

# Time-Varying Effects of Oil Supply Shocks on the US Economy\*

Christiane Baumeister  
*Bank of Canada*  
cbaumeister@bankofcanada.ca

Gert Peersman  
*Ghent University*  
gert.peersman@ugent.be

April 11, 2011

## Abstract

We investigate how the dynamic effects of oil supply shocks on the US economy have changed over time. We first document a remarkable structural change in the oil market itself, i.e. a considerably steeper, hence, less elastic oil demand curve since the mid-eighties. Accordingly, a typical oil supply shock is currently characterized by a much smaller impact on world oil production and a greater effect on the real price of crude oil, but has a similar impact on US output and inflation as in the 1970s. Second, we find a smaller role for oil supply shocks in accounting for real oil price variability over time, implying that current oil price fluctuations are more demand driven. Finally, while unfavorable oil supply disturbances explain little of the "Great Inflation", they seem to have contributed to the 1974/75, early 1980s and 1990s recessions but also dampened the economic boom at the end of the millennium.

*JEL classification:* E31, E32, Q43

*Keywords:* Oil prices, vector autoregressions, time-varying coefficients

---

\*We thank Luca Benati, Jean Boivin, Luca Gambetti, Domenico Giannone, James Hamilton, Morten Ravn, Frank Smets, Herman van Dijk and Timo Wollmershäuser as well as conference participants at the CESifo Munich, IWH Halle, EEA-ESEM Milan and seminar participants at the Erasmus University Rotterdam and the University of Kent for their useful comments and suggestions. Special thanks to Lutz Kilian for many insightful discussions and feedback on an earlier version of the paper. We acknowledge financial support from the Interuniversity Attraction Poles Programme - Belgian Science Policy [Contract No. P6/07] and the Belgian National Science Foundation. All remaining errors are ours.

# 1 Introduction

The belief that physical disruptions in the supply of crude oil in the world oil market are at the origin of stagflationary episodes has gained ground in the wake of the two big oil shocks of the 1970s.<sup>1</sup> Since that time oil price increases are associated in many people's minds with severe macroeconomic consequences in terms of higher inflation and lower economic growth. A remarkable feature of the recent, prolonged surge in oil prices however, is the relatively mild impact this seems to exert on real economic activity and the price level. This observation casts doubt on the relevance of oil shocks for the macroeconomic performance of the US economy in more recent times. In other words, the way the economy experiences oil shocks appears to have changed fundamentally. This conjecture has recently been confirmed in the empirical literature by Edelstein and Kilian (2007), Herrera and Pesavento (2009) and Blanchard and Galí (henceforth, BG 2007). In particular, these studies find a reduced impact of oil price shocks on US macroeconomic aggregates over time. The objective of this paper is to further investigate how the dynamic effects of oil supply disturbances have changed and whether global oil supply shocks can be considered an important source of economic fluctuations. We evaluate the importance of physical production shortfalls in explaining fluctuations in the real price of oil as well as their contribution to the observed movements in US real GDP growth and consumer price inflation when time variation is accounted for.

The main results that emerge from our analysis are the following. We foremost document a remarkable structural change in the global oil market. Specifically, a "typical" oil supply shock is characterized by a much smaller impact on world oil production and a greater effect on the real price of crude oil since the second half of the 1980s. Only a steeper, hence, less elastic oil demand curve can explain this stylized fact. This finding has important consequences when the macroeconomic effects of oil supply shocks are compared over time. If a comparison is based on a similar change of crude oil prices (e.g. a 10 percent rise), we currently find a more muted impact on the US economy which is consistent with the existing evidence for oil price shocks. However, such a comparison is not informative since it ignores the fact that the underlying oil supply shock is of a different nature. In particular, a constant slope of the oil demand curve is implicitly assumed in such a scenario which is at odds with our evidence. On the other hand, if an exogenous oil supply shock is measured as a similar shift in world oil production (e.g. a production shortfall of 1 percent), such a disturbance has a much greater impact on oil prices now, resulting also in stronger effects on real GDP and consumer prices compared to the 1970s and early 1980s. However, also this comparison is biased since an average oil supply shock is characterized by a disturbance in oil production of more than 2 percent in the 1970s and hardly 0.5 percent since the 1990s. When we consider a typical one standard deviation oil supply shock instead, we find a rather similar impact on US macroeconomic aggregates over time. In addition, oil supply disturbances consistently account for 15 to 20 percent of output and inflation variability.

---

<sup>1</sup>Hamilton (1983) was the first to notice that all but one of the US postwar recessions have followed sharp increases in the price of crude oil.

If the effects of average oil supply shocks have not dramatically changed since the 1970s, it is surprising that we are currently not confronted with similar macroeconomic conditions. To explain this, we demonstrate that oil supply shocks explain little of the Great Inflation, which is consistent with the propositions of Barsky and Kilian (henceforth, BK 2002, 2004). In addition, oil supply disturbances seem to have played a significant but certainly non-exclusive role in the 1974/75, early 1980s and 1990s recessions; by the same token, unfavorable oil supply shocks in 1999 made a significant negative contribution to the ongoing boom at the end of the millennium. Moreover, we show that the contribution of oil supply shocks to fluctuations in the real price of oil has decreased over time which means that current oil price movements are more demand driven. Despite a relatively constant share of oil supply shocks in explaining the variance of oil production growth (being constantly around 30 percent) and the aforementioned higher leverage effect on oil prices, the latter finding implies that also the oil supply curve is currently more inelastic. A steepening of both the oil demand and oil supply curves can be considered as a source of increased oil price volatility in more recent times.

The analysis in this paper departs from the existing literature along two dimensions. First, the paper makes use of recent methodological contributions in explicitly modeling time variation. Instabilities over time in the crude oil market and the oil-macro relationship have been widely documented in the literature.<sup>2</sup> On the one hand, the oil market itself has undergone substantial changes. Global capacity utilization rates in crude oil production have not been constant over time, with production levels being above sustainable capacity since the late 1980s, as well as in 1973/74 and 1979/80 (Kilian 2008a). In addition, the transition from a regime of administered oil prices to a market-based system of direct trading in the spot market and the collapse of the OPEC cartel in late 1985 were accompanied by a dramatic rise in oil price volatility (e.g. Hubbard 1986). Furthermore, the relative importance of the driving forces behind oil price movements has changed (e.g. BK 2002, 2004; Kilian 2009; Hamilton 2003, 2008a; Rotemberg 2007).<sup>3</sup> On the other hand, also the macroeconomic structure has changed over time which can bring about time-varying effects of oil shocks.<sup>4</sup> Prominent explanations for different macroeconomic consequences of oil shocks over time discussed in the literature are improved monetary policy (e.g. Bernanke et al. 1997; BG 2007),<sup>5</sup> more flexible labor markets (BG 2007), changes in the composition of automobile production and the overall importance of the US automobile sector (Edelstein and Kilian 2009), and variations in the role and share of oil in the economy over time (e.g. Bernanke 2006; BG 2007; Edelstein and

---

<sup>2</sup>Structural breaks in the relationship between oil prices and the macroeconomy were initially documented by Mork (1989) and Hooker (1996, 2002), when sample periods were extended to include data from the mid-1980s onwards.

<sup>3</sup>Lee et al. (1995) and Ferderer (1996) argue that the increased volatility has led to a breakdown of the empirical relationship between oil prices and economic activity in the US since the mid-eighties. Kilian (2008b, 2009) stresses that shifts in the composition of oil supply and oil demand shocks driving the price of oil help explain why regressions of macroeconomic aggregates on oil prices tend to be unstable over time and lead to the changing effects of oil price shocks.

<sup>4</sup>These structural changes are very much related to the literature on the "Great Moderation". See, for instance, McConnell and Pérez-Quirós (2000), Blanchard and Simon (2001) and Stock and Watson (2003).

<sup>5</sup>This argument has been defied by Hamilton and Herrera (2004) and Herrera and Pesavento (2009).

Kilian 2009).

All these changes suggest that a linear, constant-coefficient specification may not accurately capture the effects of oil supply shocks on the US economy. Consequently, time variation has to be allowed for in order to adequately model the interaction between oil shocks and aggregate economic activity and to explore how this relationship has evolved over time. Several studies, using a vector autoregression approach (e.g. Edelstein and Kilian 2007; Herrera and Pesavento 2009), take time variation into account by splitting the sample into two subperiods assuming a structural break sometime in the 1980s.<sup>6</sup> Alternatively, BG (2007) allow for a more gradual variation over time by estimating bivariate vector autoregressions over rolling time windows. To model time variation, we estimate a multivariate time-varying parameters Bayesian vector autoregression (TVP-BVAR) with stochastic volatility for the period 1970Q1-2006Q2 in the spirit of Cogley and Sargent (2002, 2005), Primiceri (2005) and Benati and Mumtaz (2007). The varying coefficients are meant to capture smooth transitions in the propagation mechanism of oil shocks without imposing a specific breakpoint, while the stochastic volatility component models changes in the magnitude of structural shocks and their immediate impact. The latter feature is particularly important in the present setting given the documented increased volatility in the oil market and the reduced macroeconomic volatility. By using a multivariate approach, it is also possible to learn more about potential sources of time variation.

Second, we propose a new identification strategy to isolate the unanticipated movements in the price of crude oil due to exogenous supply shocks. Most studies, including BG (2007), rely on a recursive identification scheme where all variations in oil prices are assumed to be oil supply shocks. This view, however, has been challenged in the recent literature. BK (2002), for instance, argue that the oil price rises of the seventies might be the result of expansionary monetary policy and Kilian (2008a) shows that only a small fraction of observed oil price fluctuations can be attributed to exogenous oil production disruptions. Therefore Kilian (2009), based on the assumption of a vertical short-run supply curve, identifies oil supply shocks as the only source of innovations in global oil production in a monthly VAR; thus, shocks to the demand for crude oil have an immediate effect on oil prices but not on production volumes. His identifying assumptions are, however, not designed for estimations with quarterly data such as real GDP. For that reason he uses a single-equation approach to estimate, in a second step, the impact of quarterly shocks on US real GDP and CPI inflation by averaging the monthly structural innovations over each quarter.<sup>7</sup>

To allow for an immediate effect of both oil supply and oil demand shocks on oil production and the real price of crude oil, we propose to use sign restrictions instead, to identify oil supply shocks in a quarterly VAR.<sup>8</sup> Specifically, the sign restrictions are derived from a simple supply

---

<sup>6</sup>The break date is mostly chosen in relation to the start of the Great Moderation in 1984 or the oil price plunge in early 1986.

<sup>7</sup>In this way, Kilian (2009) also demonstrates that the exact underlying source of oil price movements is crucial for the subsequent economic consequences in the US.

<sup>8</sup>See Faust (1998), Canova and De Nicoló (2002), Uhlig (2005) and Peersman and Straub (2009) for alternative applications of this identification strategy. In particular, Peersman (2005) uses sign restrictions to identify oil shocks

and demand model of the world oil market. Using both world oil production and oil price data, oil supply shocks are identified as the only disturbances that displace the oil supply curve in this market. As a consequence, our method recognizes the fact that contemporaneous movements in oil prices and oil production could also be driven by disturbances relating to the demand for crude oil. Moreover, given our time-varying framework, the magnitude of the contemporaneous impact of structural shocks on oil production and prices can vary over time and hence, changes in the relative importance of oil supply and demand shocks and the slope of both curves are captured. We also show that our conclusions do not depend on the selected methodology and underlying assumptions. In particular, robust results are found for alternative specifications, the modeling of time variation and the identification strategy.

The remainder of the paper is structured as follows. Section 2 presents the methodology and describes the identification strategy in more detail. Section 3 discusses the main empirical results and evaluates the robustness of our findings. In section 4, we discuss some potential explanations for a less elastic oil demand curve in more recent periods, and section 5 offers some concluding remarks.

## 2 Methodology

### 2.1 A VAR with time-varying parameters and stochastic volatility

We model the joint behavior of global oil production, the real refiner acquisition cost of imported crude oil,<sup>9</sup> US real GDP and US CPI as a VAR with time-varying parameters:<sup>10</sup>

$$y_t = c_t + B_{1,t}y_{t-1} + \dots + B_{p,t}y_{t-p} + u_t \equiv X_t'\theta_t + u_t \quad (1)$$

The frequency of our data is quarterly<sup>11</sup> and the overall sample covers the period 1947Q1-2006Q2. The first twenty years of data are, however, used as a training sample to generate the priors for the actual sample period.<sup>12</sup> The lag length is set to  $p = 4$  to allow for sufficient dynamics

---

but without including world oil production.

<sup>9</sup>This oil price variable measures most accurately the marginal cost of crude oil to US refiners. It is therefore the best proxy for the free market global price of imported crude oil. For different concepts of world oil prices, see Mabro (2005). We have checked the sensitivity of our results to alternative oil price measures such as the WTI spot oil price and the *composite* refiner acquisition cost; our conclusions are not altered by the choice of this variable. Results are available upon request.

<sup>10</sup>All variables are transformed to non-annualised quarter-on-quarter rates of growth by taking the first difference of the natural logarithm. The nominal refiner acquisition cost has been deflated using the US CPI. A detailed description of all the data used in this paper can be found in Appendix A.

<sup>11</sup>In Section 3.3, we investigate the robustness of our results when monthly data are used.

