

Sources of the Volatility Puzzle in the Crude Oil Market*

Christiane Baumeister

Ghent University

christiane.baumeister@ugent.be

Gert Peersman

Ghent University

gert.peersman@ugent.be

This version: January 2010

First version: February 2009

Abstract

A remarkable but unnoticed feature of the crude oil market is that the dramatic rise in oil price volatility over time has been accompanied by a substantial fall in oil production volatility. We investigate the reasons for this opposite evolution of both oil market variables. Our main finding is that the observed volatility puzzle can be rationalized by the fact that the price elasticities of both oil supply and oil demand have decreased considerably over time. This implies that small disturbances on either side of the oil market currently generate large price reactions but only modest quantity adjustments. We further document that the variance of innovations which shift oil demand and supply has even become smaller in the more recent past thereby mitigating oil price fluctuations.

JEL classification: E31, E32, Q43

Keywords: Oil prices, volatility, time variation, price elasticities

*We thank Luca Benati, Matteo Ciccarelli, Lutz Kilian, Vivien Lewis, Davide Raggi, James Smith and Paolo Surico for insightful discussions and helpful comments. We acknowledge financial support from IUAP and FWO. All remaining errors are ours.

1 Introduction

The recent rollercoaster ride of crude oil prices from values of around 50\$ per barrel at the beginning of 2007 to record highs of almost 150\$ in mid-2008, back to values as low as 40\$ at the end of the same year has attracted considerable attention of policymakers and the public, being illustrative of a dramatic rise in oil price volatility. However, sharp and erratic oil price movements are not a new phenomenon but have been a dominant feature during the last two decades.¹ A related aspect that has almost gone unnoticed in the literature is that, while oil price volatility has increased, the volatility of world oil production has decreased substantially over time. Figure 1, panel A displays the quarter-on-quarter rate of change in the nominal price of crude oil and world oil production for the period 1960Q1 to 2008Q1, whereas panel B shows the evolution of the median standard deviation of these two oil market variables over time, along with the 16th and 84th percentiles.² As is evident from the graphs, the amplitude of oil price fluctuations increased significantly in more recent periods.³ With the exception of the two sharp spikes in volatility in the 1970s, excessive swings seem a persistent feature of the nominal price of crude oil after the 1986 collapse of oil prices. World oil production on the contrary exhibited very wide fluctuations in the early part of the sample, especially during the 1970s, which gradually diminished over time. This diverging pattern of the two variables representing the global oil market is puzzling but suggestive of important transformations in the structure of the market for crude oil. It is of great importance to understand the causes of the increased oil price volatility, not only to devise suitable measures for dealing with it, but also because the underlying determinants could apply to the volatility of other commodity and asset prices as well. The goal of this paper is therefore to analyze alterations in oil market dynamics over time and to assess the factors that are at the origin of changes in the degree of volatility. More specifically, we investigate the reasons of the rise in oil price volatility and the concomitant drop in oil production volatility.

Several hypotheses could be put forward to account for the changed volatility of the oil market variables over time. Natural explanations can be sought in the evolution of

¹See Dvir and Rogoff (2009) for an empirical analysis of oil price behavior over the period 1861-2008.

²The time-varying standard deviations of the nominal price of crude oil and world oil production in panel B have been obtained from the empirical model which is presented in section 3.

³Other studies find the same pattern for oil price volatility by computing the standard deviation of log price differences over rolling time windows as an indicator of changes in volatility over time (Regnier 2007) or by estimating GARCH models over different sample periods (Yang *et al.* 2002).

the variance of shocks or the relative importance of different types of shocks affecting the oil market. Increases or decreases in the size of certain underlying shocks alone, however, cannot explain the *inverse* evolution of oil price and oil production volatility. For instance, while greater oil-specific demand shocks due to changes in inventory practices or speculative activities have the potential to account for the observed increase in oil price variability, such a hypothesis cannot explain the accompanying fall in oil production variability. Similarly, smaller oil supply disruptions in more recent periods compared to the 1970s and early 1980s could be a source of a decline in oil production volatility, but are incompatible with greater oil price fluctuations. Hence, at least a combination of different magnitudes of the underlying shocks is potentially needed to explain the oil market volatility puzzle.

Other structural changes in the oil market over time should also be considered. In particular, a fall in the price elasticity of oil supply or oil demand can rationalize an opposite movement of oil price and production volatility over time. For instance, a less elastic oil demand curve implies that similar shifts of an upward-sloping oil supply curve are characterized by smaller adjustments in oil production and larger fluctuations of oil prices. Likewise, a steeper oil supply curve could be at the origin of an increase in oil price volatility and a decrease in oil production variability for similar shocks at the demand side of the oil market. Finally, a change in the degree of flexibility of oil prices could be relevant for the volatility pattern in the crude oil market. Before the collapse of the OPEC cartel in 1985, and even more so during the 1960s, oil market transactions were mainly based on long-term contracts with predetermined oil prices. As a consequence, large production adjustments were needed to accommodate changes in the demand for crude oil, at least for the remaining period of the contract. The transition to the current market-based system of spot market trading should be conducive to more rapid translations of oil supply and demand variations into price changes. As a result, smaller shifts in global oil production would be required to clear the market.

The contribution of this paper is to evaluate the validity of any of the above hypotheses in a unifying framework and to provide empirical evidence that the increase in oil price volatility can be reconciled with a decrease in oil production volatility once we allow the price responsiveness of oil supply and oil demand as well as the variance of the underlying shocks to vary over time. To this end, in the spirit of Cogley and Sargent (2005) and Primiceri (2005), we estimate a time-varying parameter Bayesian vector autoregression

model with stochastic volatility in the innovation process over the sample period 1960Q1-2008Q1. Within this VAR framework, we identify three types of structural disturbances that drive the movements in world oil production and oil prices, namely oil supply shocks, oil demand shocks caused by shifts in economic activity and demand shocks that are specific to the crude oil market. These shocks are identified by means of sign restrictions to allow for an immediate effect of the shocks on both oil prices and oil production which can change over time.⁴

Our key finding is that the volatility puzzle in the crude oil market is mainly driven by a considerable decrease in the price responsiveness of oil supply and oil demand attaining very low levels since the mid-eighties. An important implication of these low price elasticities is that any small excess demand or supply of crude oil requires large jumps in prices to clear the global oil market. Put differently, a steepening of oil supply and oil demand curves over time is the source of higher oil price volatility accompanied by smaller oil production movements. Unlike Hubbard (1986), who views the transition from long-term oil contracts to spot market deals in the mid-eighties as causal to the rise in oil price variability, our results suggest that the substantial swings in oil prices as a result of less elastic curves could also have fostered the shift from contractual arrangements to spot market transactions. In fact, we argue that the structural transformation in the oil market is probably the result of an interplay between several features that tend to reinforce each other. On the one hand, uncertainties deriving from greater oil price volatility could have encouraged the development of derivative markets, stimulated the reliance on oil futures as risk-reduction tools and led to the introduction of crude oil options as hedging devices. On the other hand, while these financial instruments were designed to cope with the rise in oil price volatility, it is conceivable that the expansion of hedging possibilities also played a role in lowering the sensitivity of oil consumers and producers to price fluctuations thereby contributing to the steepening of the oil supply and demand curves which results in higher oil price volatility. Interestingly, if such an interplay exists, our results could extend to the volatility of other types of commodities and assets. An advantage of the oil market application is the availability of both price and quantity variables, which are necessary to

⁴Kilian (2009) disentangles a similar set of shocks by imposing short-run zero restrictions. However, a recursive identification scheme is not appropriate for our purpose. Such an identification scheme does, for instance, not allow the short-run price elasticity of oil supply to vary over time. The alternative identification strategy for several types of oil shocks can be considered as a separate contribution of this paper.

measure price elasticities.

The changed volatility of the oil market variables over time is, however, not exclusively determined by a lower price elasticity of oil supply and oil demand but also by the magnitude of disturbances affecting the oil market. By means of simple back-of-the-envelope computations, we find that the variance of all three kinds of shocks has gradually decreased. More specifically, the widespread increase in macroeconomic stability, known as the "Great Moderation", appears to have carried over to the oil market, i.e. we find smaller average shifts of the oil demand curve driven by shocks to global economic activity over time.⁵ Furthermore, consistent with expectations, the average variability of exogenous oil supply disruptions was rather low before the oil shock of 1973/74, increased notably thereafter and remained relatively high until the invasion of Kuwait in 1990, after which it declined steadily. Interestingly, the variance of an average oil-specific demand shock is also smaller in more recent times compared to the 1970s and 1980s. This is in line with Kilian (2009) who argues that precautionary oil demand shocks were also important driving forces behind oil price fluctuations in previous decades. Finally, our evidence also reveals that the transition from administered prices to a market-based regime had hardly any impact on the speed of adjustment of crude oil prices and consequently oil price volatility when quarterly data are used. In particular, even before 1985, oil prices moved almost immediately to their new long-run equilibrium value following oil supply or demand disturbances, i.e. we observe little sluggishness in the behavior of oil prices over the entire sample period.

