



**EIEF Working Paper 09/10**  
**April 2010**

**Oil and the Macroeconomy:  
A Quantitative Structural Analysis**

by

**Francesco Lippi**

**(University of Sassari and EIEF)**

**Andrea Nobili**

**(Bank of Italy)**

# Oil and the macroeconomy: A quantitative structural analysis\*

**Francesco Lippi**

University of Sassari, EIEF  
`f.lippi@uniss.it`

**Andrea Nobili**

Bank of Italy  
`andrea.nobili@bancaditalia.it`

April 27, 2010

## Abstract

We model an economy where the cost of the oil input, industrial production, and other macroeconomic variables fluctuate in response to fundamental oil supply shocks, as well as aggregate demand and supply shocks generated domestically and in the world economy. We estimate the effects of these structural shocks on US monthly data over 1973.1-2007.12, using robust sign restrictions suggested by the model. It is shown that the interplay between the oil market and the US economy goes in both ways. First, US output falls below the baseline for a prolonged period of time after a negative oil supply shock. However, oil-supply shocks explain a relatively modest part of overall output fluctuations (about 10%). Second, most variations of (real) oil prices occur in response to shocks originated in the global economy and in the US. In particular, supply shocks in the rest of the world and in the US explain more than half of the variance of oil price fluctuations. Finally, the correlation between oil prices and the US business cycle depends on the nature of the fundamental shock: a negative correlation emerges in periods when oil-supply shocks or global demand shocks occur, while a positive correlation emerges in periods of supply shocks in the global economy or the US. The unconditional correlation between oil prices and US production over a long sample period is tenuous because it blends conditional correlations with different signs.

*JEL classification:* C32; E3; F4.

*Keywords:* Business cycle; Oil prices; Structural VAR.

---

\*We thank Fernando Alvarez, Fabio Canova, Luca Dedola, Lutz Kilian, Stefano Neri and Harald Uhlig for helpful comments. We also benefited from the comments by participants to seminars at the University of Sassari, EIEF, 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 institutions with which we are affiliated. 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.<sup>1</sup>

[Barsky and Kilian \(2002\)](#) and [Barsky and Kilian \(2004\)](#) call for a structural interpretation of the reduced form correlations mentioned above. These authors challenge the view that oil price shocks tend to be driven mainly by exogenous factors in the Middle East and claim that the OPEC decisions usually respond to global macroeconomic conditions affecting the demand for oil.<sup>2</sup> Following this view [Kilian \(2009\)](#) uses a VAR model to identify 3 structural shocks in the oil market by means of “impact restrictions”. Assuming a specific economically motivated recursive structure, he identifies oil supply shocks, world aggregate demand shocks and specific oil demand innovations, interpreted as reflecting fluctuations in precautionary demand for oil driven by fears about the availability of future oil supplies. One important message is that oil price changes have different effects on the US economy depending on the underlying shock that caused them.<sup>3</sup>

Disentangling the source of oil price fluctuations is also the question studied by this paper. We model the dynamics of the oil market and the US economy by a slight modification of the three-country model of [Backus and Crucini \(2000\)](#). The model assumes two industrial economy, the US and the global economy (the “rest of the world”), who produce differentiated good using capital and the oil input. Aggregate demand and supply in both industrial countries are subject to stochastic shocks, and so is the oil-supply. The model provides a mapping between these 5 fundamental shocks and the observed responses of production and relative prices. We estimate a 5 variable VAR that includes quantities and (real) prices in the oil market, quantities and (real) prices in the US economy and a measure of the global

---

<sup>1</sup>See [Hooker \(1996\)](#), [Hamilton \(2008\)](#) and [Edelstein and Kilian \(2009\)](#). Similar conclusions are reached for European countries by [Mork et al. \(1994\)](#) and [Cunado and Perez de Gracia \(2003\)](#).

<sup>2</sup>[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.

<sup>3</sup>In a number of related papers, this scheme was used by the author to address the effects of oil price shocks on several US macroeconomic variables, see e.g. [Kilian \(2008a\)](#) for a survey.

business cycle. The estimated VAR is used to identify 5 fundamental shocks using robust predictions on the sign of the impulse responses.

A novelty of this approach is that the interplay between the oil market, the global economy and the US economy is managed within a unified framework. This differs from the analysis in [Kilian \(2009\)](#), which assumes that shocks originated in the US affect the oil market only with delays. By allowing for the simultaneous interaction between the oil market and the US economy our analysis casts light on the assumption, often used in the empirical analysis, that oil-price shocks are predetermined to the US business cycle.<sup>4</sup> The empirical estimates of the theoretical model are informative about what fundamental shocks underlie observed fluctuations in the (real) price of oil. Moreover, the framework allows us to quantify the role of oil-supply shocks, and of the other fundamental shocks, with respect to the US business cycle.

The paper has 6 sections. The next one presents the theoretical framework. [Section 3](#) describes the estimation approach, whose results are given in [Section 4](#). [Section 5](#) discusses the robustness of the empirical findings. [Section 6](#) gives a summary of results.

## 2 Theoretical frame

We present a three-country model 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 et al. \(1994\)](#) and incorporate a country that produces oil. The model features supply shocks  $z_j$  in each country  $j$ . We provide a small addition to this 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.

Two industrialized and symmetric 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

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

where  $z$  is an AR(1) stochastic productivity shock  $z_t = \rho_z z_{t-1} + \tilde{z}_t$  with iid innovation  $\tilde{z}_t$ . 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 AR(1) exogenous stochastic oil supply component

---

<sup>4</sup>The recent papers by [Anzuini et al. \(2007\)](#) and [Baumeister and Peersman \(2008\)](#) study the effect of oil-supply shocks on the US output abandoning the recursiveness assumption. Compared to these papers our study also considers oil-demand shocks and shocks to the US business cycle.

and  $(n_t^\circ)^\alpha$  the 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 ( $c$ ) using the CES aggregator

$$c(a, b, \psi_t) = [\psi_t a^{1-\mu} + (1 - \psi_t) b^{1-\mu}]^{1/(1-\mu)} \quad (2)$$

The consumption bundle is subject to stochastic AR(1) preference shocks, such that  $\psi_t \equiv s_t \cdot \psi$  with  $s_t = (1 - \rho_s) + \rho_s s_{t-1} + \tilde{s}_t$  and  $\tilde{s}_t$  is i.i.d.<sup>5</sup> An identical aggregator, with deterministic weight  $\psi$ , is used to produce the investment good,  $i$ .

Capital follows 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 the US and the RoW maximize the expected value of lifetime utility

$$\max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t U(c_t, 1 - n_t), \quad U(c, 1 - n) = \frac{[c^\theta (1 - n)^{1-\theta}]^{1-\gamma}}{1 - \gamma} \quad 0 < \gamma, \quad 0 < \theta < 1 \quad (3)$$

where  $\beta < 1$  is the intertemporal discount, and the intertemporal and intra-temporal (consumption - leisure) substitution elasticities are constant, equal to  $1/\gamma$  and 1, respectively. As in [Backus and Crucini \(2000\)](#) a different utility function is assumed for oil producers, consistent with an inelastic labor supply.<sup>6</sup> Prices and allocations are solved for a competitive equilibrium. As usual, we appeal to the first welfare theorem and compute allocations by solving a standard planning problem.

The model is used to examine the effects of supply side (productivity) and demand (preference) shocks in each economy. Since we are interested in economic implications that are robust we follow [Canova and Paustian \(2007\)](#) and [Dedola and Neri \(2007\)](#), and assess the response of endogenous variable to the different structural shocks under a range of parameterizations centered around the values used in [Backus and Crucini \(2000\)](#). We then develop Monte Carlo simulations assuming that the relevant structural parameters are uniformly and independently distributed over the range described in [Table 1](#). For each simulation the parameters are drawn from the uniform densities, and the impact and dynamic responses of all variables to fundamental shocks are recorded. Then we compute the median, the 5th and 95th percentiles of the resulting distribution of responses, point by point. This ensures that parameters' combinations that bring about extreme responses in the tails are ruled out.

---

<sup>5</sup>Similar effects are obtained by considering shocks to the intertemporal discount factor.

<sup>6</sup>Specifically the model assumes  $U(c, (1 - n)) = c^{1-\gamma}/(1 - \gamma) + \theta_L (1 - n)^{1-\xi_L}/(1 - \xi_L)$ . The separability simplifies the solution of the model; coupled with a low labor supply elasticity ( $\xi_L \approx 5$ ) this reproduces the observed low responsiveness of oil production to the relative price of oil, or production in OPEC countries (see the discussion on page 197 - 198 in [Backus and Crucini \(2000\)](#)).

The median and the percentiles of distribution of the theoretical impulse responses to each shock are shown in [Figure 1](#). Each column of the figure reports (from top to bottom) the impulse response functions of oil production, oil price, the US output level, the price of US output and the RoW output level to the structural shock indicated below the column. All prices are expressed relative to the US consumption deflator that is chosen as the numeraire, as done later in the empirical analysis. Notice that one consequence of choosing the US consumption price index as the unit of account is that the impulse responses of the (relative) price variables will appear asymmetric even though the model and the benchmark parametrization chosen are completely symmetric.<sup>7</sup>

Table 1: Parameter ranges in the model economy

<i>simulated parameters</i>		Range of values
$1/\gamma$	Intertemporal elasticity of substitution	[1.0,2.5]
$1/\mu$	Elasticity of substitution between home and foreign good	[0.5,2.0]
$1/\nu$	Elasticity of substitution between oil and capital	[4.0,12.0]
$\psi$	Preference for home good	[0.7,0.9]
$\rho_z$	Persistence of US supply shock	[0.5,0.99]
$\rho_{z^*}$	Persistence of RoW supply shock	[0.5,0.99]
$\rho_{z^o}$	Persistence of oil supply shock	[0.5,0.99]
$\rho_s$	Persistence of US demand shock	[0.5,0.99]
$\rho_{s^*}$	Persistence of RoW demand shock	[0.5,0.99]
<i>calibrated parameters (quarterly model)</i>		
$\alpha$	Labor's share in the industrial country	0.64
$\beta$	Intertemporal discount factor	0.99
$\delta$	Depreciation rate of capital	0.025
$\eta$	Oil weight in technology	0.9
$\phi$	Investment adjustment costs	0.99
$\theta$	Labour share in the industrial country	0.4
$\psi^o$	Oil share in the oil producing country	0.5
$\theta_L$	Labor share in the oil producing country	0.6
$\xi_L$	Inverse labor elasticity in the oil producing country	5.0

The first column describes the effect of a positive oil supply shock ( $z^o > 0$ ). The shock

<sup>7</sup> As an example compare the response of the oil price to a RoW demand shock with the one to a US demand shock, respectively in the third and fifth columns of [Figure 1](#). Since the model is symmetric one expects the effect of each of these shocks to be identical. Indeed the effects are identical if one looks at the real variables (e.g. the response of the oil production). But the response of the real oil-price is different across the two columns: the reason is that in both experiments the oil price ( $p_o$ ) is deflated by the US CPI ( $p_c$ ). And the dynamics of the US CPI depend on whether the shock is a US demand (in which case the US CPI increases more, hence the negative response of the real oil price) or whether it is a RoW demand (in which case the US CPI does not increase much and hence the positive response of the real oil price).

moves the real oil price and the oil quantity in opposite directions. This represents the prototype textbook case of an exogenous oil supply shock. As the relative price of the energy input falls, production in both the RoW and the US increases. The price of US output (deflated by the US CPI) is basically constant in response to the oil supply shock. The reason is that the model is symmetric so that the higher cost of oil impacts in exactly the same proportion into both foreign and home prices, leaving relative prices constant.