<sup>12</sup>We have also experimented with shorter sample periods to calibrate our priors. Given sufficiently diffuse priors, our results were not altered by the choice of the training sample.

in the system.<sup>13</sup> The time-varying intercepts and lagged coefficients are stacked in  $\theta_t$  to obtain the state-space representation of the model. The  $u_t$  of the observation equation are heteroskedastic disturbance terms with zero mean and a time-varying covariance matrix  $\Omega_t$  which can be decomposed in the following way:  $\Omega_t = A_t^{-1}H_t(A_t^{-1})'$ .  $A_t$  is a lower triangular matrix that models the contemporaneous interactions among the endogenous variables and  $H_t$  is a diagonal matrix which contains the stochastic volatilities:

$$A_t = \begin{bmatrix} 1 & 0 & 0 & 0 \\ \alpha_{21,t} & 1 & 0 & 0 \\ \alpha_{31,t} & \alpha_{32,t} & 1 & 0 \\ \alpha_{41,t} & \alpha_{42,t} & \alpha_{43,t} & 1 \end{bmatrix} \quad H_t = \begin{bmatrix} h_{1,t} & 0 & 0 & 0 \\ 0 & h_{2,t} & 0 & 0 \\ 0 & 0 & h_{3,t} & 0 \\ 0 & 0 & 0 & h_{4,t} \end{bmatrix} \quad (2)$$

The drifting coefficients are meant to capture possible nonlinearities or time variation in the lag structure of the model. The multivariate time-varying variance covariance matrix allows for heteroskedasticity of the shocks and time variation in the simultaneous relationships between the variables in the system. Allowing for time variation in both the coefficients and the variance covariance matrix, leaves it up to the data to determine whether the time variation of the linear structure comes from changes in the size of the shock and its contemporaneous impact (impulse) or from changes in the propagation mechanism (response). Given the observed instabilities in the oil-macro relationship, this approach is particularly expedient. Let  $\alpha_t$  be the vector of non-zero and non-one elements of the matrix  $A_t$  (stacked by rows) and  $h_t$  be the vector containing the diagonal elements of  $H_t$ . Following Primiceri (2005), the three driving processes of the system are postulated to evolve as follows:

$$\theta_t = \theta_{t-1} + \nu_t \quad \nu_t \sim N(0, Q) \quad (3)$$

$$\alpha_t = \alpha_{t-1} + \zeta_t \quad \zeta_t \sim N(0, S) \quad (4)$$

$$\ln h_{i,t} = \ln h_{i,t-1} + \sigma_i \eta_{i,t} \quad \eta_{i,t} \sim N(0, 1) \quad (5)$$

The time-varying parameters  $\theta_t$  and  $\alpha_t$  are modeled as driftless random walks.<sup>14</sup> In line with Primiceri (2005) we do not impose a stability constraint on the evolution of the time-varying parameters to enforce stationarity of the VAR system.<sup>15</sup> The elements of the vector of volatilities

<sup>13</sup>An appropriate lag structure is necessary to adequately capture the dynamics contained in the data. Since most studies in this literature opt for a lag order of four quarters, we also choose this for reasons of comparability. Moreover, Hamilton and Herrera (2004) also argue that too short a lag length might omit the primary effects of oil shocks. However, our findings are qualitatively similar when less lags are included.

<sup>14</sup>Canova and Gambetti (2009) have also experimented with more general autoregressive processes for the law of motion of the coefficients but report that the random walk specification was always preferred. Moreover, as has been pointed out by Primiceri (2005), the random walk assumption has the desirable property of focusing on permanent parameter shifts i.e. it incorporates the belief that changes are permanent rather than transitory which leads us to prefer this specification over a stationary process. From a computational point of view, the random walk specification avoids introducing additional free parameters that need to be estimated.

<sup>15</sup>Initially, we have included an indicator function which selected only stable draws i.e. the indicator function  $I(\theta_t) = 0$  if the roots of the associated VAR polynomial are inside the unit circle as e.g. in Cogley and Sargent (2005). However, our acceptance ratio was so high as to make this constraint obsolete.

$h_t = [h_{1,t}, h_{2,t}, h_{3,t}, h_{4,t}]'$  are assumed to evolve as geometric random walks independent of each other.<sup>16</sup> The error terms of the three transition equations are independent of each other and of the innovations of the observation equation. In addition, we impose a block-diagonal structure for  $S$  of the following form:

$$S \equiv Var(\zeta_t) = \begin{bmatrix} S_1 & 0_{1 \times 2} & 0_{1 \times 3} \\ 0_{2 \times 1} & S_2 & 0_{2 \times 3} \\ 0_{3 \times 1} & 0_{3 \times 2} & S_3 \end{bmatrix} \quad (6)$$

which implies independence also across the blocks of  $S$  with  $S_1 \equiv Var(\zeta_{21,t})$ ,  $S_2 \equiv Var([\zeta_{31,t}, \zeta_{32,t}]')$ , and  $S_3 \equiv Var([\zeta_{41,t}, \zeta_{42,t}, \zeta_{43,t}]')$  so that the covariance states can be estimated equation by equation.

We estimate the above model using Bayesian methods. An overview of the prior specifications and the estimation strategy (Markov Chain Monte Carlo algorithm) is provided in Appendix B.

## 2.2 Identification of oil supply shocks

Separating the endogenous and exogenous sources of movements in the price of crude oil has proven to be a difficult task. In most VAR-based studies (e.g. Burbidge and Harrison 1984; Jiménez-Rodríguez and Sánchez 2005; BG 2007) oil supply shocks are identified as innovations in the oil price variable.<sup>17</sup> Treating the unpredictable variations in the price of oil as exogenous with respect to the developments in the US economy (and more generally, global macroeconomic forces) finds its origin in the belief that at least the major oil price shocks were due to production shortfalls caused by political events in the Middle East. Implicit in this view is the assumption that oil price changes derive exclusively from the supply side of the oil market. However, it is now widely accepted that oil prices, especially in more recent decades, are also driven by demand conditions (see BK 2002, 2004; Kilian 2009; Rotemberg 2007). Consequently, innovations in the oil price equation of a VAR are not necessarily an adequate measure of exogenous variations in oil supply since they may also capture shifts in the demand for crude oil. In that sense, the resulting estimates only represent the economic effects of an average oil price shock determined by a combination of supply as well as demand factors which can bias the estimates. This concern is even more relevant in a time-varying context, in particular when the relative importance of oil supply versus oil demand shocks has

---

<sup>16</sup>Stochastic volatility models are typically used to infer values for unobservable conditional volatilities. The main advantage of modelling the heteroskedastic structure of the innovation variances by a stochastic volatility model as opposed to the more common GARCH specification lies in its parsimony and the independence of conditional variance and conditional mean. Put differently, changes in the dependent variable are driven by two different random variables since the conditional mean and the conditional variance evolve separately. Implicit in the random walk assumption is the view that the volatilities evolve smoothly.

<sup>17</sup>A lot of studies employ nonlinear transformations of the oil price, such as the net oil price index (Hamilton 1996, 2003), as a measure of exogenous supply shocks which already takes asymmetries into account. However, Kilian and Vigfusson (2009) show that estimates obtained with such a filtered variable are inconsistent and that implied impulse responses are incorrectly computed.

changed over time. Others (e.g. BG 2007) have made the case that this distinction does not matter since an oil price shock triggered by increased demand for oil in one country can still be experienced as a supply shock by the remaining countries. This presumption is, however, very stringent in the light of the results of Kilian (2009) and Peersman and Van Robays (2009) who show that there exist important differences in the responses of macroeconomic aggregates depending on the underlying source of the oil shock. Intuitively it is clear that an increase in the price of oil induced by favorable global economic conditions exerts a different influence on the macroeconomic performance than one due to oil supply interruptions resulting from a war (see Kilian 2009; Rotemberg 2007). Moreover, considering only oil prices also implicitly assumes that the slope of the oil demand curve remains constant over time. However, if the elasticity of the demand curve is not invariant over the sample period, a similar physical disruption in the supply of crude oil due to e.g. a military conflict, will have a very different impact on the oil price itself which complicates intertemporal comparisons. In section 3, we will show that the slope of the demand curve has indeed changed over time.

A different strand of the literature has developed measures based on physical oil supply disruptions associated with major political crises in oil-producing countries in order to extract the magnitude of exogenous oil supply shocks. Hamilton (2003) uses a quantitative version of the dummy-variable approach (Dotsey and Reid 1992) to isolate the exogenous component of oil price movements by measuring the oil supply curtailed by exogenous events which are largely political in origin.<sup>18</sup> He compares the observed production level before a military conflict started with the largest drop of oil supply in the affected countries in subsequent periods. The magnitude of the production shortfall is then defined as the exogenous oil supply shock expressed as a percentage of world oil production. Kilian (2008a) constructs an alternative measure of exogenous oil supply shocks by comparing the actual production shortfalls in the wake of a political crisis to an explicit counterfactual path of how production would have evolved in the absence of the crisis. The counterfactual production level is derived from countries which are subject to the same economic incentives but not involved in the conflict. While these methods avoid potential problems regarding the endogeneity of the oil price series, a shortcoming is that they are closely tied to a selection of relevant historical episodes so that only a sub-class of exogenous oil supply shocks related to political events is identified. Kilian (2009) allows shocks to the demand for crude oil to have a contemporaneous impact on oil prices in a monthly vector autoregression. To identify an oil supply shock, he assumes that the latter is the sole disturbance that has an immediate influence on the level of oil production. Accordingly, shifts in the demand for oil only have a delayed effect on crude oil production, i.e. a vertical short-run supply curve is assumed. This assumption, while arguably tenable in a monthly context, is however, less appropriate when quarterly data are used like in our study. In addition, oil production could already react to expected shifts in economic activity which are not fully captured in such a setting.

Elaborating on work by Faust (1998), Canova and De Nicoló (2002), Uhlig (2005), Peersman

---

<sup>18</sup>The events Hamilton (2003) considers are the Suez crisis (1956), the Arab-Israel war (1973), the Iranian revolution (1978), the Iran-Iraq war (1980-1988) and the Persian Gulf war (1990/91).

(2005), Benati and Muntaz (2007) and Peersman and Straub (2009), we propose to use sign restrictions on the estimated time-varying impulse responses to identify structural oil supply shocks. More specifically, the restrictions are derived from a simple textbook supply and demand model for the global oil market. In the spirit of Kilian (2009), we present the global oil market by world oil production and the world crude oil price. An oil supply shock is identified as any shift in the oil supply curve and hence, results in an opposite movement of oil production and the real price of crude oil.<sup>19</sup> In particular, the identifying assumptions are that after an unfavorable oil supply shock world oil production does not increase and the real price of crude oil does not decrease.<sup>20</sup> These restrictions are sufficient to uniquely identify global oil supply disturbances without having to impose zero restrictions on oil prices or production to distinguish them from other shocks. Furthermore, no condition at all constrains the responses of real output and consumer price inflation. The reactions of these variables will eventually be determined by the data. Given our interest in the time-varying effects of oil market shocks stemming from the supply side, we identify the oil supply shock only, while all other oil market shocks, i.e. shocks with an impact on oil prices and production of the same sign, are lumped together as the oil demand shock. The latter could reflect either oil-specific demand shocks or shifts in the oil demand curve resulting from changes in economic activity.<sup>21</sup> We impose the sign conditions to be binding for four quarters following the shock. Consequently, our restrictions accommodate a potential sluggishness in the adjustment of the oil market. Some robustness checks with regard to alternative restrictions and the horizon over which the sign constraints are imposed are conducted in section 3.3. The details for the computation of the time-varying impulse responses and the implementation of the sign restrictions are described in Appendix C.

## 3 Results

### 3.1 Impulse responses

Figure 1 displays the median impulse responses of world oil production, the real price of oil, US real GDP and US CPI to a one standard deviation oil supply shock for horizons up to 20 quarters at each point in time spanning the period 1970Q1 to 2006Q2.<sup>22</sup> The estimated responses have been

<sup>19</sup>For illustrative purposes of our identification assumptions, we refer the reader to Figure 3, panel A.

<sup>20</sup>These restrictions are imposed as weak inequality constraints  $\leq$  and  $\geq$ , thus including also the possibility of no change, i.e. a zero response. As a consequence of these inequality constraints, our identification scheme does not deliver exact identification. See Fry and Pagan (2007) for a discussion of this type of restrictions and potential problems.

<sup>21</sup>In Baumeister and Peersman (2009), we also estimate the impact of different oil demand shocks, in particular oil-market specific demand shocks and global shocks to economic activity.

<sup>22</sup>The 3D-graphs of the time-varying impulse responses are to be read in the following way: along the  $x$ -axis the starting quarters are aligned from 1970Q1 to 2006Q2, on the  $y$ -axis the quarters after the shock are displayed, and on the  $z$ -axis the value of the response is shown in percent. All responses have been normalized in such a way that the structural innovations raise the price of oil.

accumulated and are shown in levels. In general, a negative oil supply shock results in a permanent fall of oil production and a permanent rise of the real refiner acquisition cost of imported crude oil. Consistent with expectations, the shock is followed by a significant slowdown in real economic activity and increases in consumer prices, variables which were not constrained. This evidence emerges even more clearly in Figure 2, panel A, where the time-varying median responses of the four variables are plotted four quarters after the shock,<sup>23</sup> together with the 16<sup>th</sup> and 84<sup>th</sup> percentiles of the posterior distribution. Specifically, after a representative oil supply shock, we observe a relatively similar impact on output over time, which is statistically more significant in the second half of the sample, as well as slightly stronger inflationary effects.<sup>24</sup> We also performed some bilateral tests for time variation to assess the statistical significance of differences over time. For this purpose, we consider the joint posterior distribution of the cumulated responses at selected horizons where shifts of this distribution away from the 45-degree line indicate a systematic change across time.<sup>25</sup> For several benchmark quarters, we find a statistically significant, stronger effect of a one standard deviation oil supply shock on consumer prices in more recent periods and for some quarters even a larger long-run impact on output. This evidence is striking, given the results presented in Edelstein and Kilian (2007), Herrera and Pesavento (2009) and BG (2007), who find a reduced impact on real output and consumer prices over time.