The rest of the paper is organized as follows. In the next section, we present a small stylized model of the crude oil market to formulate the different hypotheses of the oil market volatility puzzle we wish to examine. Section 3 introduces the econometric framework, while the empirical results are reported in section 4. We briefly discuss a number of factors that might have contributed to the joint steepening of the oil supply and oil demand curves in section 5. Some final remarks in Section 6 complete the paper.

⁵See Blanchard and Simon (2001), McConnell and Pérez-Quirós (2000) and Stock and Watson (2003) for an account of the potential sources of the "Great Moderation".

2 A stylized model of the crude oil market

In this section, we set out a small time-varying model of the crude oil market, which should allow us to derive the different hypotheses to explain the changing volatility of oil production and oil prices over the sample period. In its simplest form, the crude oil market can be represented by the following demand and supply equations, measured as deviations from steady state:

$$\ln Q_t^D = -d_t \ln P_t^* + \varepsilon_t^d \quad (1)$$

$$\ln Q_t^S = s_t \ln P_t^* + \varepsilon_t^s \quad (2)$$

where oil demand Q_t^D and oil supply Q_t^S at each point in time are respectively a negative and positive function of the equilibrium price of oil P_t^* . d_t and s_t are positive values which represent the responsiveness of respectively the quantity of oil demanded and supplied to a change in the price of crude oil, i.e. the slopes of oil demand and supply curves at time t . Furthermore, the supply and demand for crude oil are driven by two mutually uncorrelated exogenous shocks: ε_t^d and ε_t^s , with $E[\varepsilon_t^d] = E[\varepsilon_t^s] = 0$, $E[\varepsilon_t^d]^2 = \sigma_{d,t}^2$, $E[\varepsilon_t^s]^2 = \sigma_{s,t}^2$ and $E[\varepsilon_t^d, \varepsilon_t^s] = 0$. In equilibrium, we can express the price and quantity variables as a linear combination of the structural shocks hitting the oil market:

$$\ln P_t^* = \frac{\varepsilon_t^d}{s_t + d_t} - \frac{\varepsilon_t^s}{s_t + d_t} \quad (3)$$

$$\ln Q_t^* = \frac{s_t \varepsilon_t^d}{s_t + d_t} + \frac{d_t \varepsilon_t^s}{s_t + d_t} \quad (4)$$

The period before 1985, however, was characterized by a regime of administered oil prices. In particular, the long-term contracts stipulated a fixed price for oil delivery over a certain period of time. Accordingly, oil producers had to adjust oil production in response to changes in the demand for crude oil until a new price was negotiated. At least in the short run, this supply behavior should be accounted for. We therefore allow the actual oil price to evolve gradually towards its equilibrium level:

$$\ln P_t = \lambda_t \ln P_t^* + (1 - \lambda_t) \ln P_{t-1} \quad (5)$$

with $0 < \lambda_t < 1$ the time-varying speed of adjustment to the new equilibrium price. If $\lambda_t = 1$, the price of oil immediately reflects its fundamental value, which is expected to be

the case in the more recent decades. The actual (short-run) price and quantity equations of oil that clear the market at each point in time are as follows:

$$\ln P_t = \frac{\lambda_t \varepsilon_t^d}{s_t + d_t} - \frac{\lambda_t \varepsilon_t^s}{s_t + d_t} \quad (6)$$

$$\ln Q_t = \frac{[s_t + (1 - \lambda_t) d_t] \varepsilon_t^d}{s_t + d_t} + \frac{\lambda_t d_t \varepsilon_t^s}{s_t + d_t} \quad (7)$$

When oil contracts are fully flexible, i.e. $\lambda_t = 1$, equations (6) and (7) are equal to their equilibrium counterparts (3) and (4).

According to this stylized model, and taking into account that oil supply and oil demand disturbances are uncorrelated, the variability of crude oil prices and oil production are respectively:

$$E [\ln P_t]^2 = \frac{\lambda_t^2 (\sigma_{d,t}^2 + \sigma_{s,t}^2)}{(s_t + d_t)^2} \quad (8)$$

$$E [\ln Q_t]^2 = \frac{[s_t + (1 - \lambda_t) d_t]^2 \sigma_{d,t}^2 + \lambda_t^2 d_t^2 \sigma_{s,t}^2}{(s_t + d_t)^2} \quad (9)$$

Relying on equations (8) and (9), we can formulate all possible hypotheses about the sources of the observed change in volatility of both oil market variables. We now discuss them one by one.

A change in the variance of oil market shocks. A first possible source of time variation in the oil market volatilities are changes in the variance of the underlying shocks. Keeping all other parameters of the model fixed, a change in the variance of oil market disturbances should have the following impact on the variability of oil prices and oil production:

$$\frac{\partial E [\ln P_t]^2}{\partial \sigma_{s,t}^2} > 0 \text{ and } \frac{\partial E [\ln Q_t]^2}{\partial \sigma_{s,t}^2} > 0 \quad (10)$$

$$\frac{\partial E [\ln P_t]^2}{\partial \sigma_{d,t}^2} > 0 \text{ and } \frac{\partial E [\ln Q_t]^2}{\partial \sigma_{d,t}^2} > 0 \quad (11)$$

Consider oil supply shocks. The 1970s are commonly perceived as a period of serious disruptions in the supply of oil due to military conflicts and political events, whereas more recent periods are rather characterized by minor disturbances on the supply side (Hamilton 2009a,b). Accordingly, a reduction in the standard deviation of oil supply shocks can be a source of reduced volatility of oil production over time, but cannot explain the opposite

evolution of oil price volatility. In contrast, smaller oil supply disturbances should also result in lower variability of crude oil prices in more recent periods, as can be seen from equation (10). Hence, this hypothesis alone does not suffice to explain the volatility puzzle.

The variance of average oil demand shocks could also have changed over time. On the one hand, the transition from the "official price" regime to a market-based system of direct trading in the spot market, which took place during the 1980s, resulted in a shift of price determination away from OPEC to the financial markets and the development of oil futures markets (Mabro 2005). This evolution is often seen as the source of the dramatic rise in oil price volatility (Hubbard 1986). Equation (11) suggests that if oil demand shocks resulting from e.g. increased speculative activities, precautionary buying or preference shifts, were indeed greater nowadays, they would have the potential to account for the observed increase in oil price variability. However, while increased competition and speculation appear to be plausible reasons for more frequent switches from price increases to decreases, also this hypothesis on its own cannot explain the concomitant drop in oil production variability.

On the other hand, at around the same time of the break in oil market volatility, a widespread increase in macroeconomic stability has taken place around the globe, commonly referred to as the "Great Moderation". Several studies indicate that this remarkable decline in volatility is not limited to output growth and inflation but also extends to other macroeconomic variables (e.g. Herrera and Pesavento 2009). As such, smaller fluctuations in oil production as a result of smaller oil demand shocks deriving from e.g. economic activity or monetary expansions, would accord well with this general phenomenon, but not the increase in oil price volatility. Since both hypotheses relating to the demand side of the oil market predict an opposite evolution of the variance of oil demand shocks over time, in our empirical analysis, we will make an explicit distinction between both types of shocks. In particular, we will identify oil-specific demand shocks and oil demand shocks which are driven by global economic activity.

Time-varying price elasticities of oil supply and oil demand. A change in the price elasticity of oil supply and oil demand could also play a role as can be inferred from the following derivations:

$$\frac{\partial E [\ln P_t]^2}{\partial d_t} < 0 \text{ and } \frac{\partial E [\ln Q_t]^2}{\partial d_t} > 0 \quad (12)$$

$$\frac{\partial E [\ln P_t]^2}{\partial s_t} < 0 \text{ and } \frac{\partial E [\ln Q_t]^2}{\partial s_t} > 0 \quad (13)$$

Baumeister and Peersman (2008) document a change in demand behavior over time. Specifically, they provide evidence of a less elastic oil demand curve since the second half of the eighties. A fall in the price elasticity of oil demand in more recent periods could indeed rationalize the oil market volatility puzzle. This evolution is in line with greater oil price volatility and smaller oil production fluctuations, as predicted by equation (12).

