Monetary Fund.

2007 (this uses the longest available production time series provided by the International Energy Agency). The period covers all the relevant episodes characterized by major oil price increases. We complete the specification by using a lag order of 6 months, as suggested by the Akaike Information Criteria (AIC).[5] The structural VAR approach sees (4) as a reduced form representation of the structural form

$$\mathcal{A}_0^{-1} y_t = A(L)\ y_{t-1} + e_t \qquad e_t \sim N(0, I) \tag{5}$$

where $A(L)$ is a lag polynomial of order p and the vector e includes the four structural innovations discussed above, assumed to be orthogonal. Identification of the structural shocks thus amounts to select a matrix $\mathcal{A}_0$ (i.e. a set of restrictions) that uniquely solves −up to an orthonormal transformation− for the following decomposition of the estimated covariance matrix $\mathcal{A}_0\mathcal{A}_0' = \Sigma$. The j-th column of the identification matrix $\mathcal{A}_0$, a_j, maps the structural innovations of the j-th structural component of e into the contemporaneous vector of responses of the endogenous variables y, $\Psi_0 = a_j$. The structural impulse responses of the endogenous variables up to the horizon k, Ψ_k, can then be computed using the $B(L)$ estimates from the reduced form VAR, $B_1,\ B_2,\ ...,\ B_p$, and the impulse vector a_j.[6]

The sign restriction approach identifies a set of structural models, the $\tilde{\mathcal{A}}_0 \in \tilde{\mathbf{A}}_\mathbf{0}$, such that the impulse responses Ψs implied by each $\tilde{\mathcal{A}}_0$ over the first k horizons are consistent with the sign restrictions derived from the theory. The approach exploits the fact that given an arbitrary identification matrix $\mathcal{A}_0$ satisfying $\mathcal{A}_0\mathcal{A}_0' = \Sigma$, any other identification matrix $\hat{\mathcal{A}}_0$ can be expressed as the product of $\mathcal{A}_0$ and an orthonormal matrix H. The set of the theory-consistent models, $\tilde{\mathbf{A}}_\mathbf{0}$, can be characterized as follows. For a given estimate of the reduced form VAR, $B(L)$

[5]The appropriate lag length was debated in previous literature, see Hamilton and Herrera (2004). Our results remain virtually unchanged if 12 lags are used.

[6]As $\Psi_s = \sum_{i=0}^{s} B_{s-i}\Psi_i$ for $s \geq 1$ and $B_{s-i} = 0$ if $s - i > p$.

and Σ, take an arbitrary identification matrix $\mathcal{A}_0$ and compute the set of candidate structural models $\hat{\mathbf{A}}_\mathbf{0} = \{\mathcal{A}_0\, H \mid HH = I\}$ by spanning the space of the orthonormal matrices H. The set $\tilde{\mathbf{A}}_\mathbf{0}$ is then obtained by removing from the set $\hat{\mathbf{A}}_\mathbf{0}$ the models that violate the desired sign restrictions. The findings can then be summarized by the properties of the resulting distribution of $\tilde{\mathcal{A}}_0$ models. The set of orthonormal theory-consistent matrices $\tilde{\mathbf{A}}_\mathbf{0}$ is computed using the efficient algorithm proposed by Rubio-Ramirez, Waggoner and Zha (2005).[7]

In the empirical analysis we restrict attention to 4 mutually orthogonal shocks: demand and supply shocks in the US economy and in the oil market, consistent with the model robust features discussed in Section 2. These restrictions are summarized in Table 1. A US demand shock is one that generates a response of the US industrial production and its deflator (relative to the CPI) of the same sign. A US supply shock is one that induces a negative correlation between the US industrial production and its deflator. These assumptions are consistent with the impulse responses reported in the fourth and fifth columns of Figure 1. The third one is an oil-supply shock, i.e. one that causes the oil production and its price (CPI deflated) to move in opposite directions, as in the first column of Figure 1. Finally we define the oil demand shock as one that, upon occurring, moves the price and the quantity of oil in the same direction. As discussed in the previous section this captures shocks to the oil market that originate from e.g. supply shocks in the rest of the world (second column in Figure 1).[8] The sign restrictions for the shocks in Table 1 are not mutually exclusive.

[7]Given the OLS estimates for $(B(L), \Sigma)$, the algorithm draws an arbitrary independent standard normal $(n \times n)$ matrix X and, using the QR decomposition of X (where Q is orthogonal and R triangular), generates impulse responses directly from A_0Q and $B(L)$. If these impulse responses do not satisfy the sign restrictions the algorithm generates a different draw for Q. Compared with Uhlig's procedure, this algorithm directly draws from a uniform distribution instead of involving a recursive column-by-column search procedure.

[8]For low values of the substitution elasticity between the US and the RoW goods the oil demand restrictions are also consistent with demand shocks originating in the RoW. However, our analysis does not attempt to distinguish supply and demand shocks in the RoW, due to the lack of a sufficiently long time series for prices and output in a country group for RoW that includes China and India.

By construction, however, each of the models in $\tilde{\mathbf{A}}_\mathbf{0}$ generates orthogonal structural shocks.

Table 1: Sign restrictions used for identification

	Structural shocks			
VAR Variables	oil-supply	oil-demand	US demand	US supply
oil production	$-$	$+$		
oil price[a]	$+$	$+$		
US industrial production			$+$	$+$
US producer's price[a]			$+$	$-$

Note: A "+" (or "−") sign indicates that the impulse response of the variable in question is restricted to be positive (negative) for 6 months after the shock. A blank entry indicates that no restrictions is imposed on the response. −[a]Price is deflated by the US CPI.

In practice we also have to decide on for how long the sign restrictions should hold. We start by imposing they hold for a period of 6 months. Similar findings are obtained when the restrictions are imposed for a period of 12 months. In all of these cases the resulting posterior distribution is made of 5,000 $\tilde{\mathcal{A}}_0$ models.

4 The estimated effects of structural shocks

This section reports summary measures from the distribution of impulse responses of the structural models in $\tilde{\mathbf{A}}_\mathbf{0}$. The distribution reflects the model uncertainty inherent to the sign restriction approach, providing a description of the range of possible outcomes consistent with the set of theoretical restrictions. In particular, in this exercise we freeze the VAR coefficients for $B(L)$ and Σ at their estimated OLS values. Therefore the distribution reflects the model uncertainty but not the sampling uncertainty that underlies the coefficient estimates. Appendix B presents the results for the case when both sampling and model uncertainty are accounted for.

4.1 Impulse responses

We follow Dedola and Neri (2007) and Uhlig (2005) and report the median (solid line), the 16th and the 84th percentiles (the dashed lines) of the distribution of impulse responses produced by the $\tilde{\mathcal{A}}_0$ models for each variable over the first 36 months.

The effects of an oil supply shock, normalized to yield a 1 per cent reduction in oil production, are displayed in Figure 2. The shock lowers the US industrial production, that reaches a through after about 12 months. The figure shows that after one year the response is negative for 76% of the models. In Table 2 we interpret this finding probabilistically, and interpret this ratio as the probability that the output response, conditional on the oil supply shock, is negative. The price of US production (CPI deflated) increases slightly in the short run, but not in the long run. The persistence of the oil-supply shock is high, well above the 6 month restriction imposed by the model selection. The output effect is also persistent.

