

OIL AND THE MACROECONOMY: A QUANTITATIVE STRUCTURAL ANALYSIS

Francesco Lippi
University of Sassari

Andrea Nobili
Bank of Italy

Abstract

We model an open economy where macroeconomic variables fluctuate in response to oil supply shocks, as well as aggregate demand and supply shocks generated domestically and abroad. We use several robust predictions of the model to identify five fundamental shocks underlying the fluctuations of the (real) oil price, the US activity and the global business cycle. The estimates show that supply shocks generated in the global economy explain the largest fraction of the oil price fluctuations, about four times more than canonical oil supply shocks. The correlation between oil prices and the US activity varies with the type of shock. (JEL: C32, E3, F4)

1. Introduction

Large fluctuations in oil prices are a recurrent feature of the macroeconomic environment. Despite oil's relatively small share as a proportion of total production costs, such dynamics raise the specter of the 1970s, worrying consumers, producers, and policy makers. Consistent with this view was the fact that nine out of ten of the US recessions since World War II were preceded by a spike in oil price (see, e.g., Hamilton 1983). This relation, however, has proved unstable over time: the correlation between oil prices and US output after the mid-1980s is much smaller.¹

Barsky and Kilian (2002, 2004) call for a structural interpretation of these reduced-form correlations. These authors challenge the view that oil supply shocks were the main driver of oil price hikes and observe that the OPEC decisions usually respond to global macroeconomic conditions affecting the demand for oil. A structural analysis appears in Kilian (2009), where a VAR model is used to identify three structural shocks assuming "zero-impact restrictions" and a recursive structure: oil supply shocks, world

The editor in charge of this paper was Fabio Canova.

Acknowledgments: We thank the editor and four anonymous referees for several stimulating comments and suggestions. We also benefited from the comments of Fernando Alvarez, David Andolfatto, Luca Dedola, Luca Deidda, Lutz Kilian, Stefano Neri, and Harald Uhlig, and from the feedback received at seminars at the University of Sassari, EIEF, the Bank of Italy, the Center for European Integration Studies (Bonn), and 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. Lippi is a research fellow at EIEF. E-mail: f.lippi@uniss.it (Lippi); andrea.nobili@bancaditalia.it (Nobili)

1. See Hooker (1996), Hamilton (2008) and Edelstein and Kilian (2009). Similar patterns emerge for European countries: see Mork, Olsen, and Mysen (1994) and Cunado and Perez de Gracia (2003).

aggregate demand shocks and oil market-specific demand innovations, interpreted as reflecting fluctuations in precautionary demand for oil. A central finding of Kilian's paper is that oil price fluctuations have historically been driven mainly by precautionary demand shocks, with a small role for traditional oil supply shocks.

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 using the three-country model of Backus and Crucini (2000). The model assumes an oil-producing country and two industrial economies, the United States and the global economy or "rest of the world", who produce differentiated goods 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 five fundamental shocks and the observed responses of production and relative prices. We estimate a five-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 business cycle. The estimated VAR is used to identify the five fundamental shocks of our theory using robust model predictions on the sign of the impulse responses.

The main novelty of this paper is that the interplay between the oil market, the global economy, and the US economy is managed within a fully specified theoretical model, which allows us to be explicit about the identifying restrictions used in the empirical analysis. We stress three points. First, we estimate the model using some robust implications of the theory, following the sign-restrictions approach pioneered by Davis and Haltiwanger (1999), Canova and De Nicolò (2002) and Uhlig (2005). The identifying assumptions are based on an explicit theoretical model which prescribes the use of novel restrictions, on, for example, the relative price of home and foreign goods and the business cycle in the rest of the world.² The method represents an alternative to the zero contemporaneous restrictions, widely used in previous works, which are difficult to reconcile with explicit dynamic models. We will present different estimates of the model, and compare them with previous results, most notably those by Kilian (2009).

Second, we allow for the simultaneous interaction between the oil market and the US economy. By doing so we cast light on the assumption, widely used in the empirical analysis, that oil prices are predetermined with respect to the US business cycle (see Leduc and Sill 2004; Kilian 2009; Blanchard and Gali 2010).³

2. This approach is novel in the macro literature on the effect of oil prices, although similar ideas have been used by, for example, Dedola and Neri (2007) to study the response of labor hours to technology shocks and by Pappa (2009) to assess the effects of fiscal shocks in the labor market.

3. The assumption of predeterminedness with respect to the US economy is particularly debated in the empirical literature. Kilian and Vega (2008) regressed cumulative changes in daily energy prices on daily news from US macroeconomic data releases and found no compelling evidence of feedback at daily and monthly horizons. In contrast, Pagano and Pisani (2009) showed that releases of US industrial production and capacity utilization data are useful in predicting oil prices. Jimenez-Rodriguez and Sanchez (2005) found that the interactions between oil prices and US macroeconomic variables is significant, with the direction of causality going in both directions. The recent papers by Anzuini, Pagano, and Pisani (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 distinguishes between shocks originated in the global economy versus the United States.

Third, we extend the seminal analysis of Kilian (2009) by decomposing generic “oil demand shocks” into four different fundamental shocks, namely aggregate demand and supply shocks in the United States and in the rest of the world (RoW henceforth). This issue is relevant because RoW-supply and RoW-demand shocks have similar dynamic effects on the oil market variables (for example, move price and quantity in the same direction) but may cause a different response of the US output. The theoretical model allows us to explicitly spell out the mechanism that underlies these shocks, and to quantify their importance.

Altogether, our analysis suggests that each of these points is useful to interpret the time series evidence on oil price fluctuations and the US business cycle (a summary of results is given in Section 6). The analysis provides a simple explanation of the unstable correlation between oil prices and the US economic activity documented in Hamilton (2008), and recently discussed by Blanchard and Gali (2010) and Blanchard and Riggi (2009).

The paper has six 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: we revisit the quantitative importance of the “oil market specific shocks”, which are central in Kilian (2009), and discuss the economic meaningfulness of the impulse responses produced by the sign-restrictions method, an issue of concern discussed in Kilian and Murphy (forthcoming). Section 6 summarizes and concludes.

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, Kehoe, and Kydland (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 United States and RoW (rest of the industrial world), produce imperfectly substitutable consumption goods, a and b , using capital (k), labor (n), and oil (o). The United States produces 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 i.i.d. 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 and $(n_t^o)^\alpha$ 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.⁴ An identical aggregator, with deterministic weight ψ , is used to produce the investment good, i .