Time variation in the price elasticity of oil supply, with a lower elasticity in the more recent past, is also a plausible hypothesis to explain the volatility puzzle in the crude oil market (see equation (13)). Kilian (2008) reports that world oil production has been close to its full capacity level since the mid-eighties, which makes it very difficult to raise production volumes in the short run when the demand for oil increases. Smith (2009) interprets this fact as the result of purposeful behavior on the part of OPEC suppliers who refrain from expanding productive capacity despite the ample oil reserves available for development in order to support cartel discipline. As a consequence, the oil market is characterized by higher oil price fluctuations accompanied by limited adjustments in oil production.

More flexible crude oil prices over time. Finally, changes in the speed of oil price adjustment subsequent to shocks are expected to affect oil market volatility. The above described shift in the pricing regime from long-term oil contracts towards a more transparent system of spot market trading and the collapse of the OPEC cartel in late 1985 could have altered the flexibility of oil prices. If a greater fraction of oil transactions is carried out on the spot market, oil supply and demand variations are expected to translate quicker into price changes. According to our stylized model, an increase in the speed of adjustment of the actual oil price to its equilibrium value influences the variability of oil prices and oil production in the following way:

$$\frac{\partial E [\ln P_t]^2}{\partial \lambda_t} > 0 \text{ and } \frac{\partial E [\ln Q_t]^2}{\partial \lambda_t} = \frac{2d_t}{(s_t + d_t)^2} \{ -[s_t + (1 - \lambda_t) d_t] \sigma_{d,t}^2 + \lambda_t d_t \sigma_{s,t}^2 \} \leq 0 \quad (14)$$

On the one hand, more flexible contracts do result in greater oil price volatility in the short run. On the other hand, the impact of a faster convergence to the equilibrium price level on short-run oil production volatility is uncertain and crucially depends on the relative variance of supply and demand shocks in combination with the price elasticities of oil

supply and demand. Intuitively, since institutional arrangements in the oil market that prevailed until the mid-eighties relied on a fixed reference price for crude oil, adjustments to new demand conditions had to be carried out by adapting production volumes leading to wide fluctuations in oil production. On the other hand, fixing the price of oil should smooth oil supply disturbances, which reduces the variability of global oil production. The net effect on oil production volatility therefore depends on the relative importance of both shocks. If oil demand shocks were relatively more important in earlier periods, an increased speed of adjustment of the crude oil price to its equilibrium value alone could explain the oil market volatility puzzle.

It is very likely that many of the potential explanations come into play simultaneously. While the theoretical demand and supply relationships are easily established, identifying them is more difficult. In the next section, we present an empirical framework that allows us to examine the different hypotheses jointly.

3 Empirical methodology

Previous empirical studies about oil price volatility (e.g. Regnier 2007) note that the surge in volatility is not a one-time event but rather a sustained development which implies that the best modelling approach is one that allows for a slow-moving but continuous change as well as for potential jumps. We therefore use a VAR framework that features time-varying coefficients and stochastic volatilities in the innovation process. This approach enables us to evaluate time variation in the variance of shocks, the price elasticities of oil supply and demand, and the speed of oil price adjustment. We disentangle the structural shocks by means of sign restrictions.⁶ In particular, we identify exogenous oil supply shocks, oil demand shocks driven by economic activity and oil-specific demand shocks.

3.1 A VAR with time-varying parameters and stochastic volatility

We consider the following VAR(p) model with time-varying parameters and stochastic volatility in the spirit of Cogley and Sargent (2005), Primiceri (2005), and Benati and Mumtaz (2007):

$$y_t = C_t + B_{1,t}y_{t-1} + \dots + B_{p,t}y_{t-p} + u_t \quad (15)$$

⁶See Krichene (2002) for an alternative way of estimating structural relationships of demand and supply for oil.

where y_t is an 3×1 vector of observed endogenous variables that contains quarterly data on global oil production, the nominal refiner acquisition cost of imported crude oil and world industrial production,⁷ C_t is a vector of time-varying intercepts, $B_{p,t}$ are 3×3 matrices of time-varying coefficients on the lags of the endogenous variables, where the number of lags is set to $p = 4$ to allow for sufficient dynamics in the system,⁸ and u_t are heteroscedastic reduced-form innovations with zero mean and a time-varying covariance matrix Ω_t . The overall sample spans the period 1947Q1-2008Q1. However, the first eleven years of data are used as a training sample to calibrate the priors for estimation over the actual sample period which starts in 1960Q1.⁹

The drifting coefficients are meant to capture possible time variation in the lag structure of the model. The multivariate time-varying covariance matrix allows for heteroscedasticity of the shocks and time variation in the simultaneous relationships between the variables in the system. Including the stochastic volatility component appears appropriate given the changes in volatility of our two key variables documented above, since ignoring heteroscedasticity of the disturbance terms could lead to fictitious dynamics in the VAR coefficients as emphasized by Cogley and Sargent (2005), i.e. movements originating from the heteroscedastic covariance structure would be picked up by the VAR coefficients leading to an upward bias. Thus, allowing for time variation in both the coefficients and the variance covariance matrix implies that time variation can derive from changes in the

⁷All variables are transformed to non-annualized quarter-on-quarter rates of growth by taking the first difference of the natural logarithm. The oil price variable is the best proxy for the free market global price of imported crude oil. Our conclusions are not altered if we use the real price of oil (deflated by US CPI) instead. World industrial production is the broadest available index to capture the state of the world economy. To test for the robustness of our findings, we also re-estimated the model with (quarterly averages of) the measure of aggregate demand for industrial commodities devised by Kilian (2009) for the period 1975Q1-2007Q4 given that his indicator only starts in 1968 and the first 5 years are used as a training sample. Our conclusions remain largely unchanged when this measure is used as an indicator of global economic activity. A detailed description of the data, which are available upon request, can be found in Appendix A.

⁸The appropriate lag length is subject to debate (see Hamilton and Herrera 2004). However, our conclusions are unaltered if shorter lag orders are used. Including more lags would lead to a proliferation of parameters which is prohibitive given the time-varying nature of our model.

⁹This starting point corresponds to the establishment of OPEC. Even though the first decade of its existence was rather uneventful, we think that it is instructive to include this period in our analysis. See also Dvir and Rogoff (2009) for the importance of a long-term view for understanding oil market dynamics. We have also experimented with shorter sample periods, which did not alter the results. Given sufficiently diffuse priors, the same applies to the choice of the length of the training sample.

magnitude of shocks and their contemporaneous impact as well as from changes in the propagation mechanism. We estimate this model using Bayesian techniques. Technical details regarding the model setup, the prior specifications and the estimation strategy (Markov Chain Monte Carlo algorithm) are provided in Appendix B.

3.2 Identification of several types of oil shocks

Baumeister and Peersman (2008) isolate oil supply shocks by means of sign restrictions to analyze the dynamic consequences for the US economy. We extend their approach to the demand side of the crude oil market. An advantage of sign restrictions is the absence of a restriction on the magnitude of the contemporaneous impact of the shocks, which is not the case for recursive identification schemes (e.g. Kilian 2009). Accordingly, this impact as well as the short-run price elasticities can vary over time.¹⁰ More specifically, we place theoretically plausible sign restrictions on the time-varying impulse responses to recover the three underlying structural shocks we postulated to drive world oil production and the price of crude oil in the model presented in Section 2. The identification restrictions are summarized in Table 1.

Table 1: Identification restrictions

	Q_{oil}	P_{oil}	Y_{world}
Oil supply shock	≤ 0	≥ 0	≤ 0
Oil-specific demand shock	≥ 0	≥ 0	≤ 0
Global demand shock	≥ 0	≥ 0	> 0

Oil supply shocks are disturbances that shift the oil supply curve and could be the result of e.g. overproduction or supply interruptions due to war-related activities or destruction of oil facilities. According to equations (6) and (7) of our stylized model, such a shock moves oil quantity and oil prices in opposite directions. After a negative oil supply disturbance, the reaction of world industrial production is also restricted to be non-positive.¹¹

¹⁰See Uhlig (2005) and Peersman (2005) for alternative applications and the implementation of this identification strategy. See Fry and Pagan (2007) for a critical assessment of the sign-restriction methodology in general.

¹¹Note that the sign conditions are imposed as weak inequality restrictions, \leq or \geq (except for the response of world industrial production after a global demand shock), which implies that a zero response is always possible.

Oil demand shocks are disturbances that displace the oil demand curve and hence, move oil production and oil prices in the same direction in our model. As already mentioned, there are two types of demand-side shocks which could matter to explain the time profile of the oil market volatilities. On the one hand, fears of future shortages of oil supplies or expectations of strong future oil demand growth which result in precautionary buying, hoarding and speculation, should have an upward impact on volatility. On the other hand, increased macroeconomic stability that possibly also translates into smaller disruptions in the demand for oil as an input factor in the production process, should reduce variability. To examine both hypotheses, we make an explicit distinction between oil demand shocks driven by economic activity and oil-specific demand shocks. In order to disentangle the two kinds of demand disturbances, we impose the restriction that favorable global demand shocks unambiguously increase world industrial production, whereas oil-specific demand shocks that are not related to the business cycle have no effect or even a negative influence on global economic activity.