The effects of an oil demand shock, normalized to yield a 1 per cent increase in oil production, are displayed in Figure 3. The oil price increase is highly persistent. The key finding from this figure is that the response of the US industrial production differs markedly from the case of the oil-supply shock. For most models, production 12 months after the shock is above the baseline. Table 2 shows that the probability that the US industrial production increases conditional on an oil demand shock is about 70 per cent after 12 months, and is greater than 50% at all horizons.

Table 2: US output response to oil market shocks

	Probability of a negative response of US output[a]				
At Horizon:	1	6	12	18	24
Oil-supply shock	0.65	0.74	0.76	0.76	0.75
Oil-demand shock	0.38	0.34	0.33	0.34	0.35

Note: –[a]Fraction of models $\tilde{\mathcal{A}}_0 \in \tilde{\mathbf{A}}_\mathbf{0}$ that yield a negative response at the given horizon.

These results suggest that in order to predict the dynamics of US business cycle conditional on observing an oil-price increase it is key to identify the fundamental shock underlying the oil-price hike. Higher oil prices are associated to an expected reduction of production conditional on a negative oil-supply shock, and to an expected rise of production conditional on a positive oil-demand shock. The fact that US industrial production increases following a positive oil-demand shock can be reproduced by the theoretical model of Section 2 under the assumption that the substitution elasticity between US and RoW goods is small, so that the higher demand in the RoW countries increases expenditure on US goods.

Figure 4 illustrates the extent to which the oil market is affected by the US shocks. An aggregate demand shock that raises the US industrial production causes a rise in oil prices. The effect on oil production is smaller, and it is centered about zero 12 months after the shock. Figure 5 shows the effects of a positive US aggregate supply shock: the response of oil quantity is positive, while the response of the oil price is small and centered about zero over the first 18 months.

4.2 Variance decomposition

What is the contribution of the different structural shocks on aggregate fluctuations and oil prices? We assess this issue by computing the percentage of the variance of the k-step ahead forecast error that is accounted for by the identified structural shocks. Table 3 reports the variance decomposition at horizons up to 24 months for the spot oil price and the US industrial production. To ensure orthogonality of the structural shocks the table entries are computed from a unique $\tilde{\mathcal{A}}_0$, chosen so as to minimize a minimum distance criterion from the median responses displayed in the previous subsection (see Appendix C for the details).[9]

[9] Qualitatively the results are similar to those produced by the median of the forecast variance posterior distribution implied by the set of $\tilde{\mathcal{A}}_0$ models (available upon request).

Table 3: Variance decomposition

k	Oil supply	Oil demand	US supply	US demand
	Spot oil price			
1	51.8	42.7	2.2	3.3
6	51.1	38.6	0.0	10.2
12	18.8	58.5	2.7	20.0
18	0.0	70.0	2.1	27.9
24	7.9	68.6	1.1	22.4
	US industrial production			
1	7.2	9.7	41.3	41.7
6	14.2	14.3	45.4	26.1
12	15.7	13.7	54.6	16.0
18	15.4	12.2	62.4	10.0
24	13.7	9.2	71.1	5.9

Notes: Entries computed by the $\tilde{\mathcal{A}}_0$ model that minimizes the distance from the median impulse response (see Appendix C). k denotes the forecast horizon (in months).

The first panel of table 3 shows that oil-demand shocks explain a large fraction of the oil price variance, between 40 and 70 per cent, over the horizons considered. Oil-supply shocks explain a large portion of the variance within the one year horizon. Another sizeable fraction of the oil-price variance, accounting for about 20 to 30 percent, is due to US aggregate demand shocks. The US supply shocks have a negligible impact on oil prices. The historical decomposition of the oil price time series, displayed in Figure 6, shows that oil-demand shocks were a key factor underlying the drop of oil prices following the Asian crisis of 1997-1998 and the strong price hike that started in 2003.

The second panel of Table 3 presents the variance decomposition of the US industrial production at horizons of up to two years. The US aggregate supply shock explains the largest share, in line with the real business cycle hypothesis and the recent contributions of Dedola and Neri (2007), and Francis and Ramey (2005). The role of US aggregate demand shocks is largest at short horizons (1 to 6 months), and smaller than the role of US supply shocks at all subsequent horizons. Concerning the role of the oil market variables on US production, both the oil-supply and the oil-demand shocks appear important, as each one explains a proportion of about 10 percent. Figure 6 presents the time-series decomposition of the US production.

4.3 Robustness

The robustness of the findings was tested along several dimensions. First, we considered whether the quantitative findings on the effects of the oil market shocks changed if we used a scheme that identifies only 2, as opposed to 4, structural shocks. This corresponds to identifying $\tilde{\mathcal{A}}_0$ matrices using the restrictions of the first two columns only of Table 1. The estimated effects of the oil market shocks on the US industrial production are virtually unchanged, as shown in Figure 7.

Second, we estimated the effects of structural shocks accounting for both model and sample uncertainty (the latter was ignored above). We also explored the critique, raised by Fry and Pagan (2007), that the sign restriction approach is flawed because the impulse response functions it generates likely violate the assumption that structural innovations are orthogonal. Details on each of these analysis are given in Appendix B and in Appendix C. Altogether, these analyses produce impulse response functions that are not significantly different from the benchmark case analyzed above.

The next section studies how the analysis is affected by extending the VAR to include a role for the monetary policy rule, for precautionary oil demand shocks, and considering the possibility that other structural changes in the economy occurred between the seventies and the last part of the century.

5 Related hypotheses in recent literature

Some recent contributions −that maintain the assumption that oil prices are exogenous to the US economy− discuss complementary mechanisms that may explain the smaller effects of oil prices on the US macroeconomy observed in recent years. Below, we explore some of them in the context of our framework.

5.1 Precautionary demand for oil

One hypothesis concerns the possibility that oil demand shocks of different nature have different effects on the US economy. For instance, the model of Section 2 showed that demand and supply shocks in the rest of the world caused an oil demand shock (comovement of oil prices and quantities), but had different consequences on US production. To analyze the possibility of oil-demand shocks of different nature, we follow Kilian (2007) who decomposes oil price dynamics into oil-supply shocks, world aggregate demand shocks and precautionary oil-demand shocks (due to e.g. fears about future oil supply availability). Kilian's shocks are identified by a recursive scheme based on short-run restrictions. His findings show that aggregate demand shocks produce effects on US real GDP qualitatively similar to the oil-demand shocks described in Figure 3, although more persistent.[10] To explore the hypothesis that the oil-demand shocks described in Figure 3 may also reflect shocks to the precautionary demand for oil discussed by Kilian, we set up a 5-variable VAR that includes a non-energy commodity price index. The oil-demand shock is now identified by assuming that the shock simultaneously increases the demand for oil quantity, as well as the demand for other non-energy commodities, thus resulting in an increase of both the oil price and the commodity price index. The impulse responses, reported in Figure 8, show that the median responses obtained with our benchmark 4-variable VAR are extremely similar to those produced by the new identification scheme.

[10]After an oil-supply shock the hike in oil price is small and transitory, smaller than what is suggested by our impulse response. Real GDP decreases on impact while inflation increases temporarily. Precautionary oil-demand shocks produce stagflationary effects which are more persistent than those implied by oil supply shocks.

5.2 Structural change

Blanchard and Gali (2007) argue that the smaller effect of oil shocks on the US economy in the recent years can be ascribed to structural changes, such as changes in the energy share used in production or real wage rigidities. These authors also discuss the possibility that other shocks had an offsetting impact on the US economy. The last hypothesis is aligned with our view that the reduced form correlation between US macro variables and oil prices reflect a variety of shocks, and the evidence presented above supports this view and shows that it is quantitatively important.