Capital follows the accumulation equation $k_{t+1} = (1 - \delta)k_t + k_t \varphi(i_t/k_t)$, where $\varphi(\cdot)$ is a concave function positing adjustment costs in capital formation as in Baxter and Crucini (1993). Consumers in the United States and the RoW maximize the expected value of lifetime utility, solving

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

where

$$U(c, 1 - n) = \frac{[c^\theta (1 - n)^{1-\theta}]^{1-\gamma}}{1 - \gamma}, \quad \gamma > 0, \quad \text{and} \quad 0 < \theta < 1.$$

Here $\beta < 1$ 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. Specifically the model assumes

$$U(c, (1 - n)) = c^{1-\gamma} / (1 - \gamma) + \theta_L (1 - n)^{1-\xi_L} / (1 - \xi_L). \quad (4)$$

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 pp. 197–198 in Backus and Crucini 2000). 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

4. Similar effects are obtained by considering shocks to the intertemporal discount factor.

TABLE 1. Parameter ranges in the model economy.

Simulated parameters		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]
ψ	Preference for home good	[0.7, 0.9]
ρ_z	Persistence of US supply shock	[0.5, 0.99]
ρ_{z^*}	Persistence of RoW supply shock	[0.5, 0.99]
ρ_{z^o}	Persistence of oil supply shock	[0.5, 0.99]
ρ_s	Persistence of US demand shock	[0.5, 0.99]
ρ_{s^*}	Persistence of RoW demand shock	[0.5, 0.99]
Calibrated parameters (quarterly model)		
α	Labor's share in the industrial country	0.64
β	Intertemporal discount factor	0.99
δ	Depreciation rate of capital	0.025
η	Oil weight in technology	0.9
φ	Investment adjustment costs	0.99
θ	Labour share in the industrial country	0.4
ψ^o	Oil share in the oil-producing country	0.5
θ_L	Labor share in the oil-producing country	0.6
ξ_L	Inverse labor elasticity in the oil-producing country	5.0

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. The full description of the model and the log-linearized first-order conditions used in the Monte Carlo simulations are shown in the Online Appendix.

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 rest of the world output level and the price of US output (CPI deflated) 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. 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

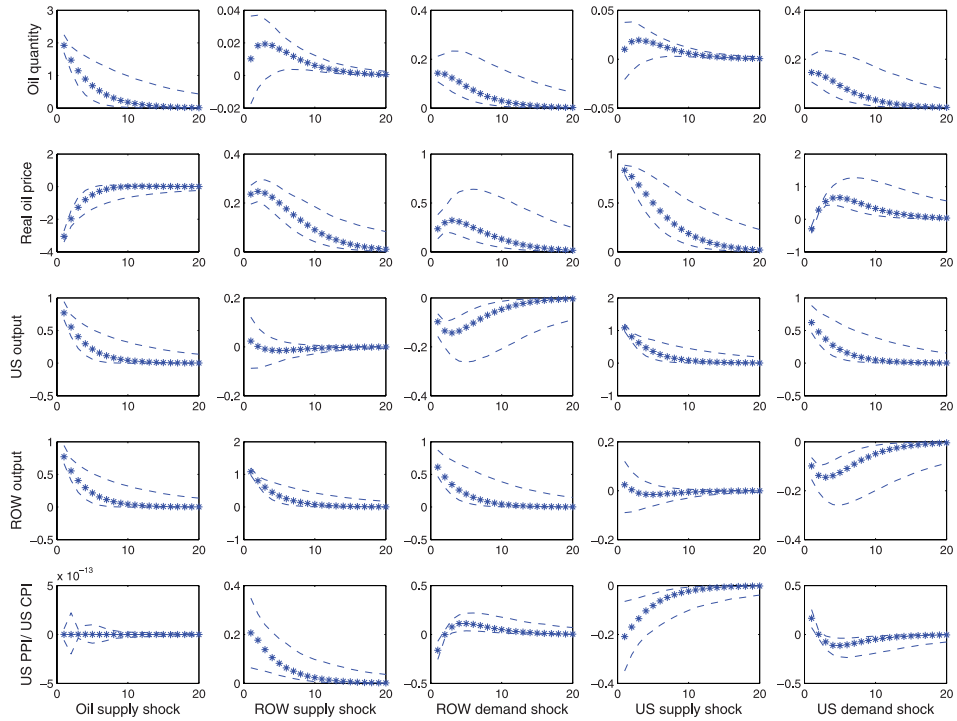


FIGURE 1. Dynamic effect of each structural shock in the model economy. The figure plots the 5th (dashed), 50th (thick) and 95th (dashed) percentiles of the distribution of the responses of each variable to the indicated structural shock at different monthly horizons. The responses are computed by carrying 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. “Real oil price” is the ratio between the oil price and the US CPI.

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).

The first column describes the effect of a positive oil supply shock ($z^o > 0$). The shock 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 United States 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 in 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—that is, these shocks appear to an observer of the oil market as “oil demand shocks”. But the figure also shows that the response of US

production to these shocks is not the same. The difference in the response is due to the fact that in a general equilibrium a positive RoW supply shock increases the total amount of resources available in the economy. With perfect risk sharing, consumption tends to increase in both countries. As the foreign and home goods are imperfect substitutes in consumption, the effect on US *production* depends on the consumption substitution elasticity between home and foreign goods. If the elasticity is unitary, then the US output stays constant. When the substitution elasticity is smaller than one, the goods are complements and US output increases after a positive RoW shock. The reverse happens for a substitution elasticity bigger than one. Now consider the effect of a positive RoW demand shock. This increases the real cost of oil and the price of imported foreign goods, so that the final effect on US consumption and on US production is an unambiguous reduction.

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 positive 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.

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 adverse oil supply shock is followed by a decrease of US production, the oil price hike caused by a positive 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 previously outlined, 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 De 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—that is, 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)

$$y_t = B(L) y_{t-1} + \varepsilon_t \quad \varepsilon_t \sim N(0, \Sigma), \quad (5)$$

where $B(L)$ is a lag polynomial of order p and y_t contains five variables (all in logs) describing the United States, 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. The fifth variable is total imports to the RoW, aiming at capturing economic activity in the RoW. 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 an indicator of cumulative global demand pressure. This measure, however, does not distinguish between increase in demand stemming from the United States and those originating in the RoW. Our results are robust to the use of Kilian's measure of world output instead of the RoW imports. Data description and the source of the variables are described in the Appendix.

Estimation of the VAR is based on monthly data spanning the period between January 1973 and 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 two months, as suggested by the Bayesian Information Criteria (BIC), which provides estimated residuals for the reduced-form VAR characterized by good white-noise properties. The appropriate lag length has been debated in the literature, see Hamilton and Herrera (2004) and Kilian (2009). Our results remain virtually unchanged if seven lags (as suggested by the Akaike Information Criteria) or if twelve lags are used.