The sign restrictions are imposed to hold for four quarters after the occurrence of the shocks which accounts for the delayed response of the oil market variables in the early period of our sample due to institutional arrangements (Barsky and Kilian 2002). In particular, the described stickiness of the nominal price of crude oil due to long-term contractual agreements and the sluggish adjustment of oil production plans in response to demand shifts should be captured. By the same token, also the global economy needs some time to adapt to new conditions in the oil market. The potential sluggishness of the oil market is allowed in response to all three shocks. In this way, the identification scheme is able to accommodate different historical settings in the crude oil market (e.g. cartel vs competitive market forces, contracts vs spot sales), the importance of which are likely to have varied over the sample period. However, similar findings are obtained when the sign constraints are imposed for shorter periods and even only contemporaneously which means that these concerns are probably of minor importance. The implementation of the sign restrictions and the computation of impulse responses are discussed in Appendix C.

4 Empirical results

4.1 Oil market dynamics over time

Figure 2, panel A displays the median responses of global oil production and the price of imported crude oil to one-standard deviation structural shocks for horizons up to 20 quarters at each point in time spanning the period 1960Q1 to 2008Q1.¹² The estimated responses have been accumulated and are shown in levels. There is evidence of considerable time variation in the dynamic responses of the oil market variables after all three types of disturbances. The decline of world oil production after an unexpected oil supply shock follows a u-shaped pattern over time. The responsiveness of oil quantity continuously increases from the mid-sixties until the early eighties and then gradually dampens to reach the same modest level in the early 1990s as during the 1960s. Oil price reactions instead get more pronounced over time ranging from barely any response during the 1960s, episodes of greater reactivity during the 1970s to consistently stronger effects from the mid-eighties onwards. The evolutionary patterns observed after the two demand-induced shocks are similar with a subdued gradual rise and subsequent decline of the effect on world oil production, whereas the impact on the oil price variable increases notably over time. The most striking regularity is the remarkable increase of the impact on crude oil prices in response to all three structural innovations with two obvious breaks in 1970 and in 1986, after which the strength of the oil price responses increases, which is suggestive of the fact that the global oil market has experienced fundamental structural transformations.

4.2 Evaluation of hypotheses

4.2.1 Speed of adjustment

To assess whether faster convergence to the equilibrium price level is at the origin of the observed volatility pattern, we first look at the pattern of the dynamic responses of the nominal oil price at three selected points in time, which are depicted in Figure 3, panel A. If sluggishness in the oil price adjustment were an important feature only in the early

¹²The 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 1960Q1 to 2008Q1, 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.

part of the sample, we would expect a gradual rise in prices but we observe that oil prices jump immediately after all three structural shocks at each date.¹³ This feature emerges even more clearly in panel B, where the time-varying median responses of the crude oil price are plotted on impact and four quarters after each shock, along with the 16th and 84th percentiles of the posterior distribution for the whole sample period. Relying on these impulse responses, we also computed half-lives which confirm our visual analysis that more than 50% of the ultimate price increase is complete on impact or after one quarter at all times indicating no change in the duration of price adjustment across time.

Thus, the degree of price flexibility appears to be more or less homogeneous over time, so that changes in the institutional structure of the oil market have not increased the speed of adjustment to changes in fundamentals as an explanation for larger oil price fluctuations.

4.2.2 Evolution of price elasticities

As noted by Hamilton (2009a), it is difficult to trace the slope of the oil demand and oil supply curves from the observed movements in oil quantity and oil price because these two variables are subject to a myriad of influences which are hard to disentangle.¹⁴ But since we have identified the structural shocks in our empirical model that induce reactions in the oil market variables by shifting either the oil demand or oil supply schedule, we can derive estimates for the short-run price elasticities of oil demand and oil supply at each point in time from the impulse responses as the percentage change in world oil production divided by the percentage change in oil prices.

Figure 4, Panel A plots the evolution of the median of the estimated slopes of the oil supply and oil demand curves on impact for the entire sample. Note that since we identify two different types of demand shocks, we trace the curvature of the oil supply schedule once

¹³It does not matter which dates are selected since this pattern is representative for the entire sample; this can also be inferred from Figure 2, panel A.

¹⁴Nevertheless, Hamilton (2009a) tries to infer the slope of the oil demand curve by selecting historical episodes in which a shift of the supply curve was the primary factor for fluctuations in oil price and production (i.e. political events or war activities) and computes elasticities from the subsequent changes in quantities and prices. However, no single episode in the oil market is an exclusive supply-side story. Hence, this way of recovering a measure for demand elasticities constitutes only a rough approximation but it conveys the same message i.e. that the price elasticity of oil demand is even smaller now than it was in 1980.

with the oil-specific demand shock and once with the shock to global economic activity. While qualitatively similar, the estimated magnitude of the price elasticity for the supply of crude oil is somewhat different for both shocks, a finding which might point to a different reaction of oil supply depending on the source of demand. The estimates clearly provide evidence in favor of the hypothesis that attributes the volatility disconnect to a decrease in the responsiveness of respectively oil demand and supply to price changes. As is evident from the graphs, the time profile of price elasticities broadly falls into three phases with a first transition point in the late 1960s and another clear break in 1986. Given the substantial differences in scale which get the results out of perspective, we display the price elasticities for the three periods separately in Figure 4, Panel B, along with the 16th and 84th percentiles of the posterior distribution.

While the oil supply curve is relatively flat in the first part of the sample period, the responsiveness abates in the late 1960s and reaches very low levels in the most recent past marked by a sharp drop in elasticity around 1986. The extreme high price elasticity of supply before 1970 is in line with an oil market which was controlled by an oligopoly of large oil companies. In order to keep the oil price stable, these companies had to fully adjust oil production to accommodate changes in the demand for oil. After 1970, a wave of nationalizations shifted the market power towards oil-exporting countries resulting in a disintegration of the oligopolistic structure in the oil market and a corresponding fall of supply elasticity (Smith 2009, Dvir and Rogoff 2009). The estimated price elasticities for the period between 1970 and 1985 are still relatively high, with values mostly greater than 1. This period has only been interrupted by the two oil episodes of the 1970s during which oil quantities supplied were not very reactive to increases in crude oil prices, i.e. the adjustment in the aftermath of the shocks has not taken place via quantities. Strikingly, since the mid-eighties, median values for the slope of the oil supply curve fall between 0.05 and 0.4, indicating that the supply for crude oil has become highly insensitive to changes in its price. In section 5, we will discuss this break in more detail.

A fall in the price elasticity of oil demand has also contributed to the opposite evolution of oil price and production volatility. The changed elasticity between the pre-1970 and 1970-85 periods should be taken with a grain of salt. In particular, the lower bands of the posterior distribution before 1970, always being between -0.5 and -1.0, clearly overlap with the 1970-85 lower bands of the posterior. Only the medians and upper bands indicate a

higher elasticity in the former period.¹⁵ Since the mid-eighties, however, the price elasticity of demand declines substantially ranging between -0.05 and -0.25. These results accord well with evidence presented in previous studies (e.g. Cooper 2003, Krichene 2002, Ryan and Plourde 2002) where estimates of price elasticities obtained with different models and econometric techniques are also quite low for recent periods.¹⁶ A striking feature of our evidence is the similarity in the evolution of the slopes for both the oil supply and demand schedules, with a sudden fall of both elasticities in the mid-eighties. In section 5, we try to explain the concurrence of these developments.

4.2.3 Evolution of shocks

Within an SVAR framework, it is technically impossible to measure the underlying volatility of shocks. Only the contemporaneous impact of a one-standard deviation shock on the variables can be measured, which is a combination of the magnitude of the shock itself and the immediate reaction of the variable to that shock. To get an approximation of the magnitude of one-standard deviation structural shocks over time, we perform some simple back-of-the-envelope calculations. More specifically, given the price elasticities recovered from the estimated impulse responses, we can compute the time-varying magnitude of average underlying shocks by means of equations (1) and (2) from our stylized model of the oil market. Given the simplified nature of the model, the results should however be interpreted with caution.

Figure 5 depicts the changes in magnitude of average structural shocks, along with the 16th and 84th percentiles of the posterior distribution.¹⁷ As emerges from the graphs, all structural oil market shocks have become smaller in size during the latter part of the

¹⁵This high uncertainty for the early period is not a surprise. Given an oil price which was kept constant by oil producers, very small measured changes in oil production automatically result in a great elasticity, a feature which also applies to the price elasticity of supply. The constant price response can be seen in Figure 2.