To address the hypothesis that important structural changes have occurred, we analyze the effects of identified oil-supply and oil-demand shocks over different sample periods. A warning is due: as our approach assumes the presence of different structural shocks, splitting the sample may lead to the selection of observations (subperiods) in which some of the shocks considered were not present. For this reason we avoid cutting the sample in very short periods. The sample split dates chosen for the analysis were suggested by previous studies: 1981 as the date in which the Federal Reserve changed operating procedures, 1984 as the conventional date used in the Great Moderation literature (Blanchard and Gali (2007)), 1987 as the beginning of the post-OPEC period in the oil market (Backus and Crucini (2000)), 1991 as the date before which major oil price shocks were caused by (exogenous) political disruptions in the Middle East (Hamilton (1983); Hamilton (1996); Bernanke, Gertler and Watson (2004)). The estimated effects of oil supply and demand shock on US variables in each subsample are summarized in Table 4.

The table shows that the probability of a negative output response to an oil supply shock is slightly larger in the post 1987 sample. The estimated effects of the oil-demand shock turns out to be very similar to the one obtained on the full sample. Overall, we no major structural breaks in the effects of structural shocks emerge from the analysis of the subsamples. One possible difference with other studies is that

Table 4: US output response to oil shocks in different sample periods

	Probability of a negative response to oil-supply shock[a]				
At Horizon:	1	6	12	18	24
1981-2007	0.56	0.63	0.56	0.49	0.46
1984-2007	0.52	0.57	0.61	0.59	0.57
1987-2007	0.70	0.78	0.77	0.75	0.71
1991-2007	0.64	0.61	0.78	0.89	0.93
	Probability of a negative response to oil-demand shock[a]				
At Horizon:	1	6	12	18	24
1981-2007	0.34	0.37	0.36	0.37	0.38
1984-2007	0.32	0.39	0.42	0.46	0.50
1987-2007	0.29	0.37	0.40	0.44	0.49
1991-2007	0.28	0.31	0.32	0.35	0.39

Notes: –[a]Fraction of models $\tilde{\mathcal{A}}_0 \in \tilde{\mathbf{A}}_\mathbf{0}$ that yield a negative response at the given horizon.

our analysis concentrates on the effects on industrial production rather than GDP. It may be that the decline in energy intensity recorded by the industrial sector was smaller than in the rest of the economy.

5.3 The monetary policy rule

Bernanke et al. (2004) and Leduc and Sill (2004) argue that a change of the monetary rule is key in explaining the time-varying effects of oil price shocks on the US economy. We focus here on the real side of the economy. Recently, Herrera and Pesavento (2007) have explored the contribution of the US monetary policy response to oil price shocks and its role in the Great Moderation. They find that the magnitude and the duration of the response of output to an oil price shock has diminished during the Volcker-Greenspan era. The contribution of systematic monetary policy to the dynamic response of most macro variables during the post-1984 period appears to be significantly smaller.

We attempt to study this issue by including the federal funds rate in the VAR. The restrictions used to identify the 4 structural shocks are as in Table 1, i.e. no restriction on the response of the federal funds rate is imposed. Figure 9 shows the estimated effects of oil supply shocks for this new specification. The Fed Fund response to an oil supply shock is hardly different from zero, especially in the more recent sub-samples. A comparison with the estimates produced by the benchmark 4 variable VAR, suggests that the response of US industrial production is very similar for most of the sample periods considered. Overall, US monetary policy does not seem to have played a role in the transmission of oil supply shocks on the real side of the US economy. The lower panel of Figure 9 describes the IRFs following oil demand shocks. Industrial production increases following the shock. The Federal Funds rate increases gradually and remains above the baseline for a three years. As a consequence of a tightening monetary policy, the response of the US industrial production remains positive in the short run but it is less persistent than in our benchmark case. This picture remains broadly unchanged in all the sample splits considered. The results also show that the response of the Federal Funds rate varies with the underlying shock: a stronger response is detected in response to oil demand shocks.[11]

6 Concluding remarks

We argued that identifying the shocks underlying oil price fluctuations is key to predict their output effect. This point was illustrated with a simple open economy model, adapted from Backus and Crucini (2000), where two industrial countries produce and consume tradeable goods using energy supplied by a third oil-producing country. In the model the effects of an oil demand shock are very different from those of an oil supply shock. This prediction is tested on the data, where oil-demand and

[11]Similar results, available on request, are obtained when the VAR includes a long-term rate. The IRFs for the US structural shocks are also available. They show that after a US demand shock, the Federal Funds rate increase, consistently with much previous evidence. In line with the effects of technological innovation estimated by Dedola and Neri (2007), the US supply shock effect on the federal funds rate is tiny, with an almost equal probability of being either positive or negative.

supply shocks are identified exploiting sign-restrictions that hold across a large class of models.

The results show that following a negative oil-supply shock the US industrial production falls, consistently with the traditional textbook analysis. In contrast, the US industrial production rises following an oil-demand shock, suggesting that growing global demand, e.g. from Europe, China or India, increases the demand for oil as well as the demand for other goods, such as US export goods. These results offer a natural explanation for the smaller impact of oil price on real economic activity in the past few years compared to the seventies, as conjectured by Blanchard and Gali (2007).

References

Backus, D. K. and Crucini, M. J. (2000), 'Oil prices and the terms of trade', *Journal of International Economics* **50**(1), 185–213.

Backus, D. K., Kehoe, P. J. and Kydland, F. E. (1994), 'Dynamics of the trade balance and the terms of trade: The J-curve?', *American Economic Review* **84**(1), 84–103.

Barsky, R. B. and Kilian, L. (2004), 'Oil and the macroeconomy since the 1970s', *Journal of Economic Perspectives* **18**(4), 115–134.

Baxter, M. and Crucini, M. J. (1993), 'Explaining saving-investment correlations', *American Economic Review* **83**(3), 416–36.

Bernanke, B. S., Gertler, M. and Watson, M. (1997), 'Systematic monetary policy and the effects of oil price shocks', *Brookings Papers on Economic Activity* pp. 91–157.

Bernanke, B. S., Gertler, M. and Watson, M. (2004), 'Oil shocks and aggregate macroeconomic behavior: The role of monetary policy: A reply', *Journal of Money, Credit and Banking* **36**(2), 286–91.

Blanchard, O. J. and Gali, J. (2007), The macroeconomic effects of oil shocks: Why are the 2000s so different from the 1970s?, NBER Working Papers # 13368.

Canova, F. and Nicolo, G. D. (2002), 'Monetary disturbances matter for business fluctuations in the G-7', *Journal of Monetary Economics* **49**(6), 1131–1159.

Cunado, J. and Perez de Gracia, F. (2003), 'Do oil price shocks matter? evidence for some european countries', *Energy Economics* **25**(2), 137–154.

Dedola, L. and Neri, S. (2007), 'What does a technology shock do? a VAR analysis with model-based sign restrictions', *Journal of Monetary Economics* **54**(2), 512–549.

Francis, N. and Ramey, V. A. (2005), 'Is the technology-driven real business cycle hypothesis dead? shocks and aggregate fluctuations revisited', *Journal of Monetary Economics* **52**(8), 1379–1399.

Fry, R. and Pagan, A. (2007), Some issues in using sign restrictions for identyfing structural VARs, NCER Working Paper Series # 14.