Supply ( $z^* > 0$ ) and demand ( $s^* > 0$ ) shocks in the RoW impact on the oil market and on the US economy. The price and quantity of oil display a positive covariance in response to both shocks, as shown by the starred lines in the second and third columns of [Figure 1](#), i.e. these shocks appear to an observer of the oil market as “oil demand shocks”. But the figure also shows that US production reacts with an opposite sign following these shocks over the first year after the shock. The different response of the US output is due to the response of the US terms of trade (see [Figure 2](#)). A positive RoW supply shock depreciates the terms of trade (less unit of the domestic good are required to buy the foreign good). At the same time, the higher demand for oil caused by the output boom in RoW raises oil prices. The net effect on the US output depends on the balancing of these two effects. If the trade links between the RoW and the US are strong, the impact of the RoW boom on oil prices is outweighed by the increased demand of US exports to the RoW, and US output increases following a supply shock in RoW. The response of US output to the demand shock in RoW is unambiguous. In this case the cost of oil goes up, and the terms of trade appreciates: both effects contribute to reducing the US output.

The fourth and fifth columns describe the effects of US shocks. 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 real oil prices and output, though the latter is not robust across the different parameterizations. Finally, a US demand shock ( $s > 0$ ) increases US production and its price (relative to the CPI). The increased demand spills over to the oil market, where production increases (see [footnote 7](#) for a discussion of the response of the real oil price).

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 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 the empirical validity of this conjecture.

### 3 The design of the empirical analysis

The empirical analysis identifies a set of structural shocks, consistent with the theory above, and studies their effects on the real oil price and US output. The identification method is based on sign restrictions, following the approach pioneered by [Davis and Haltiwanger \(1999\)](#), [Canova and Nicolo \(2002\)](#) and [Uhlig \(2005\)](#). The idea is to use some robust properties of the model, namely 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 a specific parameterization of the model is the true data generating process. Next we describe the VAR specification, the data, and the identification assumptions.

The analysis is based on the vector autoregression (VAR)

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

where  $B(L)$  is a lag polynomial of order  $p$  and  $y_t$  contains five variables (all in logs) describing the US, the RoW and the oil market. The first two are the US industrial production and the producer price index. Two additional variables describe the oil market: the (real) spot oil price and the global oil production.<sup>8</sup> The fifth variable is total imports to the RoW, aiming at capturing economic activity in the RoW.<sup>9</sup> Estimation of the VAR is based on monthly data spanning the period January 1973 - February 2009 (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, including the most recent one. We complete the specification by using a lag order of 2 months, as suggested by the Bayesian Information Criteria (BIC).<sup>10</sup>

---

<sup>8</sup>Production is in barrels per day and provided by the International Energy Agency (IEA); the spot oil price is from the International Monetary Fund, computed as the simple arithmetic average of the U.K. Brent, Dubai Fateh and West Texas Intermediate spot prices. It is expressed in real terms (e.g. deflated by the US consumer price index). In this paper we refer to this variable as our measure of real oil price. Our results are unchanged if we use a single oil price measure instead of the simple average.

<sup>9</sup>We cannot use a RoW industrial production index due to the lack of a sufficiently long time series for output in a country group for RoW that includes China and India. In [Kilian \(2009\)](#) the measure of global real economic activity is based on a global index of dry cargo single voyage freight rates (deflated by the US CPI). Increase in freight rates may be used as indicators of cumulative global demand pressures. This measure, however, does not distinguish between increase in demand stemming from the US to those originated in the RoW. Our results are robust to the use of Kilian’s measure of world output instead of the RoW imports.

<sup>10</sup>This lag order provides estimated residuals for the reduced-form VAR characterized by good white-noise properties. The appropriate lag length was debated in previous literature, see [Hamilton and Herrera \(2004\)](#) and [Kilian \(2009\)](#). Our results remain virtually unchanged if 7 lags (as suggested by the Akaike Information Criteria) or 12 lags (as in [Kilian \(2009\)](#)) are used.



The structural VAR approach sees [equation \(4\)](#) as a reduced form representation of the structural form

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

where  $A(L)$  is a lag polynomial of order  $p$  and the vector  $e$  includes the five 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 orthogonal 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  $k$ ,  $\Psi_k$ , can then be computed using the  $B(L)$  estimates from the reduced form VAR,  $B_1, B_2, \dots, B_p$ , and the impulse vector  $a_j$ .<sup>11</sup>

The sign restriction approach identifies a set of structural models, the  $\tilde{\mathcal{A}}_0 \in \tilde{\mathbf{A}}_0$ , such that the (vectors of) impulse responses  $\Psi$ s 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 orthogonal matrix  $Q$ . The set of the theory-consistent models,  $\tilde{\mathbf{A}}_0$ , can be characterized as follows. For a given estimate of the reduced form VAR,  $B(L)$  and  $\Sigma$ , take an arbitrary identification matrix  $\mathcal{A}_0$  and compute the set of candidate structural models  $\hat{\mathbf{A}}_0 = \{\mathcal{A}_0 Q \mid QQ' = I\}$  by spanning the space of the orthogonal matrices  $Q$ . The set  $\tilde{\mathbf{A}}_0$  is then obtained by removing from the set  $\hat{\mathbf{A}}_0$  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.

In practice we also have to decide on for how long the sign restrictions used for identification should hold. In this regard, [Canova and Paustian \(2007\)](#) show that sign restrictions imposed on the contemporaneous relationships among variables are robust to several types of model misspecification. Following this approach, we impose the sign restrictions only on impact. As several signs of the impulse responses depend on the model parametrization, the identification restricts attention to robust features of the contemporaneous impact responses obtained by Montecarlo simulations. The ranges for the parameters used in the simulations are given in [Table 1](#). The results of these simulations are reported in [Figure 3](#).

In the empirical analysis we restrict attention to 5 mutually orthogonal shocks: an oil supply shocks, supply and demand shocks in the RoW, and supply and demand shocks in the

---

<sup>11</sup>The vectors of responses for all horizons  $s$ ,  $\Psi_s$ , are computed recursively using the vector of contemporaneous responses,  $\Psi_0 = a_j$ , and the reduced form  $B$  matrices:  $\Psi_s = \sum_{i=0}^s B_{s-i}\Psi_i$  for  $s \geq 1$  and  $B_{s-i} = 0$  if  $s - i > p$ .

US. Next we describe the identifying assumption for each shock, consistent with the model robust properties, which are summarized in [Table 2](#). We define as an oil-supply shock one that causes the oil production and its real price (CPI deflated) to move in opposite directions, and both the RoW and the US output to decrease, as shown in the first column of the figure. We define a RoW supply shock as one that moves in the same direction the RoW output, the real oil price and the US relative price (the response of the oil quantity and US output are left unconstrained). A positive RoW demand shock raises the price and the quantity of oil, the RoW output and decreases the US industrial production. US shocks are described in the fourth and fifth columns of [Figure 1](#). A positive shock to the US supply is one that induces a negative correlation between the US industrial production and its deflator and increases the real oil price. A positive US demand shock is one that generates a positive response of the oil production, the US industrial production and its deflator (relative to the CPI), and reduces the real oil price and RoW output. The last restriction on the RoW output is useful because it allows us to distinguish a (negative) oil supply shock from a (positive) US demand shock. It is important to remark that our identification scheme defines mutually exclusive structural shocks, thus avoiding the possibility that we are confusing shocks originated in the rest of the world with US specific shocks. In this regard, [Table 2](#) shows that a RoW supply shock cannot be mixed up with a US supply shock as the US relative price variable (US CPI/PPI) responds with an opposite sign to these shocks. At the same time, a RoW demand shock cannot be confused with a US supply shock as the US output responds with opposite sign to these two shocks. Similarly, we are also able to disentangle RoW shocks from a US demand shock. Indeed, a RoW supply shock is distinguished from a US demand shock as the response of the RoW output is opposite in sign to the these two shocks; finally, a RoW demand shock is different from a US demand shock because oil production and the real price of oil comove in response to the former, while they exhibit opposite responses in sign to the latter.

## 4 The estimated effects of structural shocks

This section presents our estimates on the effect of the various structural shocks. 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  $(B(L), \Sigma)$ .

The set of theory-consistent matrices  $\tilde{\mathbf{A}}_0$  is computed using the efficient algorithm pro-

Table 2: Sign restrictions used for identification

VAR Variables	Structural shocks				
	oil-supply	RoW - supply	RoW - demand	US supply	US demand
oil production	–		+		+
oil price <sup>a</sup>	+	+	+	+	–
US output	–		–	+	+
RoW output	–	+	+		–
US output price <sup>a</sup>		+	–	–	+

Note: A “+” (or “–”) sign indicates that the impulse response of the variable in question is restricted to be positive (negative) on impact. A blank entry indicates that no restrictions is imposed on the response. –<sup>a</sup>Price is deflated by the US CPI.

posed by [Rubio-Ramirez et al. \(2005\)](#). Given an estimate for  $(B(L), \Sigma)$  and one candidate identification matrix  $A_0$  (i.e. a Choleski decomposition), the algorithm draws an arbitrary independent standard normal  $(n \times n)$  matrix  $X$  and, using the  $QR$  decomposition of  $X$  generates one orthogonal matrix  $Q$ . Impulse responses are then computed using  $A_0Q$ , the rotation of the initial identification matrix, and  $B(L)$ . If these impulse responses do not satisfy the sign restrictions the algorithm generates a different draw for  $X$ . Compared with Uhlig’s procedure, this algorithm directly draws from a uniform distribution instead of involving a recursive column-by-column search procedure. Thus, the informativeness of the sign restriction method is affected by the sampling uncertainty around the estimates regarding reduced form VAR coefficients and the covariance matrix of reduced form innovations, as well as by the model uncertainty inherent to the possible outcomes (e.g. matrices  $\tilde{\mathbf{A}}_0$ ) that are consistent with the set of theoretical restrictions.