¹⁶Ryan and Plourde (2002) explore a variety of different approaches to estimate the own-price elasticities of nontransport oil demand, among them price-decomposition techniques, systems of cost share equations and combinations thereof. Krichene (2002) instead uses a simultaneous equations model of oil supply and demand to derive short-run and long-run price elasticities and the estimates of Cooper (2003) are based on multiple regression models across 23 countries.

¹⁷Note that as in the case of the price elasticities, the variance of the oil supply disturbances are derived with both demand specifications.

sample. While the variance of oil demand shocks driven by economic activity rapidly increased in the early 1970s, it has been steadily diminishing since the mid-1980s which is in line with the onset of the "Great Moderation". Around the same time, substitution effects took hold thereby lowering the oil intensity of industrial production which might induce smaller shifts of the oil demand curve deriving from greater economic activity. However, aggregate demand shocks seem to have gained somewhat in size in the early 2000s.

The evolution of the variance of typical oil-specific demand shocks might depend on whether these shocks originate mainly from fears about future oil supply conditions, changes in inventory behavior or speculation, with the role of the latter being the most controversial. While there is limited evidence for the first two sources in more recent times, speculative and hedging activities developed since the late 1980s (Oil & Gas Journal 1989, Jan 23) and non-commercial trading in oil futures experienced an unprecedented surge in terms of market share since 2001 (Alquist and Kilian 2009) which would make speculative shocks the prime candidate for explaining the increase in oil price volatility. However, it is not straightforward to what extent trading in "paper barrels" would spill over to the physical market (Smith 2009). Indeed, we find that the variance of this kind of disturbances has decreased since the second half of the 1980s. Note, however, that the speculation-based explanation of oil price volatility is not necessarily linked to the variance of the shock itself but rather to the very low price elasticity of oil demand which, according to Hamilton (2009b), is the essential ingredient for speculation to "work".

Finally, the variance of typical oil supply disturbances has been quite large during the 1970s and 1980s but started to decline sharply in the early nineties being considerably lower since 1995. Again, estimates for the variance of oil supply innovations obtained with the oil-specific and global demand specifications differ only slightly. Consequently, smaller shocks have to some extent tempered the volatility increase of oil prices which could have been even larger had the variance of these shocks remained the same as during the 1970s.

5 Analysis of declined elasticities

So far we have documented a remarkable change in the magnitude of oil supply and demand elasticities over time. A striking feature is the similarity in the evolution of both the price elasticity of oil supply and oil demand, in particular the joint break around the mid-

eighties. In this section, we argue that a confluence of important structural changes in the oil market can account for the coincidence in timing. More specifically, we make the case that this apparent synchronicity is not coincidental but the result of specific conditions that tend to reinforce each other. Even though these proposed explanations for changes in the price elasticities are speculative in nature, it seems worthwhile discussing them in that they constitute potentially promising avenues for future research. In addition, if such an interplay exists, they may very well apply to other asset prices as well. We first discuss the interaction between the development of oil futures markets and the price elasticities, before moving to a more specific interaction between the price elasticities of supply and demand.

Oil futures markets and the price elasticity of oil supply and demand. At the same time of the substantial decline in both price elasticities in the mid-eighties, an important structural transformation occurred in the oil market. In particular, the transition from a regime of administered oil prices to a market-based system of direct trading in the spot market and an accompanying development of derivatives markets. Hubbard (1986) considers this transition the source of greater oil price volatility, which attracted speculators and fostered the development of oil futures markets. On the other hand, Smith (2009) argues that futures trading by speculators and hedgers should hardly exert any effect on the physical oil market, unless the buoyant futures market fuels expectations about spot prices, which trigger a reaction from market participants such as an accumulation of oil inventories or a fall in production that should result in a minor influence on current oil prices.

While the development of futures markets might not have a direct impact on the physical market, the availability of oil futures contracts as a risk management tool has the potential to indirectly alter the behavior of commercial traders on both sides of the oil market. More specifically, refineries and other oil consumers engage in oil futures trading to protect their business operations against unfavorable price movements by entering into a hedging contract. For instance, an airline company that wants to eliminate price risks associated with its future fuel purchases has to buy oil futures today to lock in the desired price for future delivery. Likewise, oil producers can lock in future sales revenues and profit margins and hedge themselves against declines in prices by selling oil futures contracts given their inherent long position in physical oil. As a result, both consumers

and producers will be unresponsive to price changes because physical purchases and sales of crude oil are hedged by offsetting financial positions in the oil futures market. Physical delivery is not even necessary,¹⁸ since hedgers typically use the net financial gains and losses to offset fluctuations in operating earnings. Put differently, opportunities for hedging could decrease the sensitivity of commercial dealers to oil price fluctuations in the spot market, contributing to less elastic oil supply and demand curves. The reduced price elasticities of supply and demand result in increased oil price volatility which further encourages the development of a market for derivatives. While crude oil futures were launched on the New York Mercantile Exchange (NYMEX) already in 1983, the volume of crude oil trading expanded substantially only after 1985 (Oil & Gas Journal 1989, Jan 23) which coincides with the drop in price elasticity of oil demand and oil supply uncovered in our empirical analysis.

Influence of investments and capacity constraints on supply and demand in the oil market. The uniformity in the evolution of both price elasticities can also be explained by an interplay between developments on the demand side of the crude oil market that trigger reactions on the supply side which in turn affect demand.¹⁹ According to Gately (2004), if the price elasticity of crude oil is relatively low, oil producers may deliberately refrain from increasing production capacity at a rapid pace to preserve the revenues from higher oil prices. Thus, the price elasticity of oil demand feeds back into the supply behavior of oil producers by reducing the incentives to bring new capacity on stream, which leads to less investments in infrastructure and an erosion of idle capacity.²⁰ Moreover, the geographic concentration of proven oil reserves in a limited number of OPEC countries, where investment decisions are not purely determined by economic

¹⁸According to the US Commodity Futures Trading Commission (CFTC), only about 2% to 5% of futures contracts lead to actual delivery of crude oil.

¹⁹Hamilton (2009b, p.10) exemplifies this interaction by stating that "it is a matter of conjecture whether the decline in Saudi production in 2007 should be attributed to depletion [...] or to a deliberate policy decision in response to a perceived decline in the price elasticity of demand."

²⁰We refer to Baumeister and Peersman (2008) for an overview of developments in oil-importing countries which could have triggered a reduction in the price elasticity of oil demand since the mid-eighties. For instance, rising oil prices during the 1970s induced many industries to switch away from oil to other sources of energy. As a result, the remaining amount of oil demand is absolutely necessary due to a lack of substitutes and therefore less elastic (e.g. the increased share of transportation in total oil demand).

considerations but also political factors,²¹ might have impaired the necessary investments in the oil industry. In fact, Smith (2009) advances the view that OPEC members pursue the deliberate strategy of limiting the growth of productive capacity in the interest of the common good of the cartel. This evolution is illustrated in Figure 6. Panel A displays the annual global capacity utilization rates of crude oil production over the period 1970 to 2007 derived from IMF and DoE estimates of spare oil production capacity, and panel B shows worldwide active rig counts for the period 1975M1 to 2007M12. The latter can be considered as one of the primary measures of exploratory activity in the oil industry and hence, a good indicator for investment in productive capacity. The figures clearly demonstrate that oil producers have effectively been nearing their capacity limit since the second half of the eighties, accompanied by a substantial decline in investment activities.²² When capacity constraints become binding, the flexibility of oil producers to offset unexpected oil market disturbances by raising oil supply is severely limited so that the adjustment has to take place via prices which implies a very inelastic oil supply curve.²³ In addition, the increased oil price volatility induces uncertainty which might lead to postponing investment in exploration and development needed to enhance the responsiveness of petroleum supply.

However, the presence of capacity constraints does not only affect the supply side of the world oil market but might also have the potential to induce a different behavior on the demand side. In fact, high rates of capacity utilization can put considerable strain on oil consumers in that they signal market tightness and hence raise fears about future oil scarcity, which makes market participants willing to pay a "fear premium" to shield

²¹Political impediments for the expansion of capacity can be sought in fear of expropriation, a resurgence of "resource nationalism" i.e. refusal of foreign direct investments and concerns about rapid depletion of oil resources i.e. preservation of oil reserves for future generations.

²²The maximum sustainable physical capacity is defined as "the maximum capacity that each OPEC country can produce at without damaging the reservoirs, while permitting itself long enough production life commensurate with its economic strategy" (Oil & Gas Journal 1989, Jan 9, p.29). Kilian (2008) considers capacity utilization rates close to 90% as reasonable for safeguarding the long-run productivity of an oil field. Notice also the limited response of investment to higher oil prices in more recent times compared to the late 1970s, which might be supportive of the decreased responsiveness of oil production to price changes.