We do find, however, considerable time variation in the dynamics of the oil market itself which is symptomatic of the fact that the global oil market has undergone fundamental structural changes. By all appearances the interaction between oil production and oil prices varies remarkably across time periods with an obvious declining trend in the response of oil quantity and a stronger impact on the oil price level since the mid-eighties. Also for the episodes 1973/74 and 1979/80, we observe an increased reaction of crude oil prices. Not surprisingly, a statistically significant, smaller impact on oil production and greater effect on real oil prices over time is strongly confirmed by the bilateral tests.

A natural question which emerges is why we find such a change in the oil market over time. We observe that the responses in the oil market set in immediately. Since the magnitude of the shocks are inextricably intertwined with the contemporaneous response of the variable in question, this feature makes it difficult to distinguish between dynamics and volatility.<sup>26</sup> However, a change in the underlying oil supply volatility alone cannot explain this stylized fact because the impact on

---

<sup>23</sup>This choice can be rationalized by the fact that the greatest effect on real GDP is expected to occur with a delay of 3 to 4 quarters after the shock (Hamilton 2008). However, given the persistency of the responses, the message is not altered if different horizons are selected.

<sup>24</sup>For real GDP, we also find somewhat stronger effects over time for longer horizons after the shock, as can be seen in Figure 1.

<sup>25</sup>See also Cogley et al. (2009) for similar tests in a different context. The results of this exercise are available upon request.

<sup>26</sup>This is a standard problem when VAR results are compared across different samples or estimated with time-varying parameters. Only the contemporaneous impact of a shock on a number of variables can be measured. Consequently, it is not possible to know exactly whether the shock itself (volatility) has changed or the immediate reaction (economic structure) to this shock.

oil prices and production does not change in the same direction over time (as illustrated in panel A of Figure 3). The only possible explanation for observing a smaller reaction of oil production in combination with a greater reaction of oil prices is a steepening of the oil demand curve, i.e. oil demand became less elastic over time (as shown in panel B).<sup>27</sup> Specifically, in order to push up oil prices as much as in the 2000s, a large reduction in world oil production was necessary before the mid-eighties because oil prices were much less sensitive to changes in oil supply then as they appear to be nowadays, with the exception of two episodes, namely 1973/74 and 1979/80.

The fact that the impact of a typical oil supply shock on oil quantity and prices changes dramatically every period complicates, however, the analysis of the time-varying dynamic effects of an exogenous event in the oil market on the macroeconomy. The way the normalization is done for the experiment becomes very important. Since the focus of previous research has centred on an unanticipated increase in the price of oil, consider for instance, the effect of an oil supply shock which raises the real price of oil by 10 percent on impact. The latter is used by BG (2007) as a benchmark for an intertemporal comparison. The normalized time-varying impulse responses after four quarters are shown in panel B of Figure 2. We find a more muted reaction of macroeconomic variables in the latter part of the sample, a finding which complies very well with existing empirical evidence on the time-varying effects of oil price shocks (e.g. BG 2007; De Gregorio et al. 2007). This experiment, however, implicitly assumes a constant slope of the oil demand curve over time which is clearly not the case. As a consequence, the results of such a normalization cannot be compared because a different underlying supply shock is considered. Specifically, a 10 percent rise in oil prices is currently generated by an oil production shortfall of 1 to 2 percent. To elicit the same oil price move in the 1970s, a decline in the physical supply of crude oil of up to 15 percent is required, which actually never happened within the sample period. Despite the assertion by BG (2007) that "what matters [...] to any given country is not the level of global oil production, but the price at which firms and households can purchase oil" (p.17), it appears that the volume of oil can be considered as an important input factor in the production process.<sup>28</sup>

Alternatively, an oil supply shock can also be normalized on the quantity variable rather than on the price variable. In fact, oil supply shocks have frequently been viewed as physical interruptions in the production of crude oil due to deliberate decisions by OPEC aimed at achieving a certain price level or destruction of oil facilities in the wake of war activities.<sup>29</sup> The dynamic macroeconomic effects of exogenous oil supply shocks measured as a 1 percent decrease in global oil production on impact are shown in panel C of Figure 2. The implied elasticity of the real price of crude oil with respect to a 1 percent shortfall in world oil production increases substantially over time, from an average value of 5 percent in the 1970s and 1980s to 10 percent in the 1990s up to 15 percent

---

<sup>27</sup>This observation does not exclude the possibility that the variability of oil supply shocks has also changed over time, but this alone could never explain our findings.

<sup>28</sup>Considering only the oil price might be realistic for a small country but is more problematic for the United States.

<sup>29</sup>Since OPEC only controls the quantity supplied, the world oil price is only indirectly influenced. Even if at times OPEC announced a price target or effectively "set" the reference price, it still had to regulate its production volume in order to obtain and maintain a certain price level.

in the 2000s. These dramatic oil price increases triggered by a similar reduction in oil production in the second part of the sample are in turn more disruptive to the economy and emphasize the importance of the proper definition of the oil shock concept. The accumulated loss in real GDP growth is about twice as big in the 1990s and almost three times as big in the 2000s as in the 1970s. The response of consumer prices gets more pronounced from the 1990s onwards and continues to increase considerably in the 2000s.<sup>30</sup> However, also this intertemporal comparison cannot really be made since a typical (one standard deviation) shift of the crude oil supply curve is characterized by a change in world oil production of around 2 percent in the seventies, whereas this change amounts to only 0.5 percent from the mid-nineties onwards. Given the indistinguishability of volatility and immediate impact, our analysis cannot determine whether these smaller average changes in oil production are the result of a steeper oil demand curve or also the consequence of smaller shifts in the underlying oil supply curve over time.<sup>31,32</sup>

### 3.2 Relevance of oil supply shocks

It is still a widely held belief that the oil price hikes of the 1970s and early 1980s were the underlying source of macroeconomic stagflation during that period (e.g. Hamilton 1983). Since the average effects of oil supply shocks on US output and inflation have not dramatically changed compared to the 1970s, it is surprising that the current macroeconomic conditions are so different. To shed some light on this issue, we now have a closer look at the variance and historical decompositions. Figure 4 displays the median time-varying contribution of oil supply shocks to the forecast error variance after 20 quarters of respectively world oil production growth, changes in the real price of crude oil, US real GDP growth and CPI inflation, along with the 16<sup>th</sup> and 84<sup>th</sup> percentiles of the posterior distribution. The contribution of oil supply shocks to the variance of CPI inflation and real GDP growth in the US consistently ranges between 15 and 20 percent. The share of output volatility attributable to oil supply shocks oscillates moderately over time, whereas the fraction of movements in consumer price inflation induced by unexpected oil supply disturbances exhibits a slight increase in more recent periods. The latter is not surprising given that the general volatility of CPI inflation decreased over time, while the impact of oil supply shocks on inflation did not decline. We can thus conclude that exogenous oil supply shocks are economically still relevant.

The decomposition also indicates that oil supply shocks account for approximately 30 percent of the forecast error variance in world oil production which only experiences moderate variations over time. However, the figure reveals that the contribution of oil supply shocks to the variability in

---

<sup>30</sup>This should not be too surprising since oil prices are part of the consumer price index and hence, larger oil price increases automatically lead to a higher CPI even if second-round effects are absent. However, this is probably not the main reason for our findings since, when we employ the implicit GDP deflator as the measure of inflation, we still find a substantial increase in the response of the price level over time. This result is also reported in section 3.3.

<sup>31</sup>In subsequent work (Baumeister and Peersman 2009), we show that the underlying shifts of the oil supply curve have indeed declined over time.

<sup>32</sup>Note that, in case of a vertical oil supply curve, the observed decrease in oil production changes would be fully driven by decreased oil supply volatility.

the real price of crude oil declines over time from around 30 percent in the first part of the sample to around 20 percent in later periods suggesting that they are currently a less important source of oil price movements. This evidence is consistent with Kilian (2008a, 2009) and the conventional wisdom that "demand increases rather than supply reductions have been the primary factor driving oil prices over the last several years" (Hamilton 2008, p.175). Moreover, given the almost constant proportion of non-supply shocks in world oil production and the increased contribution of these shocks to oil price variability, our results indicate that also the oil supply curve must have become more inelastic over time. A steeper oil supply curve is thus also a source of increased oil price volatility.

To evaluate whether exogenous shifts of the oil supply curve are responsible for specific episodes, the historical contribution of oil supply disturbances to the four endogenous variables are presented in Figure 5. Specifically, these figures show the baseline forecast of the variables augmented by the cumulative contribution of oil supply shocks (red line) as well as the actual time series in growth rates (blue line). The difference between both is driven by other shocks. For real GDP growth, grey bars are added to indicate periods of recessions as dated by the NBER. With regard to inflation, the graphs reveal that unfavorable oil supply shocks explain little of the Great Inflation. Despite the fact that there is a positive contribution, the bulk of excessive inflation in the 1970s is explained by other shocks. While this insight is apparently in contrast with popular perception, it can still be reconciled with several recent findings in the literature. In particular, BK (2002) suggest that other adverse shocks must have been at the origin of the stagflationary experience of the 1970s and Clarida et al. (2000) argue that bad monetary policy was the source of the Great Inflation.

The supply-shock driven contribution to the historical evolution of the real price of crude oil itself and economic activity changes from episode to episode. During the 1973/74 oil embargo a major part of the oil price increase is attributable to oil supply shocks but they cannot account completely for the price spike. The contribution of oil supply shocks to the development of the price of oil during the events of 1978-80 is more limited and clearly indicates that this was more a demand-driven price shock mainly determined by rising oil demand at a time of low spare capacity.<sup>33</sup> The latter finding confirms BK (2002) and Kilian (2009) who argue that pure oil supply shocks were never the sole driving force behind observed fluctuations in the real price of crude oil. On the other hand, a substantial share of the oil price hike after Iraq's invasion of Kuwait in August 1990 can be ascribed to unfavorable oil supply shocks. As a consequence, interruptions in the supply of crude oil did play a significant role in the 1974/75, early 1980s and 1990s recessions. However, the contribution of oil supply disturbances to the poor macroeconomic performance was certainly not exclusive since a major part of the slowdowns is still explained by other shocks.

But also in more recent times unfavorable oil supply shocks had significant effects on real economic activity. Consider the 1999 oil supply disturbance, engineered by the joint decision of OPEC and non-OPEC countries to cut oil production, which made an important contribution to

---

<sup>33</sup>Skeet (1980, p.164) reports that "only Saudi Arabia exhibited any strong wish to restrain the take-off of price, but it no longer had the production flexibility to manage this on its own" during the events of 1979.

the oil price rise. This unfavorable shock had a negative impact on real GDP growth and made the ongoing strong economic boom before the millennium turnover more subdued whereas its role in the economic downturn of 2001 was negligible. In addition, most oil price surges since 2002 were driven by shocks affecting the demand side of the oil market. The latter finding is consistent with Kilian (2009) and helps explain why these shocks were not accompanied by a major recession in the US. These evolutions could eventually have contributed to the belief that the way the economy experiences oil shocks has fundamentally changed over time.

### 3.3 Robustness analysis

**Model properties and identification.** Since our specification departs from standard models along two dimensions, namely identifying oil supply shocks by means of sign restrictions and explicitly modeling time variation, it is not entirely clear whether one of these two factors accounts for our results. We therefore check the robustness of our conclusions by looking at each aspect separately. First, we consider the changing effects of oil supply shocks identified with sign restrictions in a constant-coefficient VAR by splitting the sample into two subperiods. Given our findings with time-varying parameters, the preferred date for the sample split is 1986, even though we still observe time variation within these two subsamples, e.g. the 1973/74 and 1979/80 episodes.<sup>34</sup> Within this simpler framework, we can also assess the robustness of our results with regard to the frequency of our data. Estimating the benchmark time-varying VAR model with monthly data is not an option since this would result in a proliferation of free parameters given the number of lags we would have to include to account for dynamics. Second, following Kilian (2009), we identify oil supply shocks using zero short-run restrictions but implemented in our time-varying framework.

Figure 6, panel A, displays the median impulse responses of the four endogenous variables after an oil supply shock identified with our sign restrictions in a fixed-coefficient VAR estimated over the two subsamples 1970Q1-1985Q4 and 1986Q1-2006Q2 together with the 16<sup>th</sup> and 84<sup>th</sup> percentiles. It turns out that the response of oil production to a one standard deviation oil supply shock is indeed more muted in the second subsample, while the reaction of real oil prices is more pronounced, which confirms our finding of a less elastic oil demand curve over time. Also in line with our previous findings, there is hardly any time variation in the responses of real GDP and consumer prices across subsamples. Very similar results (not shown) are obtained for the sample split with monthly data where real GDP is replaced by US industrial production. As a consequence, our results appear robust with respect to the specific modeling of time variation as well as the choice of data frequency.

Figure 6, panel B, presents the time-varying median responses of the two oil market variables and macroeconomic aggregates to an oil supply shock identified by a recursive ordering of the variables four quarters after the shock together with 16<sup>th</sup> and 84<sup>th</sup> percentiles. More specifically, in this setting only oil supply shocks have a contemporaneous impact on all four endogenous variables,

---

<sup>34</sup>However, our results are not sensitive to the specific breakpoint.

while shocks in oil demand do not have an immediate effect on oil production. This identification strategy is consistent with Kilian (2009). Note, however, that these restrictions are more plausible in a monthly VAR and less appropriate for our quarterly specification. We nevertheless conduct the analysis as a robustness check but the results should be interpreted with more than the usual degree of caution. Also this alternative identification strategy demonstrates clearly time variation in the oil market. In particular, the impact of an oil supply shock on global oil production becomes much smaller over time and the response of oil prices stronger, implying again a steepening of the demand curve for crude oil. As before, we do not observe a reduced reaction of macroeconomic variables over time. Surprisingly, the responses of the real price of crude oil after an unfavorable oil supply shock are slightly negative in the first half of our sample, which also explains the somewhat perverse effects on real GDP and consumer prices during this period. This might be explained by the fact that an oil supply shock identified in this way is contaminated by demand factors when using quarterly data, which could underestimate the true oil demand elasticity. This observation also conforms to our concerns about applying this identification strategy in a quarterly framework in which the assumption of a short-run vertical supply curve appears to be problematic. The substantial rise in the oil price responsiveness to a typical supply shock since 1986 and the gradual reduction of its impact on oil production over time can, however, not be ignored and turns out to be a stylized fact. In sum, we can conclude that our results are not an artifact of our new identification strategy.