Operationally we use a two-step procedure. In the first step we generate 2,000 random draws from the posterior distribution of the reduced form VAR coefficients,  $B(L)$  and the covariance matrix of disturbances,  $\Sigma$ . 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 5,000 identification matrices are found. By construction, each of the models in  $\tilde{\mathbf{A}}_0$  generates orthogonal structural shocks.<sup>12</sup>

We notice that the number of theory-consistent models we choose to compute is large, so that for each draw of the reduced form VAR the simulation the algorithm finds at least one

<sup>12</sup>This algorithm differs from [Dedola and Neri \(2007\)](#), and [Uhlig \(2005\)](#), who select a priori the number of draws for the reduced form parameters and, for each of them, draw a fixed number of impulse matrices. In their approach, the number of accepted theory-consistent models is not constant.

identification matrix satisfying the sign restrictions. This helps us ensure that the posterior distribution for impulse responses that we obtain does not depend on a few selected candidate draws from the reduced form.<sup>13</sup>

## 4.1 Impulse responses

Following [Dedola and Neri \(2007\)](#) and [Uhlig \(2005\)](#), we report in [Figure 4](#) the median (solid line), the 16th and the 84th percentiles (the dashed lines) of the distribution of impulse responses produced by the algorithm discussed above for each variable over 24 months.

The effects of an oil supply shock, normalized to yield a 1 per cent reduction in oil production, are displayed in the first column of [Figure 4](#). The shock lowers the US industrial production, that reaches a trough after about 12 months. In [Table 3](#) we follow the terminology of [Dedola and Neri \(2007\)](#) and interpret this fraction as a probability. The figure shows that after one year the response is negative for 100% of the models. Notice that our identification scheme imposes a negative response of the US output only on impact, so that the persistence of the response is really a finding that is coming from the data, it is not a necessary implication of the identification assumption.

Table 3: US output response to different structural shocks

At Horizon:	“Probability” (% of estimated models) of a negative response of US output <sup>a</sup>				
	1	6	12	18	24
Oil-supply shock	1.00	1.00	1.00	1.00	1.00
RoW supply shock	0.11	0.09	0.42	0.76	0.94
RoW demand shock	1.00	1.00	1.00	1.00	1.00

Note: <sup>a</sup>Fraction of models  $\tilde{\mathcal{A}}_0 \in \tilde{\mathbf{A}}_0$  that yield a negative response at the given horizon.

The effects of a RoW supply shock, normalized to yield a 1 per cent increase in the RoW output, are displayed in the second column of [Figure 4](#). The response of the US industrial production differs markedly from the case of the oil-supply shock: production 12 months after the shock is *above* the baseline. [Table 3](#) shows that the fraction of models in which the US industrial production increases conditional on a positive RoW supply shock is about 90% after 6 months, and 60% after 1 year. The effect of a RoW demand shock on the US industrial production (normalized to produce a 1 per cent increase in the RoW output) is

<sup>13</sup> In order to disentangle the different structural shocks, we must use several robust sign-restrictions. This makes the analysis particularly severe from a computational viewpoint. [Canova and Paustian \(2007\)](#) have shown that what matters for identification is the combination of the number of restrictions and the magnitude of the variance of the shocks in the sample period considered. In particular, when a small number of identification restrictions is used the identification becomes weak and, unless, the variance of the shock very is large, results are rarely sharp.

instead negative and persistent at all horizons. The different effects of RoW demand and supply shocks for US output, that were discussed in [Section 3](#), appear important empirically. These findings suggest that the different response of the terms of trade to foreign domestic and supply shocks determines the sign of the response of US output to innovation in global business cycle, to horizons up to 12 months.

The fourth and fifth columns of [Figure 4](#) illustrate the extent to which the oil market is affected by US shocks. A positive US aggregate supply shock raises both oil quantity and prices. A positive US aggregate demand shock raises the US production and causes a small decrease in real oil prices (as the direct effect is to raise the price of domestic output more than the price of the oil input) and a significant increase in the oil production.

Our findings concerning the effect of an oil supply shock are qualitatively comparable to those in [Kilian \(2009\)](#), even though the negative response of the US output is larger and more persistent in our estimates. The main difference concerns the effect of the shocks in the RoW. In his analysis an expansion of the global business cycle (what Kilian labels an “aggregate demand shock”) causes a statistically not-significant increase in real GDP in the first year, followed by a gradual *decline* which becomes significant in the third year. Our predictions for the long run are similar, but they differ over the first year, where we find that the US output response may be either positive or negative depending on whether the fundamental innovation underlying the expansion of the global business cycle is a supply or a demand shock. Our model provides a simple explanation for the different findings: RoW demand and supply have similar dynamic effects on the oil market (e.g. move price and quantity in the same direction) but they cause effects opposite in sign on the US output at least at horizons up to one year (see [Figure 1](#) and [Figure 4](#)). This suggests that by mixing together RoW supply with RoW demand shocks, one may bias the response of US output towards zero at horizons up to one year. The variance decomposition analysis presented in the following section corroborates this hypothesis since at frequencies up to one year the contribution of these two shocks to the variation of US production is about equal.

Altogether the estimates show that identifying the fundamental shock underlying the oil-price hike is key to predict the dynamics of US output conditional on observing an oil-price increase. A higher oil price is associated to an expected reduction of US production conditional on an adverse oil-supply shock, a negative US demand shock or a positive RoW demand shock. Instead, a higher oil price is associated with an expected rise of US production conditional on a positive RoW supply shock or a US supply shock.

## 4.2 Variance decomposition

We analyze the contribution of the different structural shocks to fluctuations in the real cost of the oil input and the US output by performing a variance decomposition analysis. [Figure 5](#) reports the median (solid line), the 16th and the 84th percentiles (the dashed lines) of the distribution of the variance decomposition at horizons up to 24 months for these variables. The first row of the figure shows that supply shocks generated in the RoW explain the largest fraction of the oil price variance, between 40 to 60 per cent, over the horizons considered. This shock represents, by far, the largest source of fluctuations in oil prices. Oil-supply shocks explain a large portion of the variance within the one year horizon. The US aggregate supply shocks account for about 10 percent of the oil-price variance, with a larger share at short term horizons. The second row of [Figure 5](#) 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 recent contributions of [Dedola and Neri \(2007\)](#), and [Francis and Ramey \(2005\)](#). The role of US aggregate demand shocks is also large, though it is smaller than the role of US supply shocks at short term horizons. The other shocks affecting US production, the oil-supply and the RoW shocks, are less important than domestic shocks. The estimated median effect (marked by the continuous line) suggests that each of these shocks explains about 10 percent of total variance.

We also explored whether, as argued by [Fry and Pagan \(2007\)](#), the sign restriction approach is flawed because the impulse response functions that are generated likely violate the assumption that structural innovations are orthogonal. As suggested by [Dedola and Neri \(2007\)](#) to ensure orthogonality of the structural shocks the variance decomposition analysis could be assessed from a unique  $\tilde{\mathcal{A}}_0$ , chosen so as to minimize a minimum distance criterion from the median responses. Details on this analysis are given in the [Online Appendix A](#). To this end, we also report the percentage of the variance of the  $k$ -step ahead forecast error computed by choosing the  $\tilde{\mathcal{A}}_0$  whose impulse responses are “closest” to the median at all horizons. It corresponds to the red starred line reported in each panel of [Figure 5](#)). It provides results that are similar to those produced by the median of the forecast variance posterior distribution implied by the set of  $\tilde{\mathcal{A}}_0$  models.

## 4.3 Historical decomposition

In this Section we provide a structural interpretation for those historical episodes characterized by major oil price increases, and compare our findings with previous studies. The historical decomposition of the real oil price time series and that of the US industrial production are displayed in [Figure 6](#). This figure highlights the contribution of each structural

shock to deviations of the variables from the baseline at each point in time.

The question of whether the oil price hike recorded in the late 1973 is an oil supply shock or an oil demand shock has been widely discussed. Barsky and Kilian (2002) and Kilian (2008b) showed that it could have been a delayed consequence to a demand shock in the presence of price regulation. Figure 6 shows that this episode could have been largely driven by a positive RoW supply shock and, only to a smaller extent, by a conventional oil supply shock. Therefore, our results also seem to downplay the role of oil supply shocks.<sup>14</sup>

Oil supply shocks are important in explaining the sharp fall in the real price of oil following the collapse of the OPEC cartel in late 1985 (probably as the direct consequence of the increase in Saudi oil production) and the sharp spike in the real price of oil in 1990-1991 after the invasion of Kuwait. One robust pattern is that both RoW and US supply shocks appear to have been a key factor in many episodes characterized by oil price hikes. For example, RoW supply shocks seem to explain the rapid oil price hikes that started in 2003 and the subsequent sharp and strong reduction of oil price recorded in the second half of 2008, as well as the drop that occurred after the Asian crisis of 1997-1998. The US supply shocks are likely to have been relevant in sustaining the high level of the real price of oil during the eighties.

Regarding the fluctuations of US output the historical decomposition shows that US shocks dominate the shocks originated in the RoW. However, the distinction between RoW supply and RoW demand shocks remains relevant for the interpretation of the effects on US output. Regarding the 1990/91 episode, the real oil price increase was mainly driven by a RoW supply shock and partially by an oil supply shock; not surprisingly, this oil price shock only partially contributed to the following US recession, which was instead mainly the result of a negative US demand shock.

## 5 Robustness

The robustness of the findings was tested along several dimensions. First, we check whether our results are robust to the use of alternative identification schemes. In Figure A.2 we compare impulse responses of our benchmark model with those stemming from the same

---

<sup>14</sup>Price regulation could make the WTI price of oil as unrepresentative of the oil market during part of the sample period. We notice that the WTI price index is part of our oil price measure but not the primary oil price measure used in this study. Therefore, this issue seems to be less relevant in our analysis. However, an alternative way to completely rule out the WTI price regulation issue is to focus on the US refiners' acquisition cost, which is available since 1974. A visual inspection, suggests that this oil price measure differs from that used in this study only during the years 1979-80. However, the historical decomposition obtained with an alternative VAR including this oil price measure provides results virtually identical to those obtained with our benchmark model (see Figure A.1).



VAR but using a scheme that identifies only 3, as opposed to 5, structural shocks. This corresponds to identifying  $\tilde{\mathcal{A}}_0$  matrices using the restrictions of the first three columns only of [Table 2](#). The response of the US output to the different oil-shocks is qualitatively the same. However, the effect of an oil-supply shock is roughly one half in magnitude, while that of a RoW demand shock is smaller by one-third on impact. These results corroborates the importance of identifying both US and non-US shocks in order to avoid biased estimated response of the US industrial production to structural shocks moving the real price of oil.