²³Geroski *et al.* (1987) and Smith (2005) make the case that also the market structure plays an important role in determining the extent to which individual oil producers are willing to offset supply losses that occur elsewhere in the system and that their conduct (cooperative vs competitive) varies in function of excess capacity among other factors.

themselves from potential shortfalls in oil supplies. Put differently, each barrel of oil is of greater value to consumers given that it fulfills an insurance function against sudden dearth of crude oil delivery in the future.²⁴ This means that the share of precautionary demand in total oil demand increases when the oil sector is operating close to full sustainable capacity because agents anticipate that in case of a major oil shock, a shortfall in production volumes cannot be replaced by other producers since no idle capacity is left that could act as a buffer against abrupt interruptions. As a result, overall oil demand becomes less elastic. Hence, a more rigid demand curve reflects to some extent the degree of anxiety of oil consumers about the likelihood of future oil shortages. In the same way, expectations of growing shortages as a result of a lack of investment in productive capacity are likely to influence the current demand behavior of consumers.

6 Conclusions

In this paper, we have first documented the existence of an unnoticed puzzle in the crude oil market. In particular, an increase in oil price volatility over time which has been accompanied by a significant decline in oil production volatility. We then derived a set of potential hypotheses from a stylized demand and supply model for the crude oil market to explain this puzzle and assess their validity in a unifying empirical framework. Since the evolution of the oil market volatilities can follow from changes in the variance of structural shocks, changes in the speed of adjustment as a result of alterations in the institutional structure of the oil market and/or changes in the demand and supply behavior for crude oil, we have estimated a time-varying vector autoregression model for the period 1960Q1 to 2008Q1 that captures potential variations in the dynamic relationships and the volatility of shocks. For the identification of oil supply shocks, oil demand shocks caused by shifts in global economic activity and oil-specific demand shocks, we propose a set of sign restrictions. This specification serves both our purposes: first, to derive short-run price elasticities of oil supply and demand that are not a priori restricted to be zero on impact and second, to trace the evolution of the slopes of oil supply and demand curves,

²⁴This induced change in demand behavior (which concerns the slope of the curve) has to be clearly distinguished from oil-specific (precautionary) demand shocks; in the former case, oil consumers assign a greater value to the *same* amount of oil i.e. they pay a premium to ensure that they get this amount, whereas in the latter case, they effectively want to increase the quantity demanded (i.e. a shift of the oil demand curve) for stockbuilding.

the volatility of structural shocks and the degree of price flexibility over time.

We find that, while the variance of the shocks decreases over time and hence, changes in the shocks alone are not large enough to explain the observed swings in oil prices, the main reason for the higher oil price volatility and smaller oil production volatility in more recent times is the substantial decrease in the price elasticities of oil supply and oil demand. Put differently, both curves are so inelastic that even small disturbances generate huge price jumps but only moderate quantity adjustments. Thus, the apparent volatility puzzle is resolved once we take the steepening of the oil supply and demand curves into account which reflects alterations in the supply and demand behavior as a consequence of structural transformations in the oil market. We conjecture that the steepening of both curves since 1986 is the result of an interplay between several features. In particular, the absence of spare oil production capacity and the lack of investment in the oil industry since the mid-eighties already results in a decline of the price elasticities of oil supply and oil demand. The corresponding surge in oil price volatility fosters the deepening of oil futures markets to deal with the increased uncertainty, which by itself further reduces the sensitivity of oil supply and demand to changes in crude oil prices. The exact trigger of this interplay is not clear and deserves additional research. Another question that emerges is whether time variation in volatilities or price elasticities is also an important feature of other types of assets such as exchange rates, equity, house or commodity prices. The advantage of the oil market application is the availability of data for world oil production, a necessary condition to measure price elasticities.

A Data appendix

The world index of industrial production is taken from the United Nations Monthly Bulletin of Statistics. The index numbers are reported on a quarterly basis and span the period 1947Q1 to 2008Q1. The index covers industrial activities in mining and quarrying, manufacturing and electricity, gas and water supply. The index indicates trends in global value added in constant US dollars. The measure of value added is the national accounts concept, which is defined as gross output less the cost of materials, supplies, fuel and electricity consumed and services received. Each series is compiled using the Laspeyres formula (that is, indices are base-weighted arithmetic means). The production series of individual countries are weighted by the value added contribution, generally measured at factor costs, to gross domestic product of the given industry during the base year. For most countries the estimates of value added used as weights are derived from the results of national industrial censuses (census of production) or similar inquiries. A new set of weights is introduced every five years to account for structural changes in the composition of production in industry over time and the index series are chain-linked (by the technique of splicing) at overlapping years. These data in national currencies are converted into US dollars by means of official or free market exchange rates. The weights have been updated the last time in 2000 which is also the base year for the index (2000=100). The index has been recompiled in order to shift the whole series to this reference base. Since the (majority of) national indices have not been adjusted for fluctuations due to seasonal factors, we apply the census X12 ARIMA procedure to the reconstructed series in order to obtain a seasonally adjusted index for the entire period.

World oil production data are provided on a monthly basis by the US Department of Energy (DoE) starting in January 1973. Monthly data for global production of crude oil for the period 1953M4 to 1972M12 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).²⁵ Annual oil production data have been retrieved from World Petroleum (1947-1954). Quarterly data are averages of monthly observations.

²⁵Since 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.

The nominal refiner acquisition cost for imported crude oil is taken from the DoE database.²⁶ Since this series is only available from January 1974, it has been backdated until 1947Q1 with 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 averages over months before extrapolation. Monthly seasonally adjusted data for the US CPI (CPIAUCSL: consumer price index for all urban consumers: all items, index 1982-1984=100) are taken from the FRED database to deflate the nominal refiner acquisition cost for imported crude oil.

B A Bayesian SVAR with time-varying parameters and stochastic volatility

Model setup. The observation equation of our state space model is

$$Y_t = X_t' \theta_t + u_t \quad (16)$$

where Y_t is a 3×1 vector of observations of the dependent variables, X_t is a matrix including lags ($p = 4$) of all the dependent variables and a constant term, and θ_t is a $3(3p+1) \times 1$ vector of states which contains the time-varying parameters. The u_t of the measurement equation are heteroskedastic disturbance terms with zero mean and a time-varying covariance matrix Ω_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:

$$A_t = \begin{bmatrix} 1 & 0 & 0 \\ a_{21,t} & 1 & 0 \\ a_{31,t} & a_{32,t} & 1 \end{bmatrix} \quad H_t = \begin{bmatrix} h_{1,t} & 0 & 0 \\ 0 & h_{2,t} & 0 \\ 0 & 0 & h_{3,t} \end{bmatrix} \quad (17)$$

Let α_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

²⁶The 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.

(2005), the three driving processes of the system are postulated to evolve as follows:

$$\theta_t = \theta_{t-1} + \nu_t \quad \nu_t \sim N(0, Q) \quad (18)$$

$$\alpha_t = \alpha_{t-1} + \zeta_t \quad \zeta_t \sim N(0, S) \quad (19)$$

$$\ln h_{i,t} = \ln h_{i,t-1} + \sigma_i \eta_{i,t} \quad \eta_{i,t} \sim N(0, 1) \quad (20)$$

The time-varying parameters θ_t and α_t are modeled as driftless random walks.²⁷ The elements of the vector of volatilities $h_t = [h_{1,t}, h_{2,t}, h_{3,t}]'$ are assumed to evolve as geometric random walks independent of each other.²⁸ The error terms of the 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:

$$S \equiv \text{Var}(\zeta_t) = \text{Var} \left(\begin{bmatrix} \zeta_{21,t} \\ \zeta_{31,t} \\ \zeta_{32,t} \end{bmatrix} \right) = \begin{bmatrix} S_1 & 0_{1 \times 2} \\ 0_{2 \times 1} & S_2 \end{bmatrix} \quad (21)$$

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

Prior distributions and initial values. The priors 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 σ_i^2 . The priors are calibrated on the point estimates of a constant-coefficient VAR(4) estimated over the period 1947Q2-1958Q2.

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 α_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

²⁷As pointed out by Primiceri (2005), the random walk assumption has the desirable property of focusing on permanent parameter shifts and reducing the number of parameters to be estimated.