**Alternative specifications.** The robustness of our findings has also been analyzed by altering several features of the benchmark model. As already mentioned, the results do not depend on the number of lags included in the TVP-BVAR.<sup>35</sup> Also when assessing the sensitivity of our results with regard to the choice of the priors<sup>36</sup> and the number of periods for which the sign restrictions are imposed, the message this paper conveys is not modified. Equally, using different oil price measures such as the WTI spot oil price and the real refiner acquisition cost of *composite* crude oil does not change any of our conclusions described in the paper. Figure 7, panel A (left), shows the time-varying impulse response functions of US unemployment when this variable is included as an alternative indicator of economic activity. The responses are comparable to those of real GDP in that we find a significant rise following an unfavorable oil supply shock. Panel A (right) presents the outcome when the implicit GDP deflator is used instead of CPI inflation.<sup>37</sup> Although the responses are somewhat more subdued, the main message of the results are again not altered.

A hypothesis frequently put forward in the literature on the Great Moderation to account for the greater resilience of the economy in the face of adverse shocks is improved monetary policy (e.g. Clarida et al. 2000). Including the federal funds rate in our model to take account of the endogenous response of monetary policy to oil shocks (Bernanke et al. 1997; Hamilton and Herrera

---

<sup>35</sup>Not all the robustness checks reported in the paper are shown, but are available upon request.

<sup>36</sup>We have experimented with different prior specifications which are commonly used in the TVP-BVAR literature. We report on them in more detail in Appendix B.

<sup>37</sup>In this case the nominal price of crude oil is deflated using the GDP deflator.

2004; Herrera and Pesavento 2009) does also not change our findings about the dynamic response of the US economy to oil supply shocks and the structural change in the oil market as is evident from Figure 7, panel B.<sup>38</sup>

**Implementation of restrictions.** Since at the beginning of our sample the rise in nominal oil prices was constrained by institutional features of the oil market,<sup>39</sup> in particular long-term price agreements which were subject to revision only periodically, an obvious concern arises with regard to the timing of the restrictions imposed in the baseline case ( $t = 0$  to  $t = 4$ ) to identify oil supply shocks. When nominal prices in oil contracts are fully fixed, a positive aggregate demand shock in the real economy that raises world oil production and the general consumer price level, could then lead to a fall in the real price of crude oil unless nominal contracts are renegotiated timely to reflect the new macroeconomic conditions. The resulting opposite movement in world oil production and *real* oil prices would then imply that this shock is erroneously identified as an oil supply shock. We address this potential problem in two ways to assess the sensitivity of our results to the identification restrictions. First, we re-estimate the model with the *nominal* refiner acquisition cost of imported crude oil. The sign restrictions to identify an oil supply shock are the same as in our baseline identification scheme. Specifically, global oil production does not rise and the nominal oil price does not fall after an upward shift of the oil supply curve during the first four quarters after the shock. Note that the restrictions are imposed as  $\geq$  or  $\leq$ , so that a zero reaction of the nominal oil price is still possible and the real oil price could even temporarily fall.<sup>40</sup> Given that the bulk of real oil price movements is driven by changes in the nominal oil price rather than general inflation, it is not surprising that the results are not affected, not even at the beginning of the sample. Second, we also re-estimate the model with the real oil price, but now the sign restrictions are only binding from the fourth quarter after the shock onwards. Accordingly, the immediate reaction of oil production and the real price of crude oil are not required to conform to the expected sign. As emerges from Figure 7, panel C, where the time-varying median impulse responses of the variables representing the global oil market and the US macroeconomy are displayed, our results are again robust with respect to the time period of the sign restrictions.

In general, our results are very robust: there are no discernible differences in the evolutionary pattern of the structural features of the global oil market and the macroeconomic consequences over time compared to our benchmark model when alternative specifications are estimated.

---

<sup>38</sup>Adding an additional variable comes at the cost of having to cut down on the number of lags included in the VAR (here  $p = 2$ ); but as mentioned earlier, a shorter lag length does also not affect the main results in our baseline specification.

<sup>39</sup>Hamilton (1983, p.232) notes that "oil prices have obviously been determined under a radically different institutional regime since 1973 than before". To a lesser extent, this could also be the case for the pre-1986 period.

<sup>40</sup>It seems reasonable to assume that nominal oil contracts are revised the latest one year after the shock, especially since OPEC producers became more and more reluctant to increase oil production at artificially low prices during the early 1970s (BK 2002).

## 4 Why steepening of the oil demand curve?

In this section, we consider three important developments in the oil market which could be relevant for the substantial reduction of crude oil demand elasticity; namely, the increased flexibility of the crude oil market, a changing role of oil in the economy and oil production capacity utilization. This list is by no means exhaustive and hence, does not exclude the contribution of other factors to the steepening of the oil demand curve which could be explored in future research.

**Increased flexibility of the world oil market.** Since we observe an increased responsiveness of oil prices to a change in oil production since the mid-eighties, a natural candidate for the break is the transition from a regime of administered prices to a market-based system of direct trading in the spot market around the same time, which implied a shift of price determination away from OPEC to the financial markets exposing oil prices to greater fluctuations (Hubbard 1986; Mabro 2005).<sup>41</sup> This development had two consequences. First, oil price volatility increased considerably. Second, the increased volatility fostered the development of oil futures markets which have deepened considerably since the 1990s.

This increased flexibility of crude oil prices can, however, not explain our findings. More flexibility and increased relevance of the spot market could only affect the speed of adjustment to an oil supply shock because, in the long run, also long-term contracts should reflect the "correct" fundamental price. Figure 2 shows the impact after one year and Figure 1 presents the effect for even longer horizons. All evidence points towards a considerably stronger long-run impact on crude oil prices in the second half of the sample period. Moreover, Figure 1 indicates that not even the speed of adjustment has changed a lot over time. For the whole sample, the effect on oil prices is almost complete after 1 quarter.<sup>42</sup> In addition, increased volatility of the shocks as a result of more flexible markets should also be reflected in a stronger impact on world oil production, which is hard to reconcile with our results. Consequently, more flexibility of the crude oil market cannot explain the steepening of the oil demand curve. On the contrary, a less elastic demand curve automatically leads to greater price fluctuations after a supply shock, which could thus be a source of increased volatility. The latter could actually have fostered the development of the spot market and the abolishment of administered prices.

**Role and share of oil in the economy.** It is conceivable that there have been structural transformations in industrialized as well as developing economies which might explain why oil demand is more inelastic since the mid-eighties.

First, in response to the oil price hikes of the 1970s, the role and share of oil in the US and other industrialized economies have changed substantially. In fact, industries switched away from

---

<sup>41</sup>OPEC basically abandoned fixing the reference price in favor of a system with production quotas.

<sup>42</sup>This is a typical finding in the oil literature. A near random walk behavior of the real price of oil has also been documented by BG (2007) and Hamilton (2009).

oil to alternative sources of energy, developed more energy-efficient technologies and improved energy conservation. These efforts have been supported by the governments who reacted to the oil crises by enacting codes to reduce oil usage and increase energy awareness. The resulting gradual substitution process as well as service-biased growth (smaller share of industrial production in value added) led to a steadily falling oil intensity of economic activity, i.e. reduction in the use of oil input per unit of output as illustrated in Figure 8, panel A. Efficiency gains in the usage of oil in the production process have often been put forward as one possible explanation for the milder effects of oil shocks on the economy (e.g. BG 2007; De Gregorio et al. 2007) but also have important implications for the demand behavior and hence, the price elasticity of oil demand since they induced fundamental changes in oil consumption patterns.<sup>43</sup>

As a result of these developments, nowadays there are not as many options left for increasing energy efficiency further because most possible technical upgrades and replacements of oil-dependent capital by capital that uses alternative sources of energy<sup>44</sup> are already in place so that there is only a reduced scope for additional substitution away from oil (Dargay and Gately 1994; Ryan and Plourde 2002). More importantly, the composition of total oil demand has altered with oil consumption now being concentrated in sectors (for instance transportation) where the lack of substitutes for petroleum at all times implied a low own-price elasticity of demand. The increasing share of these sectors in total oil demand might thus have contributed to a steepening of the oil demand curve.

Second, while increased efficiency in oil use plays an important role for the declining importance of oil in industrialized economies, it is but one component of a broader concept which also takes the evolution of the real price of oil into account, namely the cost share of crude oil in US total expenditures. Following Hamilton (2009), we calculate the value share of crude oil as the ratio of the dollar value of oil expenses to nominal GDP.<sup>45</sup> Figure 8, panel B, displays the evolution of the share of oil purchases in total, economy-wide expenditures in the US economy over time. As is evident from the graph, the share of total production costs spent on oil has decreased considerably over time. One of Marshall's four rules of derived demand for factor inputs claims that there exists a direct link between the cost share of input factors in total production costs and the price elasticity of the derived demand for this factor (Marshall 1920). More specifically, the rule suggests that a smaller share of factor costs leads to a less elastic demand for that production factor if the demand elasticity for the final product is greater than the substitution elasticity between input

---

<sup>43</sup>This became apparent when falling oil prices during the 1980s did not stimulate oil demand as much as in previous periods. Dargay and Gately (1994) attribute this phenomenon to the irreversibility of technical knowledge, the durability of efficiency-improving investments and the non-abrogation of laws regarding energy-cost labeling and energy-efficiency standards. This non-reversal of efficiency improvements and energy-saving innovations could thus partly explain the decrease in price responsiveness of oil demand.

<sup>44</sup>It has to be born in mind that even if technical substitution would still be feasible, other energy sources e.g. natural gas also reach their capacity limit which restricts substitutability.

<sup>45</sup>Alternative ways to compute the cost share of crude oil have been proposed in the literature, e.g. Rotemberg and Woodford (1996), BG (2007), Edelstein and Kilian (2007). The overall pattern of our oil expenditure share complies well with these approaches.

factors (Peirson 1988).<sup>46</sup> Hence, the declining share of oil costs in total expenses over time could provide an additional clue for the decrease in oil demand elasticity. It also appears intuitively clear that a smaller share of oil in total production costs makes firms less reactive to price fluctuations, especially in an environment of increased oil price volatility where cost increases are likely to be reversed quickly. However, it is rather unlikely that the cost share remains low in light of the recent oil price developments (Hamilton 2009), i.e. oil expenditures might as well regain importance in firms' budgets which could again increase the demand elasticity.

Third, developing economies that are in the process of industrialization currently have a much higher share in global oil demand. These countries typically rely heavily on oil as an input factor and are therefore considered as being less reactive to changes in global oil prices. Specifically, as the governments of these nations are interested in fostering rapid economic growth, state-controlled oil product prices and fuel subsidies which shield consumers from the impact of rising global oil prices are prevalent features of these economies (Hang and Tu 2007).<sup>47</sup> In other words, oil demand in many developing countries is not affected by oil price hikes because price ceilings are imposed on petroleum products that keep prices significantly below market prices and hence, make consumer demand unresponsive to international price signals. This distortive pricing system could thus have important repercussions on oil demand in global markets. Given that the share of crude oil demand from the developing world is on the rise, it is possible that the responsiveness of demand will decrease further. On the other hand, the subsidies can not last forever and in some places a process of gradual dismantling of price caps has already been initiated.

**Capacity constraints in crude oil production.** The above explanations could well describe the developments since the mid-1980s but we also observe a significant fall in the elasticity of oil demand for the 1973/74 and 1979/80 episodes. To account for this observation a different reasoning is required which allows conditions on the supply side of the oil market to influence demand behavior. At times of low spare capacity in world oil production, it is possible that small supply disruptions can lead to large price increases because market participants anticipate that a loss in oil output, resulting from war activities or other incidents, cannot be replaced by other oil producers due to their operating already close to sustainable capacity. Although capacity constraints are normally a phenomenon related to the supply side, increasing capacity utilization rates also signal some tightness in the market which affects the demand behavior of consumers by raising concerns about the security of future oil supplies that motivate precautionary buying in a tightening market.<sup>48</sup> Thus, this heightened uncertainty increases the willingness of agents to pay a higher price for a barrel of oil at the margin that provides insurance against potential scarcity,

---

<sup>46</sup>For a dissenting view, see Pemberton (1989).

<sup>47</sup>"Both wholesale and retail prices of oil products in the domestic market are lower than they are in the global market" as exemplified for China by Hang and Tu (2007, p.2978). In fact, an estimate by Morgan Stanley shows that almost a quarter of the world's petrol is sold at less than the market price (The Economist, 2008).