In [Figure A.2](#), we explore the consequence of assuming a positive response of oil quantity when identifying a RoW or a US demand shock. While these assumptions are consistent with our theoretical frame, the assumption is controversial as other, e.g. [Kilian \(2009\)](#) and [Hamilton \(2009\)](#), estimate a small short-run elasticity of oil supply to oil demand changes. In order to assess how the estimated effects on US output change by relaxing this assumption, we estimate the same VAR model without a constraint on the impact response of the oil quantity following both a RoW and a US demand shock. We notice that relaxing these restrictions does not arise any issue regarding the identification of the shocks, which still remain mutually exclusive. The response of oil quantity following a RoW demand appears difficult to detect on impact, probably suggesting that the variance of the shocks is rather small in our sample, but becomes positive and significant with a delay of few months. In any event, the response of US output to each oil-shock is remarkably similar to those obtained with our benchmark identification scheme.

Second, we assess the robustness of our results by estimating a VAR model in which we treat differently the non-stationary of the variables as opposed to the inclusion of the linear trend. Global crude oil production is expressed in log first-differences while the US industrial production, the RoW output and the US relative price measure are expressed in growth rates.<sup>15</sup> The estimated impulse responses are reported in [Figure A.3](#). They suggest that our results are qualitatively robust to the way we handle the non-stationarity of the variables. Interestingly, the variance decomposition also provides considerations very similar to those obtained with the benchmark model.

Third, we assess whether our main results obtained for the US industrial production can be generalized to a broader measure of the US economic activity. The use of industrial production might have the downside that it is not necessarily the variable policymakers are most interested in and it is a measure of gross output rather than value added. This could matter for the comparison with [Kilian \(2009\)](#) to the extent that gross output responds differently to oil price shocks than value added. To this end, we repeat our analysis by

---

<sup>15</sup>In [Kilian \(2009\)](#) oil quantity is also included in log first- differences and the US real GDP and CPI inflation are also included in growth rates. Kilian’s real measure of world economic activity is instead expressed in deviation from a long-run linear trend.



replacing the US industrial production with the monthly Chicago Fed National Activity Index (CFNAI), which is commonly recognized to be a coincident indicator of the US national economic activity.<sup>16</sup> Impulse responses are reported in [Figure A.4](#) and suggest results very similar to those obtained with the US industrial production index.

Finally, our analysis might share with any empirical study the limit that it cannot consider all the relevant shocks affecting the oil market. The most natural candidates are the already mentioned “oil-specific demand shocks”, namely oil price hikes that are due to fears about future oil supply availability. [Kilian \(2009\)](#) showed that precautionary oil-demand shocks have a negative effect on the US economic activity that is more persistent than the one implied by oil-supply shocks and “aggregate demand shocks” and that they explain many historical episodes characterized by major oil price increases. To explore the hypothesis that our RoW shocks described in [Figure 4](#) may also reflect shocks to the precautionary demand for oil discussed by [Kilian \(2009\)](#), we set up a 6-variable VAR that includes a non-energy commodity price index in real terms (e.g. US-CPI deflated). Imposing sign restrictions on this variable is not obvious since our theoretical frame does not explicitly model the market for non-energy commodities. One possibility is to leave the response of the non-energy commodity price index unconstrained to all structural shocks; the other plausible option is to assume that both a RoW supply and a RoW demand shock, as capturing global business cycle fluctuations, simultaneously increase both the oil price and the commodity price index (both in real terms). In any event, the estimated impulse responses for the US output are in both cases extremely similar to those produced by our benchmark identification scheme (see for example [Figure A.5](#) which refers to the latter case).

## 6 Concluding remarks

We presented a model, adapted from [Backus and Crucini \(2000\)](#), where the cost of the oil input and US production respond to demand and supply shocks generated domestically and in the world economy. We use several robust predictions of the theoretical model to identify the nature of the shock underlying observed time series from the oil market, the US economy and the global business cycle. In the model the reduced form correlation between (real) oil prices and the US business cycle depends on the nature of the shocks. We show that this basic observation is quantitatively important to interpret the US data of the past 40 years. The main findings are as follows.

---

<sup>16</sup>The CFNAI index is a stationary principal components index constructed to summarize variation in 85 business cycle indicators including production and income, employment, unemployment and hours; personal consumption and housing; and sales, orders, and inventories.

First, the traditional view on the effects of oil-supply shocks is solid: the estimates suggest that the impact of a negative oil-supply shock on US production is negative, large and highly persistent. However, the variance decomposition analysis shows that these shocks explain only a small fraction of the US output variation (about 10%).

Second, the estimates show that oil prices respond significantly to shocks affecting the global economy and the US. In particular, supply shocks in the rest of the world and in the US explain more than half of the variance of the (real) oil price fluctuations. We view this finding as an important warning for the analyses that treat oil prices as predetermined. Reverse causality from the global economy to the cost of the oil input is evident in the data.<sup>17</sup>

Third, the estimates suggest that to assess the effect of the world business cycle on the US economy it is important to distinguish between (global) supply vs. demand shocks: both shocks increase global output and the real price of oil, but they have opposite implications concerning the US output at horizons up to one year. This finding delivers a simple interpretation of the small and unstable correlation between oil prices and the US economic activity documented in e.g. [Hamilton \(2008\)](#). Depending on the nature of the fundamental shock, a negative correlation emerges in periods when oil-supply shocks or global demand shocks occur, while a positive correlation emerges in periods of supply shocks in the rest of the industrial world or in the US. The unconditional correlation between oil prices and US production over a long sample period is tenuous because it blends conditional correlations with different signs.

---

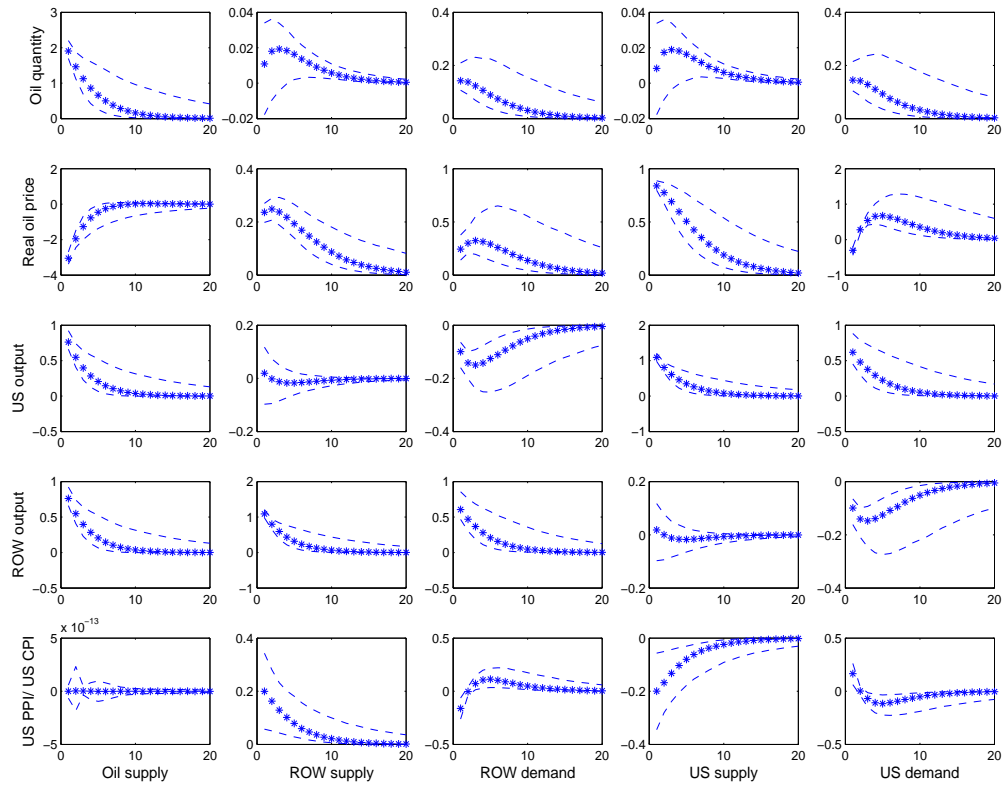
<sup>17</sup>See the recent papers by [Bodenstein et al. \(2007\)](#) and [Nakov and Pescatori \(2010\)](#) for more models where, like in ours, the price of oil is endogenous.

## References

- Anzuini, A., Pagano, P. and Pisani, M. (2007), Oil supply news in a var: Information from financial markets, Temi di discussione (Economic working papers) 632, Bank of Italy, Economic Research Department.
- 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. (2002), ‘Do we really know that oil caused the great stagflation? a monetary alternative,’, *NBER Macroeconomics Annual 2001*, B. Bernanke and K. Rogoff (eds.) Cambridge, Mit Press pp. 137–183.
- Barsky, R. B. and Kilian, L. (2004), ‘Oil and the macroeconomy since the 1970s’, *Journal of Economic Perspectives* **18**(4), 115–134.
- Baumeister, C. and Peersman, G. (2008), Time-varying effects of oil supply shocks on the us economy, Ghent University, manuscript.
- Baxter, M. and Crucini, M. J. (1993), ‘Explaining saving-investment correlations’, *American Economic Review* **83**(3), 416–36.
- 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.
- Bodenstein, M., Erceg, C. J. and Guerrieri, L. (2007), Oil shocks and external adjustment, FRB International Finance Discussion Paper No. 897.
- Canova, F. and Nicosi, G. D. (2002), ‘Monetary disturbances matter for business fluctuations in the G-7’, *Journal of Monetary Economics* **49**(6), 1131–1159.
- Canova, F. and Paustian, M. (2007), Measurement with some theory: Using sign restrictions to evaluate business cycle models, C.R.E.I., manuscript.
- Cunado, J. and Perez de Gracia, F. (2003), ‘Do oil price shocks matter? evidence for some european countries’, *Energy Economics* **25**(2), 137–154.
- Davis, S. J. and Haltiwanger, J. (1999), ‘On the driving forces behind cyclical movements in employment and job reallocation’, *American Economic Review* **89**(5), 1234–1258.
- 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.
- Edelstein, P. and Kilian, L. (2009), ‘How sensitive are consumer expenditures to retail energy prices?’, *Journal of Monetary Economics* **56**(6), 766–779.