The structural VAR approach sees (5) as a reduced-form representation of the structural form

$$\mathcal{A}_0^{-1} y_t = A(L) y_{t-1} + e_t \quad e_t \sim N(0, I), \quad (6)$$

where $A(L)$ is a lag polynomial of order p and the vector e includes the five structural innovations discussed previously, 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 , Ψ_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 .

The sign restriction approach identifies a set of structural models, the $\tilde{\mathcal{A}}_0 \in \tilde{\mathcal{A}}_0$, such that the (vectors of) impulse responses the Ψ 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 $\tilde{\mathcal{A}}_0$ can be

characterized as follows. For a given estimate of the reduced-form VAR, $B(L)$ and Σ , take an arbitrary identification matrix \mathcal{A}_0 and compute the set of candidate structural models $\hat{\mathcal{A}}_0 = \{\mathcal{A}_0 Q \mid QQ' = I\}$ by spanning the space of the orthogonal matrices Q . The set $\tilde{\mathcal{A}}_0$ is then obtained by removing from the set $\hat{\mathcal{A}}_0$ 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 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 Monte Carlo simulations. The ranges for the parameters used in the simulations are given in Table 1. The results of these simulations are reported in detail in the Online Appendix.

In the empirical analysis we restrict attention to five mutually orthogonal shocks: an oil supply shock, supply and demand shocks in the RoW, and supply and demand shocks in the United States. 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 Figure 1. We define a RoW supply shock as one that moves in the same direction as 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 oil price, the quantity of oil and the ROW output, while it 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

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 ^a	+	+	+	+	–
US output	–		–	+	+
RoW output	–	+	+		–
US output price ^a		+	–	–	+

Notes: 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.

a. Price is deflated by the US CPI.

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 confused 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 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{A}_0 is computed using the efficient algorithm proposed by Rubio-Ramirez, Waggoner, and Zha (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{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 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, Waggoner, and Zha (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 $\tilde{\mathcal{A}}_0$ generates orthogonal structural shocks.

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 algorithm finds at least one 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.

However, we used a relatively large number of robust sign restrictions in order to disentangle the structural shocks, thus making the analysis particularly severe from a computational viewpoint. This is necessary because, as shown by Canova and Paustian (2007), 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 is very large, results are rarely sharp.

4.1. Impulse Responses

Following Dedola and Neri (2007) and Uhlig (2005), we report in Figure 2 the median (solid line), and the 16th and the 84th percentiles (dashed lines) of the distribution of impulse responses produced by the algorithm discussed previously for each variable over 24 months.

The effects of an oil supply shock, normalized to yield a 1% reduction in oil production, are displayed in the first column of Figure 2. The shock lowers the US industrial production, which reaches a trough after about twelve 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.

The effects of a RoW supply shock, normalized to yield a 1% increase in the RoW output, are displayed in the second column of Figure 2. The response of the US industrial production differs markedly from the case of the oil supply shock: production twelve 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 six months, and 60% after one year. As discussed in Section 3, a positive response of the US output to a positive supply innovation in the RoW can be explained by positing that domestic and foreign goods are complements in consumption, at least for the horizons up to two years. Over the same horizons, the effect of a RoW demand shock on the US industrial production (normalized to produce a 1% increase in the RoW output) is instead negative and persistent, as shown in the

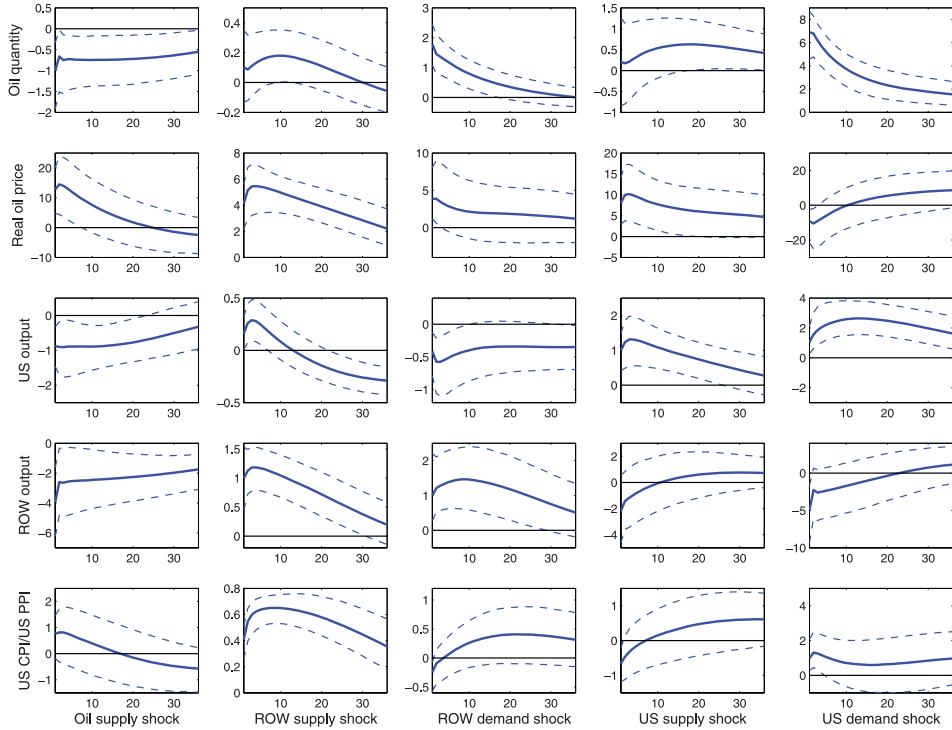


FIGURE 2. Estimated effects of fundamental shocks. The figure plots the 16th (dashed), 50th (thick), and 84th (dashed) percentiles of the distribution of responses at each monthly horizon. The “Real oil price” is the ratio between the oil price and the US CPI.

TABLE 3. US output response to different structural shocks.

“Probability” of a negative response of US output ^a					
At horizon:	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

a. Fraction of models $\tilde{A}_0 \in \tilde{A}_0$ that yield a negative response at the given horizon.

third column of Figure 2. The different effects of RoW demand and supply shocks on the US output appear important empirically.