<sup>48</sup>As has been noted by Gately (1984, p.1103), "aggravating the market tightness was an extended period of aggressive stockbuilding by the importing countries for much of 1979 and 1980. Such a stockbuilding "scramble" during a disruption was certainly perverse. It undoubtedly drove price higher than it would have gone otherwise."

i.e. they pay a risk premium.<sup>49</sup> As a result, this less-elastic precautionary fraction of total oil demand could become larger, especially if no substitutes are available in the short run.<sup>50</sup> Put differently, a certain amount of spare capacity provides a cushion which assures the market that exogenous shortfalls in production, caused intentionally or accidentally, can be compensated for; but when capacity utilization is already high before an exogenous event, this guarantee vanishes and induces market participants to alter their expectations (in anticipation of potential production disruptions) resulting in a more rigid oil demand curve. If oil supply is then indeed curtailed in such an environment, even small shocks in terms of losses in world oil supply can lead to considerable price increases because of a steeper demand curve.<sup>51</sup> Figure 9 shows global capacity utilization rates in crude oil production by year derived from IMF estimates of spare oil production capacity. A rate near 90 percent can be interpreted as an important threshold in view of the sustainability of future oil production at such high levels (Kilian 2008a). Remarkably, we find world oil production to be very close to full capacity since the second half of the eighties. On the other hand, we also note two spikes in capacity utilization rates around the time of the two oil crises of the 1970s which could provide an explanation for the estimated increased price responsiveness to oil production shortfalls around the same time.<sup>52</sup> Consequently, oil production levels which are close to full capacity might be a reason for a steeper oil demand curve during the past two decades. This evolution was further facilitated by the increased flexibility of the oil market. In such a scenario, variations in oil supply quickly translate into price changes, especially in the face of shrinking global spare capacity which leads to a process of bidding up the prices because buyers compete at the margin for limited volumes of crude oil available on the spot market.

## 5 Conclusions

In this paper, we have analyzed the time-varying effects of oil supply shocks on the US economy and the oil market from 1970 onwards. On the one hand, there are a priori many reasons to believe that the global oil market dynamics have changed over time. Consider, for instance, the transition from a regime of administered oil prices to a market-based system accompanied by a dramatic rise

---

This aggressive hoarding behavior could hint at the increased importance of less elastic precautionary demand in total oil demand in a tightening market. In fact, Adelman (2002, p.179) states that "when buyers fear damage from sudden dearth, there is also a *precautionary* motive; which may be joined to a *speculative* motive, to profit by buying sooner."

<sup>49</sup>See Alquist and Kilian (2010) for a similar argument in relation to oil futures spreads.

<sup>50</sup>Note that this is not an exogenous shift of the (precautionary) oil demand curve due to e.g. the possibility of a war, but an endogenous increase of (more inelastic) precautionary oil demand after an oil supply shock when operating close to full capacity. The former results in a shift of the oil demand curve, while the latter implies a steepening of the demand curve at higher utilization rates.

<sup>51</sup>Gately (1984, p.1103) observes that "the dynamics of the price increase (in 1979-1980), [...], were not substantially different from what would occur from a disruption in a competitive industry with *low short-run elasticities* [emphasis added]."

<sup>52</sup>Skeet (1988, p.91), for instance, notes that "in the midst of the Arab-Israeli war, the governments were far more concerned with security of supply than price."

in oil price volatility, changing capacity utilization rates in crude oil production and altering driving forces of oil prices. On the other hand, the economic structure has also changed considerably. For instance, the relative importance of oil in the production process has diminished over time, labor markets have become more flexible and monetary policy more credible.

From a methodological point of view, we depart from the existing empirical oil literature along two dimensions. To account for gradual changes over time, we have estimated multivariate structural vector autoregressions with time-varying parameters and stochastic volatility in the spirit of Cogley and Sargent (2005). Until now, time variation was analyzed via simple sample splits or the estimation of bivariate VARs with a moving window sample period. In addition, we propose a new strategy to identify oil supply shocks in a structural VAR using sign restrictions. Specifically, we identify an oil supply shock as a disturbance in the global oil market which shifts oil production and real crude oil prices in opposite directions. In contrast, all shocks that generate a positive co-movement between both variables are considered as oil demand shocks. Accordingly, we allow oil supply and demand disturbances to have an immediate impact on both oil prices and oil production. In the existing literature, oil supply shocks have been identified by imposing a zero contemporaneous impact of shifts in oil demand on either crude oil prices, oil production or both.

Surprisingly, we find that the impact of a typical one standard deviation oil supply shock on the US macroeconomy has been relatively constant over time. This finding stands in contrast to popular perception and a number of existing studies. However, this controversy can largely be explained by a remarkable structural change in the oil market itself. In particular, the oil demand curve is currently much steeper or less elastic relative to the 1970s and early 1980s which complicates comparisons over time. When the comparison is based on a standardized shift of oil prices (e.g. 10 percent rise), the impact on real GDP and inflation becomes smaller over time which is consistent with the existing evidence. This comparison, however, does not take into account that the same change in oil prices is currently characterized by a much smaller movement in oil production which can be considered as a different oil supply shock. Conversely, when an exogenous supply shock is measured as a normalized change in oil production (e.g. a fall of 1 percent), the output and inflation consequences are currently much more severe. Also this experiment is not realistic since an average oil production disturbance is currently only one fourth of a disturbance in the 1970s. Whether this reduced variability of oil production is due to the steepening of the oil demand curve and/or to smaller underlying disturbances in oil supply cannot be determined with our approach. In subsequent work (Baumeister and Peersman 2009), we provide evidence that oil supply volatility has indeed diminished over time.

We further show that the contribution of oil supply shocks to the variability of real activity and inflation is economically very relevant, being consistently between 15 and 20 percent. Also the proportion of oil supply shocks in total variability of global oil production remained more or less constant over time (approximately 30 percent). However, despite the currently stronger impact of a supply shock on real oil prices, the contribution of these shocks to crude oil price volatility has diminished considerably from 30 percent to about 20 percent. This is only possible if also the

supply curve has become more inelastic over time. Less elastic oil supply and demand curves in the global oil market both result in more variability of crude oil prices and must have contributed to the observed increase in oil price volatility. From our analysis also emerges that the Great Inflation of the 1970s cannot be explained by unfavorable oil supply shocks which confirms the propositions of BK (2002, 2004). On the other hand, there was a significant but non-exclusive contribution of oil price spikes to the recessions in 1974/75, early 1980s and 1990s. However, unfavorable oil supply disturbances also significantly reduced real activity around 1999, which made the ongoing economic boom more subdued. In addition, all more recent oil price surges can almost entirely be explained by shifts in global oil demand.

Finally, we propose some potential explanations for the steepening of the oil demand curve since the mid-eighties. In particular, after the oil price spikes of the 1970s, there has been substantial substitution to alternative sources of energy and more energy-efficient technologies were developed. The remaining amount of oil needed is thus an absolute necessity which could result in more inelastic oil demand. Moreover, due to insufficient investments in the oil industry, world oil production has been operating close to full capacity since the second half of the 1980s. Accordingly, physical shortfalls in oil production due to e.g. war activities, are difficult to replace by increased production elsewhere. As a consequence, less elastic precautionary oil demand becomes more important in total oil demand at current high levels of capacity utilization. Which of these structural changes dominates or whether there exist alternative explanations should be explored in future research.

## A Data appendix

World oil production data are available on a monthly basis from January 1973 onwards from the US Department of Energy (DoE). Monthly data for global production of crude oil for the period 1953M4 to 1972M12 (which are available upon request) have been taken from the Oil & Gas Journal (issue of the first week of each month). For the period 1947M1 to 1953M3 monthly data have been obtained by interpolation of yearly oil production data with the Litterman (1983) methodology using US monthly oil production as an indicator variable (available at DoE).<sup>53</sup> Annual oil production data have been retrieved from World Petroleum (1947-1954), the Oil & Gas Journal (end-of-year issues, 1954-1960) and the Annual Energy Review 2005 (1960-2005). Consistency between these different data sources has been checked at overlapping periods. Quarterly data are averages of monthly observations.

US petroleum consumption data are available on a monthly basis from January 1982 onwards in the Monthly Energy Review June 2008. For the period 1970M1 to 1981M12 monthly data have been reconstructed with the growth rate of total monthly supply of crude oil in the US calculated as the sum of domestic oil production and net oil imports. Monthly data for US crude oil field production (MCRFPUS1), total US crude oil imports (MCRIMUS1) and US crude oil exports (MCREXUS1) for the period 1970M1 to 2006M12 have been obtained from the US Energy Information Administration (EIA). The monthly oil consumption data have been aggregated to quarterly frequency to obtain the number of barrels of crude oil consumed in the US economy over the period 1970Q1-2006Q4. Data on annual petroleum consumption by sector have been taken from the Annual Energy Review 2007 (Table 2.1c-2.1e) for the period 1970-2006. Primary and secondary uses of petroleum in the industrial, commercial and part of the transportation sector (buses and heavy trucks) have been added up to obtain total oil consumption in the production process. Data on highway transportation energy consumption by mode (1970-2006) have been obtained from the Transportation Energy Data Book (Table 2.7).

The refiner acquisition cost of imported crude oil is taken from the DoE database.<sup>54</sup> Since this series is only available from January 1974, it has been backdated until 1947Q1 using the (quarterly) growth rate of the producer price index (PPI) for crude oil from the BLS database (WPU056). Data have been converted to quarterly frequency by taking monthly averages before the extrapolation. For our robustness checks with regard to the choice of the oil price variable, we use the quarterly average of the West Texas Intermediate (WTI) spot oil price obtained from the FRED, the St.Louis FED, database and the refiner acquisition cost of composite<sup>55</sup> crude oil from the DoE database.

---

<sup>53</sup>Since this part of the data is only needed for the training sample to initialize the priors based on the estimation of a fixed-coefficient VAR, the use of interpolated data as opposed to actual ones is of minor importance.

<sup>54</sup>The refiner acquisition cost of imported crude oil (IRAC) is a volume-weighted average price of all kinds of crude oil imported into the US over a specified period. Since the US imports more types of crude oil than any other country, it may represent the best proxy for a true “world oil price” among all published crude oil prices. The IRAC is also similar to the OPEC basket price.

<sup>55</sup>The entitlement system in force during the 1970s in the US required buyers to purchase foreign and domestic oil in fixed proportions so that the aggregate price was a linear average of these two kinds of oil.

The latter has been adjusted for price controls on domestic oil production during the 1970s as described in Mork (1989) and reconstructed back in time in the same way as the imported refiner acquisition cost series.

Quarterly seasonally adjusted series for US real and nominal GDP (GDPC1: real gross domestic product, billions of chained 2000 dollars; GDP: gross domestic product, billions of dollars) and quarterly seasonally adjusted data for the US GDP deflator (GDPDEF: gross domestic product implicit price deflator) have been obtained from the U.S. Department of Commerce: Bureau of Economic Analysis. Monthly seasonally adjusted data for the US CPI (CPIAUCSL: consumer price index for all urban consumers: all items, index 1982-1984=100) and US industrial production (INDPRO: industrial production index, index 2002=100) are taken from the FRED database, where the former has been converted to quarterly frequency by taking monthly averages. Monthly data for the civilian unemployment rate, 16 years and older, seasonally adjusted have been retrieved from the BLS database (UNRATE) and have been averaged over quarters.

## B Priors and Estimation

**Prior distributions and initial values.** The priors for the initial states of the regression coefficients, the covariances and the log volatilities,  $p(\theta_0)$ ,  $p(\alpha_0)$  and  $p(\ln h_0)$  respectively, are assumed to be normally distributed, independent of each other and independent of the hyperparameters which are the elements of  $Q$ ,  $S$ , and the  $\sigma_i^2$ . The priors are calibrated on the point estimates of a constant-coefficient VAR(4) estimated over the period 1947Q2-1967Q2.

We set  $\theta_0 \sim N[\hat{\theta}_{OLS}, \hat{P}_{OLS}]$  where  $\hat{\theta}_{OLS}$  corresponds to the OLS point estimates of the training sample and  $\hat{P}_{OLS}$  to four times the covariance matrix  $\hat{V}(\hat{\theta}_{OLS})$ . With regard to the prior specification of  $\alpha_0$  and  $h_0$  we follow Primiceri (2005) and Benati and Mumtaz (2007). Let  $P = AD^{1/2}$  be the Choleski factor of the time-invariant variance covariance matrix  $\hat{\Sigma}_{OLS}$  of the reduced-form innovations from the estimation of the fixed-coefficient VAR(4) where  $A$  is a lower triangular matrix with ones on the diagonal and  $D^{1/2}$  denotes a diagonal matrix whose elements are the standard deviations of the residuals. Then the prior for the log volatilities is set to  $\ln h_0 \sim N(\ln \mu_0, 10 \times I_4)$  where  $\mu_0$  is a vector that contains the diagonal elements of  $D^{1/2}$  squared and the variance covariance matrix is arbitrarily set to ten times the identity matrix to make the prior only weakly informative. The prior for the contemporaneous interrelations is set to  $\alpha_0 \sim N[\tilde{\alpha}_0, \tilde{V}(\tilde{\alpha}_0)]$  where the prior mean for  $\alpha_0$  is obtained by taking the inverse of  $A$  and stacking the elements below the diagonal row by row in a vector in the following way:  $\tilde{\alpha}_0 = [\tilde{\alpha}_{0,21}, \tilde{\alpha}_{0,31}, \tilde{\alpha}_{0,32}, \tilde{\alpha}_{0,41}, \tilde{\alpha}_{0,42}, \tilde{\alpha}_{0,43}]'$ . The covariance matrix,  $\tilde{V}(\tilde{\alpha}_0)$ , is assumed to be diagonal with each diagonal element arbitrarily set to ten times the absolute value of the corresponding element in  $\tilde{\alpha}_0$ . While this scaling is obviously arbitrary, it accounts for the relative magnitude of the elements in  $\tilde{\alpha}_0$  as has been noted by Benati and Mumtaz (2007).