²⁸Stochastic 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 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.

matrix $\widehat{\Sigma}_{OLS}$ 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 $\ln h_0 \sim N(\ln \mu_0, 10 \times I_3)$ where μ_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 α_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: $\tilde{\alpha}_0 = [\tilde{\alpha}_{0,21}, \tilde{\alpha}_{0,31}, \tilde{\alpha}_{0,32}]'$. 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 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 (11 years of quarterly data) to guarantee a proper prior. Following Cogley and Sargent (2005), we adopt a relatively conservative prior for the time variation in the parameters setting the scale matrix to $\bar{Q} = (0.01)^2 \cdot \widehat{V}(\widehat{\theta}_{OLS})$ 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. 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 the 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, $\bar{Q} = (0.003)^2 \cdot \widehat{V}(\widehat{\theta}_{OLS})$, which puts as little weight as possible on our prior belief about the drift in θ_t . We have also investigated the opposite assumption by choosing $\bar{Q} = 0.01 \cdot \widehat{V}(\widehat{\theta}_{OLS})$ which postulates a substantial amount of time variation in the parameters. Our results are not affected by different choices for the initial values of this prior. The two 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: $S_1 \sim IW(\bar{S}_1^{-1}, 2)$ and

$S_2 \sim IW(\bar{S}_2^{-1}, 3)$. 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 θ^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 σ_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 θ^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 θ^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 ε_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:

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

Step 3: Drawing volatility states

Conditional on θ^T , A^T , and Y^T , the orthogonalized innovations $\varepsilon_t \equiv A_t (y_t - X_t' \theta_t)$, with $Var(\varepsilon_t) = H_t$, 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, Polson, and Rossi (1994) that draws the volatility states $h_{i,t}$ one at a time.²⁹

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 θ^T , A^T , H^T and Y^T .

We perform 50,000 iterations of the Gibbs sampler but keep only every 10^{th} 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

²⁹As opposed to Primiceri (2005) who uses the method proposed by Kim, Shephard, and Chib (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.

usual set of convergence tests (see Primiceri 2005; Benati and Mumtaz 2007).³⁰ 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 impulse response functions to our three structural shocks. In the spirit of Koop, Pesaran, and Potter (1996) we compute the generalized impulse responses as the difference between two conditional expectations with and without exogenous shocks:

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

where y_{t+k} contains the forecasts of the endogenous variables at horizon k , ω_t represents the current information set and ε_t is a vector of current disturbance terms. 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.³¹ 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.

³⁰The results of these convergence diagnostics are available upon request.

³¹Alternatively, 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.

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 covariance matrix Ω_t at time t . Draw an $N \times N$ matrix, 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 ε_t contains three 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 disturbances during the propagation of the shocks 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 whole set of sign restrictions, i.e. jointly display the effects on the endogenous variables associated with the structural shocks we wish to identify; all others are discarded. Within this loop, we also compute the price elasticities of oil supply and oil demand from all accepted draws of the impulse responses.

We repeat this procedure until 100 iterations fulfil the sign restrictions and then calculate the mean responses of our three endogenous variables over these accepted simulations as well as the average price elasticities. For each point in time, we randomly draw 200 current states of the economy which provide the distribution of mean impulse responses and elasticities 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. The same applies for the price elasticities.

References

- [1] Alquist, R. and L. Kilian (2009), "What Do We Learn from the Price of Crude Oil Futures?", *Journal of Applied Econometrics*, forthcoming.
- [2] Barsky, R.B. and L. Kilian (2002), "Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative", in: *NBER Macroeconomics Annual 2001*, B.S. Bernanke and K. Rogoff (eds.), MIT Press, Cambridge, MA, p 137-183.
- [3] Baumeister, C. and G. Peersman (2008), "Time-Varying Effects of Oil Supply Shocks on the US Economy", *Ghent University Working Paper 515*.
- [4] Benati, L. and H. Mumtaz (2007), "U.S. Evolving Macroeconomic Dynamics: A Structural Investigation", *ECB Working Paper 476*.
- [5] Blanchard, O.J. and J. Simon (2001), "The Long and Large Decline in U.S. Output Volatility", *Brookings Papers on Economic Activity*, 1, p135-174.
- [6] Carter, C.K. and R.P. Kohn (2004), "On Gibbs Sampling for State Space Models", *Biometrika*, 81, p 541-553.
- [7] Cogley, T. and T.J. Sargent (2005), "Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S.", *Review of Economic Dynamics*, 8(April), p 262-302.
- [8] Cooper, J.C.B. (2003), "Price Elasticity of Demand for Crude Oil: Estimates for 23 Countries", *OPEC Review*, 27(1), p 1-8.
- [9] Dvir, E. and K.S. Rogoff (2009), "Three Epochs of Oil", *NBER Working Paper 14927*.
- [10] Fry, R.A. and A.R. Pagan (2007), "Some Issues in Using Sign Restrictions for Identifying Structural VARs", *NCER Working Paper 14*.
- [11] Gately, D. (1984), "A Ten-Year Retrospective: OPEC and the World Oil Market", *Journal of Economic Literature*, XXII (September), p 1100-1114.
- [12] Gately, D. (2004), "OPEC's Incentives for Faster Output Growth", *Energy Journal*, 25(2), p 75-96.
- [13] Geroski, P.A., A.M. Ulph, and D.T. Ulph (1987), "A Model of the Crude Oil Market in which Market Conduct Varies", *Economic Journal*, 97, p 77-86.


- [14] Hamilton, J.D. (2009a), "Understanding Crude Oil Prices", *Energy Journal*, 30(2), p 179-206.
- [15] Hamilton, J.D. (2009b), "Causes and Consequences of the Oil Shock of 2007-08", *Brookings Papers on Economic Activity*, forthcoming.
- [16] Hamilton, J.D. and A.M. Herrera (2004), "Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy: Comment", *Journal of Money, Credit, and Banking*, 36(2), p 265-286.
- [17] Herrera, A.M. and E. Pesavento (2009), "Oil Price Shocks, Systematic Monetary Policy and the 'Great Moderation'", *Macroeconomic Dynamics*, 13(1), p 107-137.
- [18] Hubbard, R.G. (1986), "Supply Shocks and Price Adjustment in the World Oil Market", *Quarterly Journal of Economics*, 101(1), p 85-102.
- [19] Jacquier, E., N.G. Polson, and P. Rossi (1994), "Bayesian Analysis of Stochastic Volatility Models", *Journal of Business and Economic Statistics*, 12, p 371-418.
- [20] Kilian, L. (2008), "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), p 216-240.
- [21] Kilian, L. (2009), "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market", *American Economic Review*, 99(3), p 1053-1069.
- [22] Kim, C.-J. and C. Nelson (1999), *State Space Models with Regime Switching*. Cambridge, Mass., The MIT Press.
- [23] Kim, S., N. Shephard, and S. Chib (1998), "Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models", *Review of Economic Studies*, 65(3), p 361-393.
- [24] Koop, G., M.H. Pesaran, and S.M. Potter (1996), "Impulse Response Analysis in Nonlinear Multivariate Models", *Journal of Econometrics*, 74, p 119-147.
- [25] Krichene, N. (2002), "World Crude Oil and Natural Gas: A Demand and Supply Model", *Energy Economics*, 24, p 557-576.

- [26] Litterman, R.B. (1983), "A Random Walk, Markov Model for the Distribution of Time Series", *Journal of Business and Economic Statistics*, 1(2), p 169-173.
- [27] Mabro, R. (2005), "The International Oil Price Regime: Origins, Rationale and Assessment", *Journal of Economic Literature*, 11(1), p 3-20.
- [28] McConnell, M.M. and G. Pérez-Quirós (2000), "Output Fluctuations in the United States: What Has Changed Since the Early 1980's?", *American Economic Review*, 90(5), p 1464-1476.
- [29] *Oil & Gas Journal*, weekly, various issues since 1954, Tulsa, Oklahoma: Pennwell Corporation.
- [30] Peersman, G. (2005), "What Caused the Early Millennium Slowdown? Evidence Based on Vector Autoregressions", *Journal of Applied Econometrics*, 20, p 185-207.
- [31] Primiceri, G.E. (2005), "Time Varying Structural Vector Autoregressions and Monetary Policy", *Review of Economic Studies*, 72, p 821-852.
- [32] Regnier, E. (2007), "Oil and Energy Price Volatility", *Energy Economics*, 29, p 405-427.
- [33] Ryan, D.L. and A. Plourde (2002), "Smaller and Smaller? The Price Responsiveness of Nontransport Oil Demand", *Quarterly Review of Economics and Finance*, 42, p 285-317.
- [34] Smith, J.L. (2005), "Inscrutable OPEC? Behavioral Tests of the Cartel Hypothesis", *Energy Journal*, 26(1), p 51-82.
- [35] Smith, J.L. (2009), "World Oil: Market or Mayhem?", *Journal of Economic Perspectives*, 23(3), p 145-164.
- [36] Stock, J.H. and M.W. Watson (2003), "Has the Business Cycle Changed and Why?", in: *NBER Macroeconomics Annual 2002*, M. Gertler and K. Rogoff (eds.), MIT Press, Cambridge, MA, p 159-218.
- [37] Uhlig, H. (2005), "What are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure", *Journal of Monetary Economics*, 52(2), p 381-419.