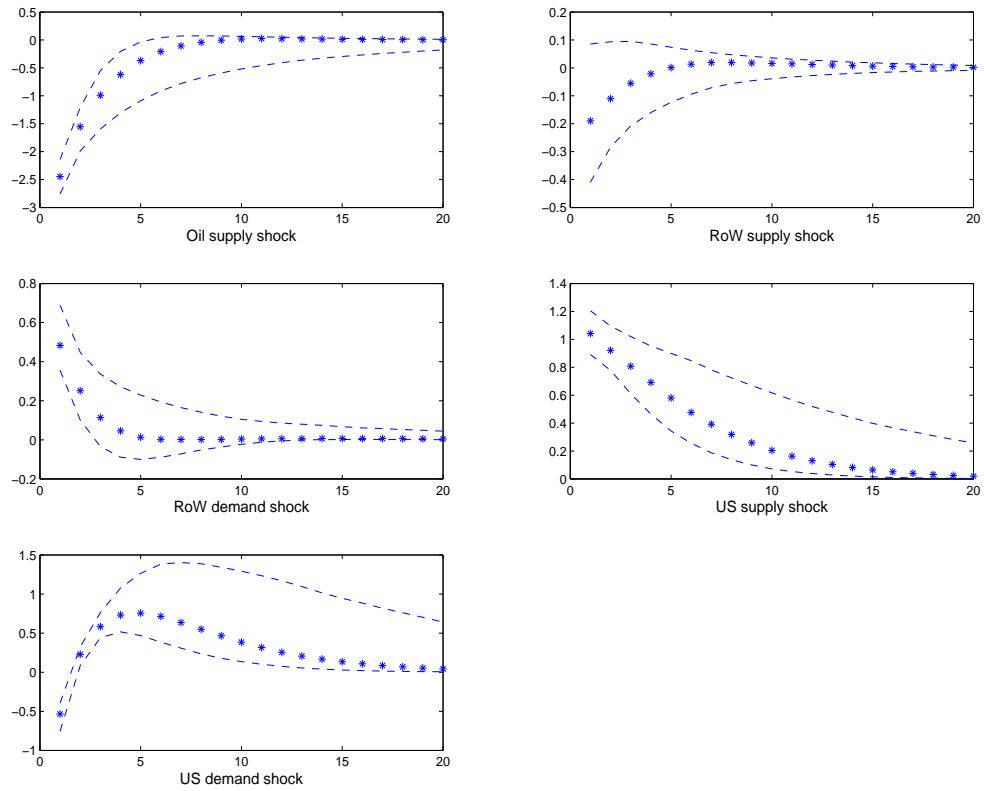
- Francis, N. and Ramey, V. A. (2005), ‘Is the technology-driven real business cycle hypothesis dead? shocks and aggregate fluctuations revisited’, *Journal of Monetary Economics* **52**(8), 1379–1399.
- Fry, R. and Pagan, A. (2007), Some issues in using sign restrictions for identifying structural VARs, 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. (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. (2009), Causes and consequences of the oil shock of 2007-08, NBER Working Papers 15002, National Bureau of Economic Research, Inc.
- 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.
- Hooker, M. A. (1996), ‘What happened to the oil price-macroeconomy relationship?’, *Journal of Monetary Economics* **38**(2), 195–213.
- Kilian, L. (2008a), ‘The economic effects of energy price shocks’, *Journal of Economic Literature* **46**(4), 871–909.
- Kilian, L. (2008b), ‘Exogenous oil supply shocks: How big are they and how much do they matter for the u.s. economy?’, *The 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), 1053–69.
- Mork, K. A., Olsen, . and Mysen, H. T. (1994), ‘Macroeconomic responses to oil price increases and decreases in seven OECD countries’, *Energy Journal* **15**(4), 19–35.
- Nakov, A. and Pescatori, A. (2010), ‘Oil and the great moderation’, *Economic Journal* **120**(543), 131–156.
- 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.

Figure 1: Dynamic effect of each structural shock in the model economy



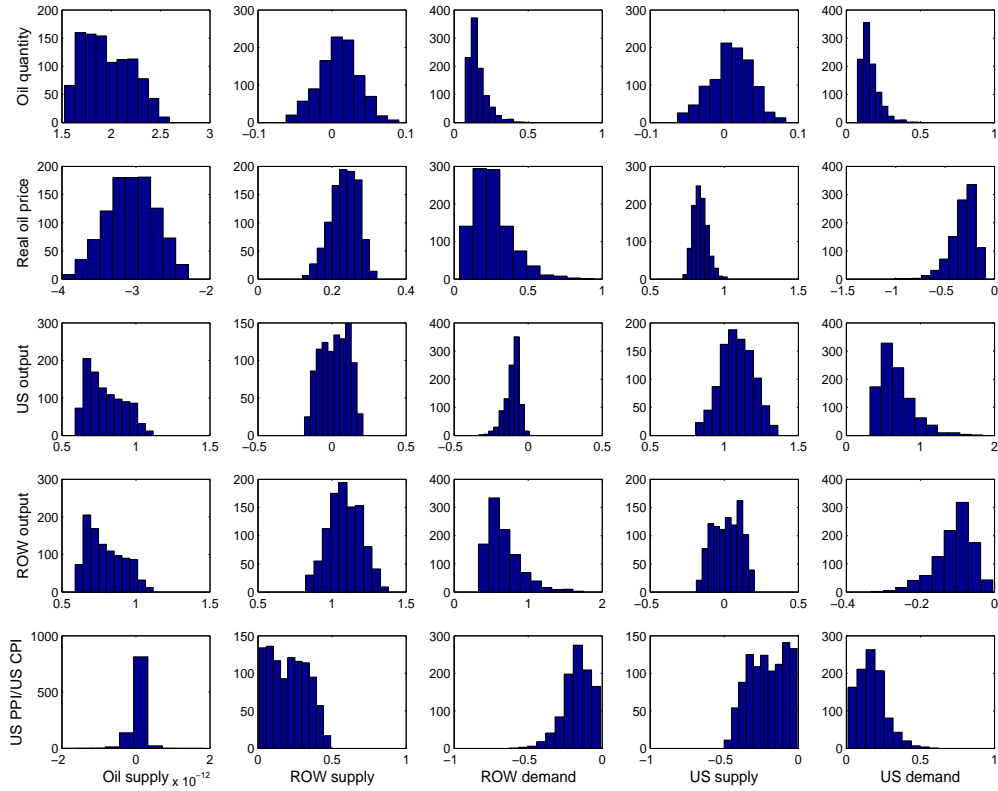
Note: The figure plots the 5th (dashed), 50th (thick) and 95th (dashed) percentiles of the distribution of the responses of each variable to the indicated structural shock at different horizons. The responses are computed by carrying out a Monte Carlo simulation on the parameters of the theoretical model. The simulation is based on 1,000 draws. The model parameters are allowed to vary over the ranges reported in [Table 1](#). The “Real oil price” is the ratio between the oil price and the US CPI.

Figure 2: Dynamic effect on the US terms of trade in the model economy



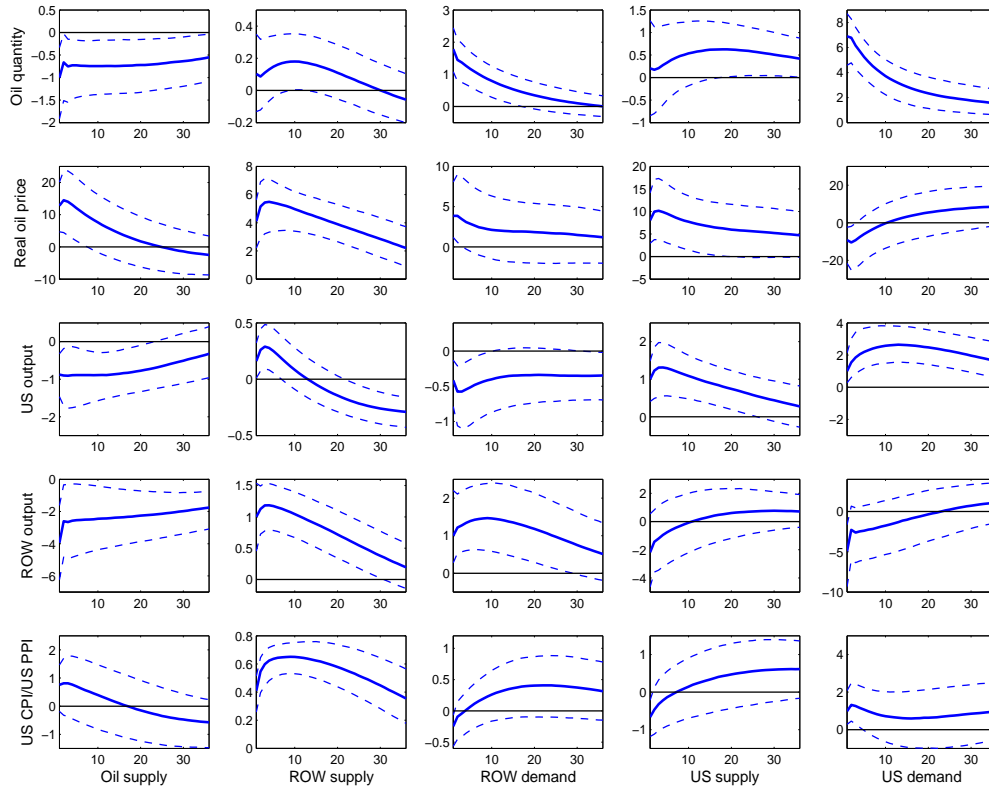
Note: The figure plots the 5th (dashed), 50th (thick) and 95th (dashed) percentiles of the distribution of the responses of the US terms of trade to the indicated structural shock at different horizons. The US terms of trade is defined as the import price index divided by the export price index. The import price index is computed as the weighted average of the response of the oil price and the RoW output price index. The responses are computed by carrying out a Monte Carlo simulation on the parameters of the theoretical model. The simulation is based on 1,000 draws. The model parameters are allowed to vary over the ranges reported in [Table 1](#).

Figure 3: Impact effect of each structural shock in the model economy



Note: The figure plots the distribution of the responses of each variable to the indicated structural shock at the impact horizon. The responses are computed by means of a Monte Carlo simulation on the parameters of the model. The simulation is based on 1,000 draws. The model parameters are allowed to vary over the ranges reported in [Table 1](#). The “Real oil price” is the ratio between the oil price and the US CPI.