The fourth and fifth columns of Figure 2 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.⁵

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 nonsignificant 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 of up to one year (see Figures 1 and 2). 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 predicting the dynamics of US output conditional on observing an oil price increase. A higher oil price is associated with an expected reduction of US production conditional on an adverse oil supply shock, a negative US demand shock, or a positive RoW demand shock. However, 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 3 reports the median (solid line), and 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

5. For more insight into the mechanism at work, we assessed the importance of the US labor market in the transmission of the shocks. The theory suggests that the median response of US labor hours is positive following a positive supply shock in the RoW, while it is negative following a positive RoW demand shock, albeit not robust to all specifications). The inclusion of the US hours worked in our benchmark VAR suggests a significant and positive short-run response in the US hours worked after a positive RoW supply shock, and a persistent decline following a positive RoW demand shock. A negative oil supply shock also has a negative effect on this variable. The response of the US hours worked following a US supply and a US demand shock are instead positive and consistent with both theory and the empirical literature. See Christiano, Eichenbaum, and Vigfusson (2004), Dedola and Neri (2007), and Peersman and Straub (2004, 2009).

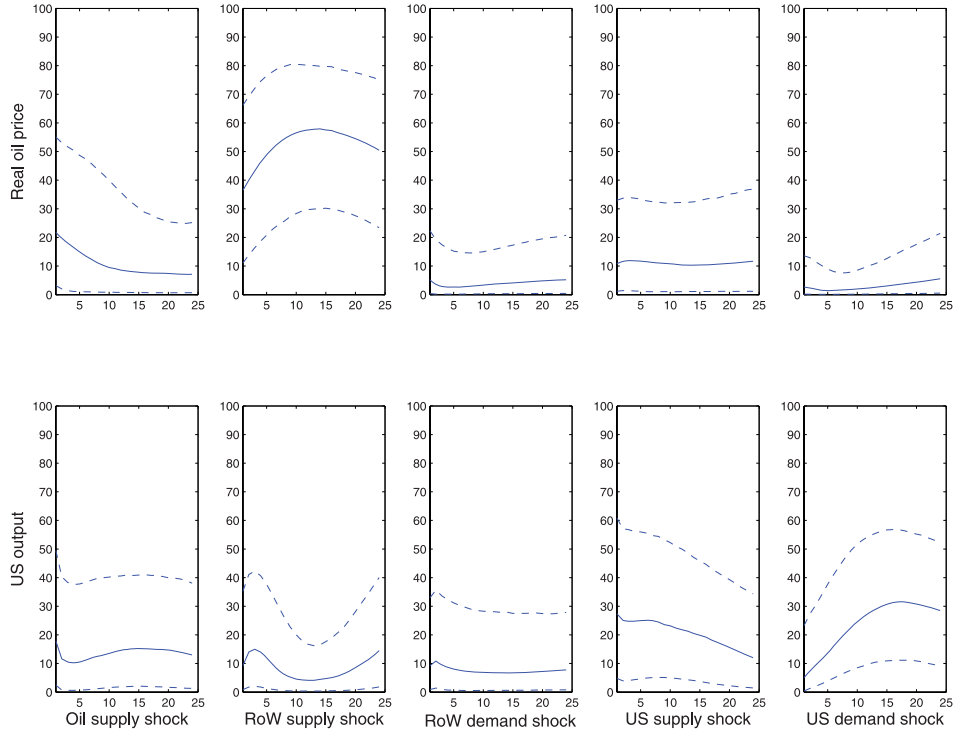


FIGURE 3. Variance decomposition. The figure plots the 16th (dashed), 50th (solid) and 84th (dash-dot) percentiles of the distribution of the variance explained by each structural shock at each monthly horizon. The “Real oil price” is the ratio between the oil price and the US CPI.

generated in the RoW explain the largest fraction of the oil price variance, between 40% and 60%, over the horizons considered. This shock represents, by far, the largest source of fluctuations in oil prices. Oil supply shocks explain 10% of the variance within the one-year horizon. The US aggregate supply shocks account for about 10% of the oil price variance, with a larger share at short-term horizons.

The second row of Figure 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 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 the short-term horizons. The other shocks affecting US production, the oil-supply and the RoW shocks, are less important than domestic shocks. However, the estimated median effect (marked by the continuous line) suggests that oil supply shocks explain about 10% of total output 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. To ensure

orthogonality of the structural shocks we follow their suggestion and select 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. The results, as measured by the impulse responses and the variance decomposition, 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 4. This figure highlights the contribution of each structural shock to deviations of the variables from the baseline at each point in time.

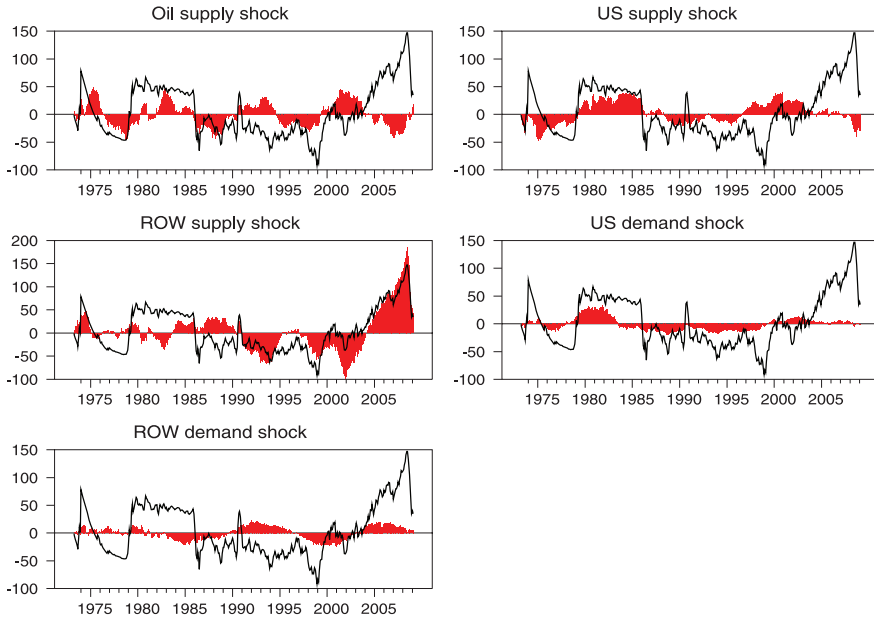
The question of whether the oil price hike recorded in late 1973 is an oil supply shock or an oil demand shock has been widely discussed. Barsky and Kilian (2002) and Kilian (2008) showed that it is explained to a large extent by a delayed consequence to a demand shock in the presence of price regulation. Figure 4 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.⁶

Oil supply shocks are important in explaining some increases in the real price of oil in the 1980s, 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), only partly the sharp spike in the real price of oil in 1990–1991 after the invasion of Kuwait, and the sharp rise in the real oil price in 1999–2000. 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 sudden increase in the real oil price in 1990–1991, 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 1980s.

