Macroeconomic Uncertainty and the Impact of Oil Shocks

Ine Van Robays*

February 2013

Abstract

This paper evaluates whether macroeconomic uncertainty changes the responsiveness of oil prices to shocks in oil demand and supply. Using a threshold VAR model, we identify different regimes of uncertainty in which we estimate the effects of structural oil demand and supply shocks. The results show that higher macroeconomic uncertainty, as measured by higher world industrial production volatility, significantly increases the oil price response to oil demand and supply changes. This implies a lower price elasticity of oil demand and supply in uncertain times, or in other words, that both oil curves become steeper when uncertainty is high. The difference in oil demand elasticities is both statistically and economically meaningful. Accordingly, varying uncertainty about the macroeconomy can explain why the oil price elasticity, and therefore also oil price volatility, changes over time. Also the impact of oil shocks on economic activity appears to be significantly stronger in uncertain times. Among other possible channels, this effect might be explained by irreversibility of investment under uncertainty.

JEL classification: E31, E32, Q41, Q43

Keywords: Oil prices, uncertainty, price elasticity, threshold VAR, sign restrictions

*European Central Bank, Kaiserstrasse 29, 60311 Frankfurt am Main, Germany, ine.van-robays@ecb.europa.eu. I am grateful to Nathan Balke for providing his computer codes on which my estimations are based, and to Ron Alquist, Hilde Bjørnland, Thierry Bracke, Julio Carrillo, Selien De Schryder, Gerdie Everaert, Gert Peersman, Joris Wauters, two anonymous referees for the ECB Working Paper Series and the participants at the CESifo conference for useful comments and suggestions. I acknowledge financial support from the Interuniversity Attraction Poles Programme - Belgian Science Policy [Contract No. P6/7] and Belgian National Science Foundation.

1 Introduction

The remarkable increase in oil price volatility over the past decade sparked an intensive debate about its driving factors. Many studies argue that the stronger oil price fluctuations can be explained by sharp movements in fundamental oil supply and demand-side factors (Baumeister and Peersman 2012a,b; Hamilton 2009, Kilian and Murphy 2013). Others claim that changes in fundamentals are not sufficient to fully explain recent oil price volatility, and argue that also financial speculation played some role (Lombardi and Van Robays 2011, Tang and Xiong 2011, Singleton 2012). A factor which has been overlooked in this debate is the role of *uncertainty*. In periods of high uncertainty about the macroeconomy, oil price volatility is typically very strong. It is well documented that increased uncertainty can influence the decision behavior of economic agents (Bernanke 1983, Pindyck 1991, Litzenberger and Rabinowitz 1995, Bloom et al. 2007); higher uncertainty causes a delay in the production or consumption decision, thereby lowering the quantity response and increasing the price impact of shocks, which leads to a higher price volatility. Similarly, uncertainty could increase the responsiveness of oil prices to oil demand and supply shocks, and thereby increase oil price volatility.

In this paper, we evaluate whether the impact of oil supply and demand shocks differs in times of higher uncertainty, and more specifically, whether macroeconomic uncertainty lowers the price elasticity of oil demand and supply. We define macroeconomic uncertainty as volatility in world industrial production growth. Using a monthly threshold vector autoregressive (TVAR) model that we estimate over the period 1986:01-2011:07, we endogenously identify high and low uncertainty regimes based on our measure of macroeconomic volatility crossing an estimated threshold. Conditional on being in a particular regime, we quantify the impact of different types of oil shocks on oil prices, oil production and economic activity. We identify three types of oil shocks using sign restrictions; oil supply shocks, oil demand shocks driven by economic activity, and oil-specific demand shocks, similar to Peersman and Van Robays (2009, 2012), Baumeister, Peersman and Van Robays (2010), Baumeister and Peersman (2012b) and Kilian and Murphy (2012). The aim of this paper is to establish some stylized facts on the interaction between uncertainty and oil price volatility that seem worthwhile exploring further in general equilibrium models.