Hamilton, J. D. (1983), 'Oil and the macroeconomy since world war II', *Journal of Political Economy* **91**(2), 228–48.

Hamilton, J. D. (1996), 'This is what happened to the oil price-macroeconomy relationship', *Journal of Monetary Economics* **38**(2), 215–220.

Hamilton, J. D. (2003), 'What is an oil shock?', *Journal of Econometrics* **113**(2), 363–398.

Hamilton, J. D. (2008), 'Oil and the macroeconomy', *The New Palgrave Dictionary of Economics. Second Edition.* S. Durlauf and L. Blume (Eds), Palgrave MacMillan Ltd.

Hamilton, J. D. and Herrera, A. M. (2004), 'Oil shocks and aggregate macroeconomic behavior: The role of monetary policy: Comment', *Journal of Money, Credit and Banking* **36**(2), 265–86.

Herrera, A. M. and Pesavento, E. (2007), Oil price shocks, systematic monetary policy and the great moderation, Michigan State Univeristy, Emory Univeristy, mimeo.

Hooker, M. A. (1996), 'What happened to the oil price-macroeconomy relationship?', *Journal of Monetary Economics* **38**(2), 195–213.

Jimenez-Rodriguez, R. and Sanchez, M. (2005), 'Oil price shocks and real GDP growth: empirical evidence for some OECD countries', *Applied Economics* **37**(2), 201–228.

Kilian, L. (2005), 'Exogenous oil supply shocks: How big are they and how much do they matter for the US economy?', *Review of Economics and Statistics, forthcoming.*

Kilian, L. (2007), The economic effects of energy price shocks, CEPR Discussion Papers # 6559.

Kilian, L. (2008), 'A comparison of the effects of exogenous oil supply shocks on output and inflation in the G7 countries', *Journal of the European Economic Association* **6**(1), 78–121.

Leduc, S. and Sill, K. (2004), 'A quantitative analysis of oil-price shocks, systematic monetary policy, and economic downturns', *Journal of Monetary Economics* **51**(4), 781–808.

Mork, K. A., Olsen, y. and Mysen, H. T. (1994), 'Macroeconomic responses to oil price increases and decreases in seven OECD countries', *Energy Journal* **15**(4), 19–35.

Rubio-Ramirez, J. F., Waggoner, D. and Zha, T. (2005), Markov switching structural vector autoregressions: Theory and application, Federal Reserve Bank of Atlanta Working Paper # 27.

Uhlig, H. (2005), 'What are the effects of monetary policy on output? results from an agnostic identification procedure', *Journal of Monetary Economics* **52**(2), 381–419.

Appendices

A The model economy

By the first welfare theorem the competitive equilibrium solves the planner's problem:

$$\max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \{ \frac{\left[c_t^{\theta}(1-n_t)^{1-\theta}\right]^{1-\gamma}}{1-\gamma} + \frac{\left[(c_t^*)^{\theta}(1-n_t^*)^{1-\theta}\right]^{1-\gamma}}{1-\gamma} + \left(\frac{(c_t^o)^{1-\gamma}}{1-\gamma} + \theta_L \frac{(1-n_t^o)^{1-\xi_L}}{1-\xi_L} \right)$$

$$+ p_{c,t}\left[\left(s_t\psi\left(a_{c,t}\right)^{1-\mu} + (1-s_t\psi)\left(b_{c,t}\right)^{1-\mu}\right)^{\frac{1}{1-\mu}} - c_t\right]$$

$$+ p_{c^*,t}\left[\left(s_t^*\psi\left(b_{c,t}^*\right)^{1-\mu} + (1-s_t^*\psi)\left(a_{c,t}^*\right)^{1-\mu}\right)^{\frac{1}{1-\mu}} - c_t^*\right]$$

$$+ p_{c^o,t}\left[\left(\psi_o\left(a_{c,t}^o\right)^{1-\mu} + (1-\psi_o)\left(b_{c,t}^o\right)^{1-\mu}\right)^{\frac{1}{1-\mu}} - c_t^o\right]$$

$$+ p_{y,t}\left[z_t\ n_t^{\alpha}\left(\eta\ k_t^{1-\nu} + (1-\eta)\ \ o_t^{1-\nu}\right)^{\frac{1-\alpha}{1-\nu}} - a_{c,t} - a_{i,t} - a_{c,t}^* - a_{i,t}^* - a_{c,t}^o\right]$$

$$+ p_{y^*,t}\left[z_t^*\ n_t^{*\alpha}\left(\eta k_t^{*1-\nu} + (1-\eta)\ \ o_t^{*\ 1-\nu}\right)^{\frac{1-\alpha}{1-\nu}} - b_{c,t} - b_{i,t} - b_{c,t}^* - b_{i,t}^* - b_{c,t}^o\right]$$

$$+ p_{o,t}\left[z^o + (n^o)^{\alpha} - o_t - o_t^*\right]$$

$$+ q_t\left[(1-\delta)\,k_t + \left(\frac{\left(\psi\left(a_{i,t}\right)^{1-\mu} + (1-\psi)\left(b_{i,t}\right)^{1-\mu}\right)^{\frac{1}{1-\mu}}}{k_t}\right)^{\phi} k_t - k_{t+1}\right]$$

$$+ q_t^*\left[(1-\delta)\,k_t^* + \left(\frac{\left(\psi\left(b_{i,t}^*\right)^{1-\mu} + (1-\psi)\left(a_{i,t}^*\right)^{1-\mu}\right)^{\frac{1}{1-\mu}}}{k_t^*}\right)^{\phi} k_t^* - k_{t+1}^*\right]\}$$

There are 20 choice variables (all chosen at time t):

$$c\ ,\ c^*,\ c^o,\ n\ ,\ n^*,\ n^o, a_c\ ,\ a_i\ ,\ a_c^*\ ,\ a_i^*\ ,\ a_c^o,\ b_c\ ,\ b_i\ ,\ b_c^*\ ,\ b_i^*\ ,\ b_c^o,\ o\ ,\ o^*, k_{t+1}\ ,\ k_{t+1}^*$$

and 16 endogenous variables:

$$y\ ,\ y^*,\ y^o,\ i\ ,\ i^*\ ,\ z,\ z^*,\ z^o, p_c\ ,\ p_{c^*}, p_{c^o}, p_y\ ,\ p_{y^*}, p_o, q^*\ ,\ q$$

The parameters used to construct the benchmark IRFs presented in figure 1 are given in table 5.

Table 5: Benchmark parametrization for the model economy

α	β	γ	δ	η	μ	ν	ψ	ϕ
0.64	0.99	1/0.5	0.025	0.9	2	11	0.8	0.99
ρ_z	ρ_{z^*}	ρ_{z^o}	θ	ψ^o	ξ_L	θ_L	ρ_χ	
0.5	0.5	0.5	0.4	0.5	5	0.6	0.8	

B Analysis with sampling and model uncertainty

This Appendix describes how we derive impulse response analysis when the model is subject to both sampling and model uncertainty. In this case, the informativeness of the sign restriction method is also affected by the uncertainty around the OLS estimates regarding reduced form VAR coefficients and the covariance matrix of reduced form innovations.

The empirical distribution for the impulse responses are derived in a Bayesian framework. As shown by Uhlig (2005) under a standard diffuse prior for $(B(L), \Sigma)$ and a Gaussian likelihood for the data sample, the posterior density for the reduced-form VAR parameters with sign restrictions is proportional to a standard Normal-Wishart. Thus one can simply draw from the Normal-Wishart posterior for $(B(L),\Sigma)$, and then use the algorithm by Rubio-Ramirez et al. (2005) to find an orthonormal theory-consistent identification matrix.