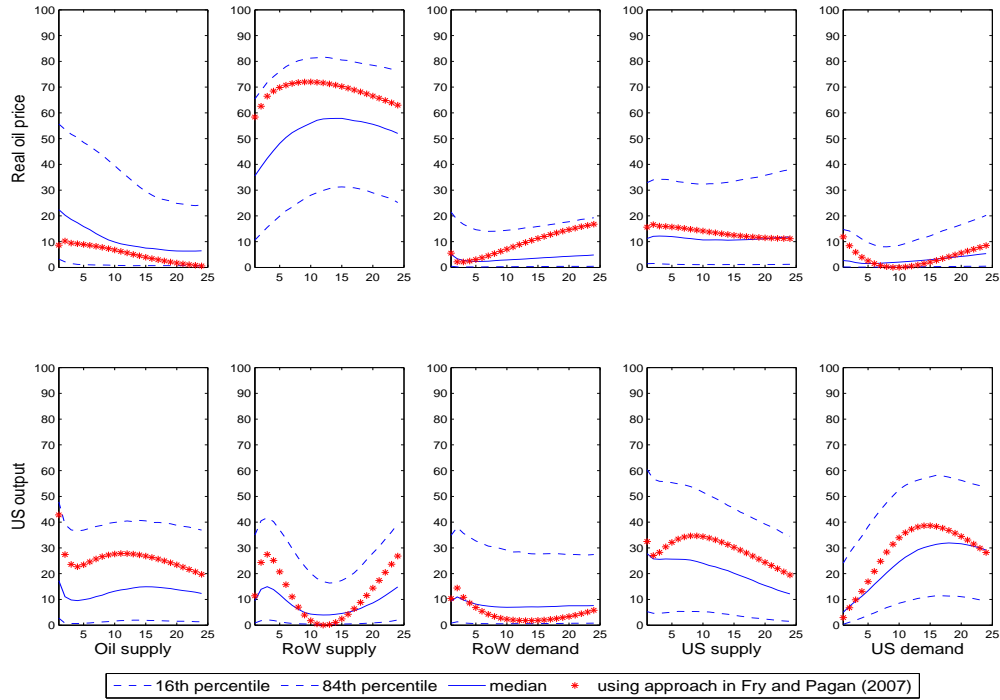
Figure 4: Estimated Effects of Fundamental Shocks



Note: The figure plots the 16th (dashed), 50th (thick) and 84th (dashed) percentiles of the distribution of responses at each horizon. The “Real oil price” is the ratio between the oil price and the US CPI.



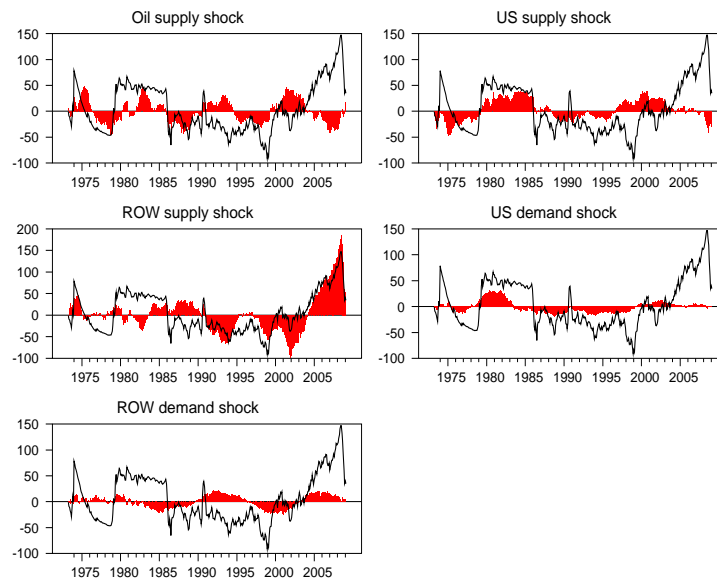
Figure 5: Variance decomposition



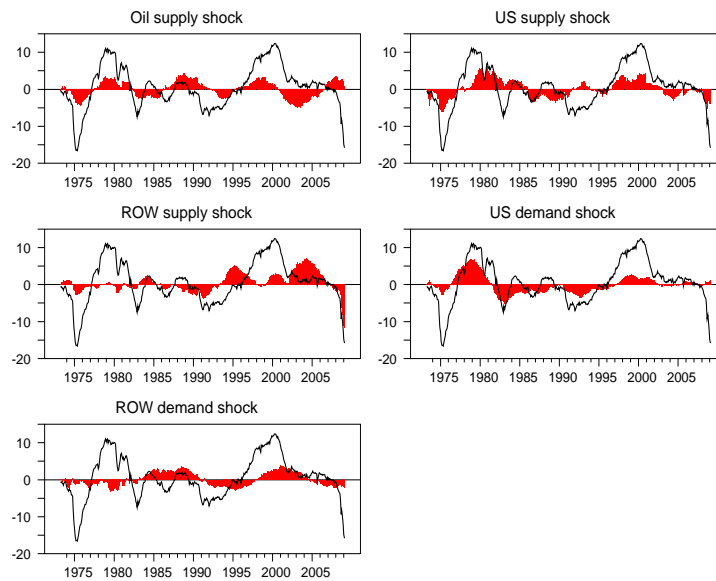
Note: The figure plots the 16th (blue dashed), 50th (blue thick) and 84th (blue dashed) percentiles of the distribution of the variance explained by each structural shock at each horizon. The red line in each panel denotes the corresponding explained variance computed by the  $\tilde{A}_0$  model that minimizes the distance from the median impulse response as in Fry and Pagan (2007). See Online Appendix A. The “Real oil price” is the ratio between the oil price and the US CPI.

Figure 6: Historical decomposition

Oil Price (CPI deflated)



US industrial production



Note: The thin line denotes the real oil price (or the US industrial production), in deviation from the baseline. The red bars in each panel denote the component of the series accounted for by each structural shock. The “Real oil price” is the ratio between the oil price and the US CPI.

Online Appendices - Not for publication

Oil and the Macroeconomy:  
A quantitative structural analysis

*Francesco Lippi (U. Sassari, Eief)*  
*Andrea Nobili (Bank of Italy)*

April 2010

## A Fry and Pagan’s 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.

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

$$y_t = C(L)e_t \tag{A.1}$$

where

$$C(L) = [I - B(L)L]^{-1} \tilde{\mathcal{A}}_0 \tag{A.2}$$

where  $C(L)$  contains the matrices of the estimated impulse responses to structural shocks. Let  $C_{i,j,h}^{(k)}$  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  $h'$ ,  $(\text{med}(C_{i,j,h}^{(k)}), \text{med}(C_{i,j,h'}^{(k)}))$  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  $\tilde{\mathcal{A}}_0 \in \tilde{\mathcal{A}}_0$ , 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

$$Z_{ij,h}^{(k)} = [C_{ij,h}^{(k)} - \text{med}(C_{ij,h}^{(k)})] / [\text{stdev}(C_{ij,h}^{(k)})] \tag{A.3}$$

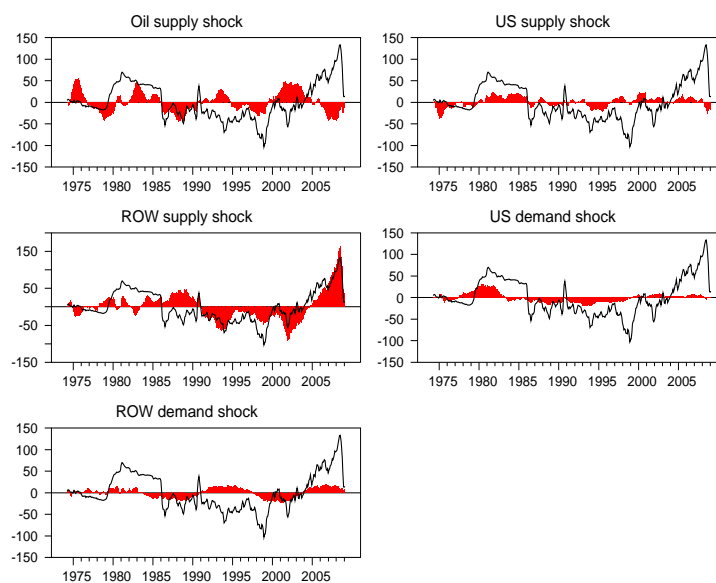
The  $Z_{ij,h}^{(k)}$  are then collected into a vector  $\phi^k$  of dimension  $(n \cdot h \cdot s \times 1)$ , where  $n$  is the number of variables in the VAR,  $h$  the horizons over which the impulse responses is computed, and  $s \leq n$  the number of identified structural shocks. In our case we have  $n = 5$   $h = 36$  and  $s = 5$ , therefore,  $\phi^k$  is a vector  $(900 \times 1)$ . Finally, we choose the value of  $k$  that minimizes  $\phi^k \phi^{k'}$ , and use it to derive the estimated impulse responses.

Figure A.6 compares the median impulse responses of Figure 4 (dashed lines) with the one produced by the model satisfying the criterion describe above (solid lines). The

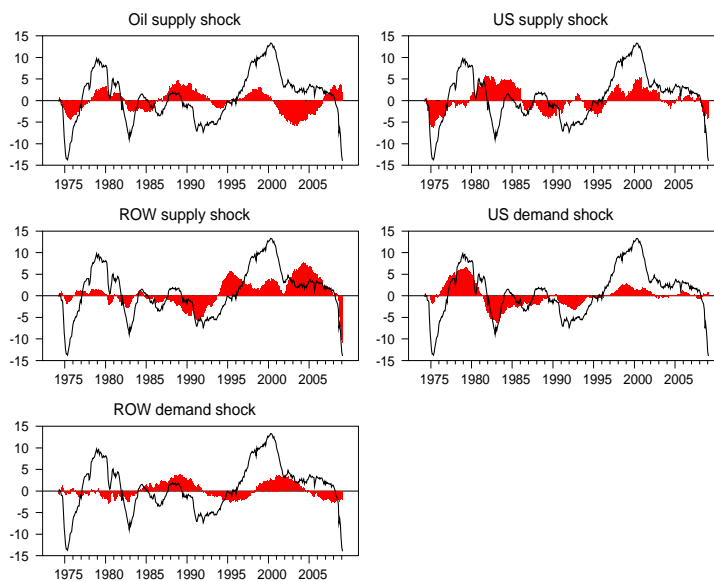
dynamic effects of structural shocks are very similar, even if there are some differences in the magnitude of the responses. In particular, after a supply oil shock the negative response of the US industrial production appears to be more pronounced. After a RoW supply shock the hike in the oil price is even larger, leading to a magnified increase in the US industrial production.