Regarding the fluctuations of US output, the historical decomposition shows that US shocks dominate the shocks originated in the RoW. However, the distinction

6. Price regulation could make the WTI price of oil unrepresentative of the oil market during part of the sample period. 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 1979–1980. However, as shown in the Online Appendix, the historical decomposition obtained with an alternative VAR including this oil price measure provides results virtually identical to those obtained with our benchmark model.

Real oil price



US industrial production.

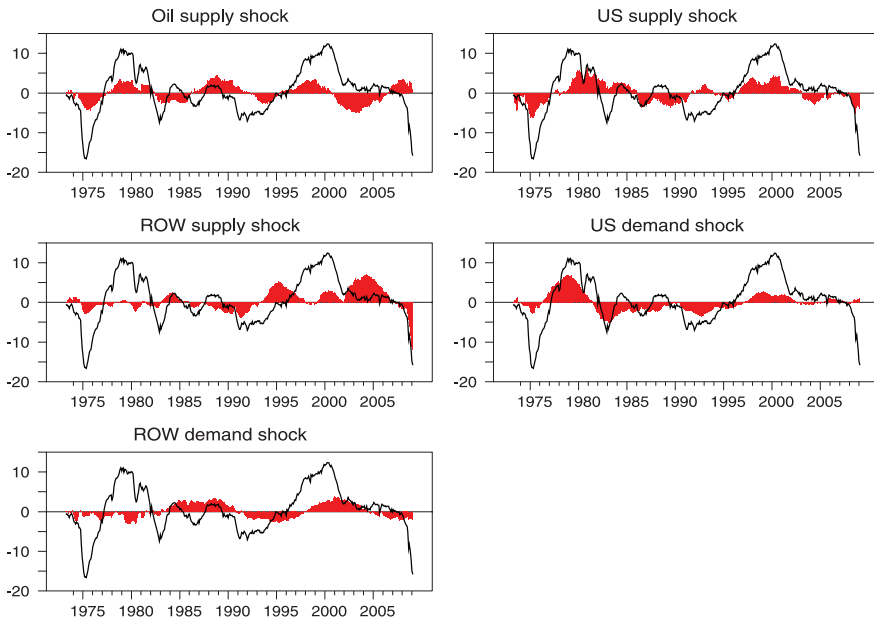


FIGURE 4. Historical decomposition. The thin line denotes the real oil price (or the US industrial production), in deviation from the baseline. The bars in each panel denote the component of the series accounted for by each structural shock. The “Real oil price” is the ratio between the oil price and the US CPI.

between oil supply shocks and RoW supply shocks remains relevant for the interpretation of the effects on US output. Only oil supply shocks contributed significantly to subsequent US economic slowdown or recessions: this happened indeed in 1983, to some extent in the 1990/1991 episode and in 2000–2001. In contrast, oil price hikes generated by RoW supply shocks played a negligible role in explaining historical episodes characterized by a fall in the US output.

5. Robustness

The robustness of the findings was tested on several dimensions.

5.1. On “Oil-Market-Specific” Demand Shocks

Kilian (2009) argues that “oil market-specific demand shocks”, namely oil price hikes related to concerns about future oil supply shortfalls, have a negative effect on the US economic activity which is more persistent than the one implied by oil supply shocks and by “aggregate demand shocks”, and that they explain many historical episodes characterized by major oil price fluctuations, such as the sharp fall of the real oil price in late 1985 and its sudden spike in 1990–1991. The role of these shocks has been also assessed empirically in Kilian and Murphy (forthcoming) and Alquist and Kilian (2010).

Our simple model economy does not allow for such “expectational shocks”. To explore the hypothesis that the estimates described in Figure 2 may also reflect the oil-market-specific demand shocks, we abandon (temporarily) the internal consistency of our model economy and consider the following modification of our benchmark VAR model. We replace the RoW demand shock with an oil-market-specific demand shock, whose identification scheme is presented in Table 4 below. In this scheme, the oil-market-specific demand shock is different from the RoW demand shock because the RoW output is assumed to respond negatively, rather than positively, to a positive

TABLE 4. Sign restrictions with “oil-market-specific” demand shock.

VAR variables	Structural shocks				
	oil supply shock	RoW supply shock	oil-market-specific demand shock	US supply demand shock	US demand shock
Oil production	–		+		+
Oil price ^a	+	+	+	+	–
US output	–		–	+	+
RoW output	–	+	–		–
US output price ^a		+		–	+

Notes: 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.

a. Price is deflated by the US CPI.

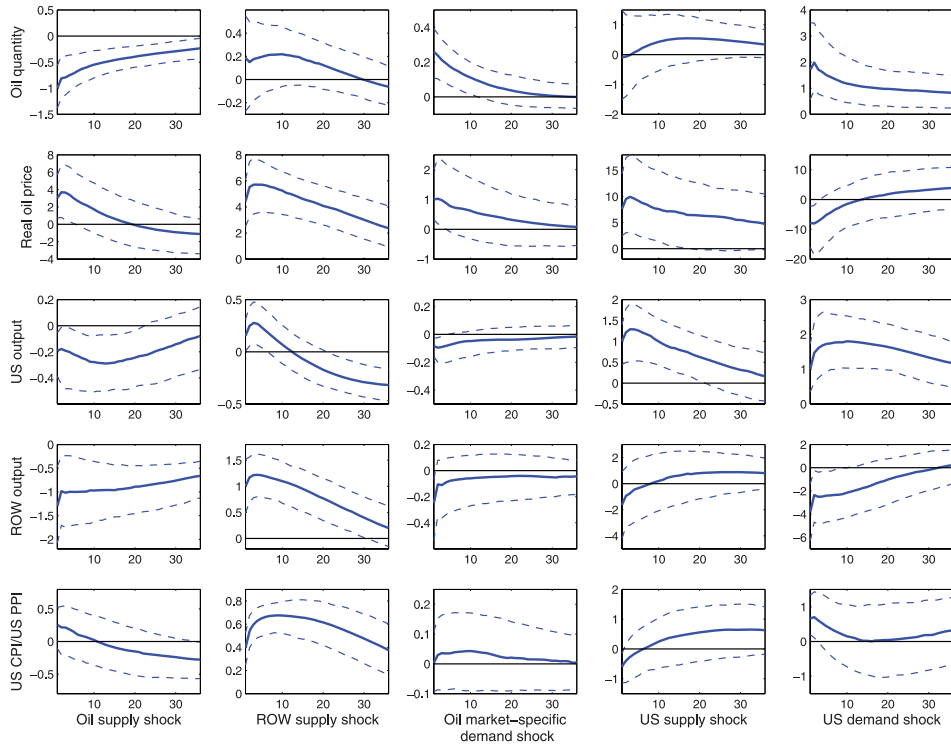


FIGURE 5. Impulse response in a scheme that allows for oil-market-specific demand shocks, showing the 16th (dashed), 50th (solid), and 84th (dashed) percentiles of the IRFs distribution. The “Real oil price” is the ratio between the oil price and the US CPI.

innovation. The US relative price is left unconstrained. This scheme generates mutually orthogonal shocks, and nests the scheme used by Kilian and Murphy (forthcoming) for their smaller-scale VAR model, where the sign restrictions imposed on the oil quantity, the real oil price, and the global economic activity are consistent with those used in Table 4.