Operationally, our simulation is based on a two-step procedure. In the first step, we derive 1,000 random draws from the posterior distribution of the reduced form VAR coefficients, $B(L)$ and the covariance matrix of disturbance, Σ. In the second step, the procedure runs a loop. It starts by randomly selecting one draw from the posterior distribution of the reduced form VAR and, conditionally on it, uses the QR decomposition by Rubio-Ramirez et al. (2005) to find an impulse matrix satisfying the sign restrictions. Then, it selects an alternative draw

The loop ends when we obtain 10,000 identification matrices.[12] We notice that the number of theory-consistent models we choose to shape impulse responses should be enough large to avoid an important drawback. First, for each draw of the reduced form VAR the simulation should finds at least one orthonormal impulse matrix satisfying the sign restrictions. This allows us to derive posterior distribution for impulse responses which are not too dependent from few selected candidate draws of the reduced form.

In Figure 11 we report the resulting impulse responses of US output and CPI to both an oil supply and an oil demand shock allowing a comparison with the corresponding ones obtained by taking into account only the model uncertainty. The median estimated effects are almost identical, while the upper and lower confidence bands appear to be slightly wider. This evidence may suggest a relatively low uncertainty around OLS estimates of our parsimonious reduced form VAR.

[12] Our simulation works partly differently than those in Dedola and Neri (2007), and Uhlig (2005). In these papers the authors select a priori the number of draws for reduced form parameters and for each of them they draw a fixed number of impulse matrices. The resulting posterior distribution for impulse responses. Therefore, the number of accepted theory-consistent models is not fixed.

C Fry and Pagan critique

Fry and Pagan (2007) argue that the sign restriction approach is in principle flawed because it ends up reporting impulse responses drawn from different $\tilde{\mathcal{A}}_0$ models, thus possibly violating the assumption that the structural innovations are orthogonal. These authors recommend to check the robustness of the results by comparing them with the impulse responses drawn from a single $\tilde{\mathcal{A}}_0$, chosen to e.g. minimize some distance criterion from the median response. Appendix C shows that the results displayed above remain virtually unchanged when this prescription is followed.

This appendix explores the sensitivity of the impulse responses reported in Section 4 to the critique raised by Fry and Pagan (2007) to the sign restriction approach. These authors note that the practice of reporting selected statistics from the posterior distribution of, say, the magnitude of impulse responses is subject to a potential methodological flaw. While the model uncertainty captured by this distribution squares nicely with a Bayesian view of model uncertainty, it is important to realize that under this approach different models are used in representation of the e.g. median impulse response to a given shock. To see this consider the moving average representation for the VAR reduced form

$$y_t = C(L)e_t \tag{6}$$

where

$$C(L) = [I - B(L)L]^{-1}\tilde{\mathcal{A}}_0 \tag{7}$$

where $C(L)$ contains the matrices of the estimated impulse responses to structural shocks. Let $C^{(k)}_{i,j,h}$ denote the response of variable i to shock j at horizon h, where k indexes the value of the estimated response in the set of the theory-consistent models. It is straightforward to notice that there is no guarantee that the median response of variable i with respect to shock j at two different horizons, h and h', $(med(C^{(k)}_{i,j,h}), med(C^{(k)}_{i,j,h'}))$ is generated by the same model $\hat{k}$. This issue also arises in comparison across all variables, shocks and for any quantile of the impulse response distribution. The ensuing violation of the shocks orthogonality may cast doubt on the results of the effects of structural shocks.

One way to tackle this problem, suggested by Fry and Pagan (2007), is to perform the structural analysis using a single structural model $\tilde{\mathcal{A}}_0 \in \tilde{\mathbf{A}}_\mathbf{0}$, choosing the one whose impulse responses are "closest" to the median at all horizons. This strategy preserves the view that the median is an appealing way of summarizing the estimated effects of structural shocks at all horizons while ensuring the orthogonality of the shocks. Implementing this strategy requires us to define what we mean by "close". As impulse responses are not unit free, we first standardize them as

$$Z^{(k)}_{ij,h} = [C^{(k)}_{ij,h} - med(C^{(k)}_{ij,h})]/[stdev(C^{(k)}_{ij,h})] \tag{8}$$

The $Z^{(k)}_{ij,h}$ are then collected into a vector ϕ^k of dimension $(n{\cdot}h{\cdot}s{\times}1)$, where n is the number of variables in the VAR, h the horizons over which the impulse responses is computed, and $s <= n$ the number of identified structural shocks. In our case we have $n = 6$ $h = 36$ and $s = 4$, therefore, ϕ^k is a vector (846×1). Finally, we choose the value of k that minimizes $\phi^k\phi^{k\,\prime}$, and use it to derive the estimated impulse responses.

Figure 10 compares the median impulse responses of Figure 3 (dashed lines) with the one produced by the model satisfying the criterion describe above (solid lines). The dynamic effects of structural shocks are very similar, even if there are some differences in the magnitude of the responses. In particular, after a supply oil shock the negative response of the US industrial production appears to be more pronounced. After an oil demand shock the hike in the oil price is even larger, leading to a magnified increase in the US industrial production.

Figure 2: Effects of an oil-supply shock

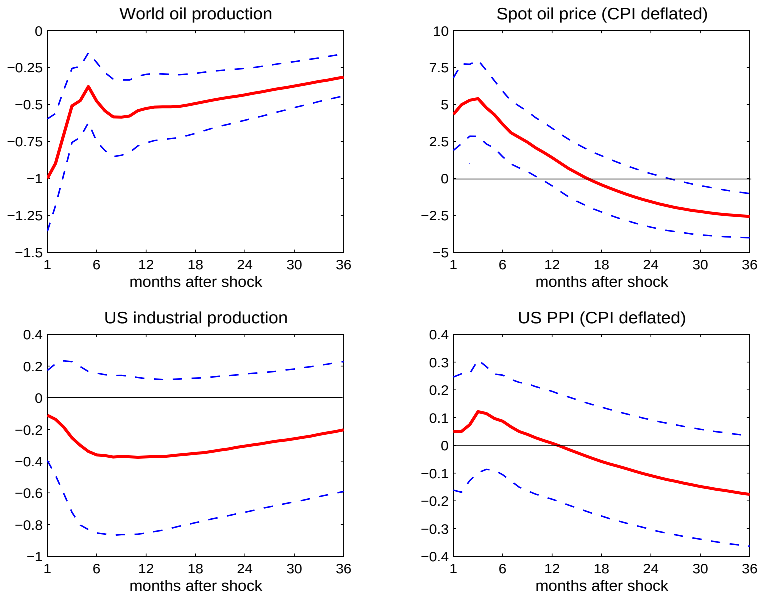

Figure 3: Effects of an oil-demand shock

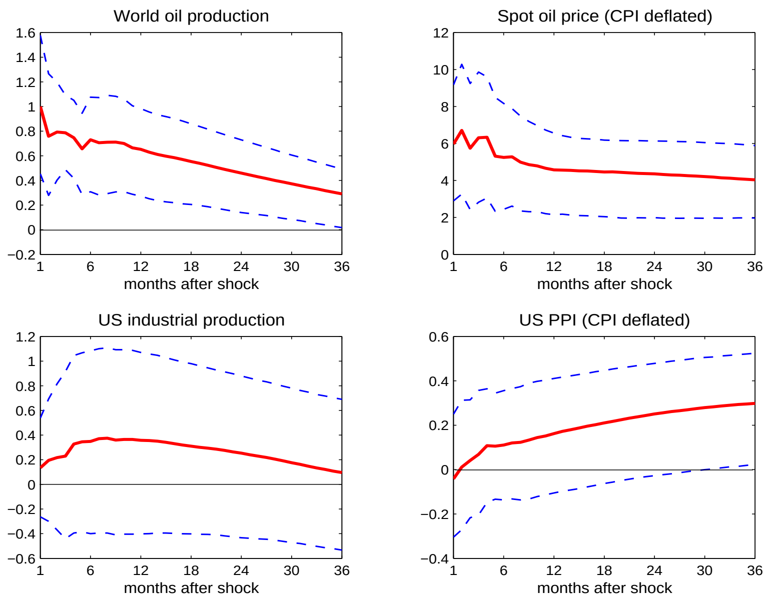

Note: The figures report the 16th, 50th and 84th percentiles of the IRFs distribution.