Figure A.1: Robustness: Historical decomposition using refiners' acquisition cost

### Oil Price (CPI deflated)

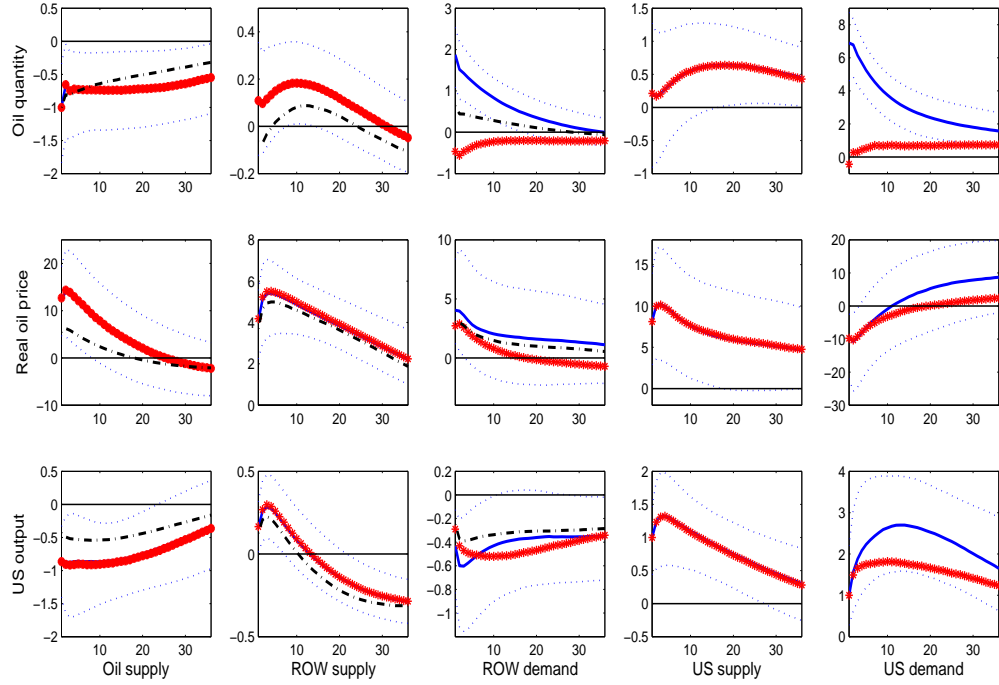


### US industrial production



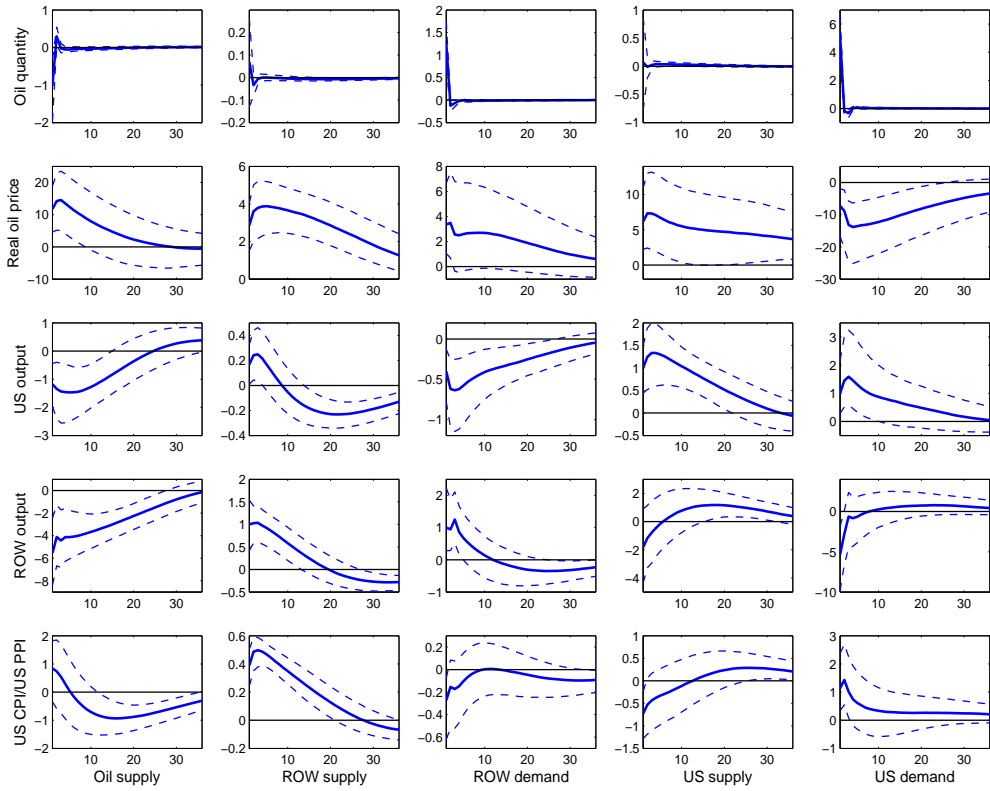
Note: The thin line denotes the real oil price (or the US industrial production), in deviation from the baseline. The red bars in each panel denote the component of the series accounted for by each structural shock.

Figure A.2: Robustness: alternative identification scheme



Note: The figure reports the 16th (blue dashed), 50th (blue tick) and 84th (blue dashed) percentiles of the IRFs distribution produced by the benchmark model. The black dashdot lines are produced by a VAR model with only 3 identified shocks (oil supply shocks, RoW supply and RoW demand shocks); the red lines with stars are those of a model in which no sign restrictions are imposed on the response of oil quantity when the RoW and US demand shocks are identified.

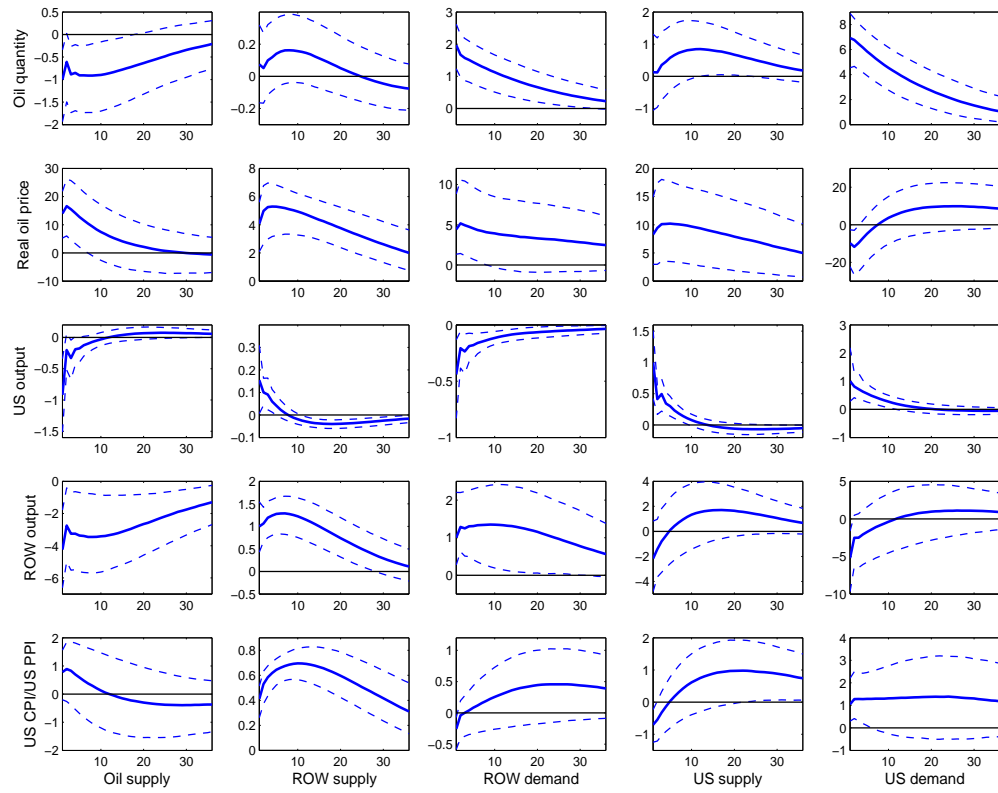
Figure A.3: Robustness: using stationary variables in the VAR



Note: The figure plots the 16th (dashed), 50th (solid) and 84th (dashed) percentiles of the distribution of responses at each horizon.