With regard to the hyperparameters, we make the following assumptions along the lines of

Benati and Mumtaz (2007). We postulate that  $Q$  follows an inverted Wishart distribution:  $Q \sim IW(\bar{Q}^{-1}, T_0)$ , where  $T_0$  are the prior degrees of freedom which are set equal to the length of the training sample which is sufficiently long (20 years of quarterly data) to guarantee a proper prior. Following Cogley and Sargent (2002), we adopt a relatively conservative prior for the time variation in the parameters setting the scale matrix to  $\bar{Q} = (0.01)^2 \cdot \hat{V}(\hat{\theta}_{OLS})$  multiplied by the prior degrees of freedom. This is a weakly informative prior and the particular choice for its starting value is not expected to influence the results substantially since the prior is soon to be dominated by the sample information as time moves forward adding more time variation. We have experimented with different initial conditions inducing a different amount of time variation in the coefficients to test whether our results are sensitive to the choice of prior specification. We follow Primiceri (2005) in setting the prior degrees of freedom alternatively to the minimum value allowed for the prior to be proper,  $T_0 = \dim(\theta_t) + 1$ , together with a smaller value of the scale matrix,  $\bar{Q} = (0.003)^2 \cdot \hat{V}(\hat{\theta}_{OLS})$ , which puts as little weight as possible on our prior belief about the drift in  $\theta_t$ . Our results are not affected by different choices for the initial values of the prior. The three blocks of  $S$  are postulated to follow inverted Wishart distributions, with the prior degrees of freedom set equal to the minimum value required for the prior to be proper:  $S_1 \sim IW(\bar{S}_1^{-1}, 2)$ ,  $S_2 \sim IW(\bar{S}_2^{-1}, 3)$  and  $S_3 \sim IW(\bar{S}_3^{-1}, 4)$ . As for the scale matrices, they are calibrated on the absolute values of the respective elements in  $\tilde{\alpha}_0$  as in Benati and Mumtaz (2007). Given the univariate feature of the law of motion of the stochastic volatilities, the variances of the innovations to the univariate stochastic volatility equations are drawn from an inverse-Gamma distribution as in Cogley and Sargent (2005):  $\sigma_i^2 \sim IG\left(\frac{10^{-4}}{2}, \frac{1}{2}\right)$ .

**MCMC algorithm (Metropolis within Gibbs sampler): Simulating the Posterior Distribution.** Since sampling from the joint posterior is complicated, we simulate the posterior distribution by sequentially drawing from the conditional posterior of the four blocks of parameters: the coefficients  $\theta^T$ , the simultaneous relations  $A^T$ , the variances  $H^T$ , where the superscript  $T$  refers to the whole sample, and the hyperparameters – the elements of  $Q$ ,  $S$ , and the  $\sigma_i^2$  – collectively referred to as  $M$ . Posteriors for each block of the Gibbs sampler are conditional on the observed data  $Y^T$  and the rest of the parameters drawn at previous steps.

### Step 1: Drawing coefficient states

Conditional on  $A^T$ ,  $H^T$ ,  $M$  and  $Y^T$ , the measurement equation is linear and has Gaussian innovations with known variance. Therefore, the conditional posterior is a product of Gaussian densities and  $\theta^T$  can be drawn using a standard simulation smoother (see Carter and Kohn 1994; Cogley and Sargent 2002) which produces a trajectory of parameters:

$$p(\theta^T | Y^T, A^T, H^T) = p(\theta_T | Y^T, A^T, H^T) \prod_{t=1}^{T-1} p(\theta_t | \theta_{t+1}, Y^T, A^T, H^T)$$

From the terminal state of the forward Kalman filter, the backward recursions produce the required smoothed draws which take the information of the whole sample into account. More

specifically, the last iteration of the filter provides the conditional mean  $\theta_{T|T}$  and variance  $P_{T|T}$  of the posterior distribution. A draw from this distribution provides the input for the backward recursion at  $T - 1$  and so on until the beginning of the sample according to:

$$\begin{aligned}\theta_{t|t+1} &= \theta_{t|t} + P_{t|t}P_{t+1|t}^{-1}(\theta_{t+1} - \theta_t) \\ P_{t|t+1} &= P_{t|t} - P_{t|t}P_{t+1|t}^{-1}P_{t|t}\end{aligned}$$

### Step 2: Drawing covariance states

Similarly, the posterior of  $A^T$  conditional on  $\theta^T$ ,  $H^T$ , and  $Y^T$  is a product of normal densities and can be calculated by applying the same algorithm as in step 1 thanks to the block diagonal structure of the variance covariance matrix  $S$ . More specifically, a system of unrelated regressions based on the following relation:  $A_t u_t = \varepsilon_t$ , where  $\varepsilon_t$  are orthogonalized innovations with known time-varying variance  $H_t$  and  $u_t = y_t - X_t' \theta_t$  are observable residuals, can be estimated to recover  $A^T$  according to the following transformed equations where the residuals are independent standard normal:

$$\begin{aligned}u_{1,t} &= \varepsilon_{1,t} \\ \begin{pmatrix} h_{2,t}^{-\frac{1}{2}} u_{2,t} \end{pmatrix} &= -\alpha_{2,1} \begin{pmatrix} h_{2,t}^{-\frac{1}{2}} u_{1,t} \end{pmatrix} + \begin{pmatrix} h_{2,t}^{-\frac{1}{2}} \varepsilon_{2,t} \end{pmatrix} \\ \begin{pmatrix} h_{3,t}^{-\frac{1}{2}} u_{3,t} \end{pmatrix} &= -\alpha_{3,1} \begin{pmatrix} h_{3,t}^{-\frac{1}{2}} u_{1,t} \end{pmatrix} - \alpha_{3,2} \begin{pmatrix} h_{3,t}^{-\frac{1}{2}} u_{2,t} \end{pmatrix} + \begin{pmatrix} h_{3,t}^{-\frac{1}{2}} \varepsilon_{3,t} \end{pmatrix} \\ \begin{pmatrix} h_{4,t}^{-\frac{1}{2}} u_{4,t} \end{pmatrix} &= -\alpha_{4,1} \begin{pmatrix} h_{4,t}^{-\frac{1}{2}} u_{1,t} \end{pmatrix} - \alpha_{4,2} \begin{pmatrix} h_{4,t}^{-\frac{1}{2}} u_{2,t} \end{pmatrix} - \alpha_{4,3} \begin{pmatrix} h_{4,t}^{-\frac{1}{2}} u_{3,t} \end{pmatrix} + \begin{pmatrix} h_{4,t}^{-\frac{1}{2}} \varepsilon_{4,t} \end{pmatrix}\end{aligned}$$

### Step 3: Drawing volatility states

Conditional on  $\theta^T$ ,  $A^T$ , and  $Y^T$ , the orthogonalized innovations  $\varepsilon_t \equiv A_t (y_t - X_t' \theta_t)$ , with  $\text{Var}(\varepsilon_t) = H_t$ , are observable. However, drawing from the conditional posterior of  $H^T$  is more involved because the conditional state-space representation for  $\ln h_{i,t}$  is not Gaussian. The log-normal prior on the volatility parameters is common in the stochastic volatility literature but such a prior is not conjugate. Following Cogley and Sargent (2005, Appendix B.2.5) and Benati and Mumtaz (2007), we apply the univariate algorithm by Jacquier et al. (1994) that draws the volatility states  $h_{i,t}$  one at a time.<sup>56</sup>

### Step 4: Drawing hyperparameters

The hyperparameters  $M$  of the model can be drawn directly from their respective posterior distributions since the disturbance terms of the transition equations are observable given  $\theta^T$ ,  $A^T$ ,  $H^T$  and  $Y^T$ .

---

<sup>56</sup>An alternative way has been applied by Primiceri (2005) who uses the method proposed by Kim et al. (1998) which consists of transforming the non-Gaussian state-space form into an approximately Gaussian one by using a discrete mixture of normals. This linear transformation then allows to apply a standard simulation smoother conditional on a member of the mixture.

We perform 50,000 iterations of the Bayesian Gibbs sampler but keep only every 10<sup>th</sup> draw in order to mitigate the autocorrelation among the draws. After a "burn-in" period of 50,000 iterations, the sequence of draws of the four blocks from their respective conditional posteriors converges to a sample from the joint posterior distribution  $p(\theta^T, A^T, H^T, M | Y^T)$ . We ascertain that our chain has converged to the ergodic distribution by performing the usual set of convergence tests (see Primiceri 2005; Benati and Mumtaz 2007). The results of these convergence diagnostics are available upon request. In total, we collect 5000 simulated values from the Gibbs chain on which we base our structural analysis.

## C Impulse responses and sign restrictions

Here we describe the Monte Carlo integration procedure we use to compute the path of structural impulse response functions to an oil supply shock. In the spirit of Koop et al. (1996) we compute the generalized impulse responses as the difference between two conditional expectations with and without the exogenous shock:

$$IRF_{t+k} = E[y_{t+k} | \varepsilon_t, \omega_t] - E[y_{t+k} | \omega_t]$$

where  $y_{t+k}$  contains the forecasts of the endogenous variables at horizon  $k$ ,  $\omega_t$  represents the current information set and  $\varepsilon_t$  is the current disturbance term. At each point in time the information set we condition upon contains the actual values of the lagged endogenous variables and a random draw of the model parameters and hyperparameters. More specifically, in order to calculate the conditional expectations we simulate the model in the following way: We randomly draw one possible state of the economy at time  $t$  from the Gibbs sampler output represented by the time-varying lagged coefficients and the elements of the variance covariance matrix. Starting from this random draw from the joint posterior including hyperparameters, we stochastically simulate the future paths of the coefficient vector as well as the (components of the) variance covariance matrix based on the transition laws for 20 quarters into the future.<sup>57</sup> By projecting the evolution of the system into the future in this way, we account for all the potential sources of uncertainty deriving from the additive innovations, variations in the lagged coefficients and changes in the contemporaneous relations among the variables in the system.

Given the current state of the economy, let  $\Omega_t = P_t D_t P_t'$  be the eigenvalue-eigenvector decomposition of the VAR's time-varying variance covariance matrix  $\Omega_t$  at time  $t$ . Draw an  $N \times N$  matrix,

---

<sup>57</sup>Alternatively, one could draw the entire time-varying path of current and future coefficients and covariances from the Gibbs sampler for the horizon  $k$  over which one wants to study the dynamics of the system. However, in order to be able to analyse the system dynamics also for the last years of the sample, one would have to extend the coefficient vector as well as the components of the variance covariance matrix since posterior information for the parameters of the VAR is only available up to the last date in the sample. Even though the last observations of these elements would constitute the best forecast when the evolution of the parameters are modeled as random walks, imposing constant parameters on the last part of the sample appears to be overly restrictive and might omit important dynamics deriving from future parameter variation.

$K$ , from the  $N(0, 1)$  distribution, take the  $QR$  decomposition of  $K$  where  $Q$  is a matrix whose columns are orthogonal to each other and compute the time-varying structural impact matrix as  $B_{0,t} = P_t D_t^{\frac{1}{2}} Q'$ . Given this contemporaneous impact matrix, we compute the reduced-form innovations based on the relationship  $u_t = B_{0,t} \varepsilon_t$ , where  $\varepsilon_t$  contains four structural shocks obtained by drawing from a standard normal distribution. Impulse responses are then computed by comparing the effects of a shock on the evolution of the endogenous variables to the benchmark case without shock, where in the former case the shock is set to  $\varepsilon_{i,t} + 1$ , while in the latter we only consider  $\varepsilon_{i,t}$ . The reason for this is to allow the system to be hit by other shocks during the propagation of the shock of interest. From the set of impulse responses derived in this way, we select only those impulse responses which at horizons  $t + k, k = 0, 1, \dots, 4$ , satisfy the sign restrictions, i.e. display the effects on the endogenous variables associated with the structural shock we wish to identify; all others are discarded.

We repeat this procedure until 100 iterations fulfil the sign restrictions and then calculate the mean responses of our four endogenous variables over these accepted simulations. For each point in time, we randomly draw 500 current states of the economy which provide the distribution of impulse responses taking into account possible developments of the structure of the economy. The representative impulse response function for each variable at each date is the median of this distribution.

## References

- [1] Adelman, Morris A., "World Oil Production & Prices 1947-2000", *The Quarterly Review of Economics and Finance* 42:2 (2002), 169-191.
- [2] Alquist, Ron, and Lutz Kilian, "What Do We Learn from the Price of Crude Oil Futures?", *Journal of Applied Econometrics* 25:4 (2010), 539-573.
- [3] Barsky, Robert B., and Lutz Kilian, "Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative" (pp. 137-183), in B.S. Bernanke and K. Rogoff (Eds.), *NBER Macroeconomics Annual 2001* (MIT Press, Cambridge, MA).
- [4] Barsky, Robert B., and Lutz Kilian, "Oil and the Macroeconomy since the 1970s", *Journal of Economic Perspectives* 18:4 (2004), 115-134.
- [5] Baumeister, Christiane, and Gert Peersman, "Sources of the Volatility Puzzle in the Crude Oil Market", *mimeo* (2009), Ghent University.
- [6] Bernanke, Ben S., "Energy and the Economy", speech to the Economic Club of Chicago, June 15, 2006.
- [7] Bernanke, Ben S., Mark Gertler, and Mark W. Watson, "Systematic Monetary Policy and the Effect of Oil Price Shocks", *Brookings Papers on Economic Activity* 1 (1997), 91-157.
- [8] Benati, Luca, and Haroon Mumtaz, "U.S. Evolving Macroeconomic Dynamics: A Structural Investigation", *ECB Working Paper 746* (2007).
- [9] Blanchard, Olivier J., and Jordi Galí, "The Macroeconomic Effects of Oil Price Shocks: Why are the 2000s so different from the 1970s?", in J. Galí and M. Gertler (Eds.), *International Dimensions of Monetary Policy* (2007).
- [10] Blanchard, Olivier J., and John A. Simon, "The Long and Large Decline in U.S. Output Volatility", *Brookings Papers on Economic Activity* 1 (2001), 135-174.
- [11] Burbidge, John, and Alan Harrison, "Testing for the Effects of Oil-Price Rises Using Vector Autoregressions", *International Economic Review* 25:2 (1984), 459-484.
- [12] Canova, Fabio, and Gianni De Nicoló, "Monetary Disturbances Matter for Business Fluctuations in the G-7", *Journal of Monetary Economics* 49:6 (2002), 1131-1159.
- [13] Canova, Fabio, and Luca Gambetti, "Structural Changes in the U.S. Economy: Is There a Role for Monetary Policy?", *Journal of Economic Dynamics and Control* 33:2 (2009), 477-490.
- [14] Carter, Christopher K., and Robert P. Kohn, "On Gibbs Sampling for State Space Models", *Biometrika*, 81:3 (2004), 541-553.