Figure 4: Effects of a US aggregate demand shock

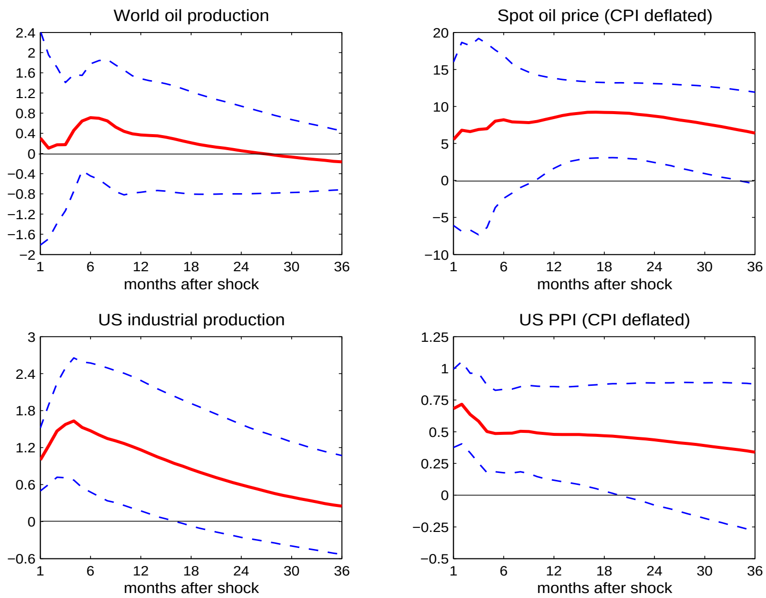

Figure 5: Effects of a US aggregate supply shock

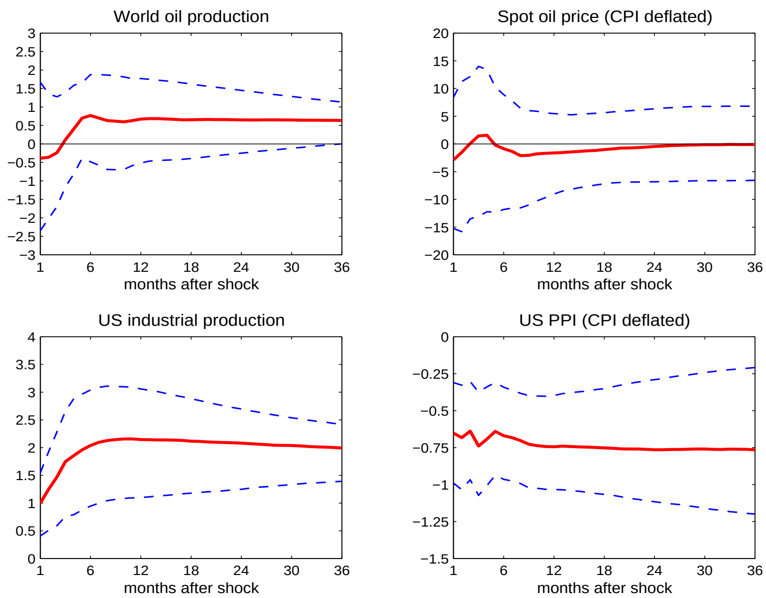

Note: The figures report the 16th, 50th and 84th percentiles of the IRFs distribution.