The impulse responses, presented in Figure 5, suggest that the effects of the oil-market-specific demand shock on the US industrial production are small, and significantly different from zero only over the very short run. In contrast, the effects of a RoW demand shock produced by our benchmark model are larger and more persistent. A similar pattern (i.e. smaller responses) emerges concerning the effect of oil-market-specific demand shocks on the real oil price. One possible explanation of the small effect of precautionary shocks is that they are less persistent than the RoW demand shocks. By contrast, the estimated effect of the other shocks is similar to the one produced by our benchmark model. Likewise, the historical decomposition shows that oil-market-specific demand shocks explain a small fraction of the real oil price fluctuations and do not seem to have been an important driver of the US recessions in the past.

TABLE 5. Median identification matrix produced by the benchmark model.

VAR variables	Structural shocks				
	Oil supply shock	RoW supply shock	RoW demand shock	US supply shock	US demand shock
Oil production	-0.0032	0.0013	0.0089	0.0007	0.0105
Oil price ^a	0.0411	0.0511	0.0192	0.0287	-0.0140
US output	-0.0028	0.0020	-0.0021	0.0036	0.0015
RoW output	-0.0126	0.0124	0.0050	-0.0078	-0.0075
US output price ^a	0.0024	0.0051	-0.0012	-0.0024	0.0015
Elasticity of oil quantity to oil price	-0.0779	0.0254	0.4635	0.0244	-0.7500

The values are computed as the median of all matrices satisfying the sign restrictions.

One simple explanation of the difference with Kilian (2009) findings is that the shock that he labels “oil-market specific” is capturing the effects of a RoW aggregate demand shock. Indeed, a positive oil-market-specific demand shock generates an increase in the real price of oil, a positive short-run response of global output, albeit small, and a negative response of the US real GDP. This is fully consistent with the effects of a RoW aggregate demand shock, see Table 2. Moreover, as pointed out also by Kilian and Murphy (forthcoming), some features of the oil-market-specific shock are “mildly implausible”: a positive oil-market-specific demand shock should have a negative effect on both the global output and the US real GDP. But in the data the output response of the US and the RoW is different. By contrast, these effects are consistent with our identified RoW demand shock, which also generates an increase in the real oil price, a positive effect on the RoW output, and a negative effect on the US output.

5.2. On the Estimated Elasticity Of Oil Supply Curve

Kilian and Murphy (forthcoming) warn about the adoption of the sign restriction approach when applied to empirical studies for the oil market. In particular, they argue that the use of the median to report the impulse response functions may be inconsistent with desirable economic features, such as a rather small price elasticity of the oil supply (Kilian 2009; Hamilton 2009), thus raising concerns about the estimated effects on the other variables of interest in the VAR.⁷

Table 5 reports the identification matrix obtained with the median estimates of our empirical analysis. We find it interesting that our median estimate for the price

7. Kilian and Murphy (forthcoming) showed that the elasticity of the oil supply curve to an aggregate demand shock generated with the sign restriction approach is 1.89. They suggest combining sign restrictions with empirically plausible bounds on the magnitude of the short-run oil supply elasticity. The bound imposed to these fundamental structural parameters is 0.0257, derived by considering the changes in the crude oil production and in the real oil price that occurred during a single particular episode (the outbreak of the Persian Gulf War on August 1990).

elasticity of the oil supply following a RoW supply or a US supply shock is very close to the bound discussed by Kilian (2009), around 3%. Therefore, for those shocks that explain the bulk of fluctuations, as discussed in Section 4.2, our identification scheme does not seem to generate economically implausible estimated elasticities of the oil supply curve, at least judging by the criterion proposed by Kilian and Murphy (forthcoming).

The estimated oil supply elasticity to a RoW demand shock is instead about 0.5, which is still much smaller than the high value of 1.89 found by Kilian and Murphy (forthcoming). While it is hard to assess whether this computed elasticity is economically plausible, to further our analysis we estimated the model by relaxing the assumption of a positive response of the oil quantity when identifying a RoW demand (or a US demand) shock. Notice that Table 2 shows that relaxing these restrictions does not impede the identification of the shocks, which remain mutually exclusive. We thus estimate the same benchmark VAR model without constraining the impact response of the oil quantity following a RoW and a US demand shock. The estimated price elasticity of the oil supply remain identical for both the RoW supply and the US supply shock; instead, the price elasticity of the oil supply following a RoW demand shock becomes nonsignificant. The main results of our paper remain unchanged (more details are in the Online Appendix). In particular, following a RoW demand shock, the increase in the real oil price becomes slightly smaller and less persistent, while the response of US output is remarkably similar to that obtained with the benchmark identification scheme.

We conclude that the sign restrictions we imposed in our benchmark identification scheme on oil production do not seem to be crucial for the assessment of the median responses on the variables of interest, notably the US output. The impact sign restrictions imposed on the other variables are instead more important for a reliable identification of the shocks and for the assessment of their effects on the US output.

5.3. *Other Econometric Issues*

We briefly summarize the findings of a few other issues that have been explored. As the Online Appendix shows, first, we compare the impulse responses of our benchmark model with those stemming from the same VAR but using a scheme that identifies only three, as opposed to five, structural shocks. This corresponds to identifying \tilde{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 corroborate the importance of identifying simultaneously both US-specific shocks, as well as shocks originating outside the US economy, in order to avoid biased estimated response of the US industrial production to structural shocks moving the real price of oil.

Second, we re-estimated the VAR where we use first-difference of the non-stationary variables, instead of including 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. The estimated impulse responses suggest that our results are qualitatively robust to the way we handle the non-stationarity of the variables. Interestingly, the variance decomposition analysis is also consistent with the one previously discussed.

Third, we investigated whether our results 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 does 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 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. Impulse responses suggest results very similar to those obtained with the US industrial production index.