[38] United Nations. *Monthly Bulletin of Statistics*. 1947-2008.

[39] Yang, C.W., M.J. Hwang, and B.N. Huang (2002), "An Analysis of Factors Affecting Price Volatility of the US Oil Market", *Energy Economics*, 24, p 107-119.

PANEL A



PANEL B




Figure 1: Volatility measures.

Panel A: Volatility of the nominal refiner acquisition cost of imported crude oil and of world oil production.

Panel B: Median time-varying unconditional standard deviation of the nominal refiner acquisition cost of imported crude oil and world oil production together with 16th and 84th percentiles of the posterior distribution.


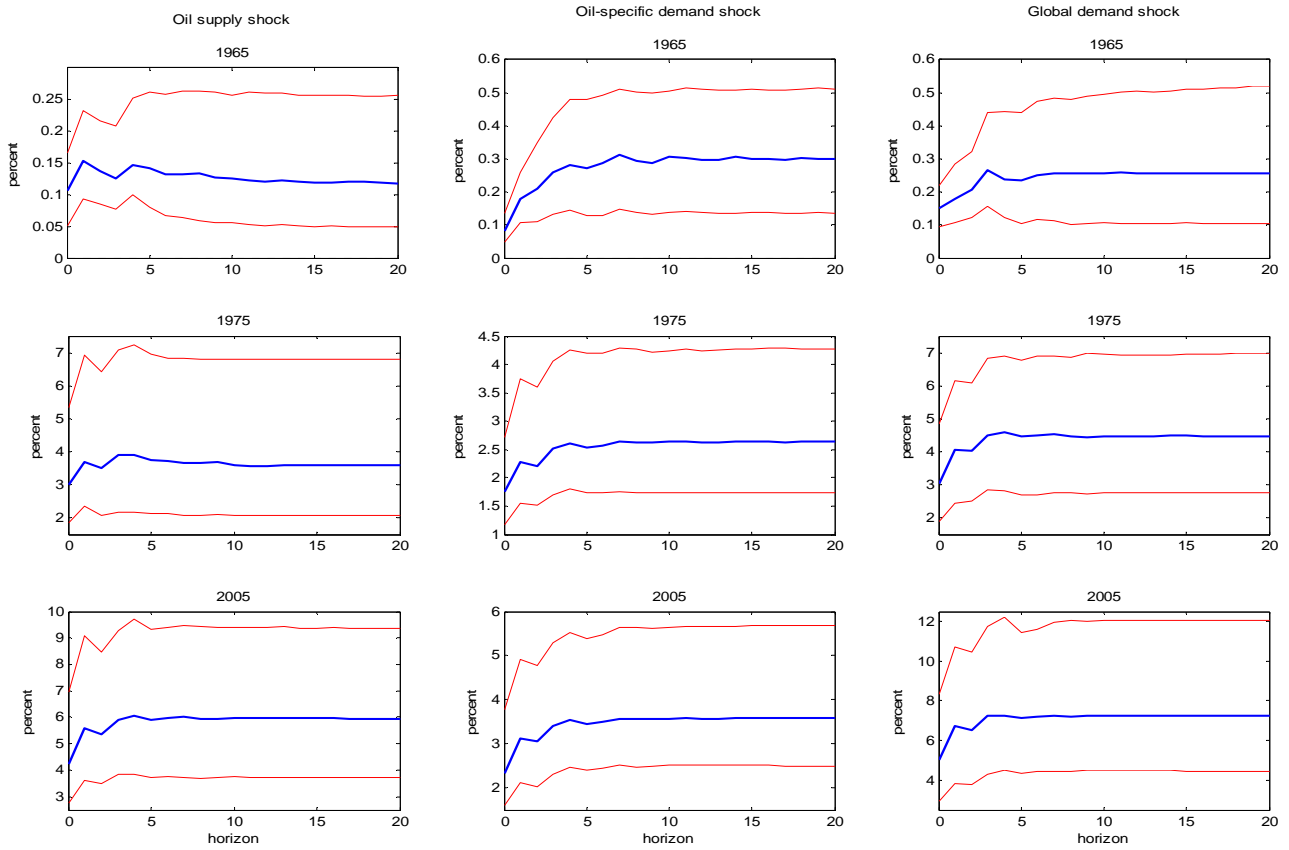


Figure 2: Time-varying median impulse response functions of world oil production and the nominal price of crude oil after an oil supply shock (first row), an oil-specific demand shock (second row) and a global demand shock (third row).

PANEL A



PANEL B





Figure 3: Evolution of the speed of adjustment.

Panel A: Median impulse responses of the price of crude oil with 16th and 84th percentiles after all three shocks at selected dates – 1965Q1, 1975Q1 and 2005Q1.

Panel B: Time-varying median effect of all three shocks on the price of crude oil on impact (black lines) and after 4 quarters (blue lines) with 16th and 84th percentiles.

PANEL A



PANEL B

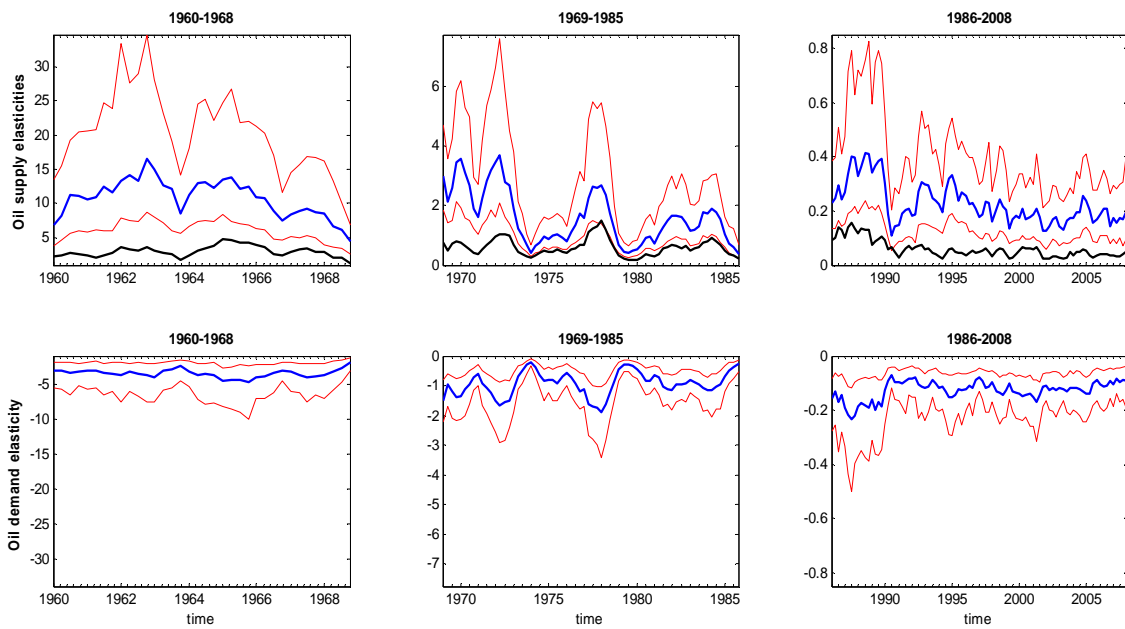


Figure 4: Short-run price elasticities of oil supply and oil demand.

Panel A: Evolution of the median oil supply and demand elasticities on impact. The supply elasticities are derived with oil-specific demand shocks (blue line) and with global demand shocks (black line).

Panel B: Evolution of the median oil supply and demand elasticities on impact over three subperiods with 16th and 84th percentiles of the posterior distribution. Supply elasticities are derived with oil-specific demand shocks (blue line and quantiles) and global demand shocks (black line).





Figure 5: Magnitude of the median structural shocks on impact together with 16th and 84th percentiles derived with back-of-the-envelope calculation. The oil supply shock is derived with elasticities obtained once with oil-specific demand shocks (black line and percentiles) and once with the global demand shocks (blue line).

PANEL A



PANEL B




Figure 6: Capacity constraints and investments in the oil sector.

Panel A: Global capacity utilization rates in crude oil production by year.

Panel B: Monthly worldwide oil rig counts.

Notes: Estimates of global spare oil production capacity are obtained from the IMF World Economic Outlook (August 2006) and the DoE Short-Term Energy Outlook (January 2009). 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. Data on worldwide rig counts are obtained from Baker & Hughes Inc.