- [15] Clarida, Richard, Jordi Galí, and Mark Gertler, "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory", *Quarterly Journal of Economics* 115:1 (2000), 147-180.
- [16] Cogley, Timothy, Giorgio E. Primiceri, and Thomas J. Sargent, "Inflation-Gap Persistence in the U.S.", *American Economic Journal: Macroeconomics* (2009), forthcoming.
- [17] Cogley, Timothy, and Thomas J. Sargent (2002), "Evolving Post-WWII U.S. Inflation Dynamics" (pp. 331-347), in B.S. Bernanke and K. Rogoff (Eds.), *NBER Macroeconomics Annual 2001* (MIT Press, Cambridge, MA).
- [18] Cogley, Timothy, and Thomas J. Sargent, "Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII US", *Review of Economic Dynamics* 8:2 (2005), 262-302.
- [19] Dargay, Joyce, and Dermot Gately, "Oil Demand in the Industrialized Countries", *Energy Journal*, Special Issue: The Changing World Petroleum Market (1994), 69-85.
- [20] De Gregorio, José, Oscar Landerretche, and Christopher Neilson, "Another Passthrough Bites the Dust? Oil Prices and Inflation", *Economía* 7:2 (2007), 155-196.
- [21] Dotsey, Michael, and Max Reid, "Oil Shocks, Monetary Policy and Economic Activity", *Federal Reserve Bank of Richmond Economic Review*, July/August (1992), 14-27.
- [22] Edelstein, Paul, and Lutz Kilian, "The Response of Business Fixed Investment to Changes in Energy Prices: A Test of Some Hypotheses About the Transmission of Energy Price Shocks", *The B.E. Journal of Macroeconomics* 7:1 (2007) (contributions).
- [23] Edelstein, Paul, and Lutz Kilian, "How Sensitive are Consumer Expenditures to Retail Energy Prices?", *Journal of Monetary Economics* 56:6 (2009), 766-779.
- [24] Faust, Jon, "The Robustness of Identified VAR Conclusions About Money", *Carnegie-Rochester Conference Series on Public Policy* 49:1 (1998), 207-244.
- [25] Ferderer, J. Peter, "Oil Price Volatility and the Macroeconomy", *Journal of Macroeconomics*, 18:1 (1996), 1-26.
- [26] Fry, Renée A., and Adrian R. Pagan, "Some Issues in Using Sign Restrictions for Identifying Structural VARs", *NCER Working Paper 14* (2007).
- [27] Gately, Dermot, "A Ten-Year Retrospective: OPEC and the World Oil Market", *Journal of Economic Literature* XXII September (1984), 1100-1114.
- [28] Hamilton, James D., "Oil and the Macroeconomy Since World War II", *Journal of Political Economy* 91:2 (1983), 228-248.
- [29] Hamilton, James D., "A Neoclassical Model of Unemployment and the Business Cycle", *Journal of Political Economy* 96:3 (1988), 593-617.

- [30] Hamilton, James D., "This is What Happened to the Oil Price–Macroeconomy Relationship", *Journal of Monetary Economics* 38:2 (1996), 215-220.
- [31] Hamilton, James D., "What is an Oil Shock?", *Journal of Econometrics* 113:2 (2003), 363-398.
- [32] Hamilton, James D., "Oil and the Macroeconomy", in S. Durlauf and L. Blume (Eds.), *The New Palgrave Dictionary of Economics* (Palgrave Macmillan, 2<sup>nd</sup> edition, 2008).
- [33] Hamilton, James D., "Understanding Crude Oil Prices", *Energy Journal* 30:2 (2009), 179-207.
- [34] Hamilton, James D. and Ana María Herrera, "Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy: Comment", *Journal of Money, Credit, and Banking* 36:2 (2004), 265-286.
- [35] Hang, Leiming, and Meizeng Tu (2007), "The Impact of Energy Prices on Energy Intensity: Evidence from China", *Energy Policy* 35 (2007), 2978-2988.
- [36] Herrera, Ana María, and Elena Pesavento, "Oil Price Shocks, Systematic Monetary Policy and the ‘Great Moderation’", *Macroeconomic Dynamics* 13:1 (2009), 107-137.
- [37] Hooker, Mark A., "What Happened to the Oil Price–Macroeconomy Relationship?", *Journal of Monetary Economics* 38:2 (1996), 195-213.
- [38] Hooker, Mark A., "Are Oil Shocks Inflationary? Asymmetric and Nonlinear Specifications versus Changes in Regime", *Journal of Money, Credit, and Banking* 34:2 (2002), 540-561.
- [39] Hubbard, R. Glenn, "Supply Shocks and Price Adjustment in the World Oil Market", *Quarterly Journal of Economics* 101:1 (1986), 85-102.
- [40] Jacquier, Eric, Nicholas G. Polson, and Peter E. Rossi, "Bayesian Analysis of Stochastic Volatility Models", *Journal of Business and Economic Statistics* 12:4 (1994), 371-418.
- [41] Jiménez-Rodríguez, Rebeca and Marcelo Sánchez, "Oil Price Shocks and Real GDP Growth: Empirical Evidence for Some OECD Countries", *Applied Economics* 37:2 (2005), 201-228.
- [42] Kilian, Lutz, "The Economic Effects of Energy Price Shocks", *Journal of Economic Literature* 46:4 (2008a), 871-909.
- [43] Kilian, Lutz, "Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the US Economy?", *Review of Economics and Statistics* 90:2 (2008b), 216-240.
- [44] Kilian, Lutz, "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market", *American Economic Review* 99:3 (2009), 1053-1069.
- [45] Kilian, Lutz, and Robert Vigfusson, "Pitfalls in Estimating Asymmetric Effects of Energy Price Shocks", *mimeo* (2009), University of Michigan.

- [46] Kim, Chang-Jin, and Charles R. Nelson, *State Space Models with Regime Switching* (Cambridge, MA: MIT Press, 1999).
- [47] Kim, Sangjoon, Neil Shephard, and Siddharta Chib, "Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models", *Review of Economic Studies* 65:3 (1998), 361-393.
- [48] Koop, Gary, M. Hashem Pesaran, and Simon M. Potter, "Impulse Response Analysis in Non-linear Multivariate Models", *Journal of Econometrics* 74:1 (1996), 119-147.
- [49] Lee, Kiseok, Shawn Ni, and Ronald A. Ratti, "Oil Shocks and the Macroeconomy: The Role of Price Variability", *Energy Journal* 16:4 (1995), 39-56.
- [50] Litterman, Robert B., "A Random Walk, Markov Model for the Distribution of Time Series", *Journal of Business and Economic Statistics* 1:2 (1983), 169-173.
- [51] Mabro, Robert, "The International Oil Price Regime: Origins, Rationale and Assessment", *Journal of Energy Literature* 11:1 (2005), 3-20.
- [52] Marshall, Alfred, *Principles of Economics* (London: Macmillan, 8<sup>th</sup> edition, 1920).
- [53] McConnell, Margaret M., and Gabriel Pérez-Quirós, "Output Fluctuations in the United States: What Has Changed Since the Early 1980's?", *American Economic Review* 90:5 (2000), 1464-1476.
- [54] Mork, Knut, "Oil and the Macroeconomy When Prices Go Up and Down: An Extension of Hamilton's Results", *Journal of Political Economy* 97:3 (1989), 740-744.
- [55] *Oil and Gas Journal*, weekly, various issues since 1954 (Tulsa, Oklahoma: Pennwell Corporation).
- [56] Peersman, Gert, "What Caused the Early Millennium Slowdown? Evidence Based on Vector Autoregressions", *Journal of Applied Econometrics* 20:2 (2005), 185-207.
- [57] Peersman, Gert, and Roland Straub, "Technology Shocks and Robust Sign Restrictions in a Euro Area SVAR", *International Economic Review* 50:3 (2009), 727-750.
- [58] Peersman, Gert, and Ine Van Robays, "Oil and the Euro Area Economy", *Economic Policy* 24:60 (2009), 603-651.
- [59] Peirson, John, "The Importance of Being Unimportant: Marshall's Third Rule of Derived Demand", *Scottish Journal of Political Economy* 35:2 (1988), 105-114.
- [60] Pemberton, James, "Marshall's Rules for Derived Demand: A Critique and a Generalization", *Scottish Journal of Political Economy* 36:4 (1989), 396-405.
- [61] Primiceri, Giorgio E., "Time Varying Structural Vector Autoregressions and Monetary Policy", *Review of Economic Studies* 72:3 (2005), 821-852.

- [62] Rotemberg, Julio J. (2007), "Comment on Blanchard-Galí: The Macroeconomic Effects of Oil Price Shocks: Why are the 2000s so different from the 1970s?" (forthcoming), in J. Galí and M. Gertler (Eds.), *International Dimensions of Monetary Policy*.
- [63] Rotemberg, Julio J., and Michael Woodford, "Imperfect Competition and the Effects of Energy Price Increases on Economic Activity", *Journal of Money, Credit, and Banking* 28:4, Part1 (1996), 549-577.
- [64] Ryan, David L., and André Plourde, "Smaller and Smaller? The Price Responsiveness of Nontransport Oil Demand", *Quarterly Review of Economics and Finance* 42 (2002), 285-317.
- [65] Skeet, Ian, *OPEC: Twenty-five Years of Prices and Politics* (Cambridge: Cambridge University Press, 1988).
- [66] Stock, James H., and Mark W. Watson, "Has the Business Cycle Changed and Why?" (pp. 159-218), in M. Gertler and K. Rogoff (Eds.), *NBER Macroeconomics Annual 2002* (MIT Press, Cambridge, MA).
- [67] *The Economist*, "Crude Measures", May 29, 2008.
- [68] Uhlig, Harald, "What are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure", *Journal of Monetary Economics* 52:2 (2005), 381-419.

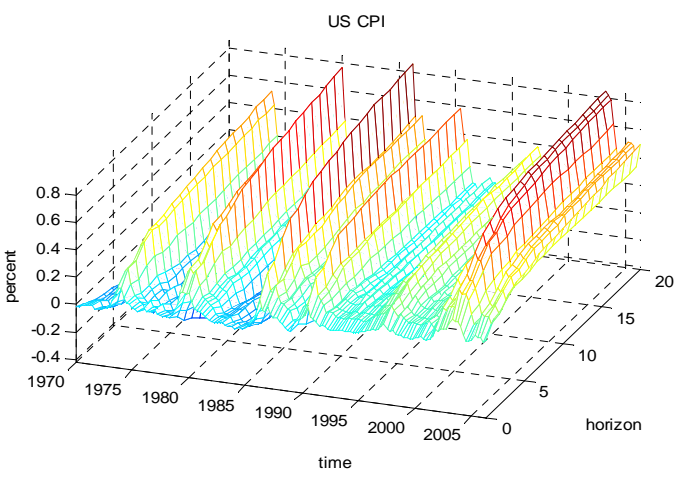
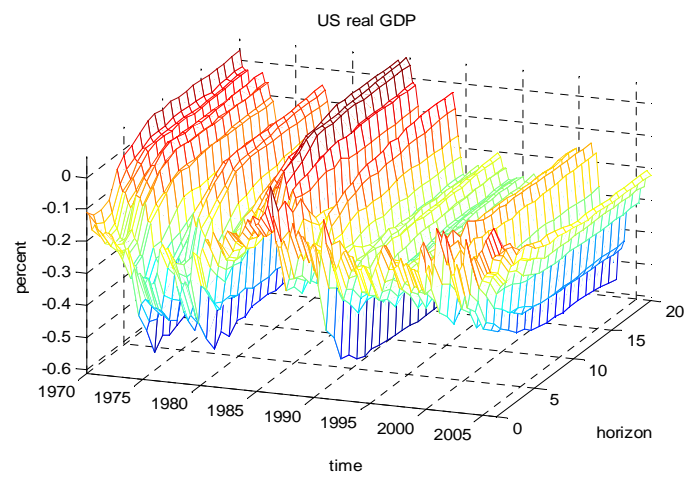
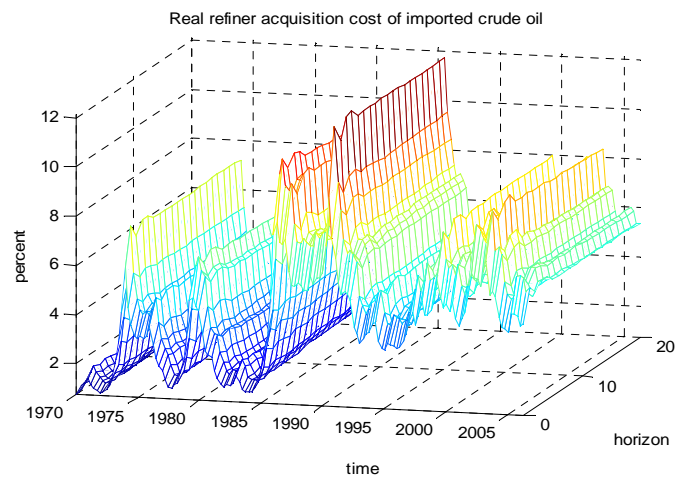
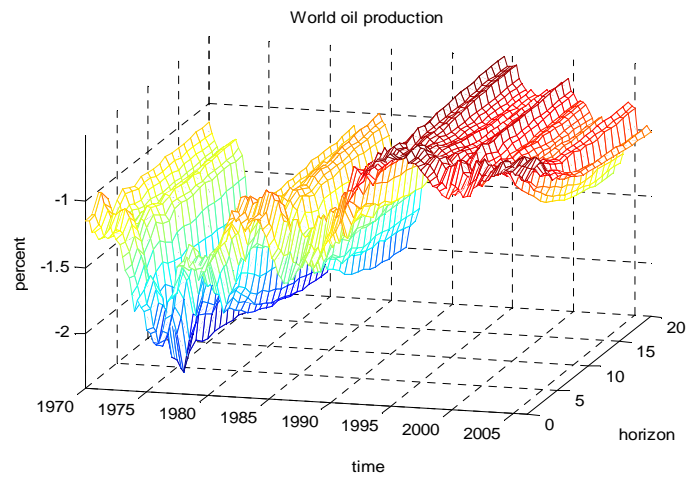