Finally, we explored the consequence of applying the sign restrictions on the US relative prices in later periods than the impact one. This might be important if prices adjust slowly to shocks. In particular, we check the results obtained by imposing sign restrictions on the US relative prices after three and six months, leaving unconstrained the response in the previous periods. Results are very similar to those obtained with the benchmark identification scheme. An exception is a stronger response of the US output following a RoW demand shock.

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 fundamental shocks underlying observed time series from the oil market, the US economy, and the global business cycle.

We summarize the main findings of the paper as follows. First, the variance decomposition analysis shows that about one-half of the (real) oil price fluctuations are explained by shocks to the RoW business cycle. The oil price is also shown to respond to shocks originated in the US economy, which explain a fraction of variance comparable of the one stemming from canonical oil supply shocks. This finding is novel and highlights the importance of not assuming that oil market variables are predetermined with respect to the US economy. The reverse causality, from aggregate demand and supply shocks to the oil price, supports the effort in building models where the oil price, like all other prices, responds to business cycle shocks (see the recent papers by Bodenstein, Erceg, and Guerrieri (2007) and Nakov and Pescatori (2010b).

Second, the estimates qualify the results in Kilian (2009) by showing that not all oil demand shocks are alike. In particular, positive innovations to RoW-demand or RoW-supply shocks increase global output and the real price of oil, but have opposite

implications concerning the US industrial production at horizons of up to one year. Therefore, shocks that appear as “oil-demand” may have very different implications for the US depending on whether they are originated in the US or in the RoW. The theory offers a simple explanation for this finding: a supply shock in the rest of the world increases world income and consumption and, provided the substitution elasticity between home and foreign goods is small enough, it also increases the production of US goods. Instead, an aggregate demand shock in the RoW does not increase the world income, it increases the cost of US imported goods, and hence reduces US demand and production.

Third, 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 role of oil supply shocks with respect to US output fluctuations is limited, explaining about 10% of the total variance. This is due to the fact that the variance of these shocks is small compared to the variance of domestic aggregate demand and supply shocks.

These findings offer a simple interpretation of the small and unstable correlation between oil prices and the US economic activity documented in, for example, 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 United States. The unconditional correlation between oil prices and US production over a long sample period is tenuous because it blends conditional correlations with different signs. Our explanation does not appeal to “structural change”. In this sense it is different from the hypothesis recently put forth by Blanchard and Gali (2010) who maintain the assumption that oil prices are predetermined to the US economy, and argue that the smaller effect of oil shocks on the US economy in recent years can be ascribed to structural changes, such as changes in real and nominal wage rigidity, or the energy share of production.

We see several interesting questions for future research. For instance, the simple structural model we considered abstracts from other mechanisms that might have affected oil prices, such as “speculation” or “precautionary” oil shocks (Alquist and Kilian 2010; Dvir and Rogoff 2009). Our tentative analysis on the role of these shocks, in Section 5.1, seemed to downplay their importance. But we think that a more rigorous analysis is necessary to precisely identify these shocks and their effects.

Appendix: Data Description and Source of the Variables

US Output. US Industrial Production Index (index 2007 = 100), seasonally adjusted, measured in logarithms. Source: Board of Governors of the Federal Reserve System.

US CPI. US Consumer Price Index for All Urban Consumers (index 1982–1984 = 100), all items, seasonally adjusted. Source: US Department of Labor, Bureau of Labor Statistics.

US Relative Prices. US Producer Price Index (index 1982=100), all commodities, not seasonally adjusted. It is expressed in real terms (e.g. deflated by the US Consumer Price Index) and measured in logarithms. Source: authors' calculation based on data from US Department of Labor, Bureau of Labor Statistics.

Oil Quantity. Global oil production in barrels per day, measured in logarithms. Source: International Energy Agency.

Oil Price. The nominal spot oil price is the simple arithmetic average of the UK Brent, Dubai Fateh, and West Texas Intermediate spot prices, in dollars per barrel. It is expressed in real terms (e.g. deflated by the US CPI) and measured in logarithms. Source: authors' calculations based on data from International Monetary Fund (IMF) and US Department of Labor, Bureau of Labor Statistics.

RoW Output. Global exports to World (IMF code: WDI7D0WDA) net of global imports from United States (IMF code: WDI7D1USA) and global imports from oil-exporting countries (IMF code: WDI7D1OPA), seasonally adjusted by TRAMO-SEATS. It is expressed in real terms (e.g. deflated by the US CPI) and measured in logarithms. Source: authors' calculations based on data from Thomson Reuters Datastream and US Department of Labor, Bureau of Labor Statistics.

CFNAI Index. Chicago Fed National Activity Index. Source: Federal Reserve Bank of Chicago.

Kilian's Measure of World Output. Global index of dry cargo single-voyage freight rates, deflated by the US Consumer Price Index, expressed in deviation from a long-run linear trend. Source: <http://www-personal.umich.edu/lkilian/rea.txt>.

US Refiners' Acquisition Cost of Crude Oil. US crude oil composite (domestic and imported) acquisition cost by refiners, expressed in dollars per barrel and measured by logarithms. Source: US Energy Information Administration.

Supporting Information

Additional Supporting Information may be found in the online version of this article:

Appendix S1. Fry and Pagan's critique (pdf file)

Appendix S2. Data sets and software for replication of the analysis in the paper (zip file)

Please note: Blackwell Publishing are not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.