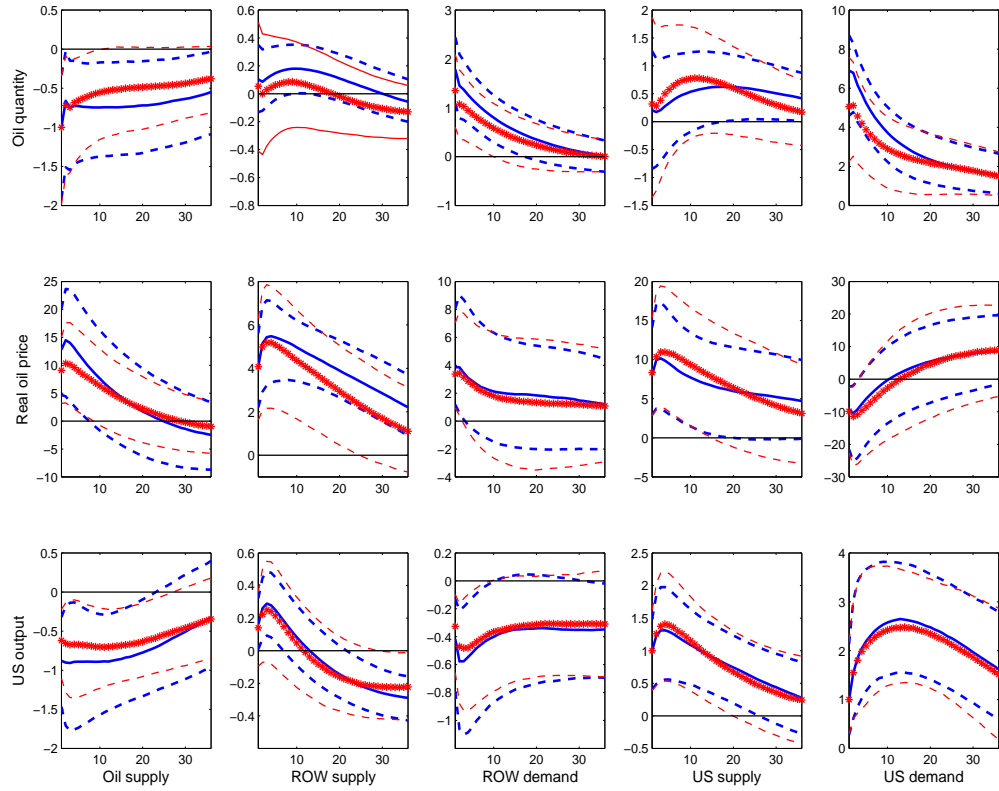


Figure A.4: Robustness: using the CFNAI index for US output



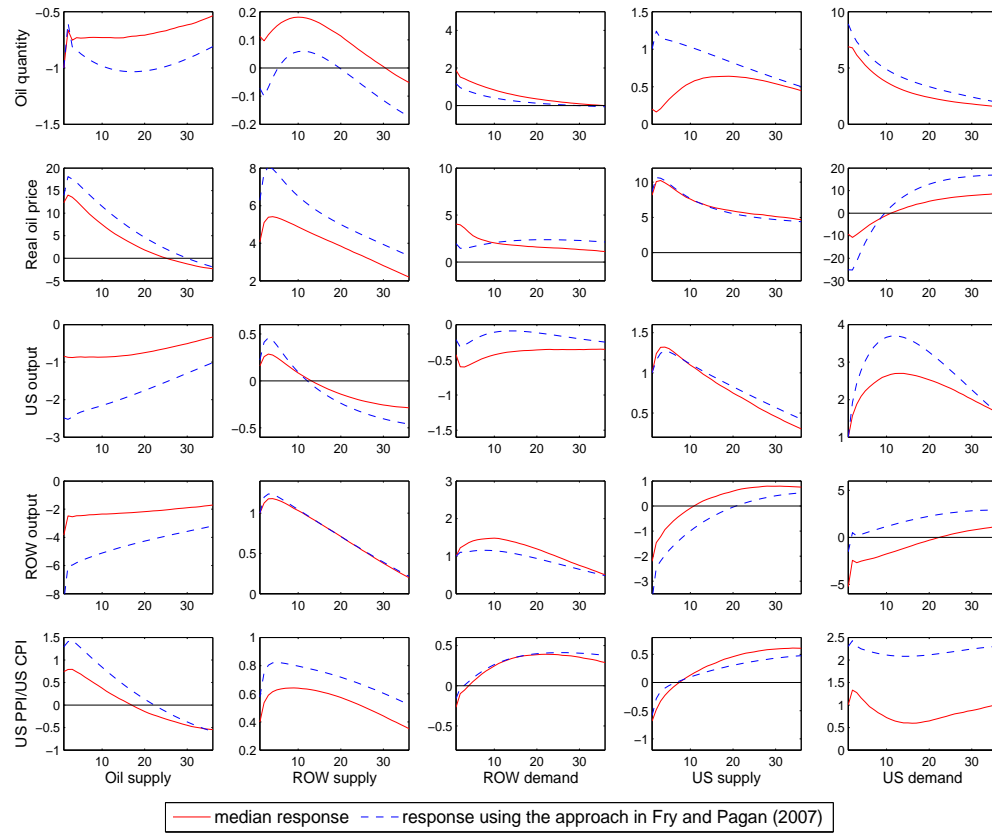
Note: The figure reports the 16th (dashed), 50th (solid) and 84th (dashed) percentiles of the IRFs distribution. In the VAR model the US output is measured by the CFNAI index of economic activity.

Figure A.5: Robustness: avoiding precautionary demand shocks



Note: The figure reports the 16th (dashed), 50th (solid) and 84th (dashed) percentiles of the IRFs distribution. The blue line are produced by the benchmark VAR model; the red lines are produced by a VAR model also including the non-energy price index (US-CPI deflated) in order to avoid shocks in the real price of oil due to unexpected shifts in the precautionary demand for oil discussed in [Kilian \(2009\)](#).

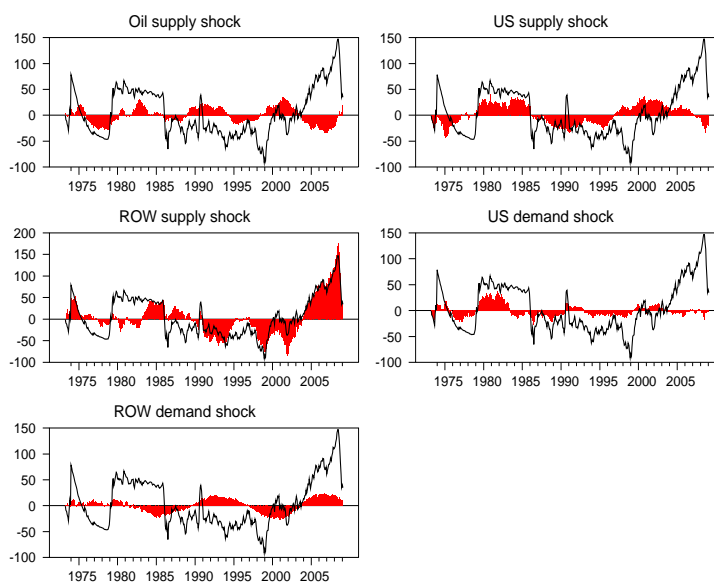
Figure A.6: Robustness: Results based on Fry and Pagan method



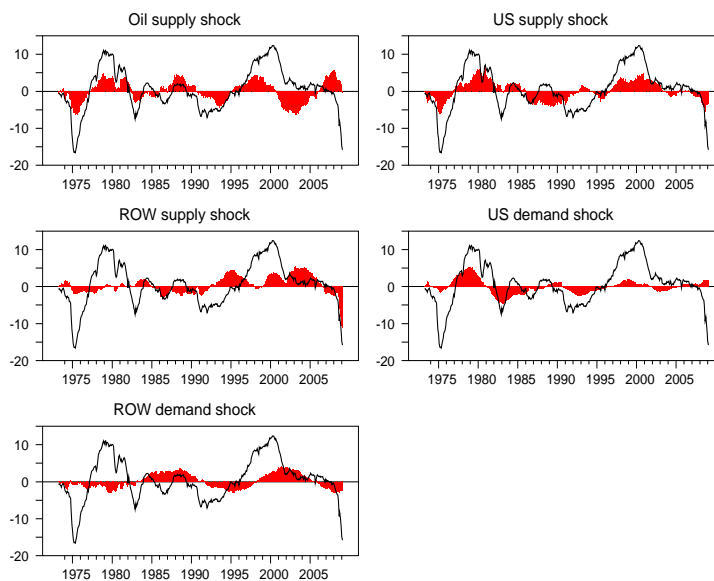
Note: The solid line is the response obtained applying the algorithm proposed by Fry and Pagan (see [Appendix A](#)); the dashed line gives the median response of our benchmark VAR.

Figure A.7: Robustness: Historical decomposition based on Fry and Pagan method

### Oil Price (CPI deflated)

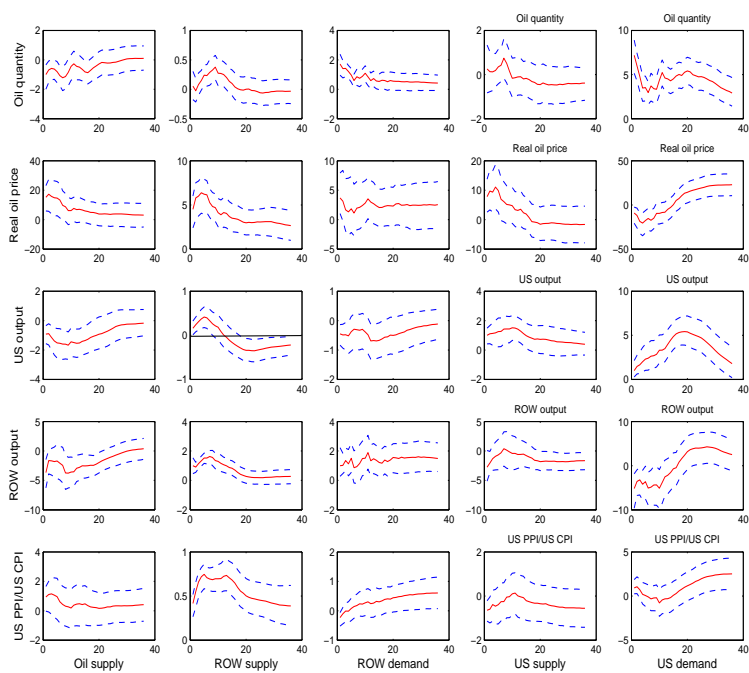


### US industrial production



Note: The thin line denotes the real oil price (or the US industrial production), in deviation from the baseline. The red bars in each panel denote the component of the series accounted for by each structural shock. All lines are obtained applying the algorithm proposed by Fry and Pagan (see [Appendix A](#)).

Figure A.8: Robustness: using 12 lags in the VAR model



Note: The figure reports the 16th, 50th and 84th percentiles of the IRFs distribution.

## B The model economy

By the first welfare theorem the competitive equilibrium solves the planning problem:

$$\begin{aligned}
& \max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ \frac{[c_t^\theta (1-n_t)^{1-\theta}]^{1-\gamma}}{1-\gamma} + \frac{[(c_t^*)^\theta (1-n_t^*)^{1-\theta}]^{1-\gamma}}{1-\gamma} + \left( \frac{(c_t^o)^{1-\gamma}}{1-\gamma} + \theta_L \frac{(1-n_t^o)^{1-\xi_L}}{1-\xi_L} \right) \right. \\
& + p_{c,t} \left[ \left( s_t \psi (a_{c,t})^{1-\mu} + (1-s_t \psi) (b_{c,t})^{1-\mu} \right)^{\frac{1}{1-\mu}} - c_t \right] \\
& + p_{c^*,t} \left[ \left( s_t^* \psi (b_{c,t}^*)^{1-\mu} + (1-s_t^* \psi) (a_{c,t}^*)^{1-\mu} \right)^{\frac{1}{1-\mu}} - c_t^* \right] \\
& + p_{c^o,t} \left[ \left( \psi_o (a_{c,t}^o)^{1-\mu} + (1-\psi_o) (b_{c,t}^o)^{1-\mu} \right)^{\frac{1}{1-\mu}} - c_t^o \right] \\
& + p_{y,t} \left[ z_t n_t^\alpha (\eta k_t^{1-\nu} + (1-\eta) o_t^{1-\nu})^{\frac{1-\alpha}{1-\nu}} - a_{c,t} - a_{i,t} - a_{c,t}^* - a_{i,t}^* - a_{c,t}^o \right] \\
& + p_{y^*,t} \left[ z_t^* n_t^{*\alpha} (\eta k_t^{*1-\nu} + (1-\eta) o_t^{*1-\nu})^{\frac{1-\alpha}{1-\nu}} - b_{c,t} - b_{i,t} - b_{c,t}^* - b_{i,t}^* - b_{c,t}^o \right] \\
& + p_{o,t} [z^o + (n^o)^\alpha - o_t - o_t^*] \\
& + q_t \left[ (1-\delta) k_t + \left( \frac{\left( \psi (a_{i,t})^{1-\mu} + (1-\psi) (b_{i,t})^{1-\mu} \right)^{\frac{1}{1-\mu}}}{k_t} \right)^\phi k_t - k_{t+1} \right] \\
& + q_t^* \left[ (1-\delta) k_t^* + \left( \frac{\left( \psi (b_{i,t}^*)^{1-\mu} + (1-\psi) (a_{i,t}^*)^{1-\mu} \right)^{\frac{1}{1-\mu}}}{k_t^*} \right)^\phi k_t^* - k_{t+1}^* \right] \left. \right\}
\end{aligned}$$

The AR(1) laws of motion (identical across countries) for the oil and preference shocks complete the model:  $z_t = \rho_z z_{t-1} + \tilde{z}_t$  and  $s_t = \rho_s s_{t-1} + \tilde{s}_t$ . There are 20 choice variables (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:

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