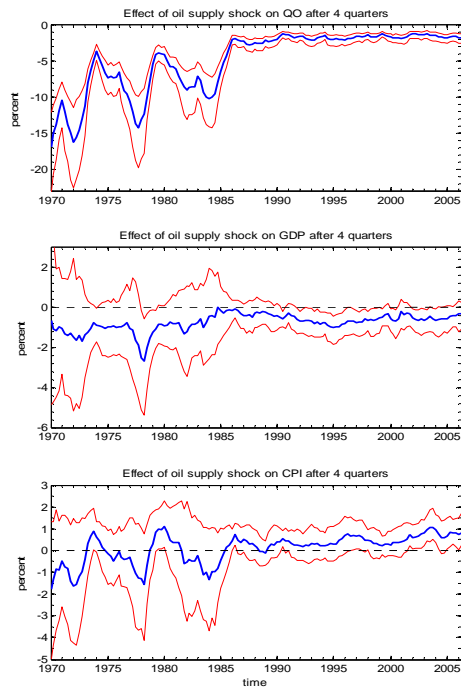


Figure 1: Time-varying median impulse response functions to a one standard deviation oil supply shock.

## Panel A



## Panel B



## Panel C

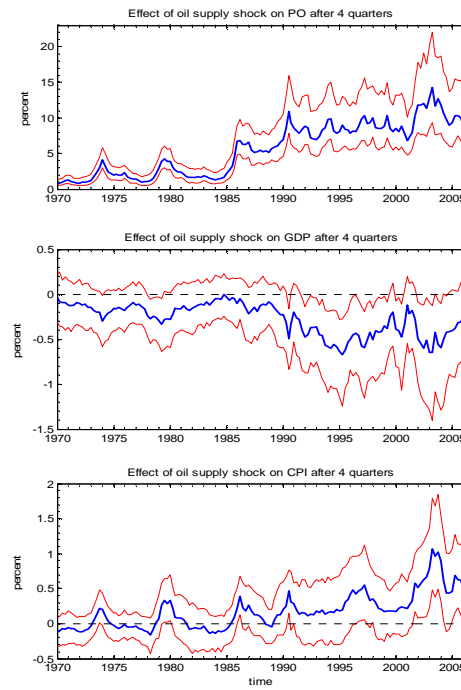


Figure 2: Median effect of an oil supply shock four quarters after the shock with 16<sup>th</sup> and 84<sup>th</sup> percentile confidence bands at each point in time.

Panel A: Responses to a one standard deviation shock.

Panel B: Responses after a 10% increase in the real price of oil.

Panel C: Responses after a 1% shortfall in world oil production.




Figure 3: Structural changes in the global crude oil market. Panel A: Change in the volatility of oil supply shocks. Panel B: Steeper oil demand curve.







Figure 4: Contribution of oil supply shocks to the forecast error variance of all four endogenous variables (in growth rates) after 20 quarters with 16<sup>th</sup> and 84<sup>th</sup> percentiles.


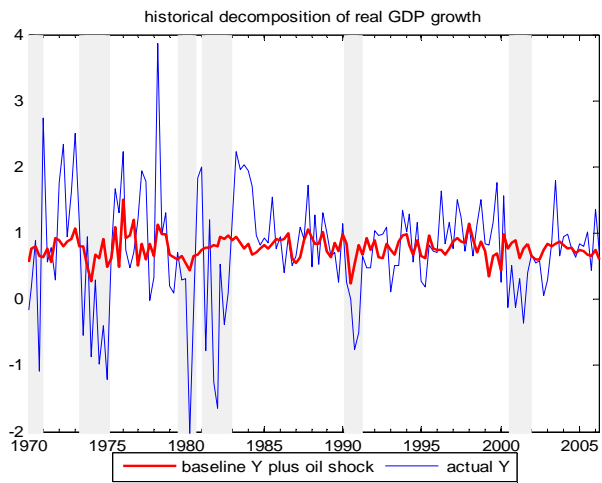

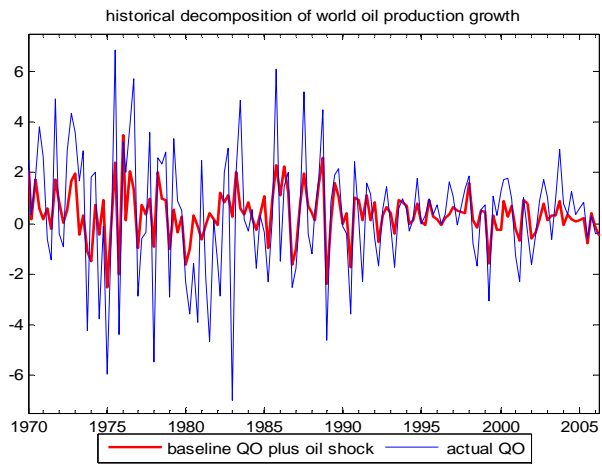



Figure 5: Historical contribution of oil supply shocks to world oil production growth, changes in the real price of crude oil, real GDP growth and CPI inflation.

### Panel A



### Panel B

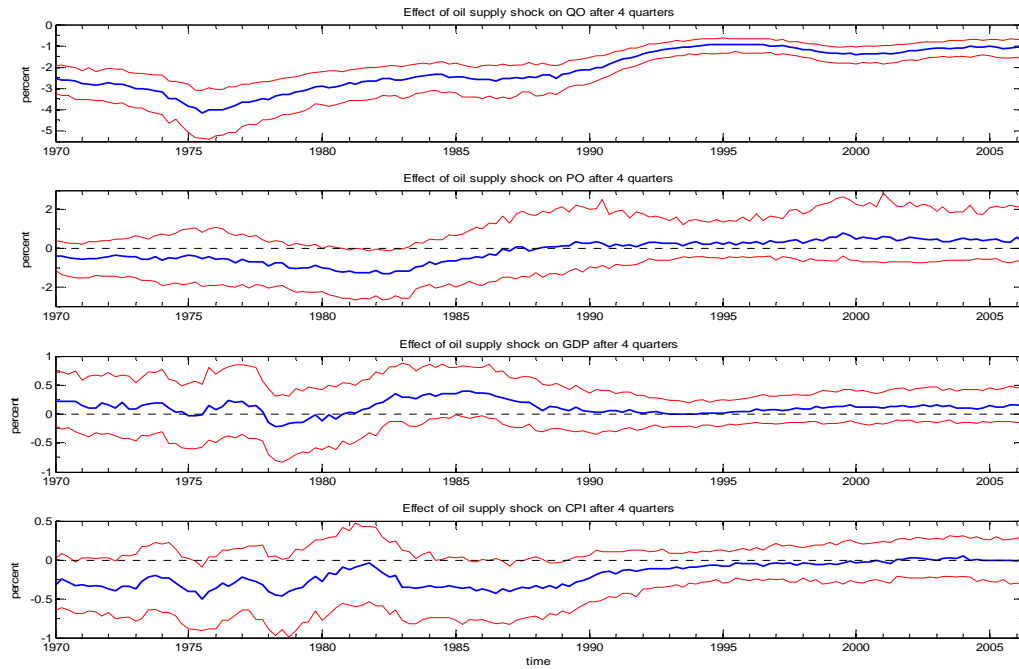
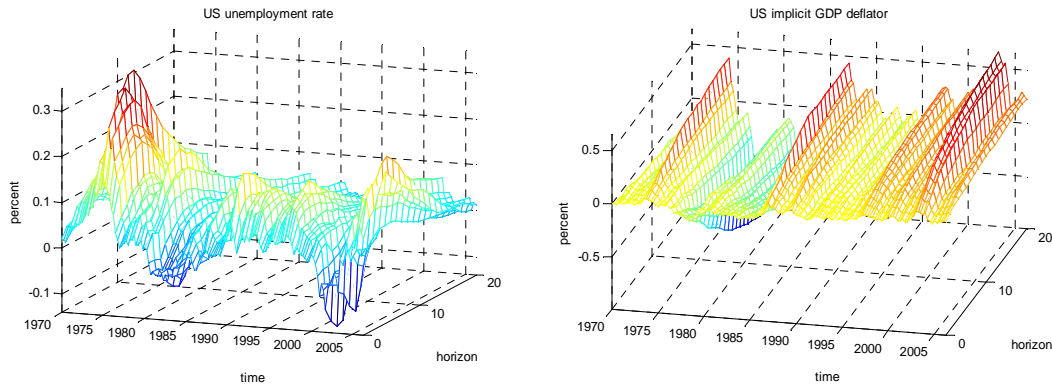


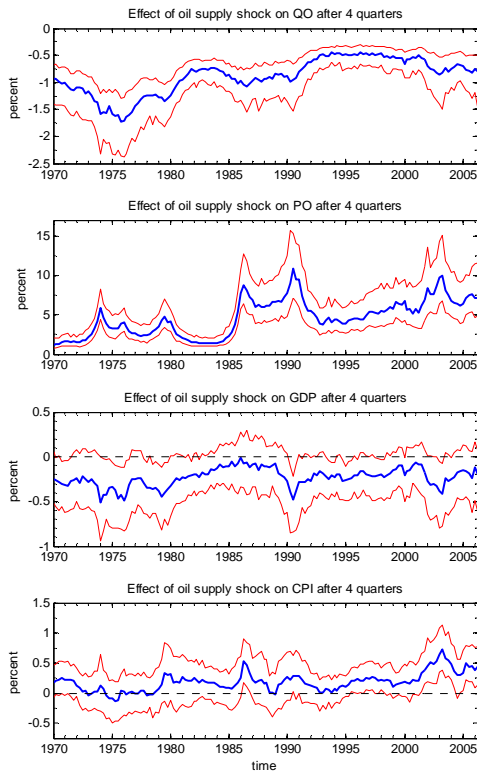
Figure 6: Robustness checks. Panel A: Median impulse responses after an oil supply shock identified with sign restrictions in a constant-coefficient VAR estimated over the two subsamples 1970Q1-1985Q4 and 1986Q1-2006Q2 together with 16<sup>th</sup> and 84<sup>th</sup> percentiles.

Panel B: Time-varying median impulse responses after an oil supply shock identified with zero short-run restrictions.

Panel A



Panel B



Panel C

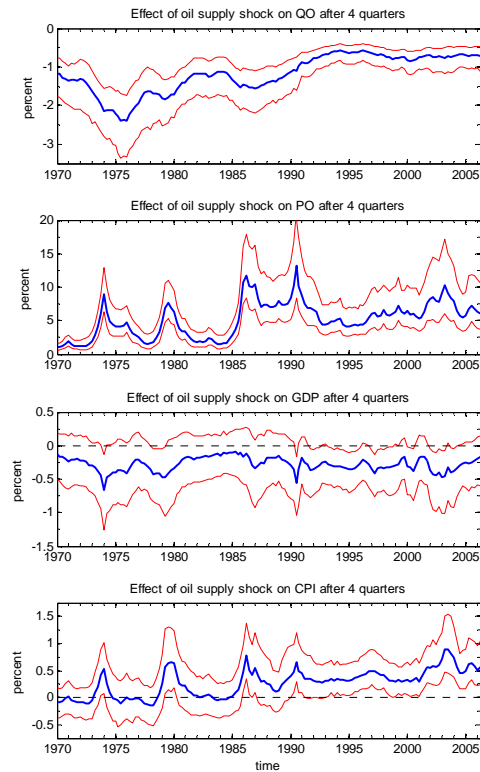



Figure 7: Robustness checks. Panel A: Median responses of US unemployment rate (left) and US implicit GDP deflator (right) to an oil supply shock. Panel B: Median responses to an oil supply shock after 4 quarters with 16<sup>th</sup> and 84<sup>th</sup> percentiles when federal funds rate is included. Panel C: Median responses to an oil supply shock after 4 quarters with 16<sup>th</sup> and 84<sup>th</sup> percentiles when sign restrictions are imposed from  $t=4$  to  $t=8$ .

Panel A



Panel B




Figure 8: Role of oil in the US economy. Panel A: US oil intensity of production.  
Panel B: Share of crude oil in US total expenditures.

Notes: The former has been computed as the annual petroleum consumption in British thermal units (BTU) by the industrial, commercial and part of the transportation sector per dollar of US real GDP, the latter has been calculated as number of barrels of crude oil consumed in the US economy each quarter times the nominal refiner acquisition cost of imported crude oil divided by US nominal GDP.




Figure 9: Global capacity utilization rates in crude oil production by year.

Notes: Estimates of global spare oil production capacity are obtained from the IMF World Economic Outlook (August 2006). Spare capacity refers to production capacity that can be brought online within 30 days and sustained for 90 days. Global capacity utilization rates are calculated as percentage of total potential annual world oil production.