References

- Alquist, Ron, and Lutz Kilian (2010). "What Do We Learn from the Price of Crude Oil Futures?" *Journal of Applied Econometrics*, 25, 539–573.
- Anzuini, Alessio, Patrizio Pagano, and Massimiliano Pisani (2007). "Oil Supply News in a VAR: Information from Financial Markets." Temi di Discussione (Economic Working Papers) No. 632, Economic Research Dept, Bank of Italy. Available at http://ideas.repec.org/p/bdi/wptemi/td_632_07.html.
- Backus, David K., and Mario J. Crucini (2000). "Oil Prices and the Terms of Trade." *Journal of International Economics*, 50, 185–213.
- Backus, David K., Patrick J. Kehoe, and Finn E. Kydland (1994). "Dynamics of the Trade Balance and the Terms of Trade: The J-Curve?" *American Economic Review*, 84(1), 84–103.
- Barsky, Robert B., and Lutz Kilian (2002). "Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative." In *NBER Macroeconomics Annual 2001*, edited by B. Bernanke and K. Rogoff. MIT Press, pp. 137–183.
- Barsky, Robert B., and Lutz Kilian (2004). "Oil and the Macroeconomy since the 1970s." *Journal of Economic Perspectives*, 18(4), 115–134.
- Baumeister, Christiane, and Gert Peersman (2008). "Time-Varying Effects of Oil Supply Shocks on the US Economy." Working Paper No. 08/515, Ghent University.
- Baxter, Marianne, and Mario J. Crucini (1993). "Explaining Saving Investment Correlations." *American Economic Review*, 83(3), 416–436.
- Blanchard, Olivier, and Jordi Gali (2010). "Labor Markets and Monetary Policy: A New Keynesian Model with Unemployment." *American Economic Journal: Macroeconomics*, 2, 1–30.
- Blanchard, Olivier J., and Marianna Riggi (2009). "Why are the 2000s so Different from the 1970s? A Structural Interpretation of Changes in the Macroeconomic Effects of Oil Prices." NBER Working Paper No. 15467. Available at <http://ideas.repec.org/p/nbr/nberwo/15467.html>.
- Bodenstein, Martin, Christopher J. Erceg, and Luca Guerrieri (2007). "Oil Shocks and External Adjustment." FRB International Finance Discussion Paper No. 897.
- Canova, Fabio, and Gianni De Nicolò (2002). "Monetary Disturbances Matter for Business Fluctuations in the G-7." *Journal of Monetary Economics*, 49, 1131–1159.
- Canova, Fabio, and Matthias Paustian (2007). "Measurement with Some Theory: Using Sign restrictions to Evaluate Business Cycle Models." Working paper, CREI.
- Christiano, Lawrence J., Martin Eichenbaum, and Robert Vigfusson (2004). "The Response of Hours to a Technology Shock: Evidence Based on Direct Measures of Technology." *Journal of the European Economic Association*, 2, 381–395.
- Cunado, Juncal, and Fernando Perez de Gracia (2003). "Do oil price shocks matter? Evidence for some European countries." *Energy Economics*, 25, 137–154.
- Davis, Steven J., and John Haltiwanger (1999). "On the Driving Forces behind Cyclical Movements in Employment and Job Reallocation." *American Economic Review*, 89(5), 1234–1258.
- Dedola, Luca, and Stefano Neri (2007). "What Does a Technology Shock do? A VAR Analysis with Model-Based Sign Restrictions." *Journal of Monetary Economics*, 54, 512–549.
- Dvir, Eyal, and Kenneth S. Rogoff (2009). "Three Epochs of Oil." NBER Working Paper No. 14927. Available at <http://ideas.repec.org/p/nbr/nberwo/14927.html>.
- Edelstein, Paul, and Lutz Kilian (2009). "How Sensitive are Consumer Expenditures to Retail Energy Prices?" *Journal of Monetary Economics*, 56, 766–779.
- Francis, Neville, and Valerie A. Ramey (2005). "Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited." *Journal of Monetary Economics*, 52, 1379–1399.
- Fry, Renee, and Adrian Pagan (2007). "Some Issues in Using Sign Restrictions for Identifying Structural VARs." NCEER Working Paper No. 14.
- Hamilton, James D. (1983). "Oil and the Macroeconomy since World War II." *Journal of Political Economy*, 91, 228–248.

- Hamilton, James D. (2008). "Oil and the Macroeconomy." In *The New Palgrave Dictionary of Economics*, 2nd ed, edited by S. Durlauf and L. Blume. Palgrave MacMillan.
- Hamilton, James D. (2009). "Causes and Consequences of the Oil Shock of 2007–08." NBER Working Paper No. 15002. Available at <http://ideas.repec.org/p/nbr/nberwo/15002.html>.
- Hamilton, James D, and Ana Maria Herrera (2004). "Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy: Comment." *Journal of Money, Credit and Banking*, 36, 265–86.
- Hooker, Mark A. (1996). "What Happened to the Oil Price–Macroeconomy Relationship?" *Journal of Monetary Economics*, 38, 195–213.
- Jimenez-Rodriguez, Rebeca, and Marcelo Sanchez (2005). "Oil Price Shocks and Real GDP Growth: Empirical Evidence for some OECD Countries." *Applied Economics*, 37, 201–228.
- Kilian, Lutz (2008). "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, 216–240.
- Kilian, Lutz (2009). "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market." *American Economic Review*, 99(3), 1053–1069.
- Kilian, Lutz, and Dan Murphy (forthcoming). "Why Agnostic Sign Restrictions Are Not Enough: Understanding the Dynamics of Oil Market VAR Models." *Journal of the European Economic Association*.
- Kilian, Lutz, and Clara Vega (2008). "Do Energy Prices Respond to U.S. Macroeconomic News? A Test of the Hypothesis of Predetermined Energy Prices." CEPR Discussion Paper No. 7015. Available at <http://ideas.repec.org/p/cpr/ceprdp/7015.html>.
- Leduc, Sylvain, and Keith Sill (2004). "A Quantitative Analysis of Oil-Price Shocks, Systematic Monetary Policy, and Economic Downturns." *Journal of Monetary Economics*, 51, 781–808.
- Mork, Knut A., Oystein Olsen, and Hans Terje Mysen (1994). "Macroeconomic Responses to Oil Price Increases and Decreases in Seven OECD Countries." *Energy Journal*, 15, 19–35.
- Nakov, Anton, and Andrea Pescatori (2010a). "Monetary Policy Trade Offs with a Dominant Oil Producer." *Journal of Money, Credit and Banking*, 42, 1–32.
- Nakov, Anton, and Andrea Pescatori (2010b). "Oil and the Great Moderation." *Economic Journal*, 120, 131–156.
- Pagano, Patrizio, and Massimiliano Pisani (2009). "Risk-Adjusted Forecasts of Oil Prices." *The B.E. Journal of Macroeconomics*, 9(1).
- Pappa, Evi (2009). "The Effects Of Fiscal Shocks On Employment And The Real Wage." *International Economic Review*, 50, 217–244.
- Peersman, Gert, and Roland Straub (2004). "Technology Shocks and Robust Sign Restrictions in a Euro Area SVAR." Working Paper No. 373, European Central Bank.
- Peersman, Gert, and Roland Straub (2009). "Technology Shocks And Robust Sign Restrictions In A Euro Area Svar." *International Economic Review*, 50, 727–750.
- Rubio-Ramirez, Juan F., Daniel Waggoner, and Tao Zha (2005). "Markov Switching Structural Vector Autoregressions: Theory and Application." Federal Reserve Bank of Atlanta Working Paper No. 27.
- Uhlig, Harald. (2005). "What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure." *Journal of Monetary Economics*, 52, 381–419.