Figure 6: Historical decomposition

Oil Price

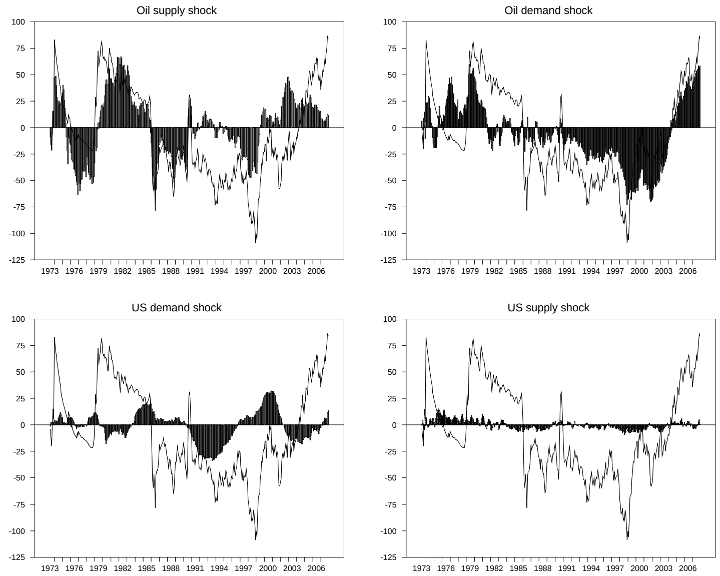

US industrial production

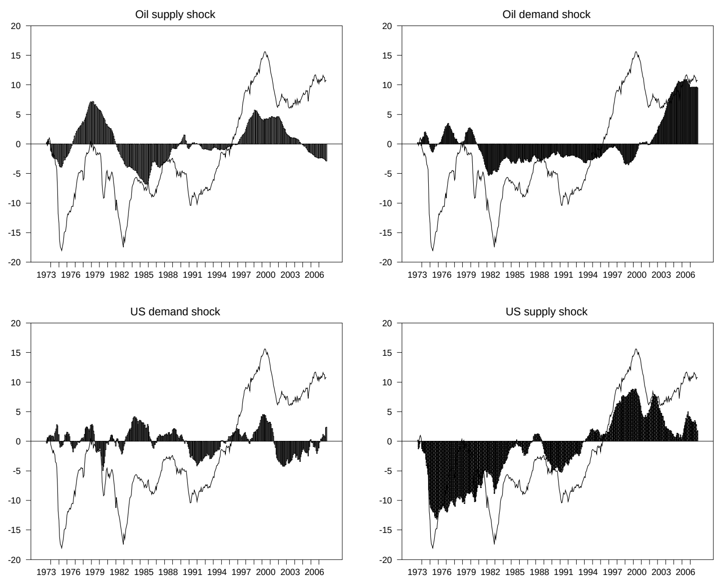

Figure 7: Robustness: Two vs. Four shocks

Oil supply shock

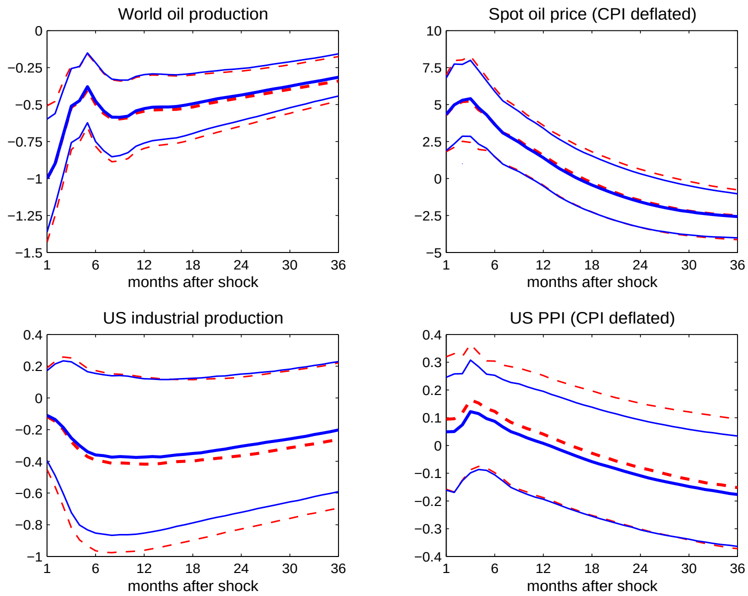

Oil demand shock

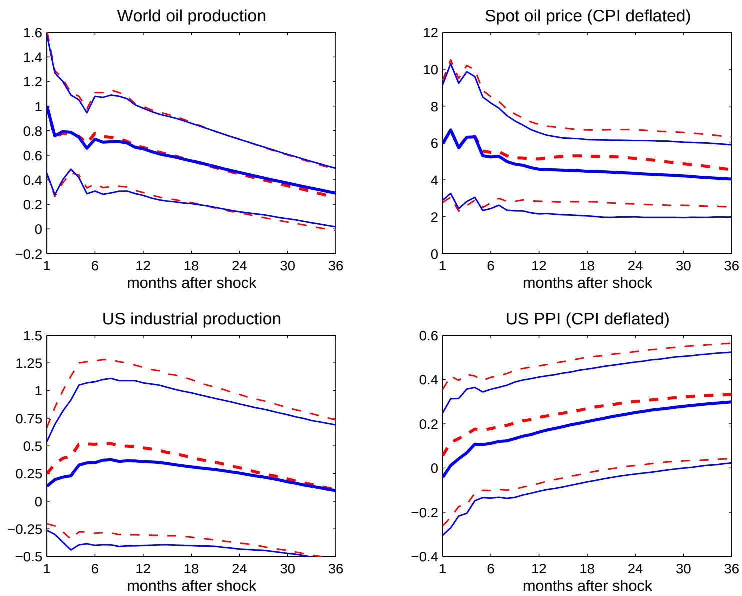

Note: The figure reports the 16th, 50th and 84th percentiles of the IRFs distribution. The solid lines are produced by the model with 2 identified shocks; the dashed lines are those of the benchmark model (with 4 identified shocks).

Figure 8: Effects of an oil demand shock using Kilian's Hypothesis

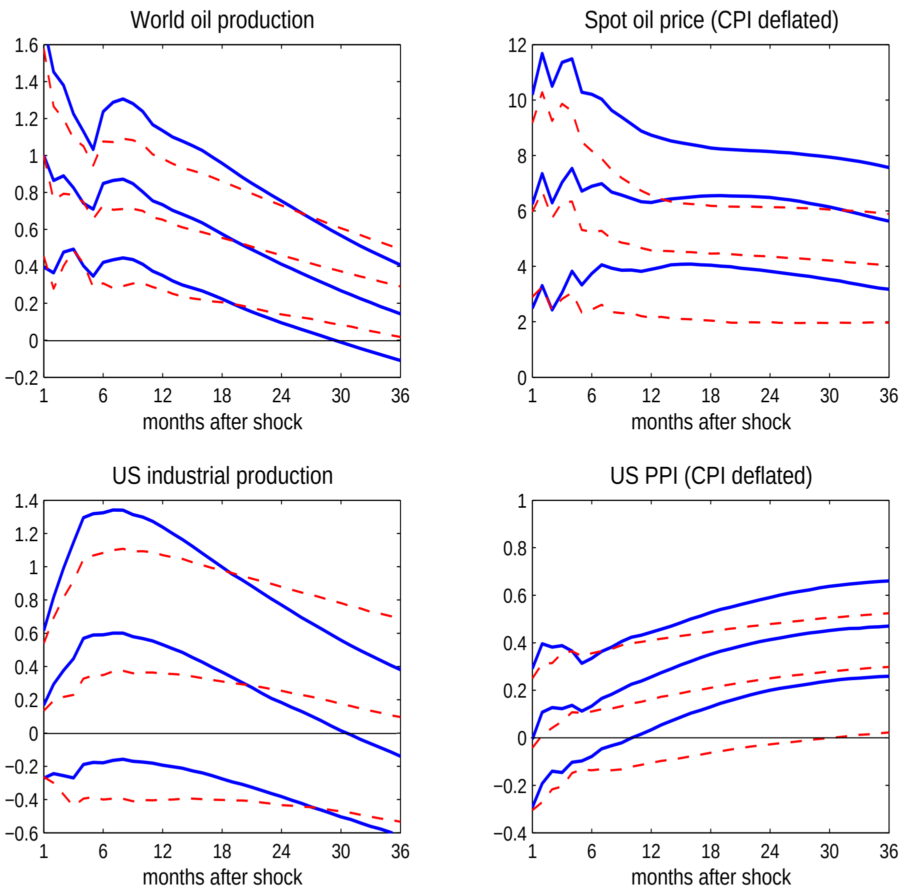

Note: The figure reports the 16th, 50th and 84th percentiles of the IRFs distribution. The solid lines are produced by the VAR including non-energy commodity price index; the dashed lines are those of the benchmark VAR (excluding non-energy commodity price index).

Figure 9: Robustness: VAR with Fed Fund rate

Oil supply shock

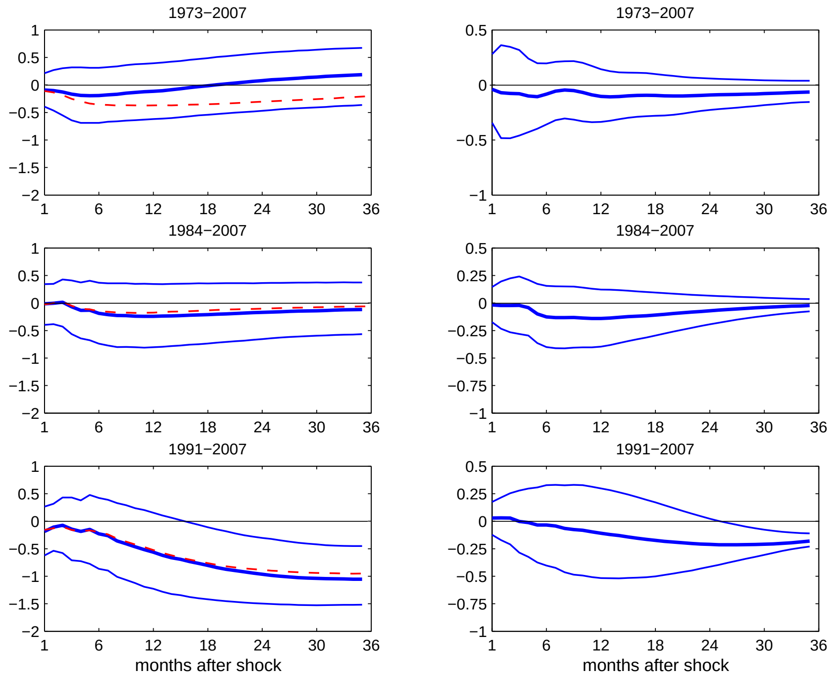

Oil demand shock

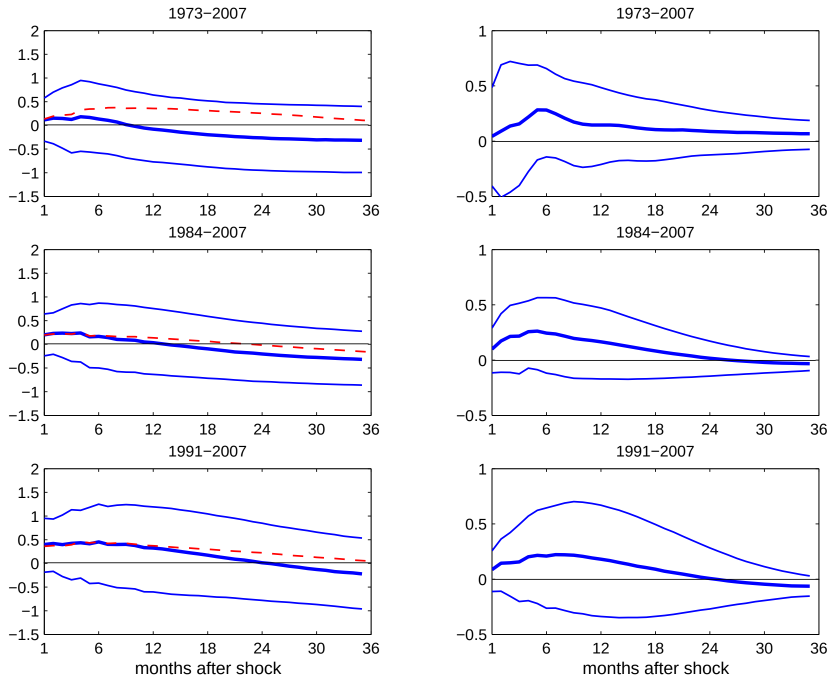

Note: The figure reports the 16th, 50th and 84th percentiles of the IRFs distribution. The solid lines are produced by the VAR including the Fed Fund rate; the dashed lines are those of the benchmark VAR.

Figure 10: Robustness: Results based on Fry and Pagan method

Oil supply shock

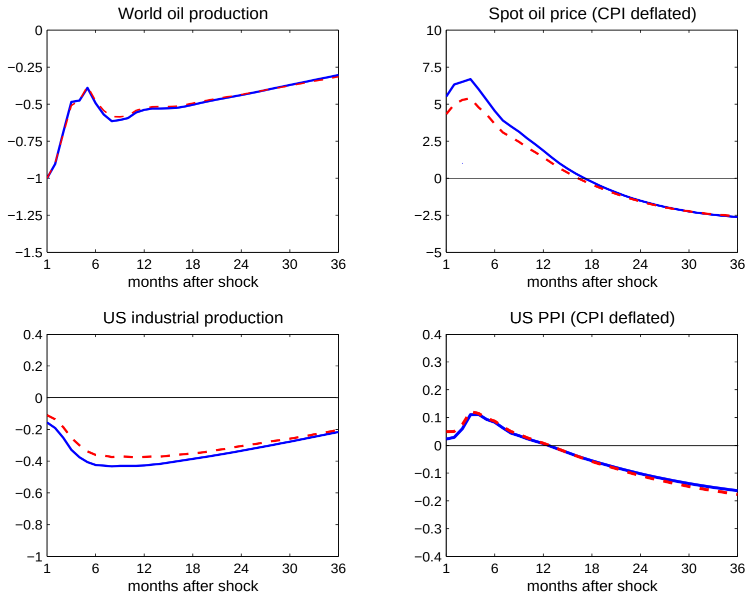

Oil demand shock

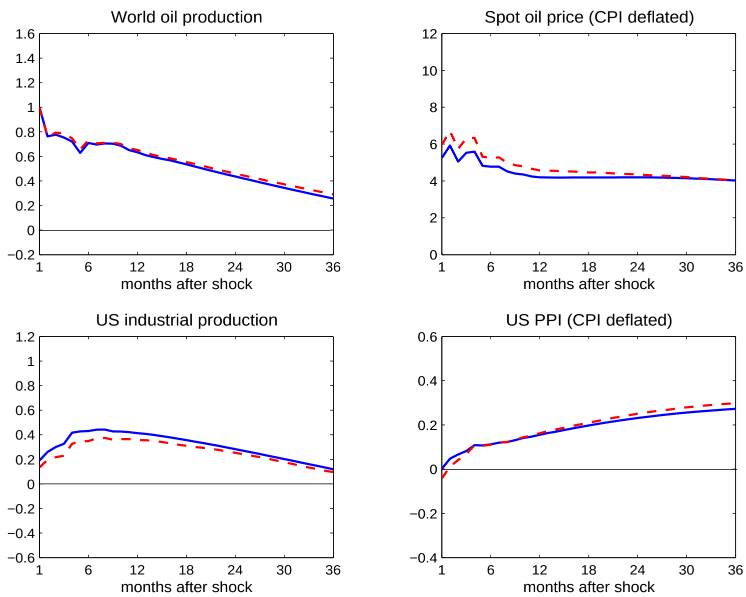

Note: The solid line is the response obtained applying the algorithm of Fry and Pagan (see Appendix C); the dashed line gives the median response of our benchmark VAR.

Figure 11: Response of US production with model and sampling uncertainty

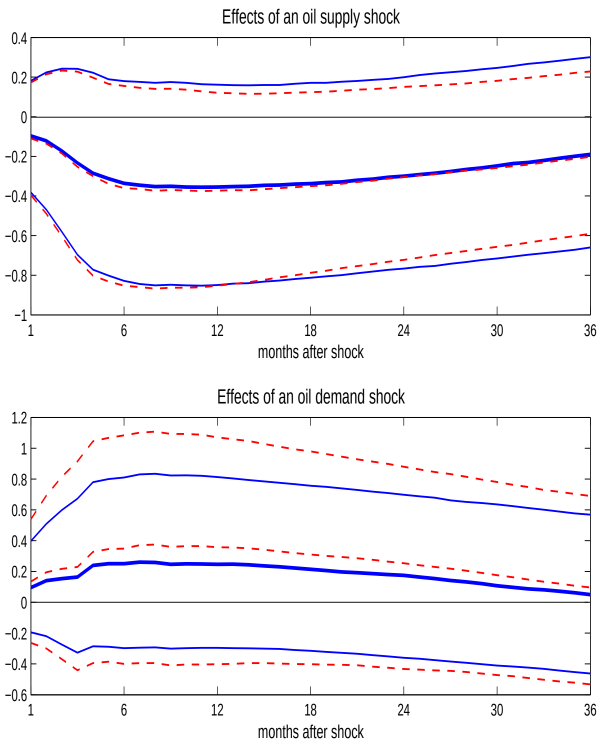

Note: The figure reports the 16th, 50th and 84th percentiles of the IRFs distribution. The solid lines are produced assuming model and sampling uncertainty; the dashed line assuming only model uncertainty.