Our results show that the impact of oil demand and supply shocks tends to differ substantially when macroeconomic uncertainty is high. Oil demand and supply shocks have a significantly stronger effect on oil prices for a given change in oil production, implying that the price elasticity of oil demand and supply is lower in the high uncertainty regime. In other words, the oil demand and oil supply curve become steeper in uncertain times. We estimate the impact oil demand elasticity to decline from a range of -0.52 to -0.15 when uncertainty is low, to -0.36 to -0.11 when uncertainty is high. The oil supply elasticity, on the other hand, drops from a range of 0.21 to 0.03 to a number in between 0.15 and 0.02 conditional on a highly uncertain environment. Although there is some overlap across the regimes, the difference in estimated elasticity across regimes is statistically significant. The difference is also economically significant, as the price impact of a similar oil shock might even double when it hits the economy in uncertain times. Hence, we show that different levels of macroeconomic uncertainty over time can explain time variation in the price elasticity of oil, and therefore in oil price volatility. Hamilton (2009), Kahn (2009) and Baumeister and Peersman (2012b) argue that a lower price elasticity could explain why fundamental oil supply and demand shocks impacted more strongly on oil prices over the last decade, and we empirically demonstrate that this could have been the case because of higher uncertainty. Moreover, not only oil prices and oil production react differently, but also economic activity reacts more aggressively to oil shocks when macroeconomic volatility is already high.

As far as we are aware, this is the first paper which endogenously explains why oil price elasticities, and hence also oil price volatility, changes over time. Using a time-varying VAR model, Baumeister and Peersman (2012b) explain time variation in oil price volatility by showing that the elasticity of oil demand and supply varies substantially over time. However, in contrast to our approach, their model does not allow to empirically test which factors are responsible for explaining the time-variation in oil price elasticities. We built upon their findings and show that variation in the elasticity of both oil demand and supply can be explained by changes in macroeconomic uncertainty. Several papers have touched upon the relationship between uncertainty and oil prices, but mostly they focus on uncertainty with respect to the oil price itself, i.e. oil price volatility instead of macroeconomic volatility more generally (Bredin et al. 2011, Elder and Serletis 2010, Ferderer 1996, Kellogg 2010, Lee, Ni and Ratti 1995, Pindyck 2004).[1] The main aim

[1] Two exceptions to this are Pindyck (1980) and Litzenberger and Rabinowitz (1995), although their focus is different. Pindyck (1980) concentrates on the theoretical effect of demand and oil reserves uncertainty on expected oil price behavior, and Litzenberger and Rabinowitz (1995) focus on explaining backwardation in oil futures markets.

of this paper is not to explain the economic effects of higher oil price volatility, but to understand better why oil prices behave more volatile in some periods compared to others.

The remainder of this paper is organized as follows. In the next section, we provide some intuition and evidence on why uncertainty could matter for the impact of oil shocks. In Section 3, we describe the threshold VAR model and its specification, test for threshold effects and explain the identification strategy. The empirical results are discussed in Section 4 and Section 5 briefly evaluates the robustness of the results. Section 6 concludes.

2 How can uncertainty affect the oil market?

A lower price elasticity of oil demand and supply during uncertain economic times means that shocks hitting the oil market generate larger responses in prices but smaller responses in quantities compared to more certain times. In this section, we discuss several possible ways in which macroeconomic uncertainty can negatively impact on the price elasticity of oil demand and supply. Most of these channels are well documented in the literature. The explanations are not mutually exclusive and mainly serve to provide intuition behind the empirical results of this paper. Exploring the relevance of these transmission channels is an interesting and promising avenue for further research.

First, both oil demand and oil supply could be less responsive because of an *option value to wait*. Under the condition that the action to be decided on is irreversible, uncertainty creates an option value to wait through which investors are willing to forego current returns in order to gain from more information in the future. In other words, uncertainty over future demand reduces current investment. There exists a large literature providing both theoretical and empirical evidence on this link. Bernanke (1983) relies on this concept to explain cyclical fluctuations in investment, and in more recent work, Bloom et al. (2007) and Bloom (2009) confirm that firms delay investment and hiring decisions because of higher uncertainty about future demand.[2] This theory or irreversible investment under uncertainty has also been applied to the oil market. Using a duration model of irreversible investment, Favero, Pesaran and Sharma (1994) show both theoretically and empirically that increased oil price uncertainty causes oil firms to postpone production. As oil prices are endogenous to the macroeconomy, macroeconomic uncertainty will lead to oil price

[2] Other examples are Arrow (1968), Henry (1974a,b), Pindyck (1991), Brennan and Schwartz (1985), Majd and Pindyck (1987), Elder and Serletis (2010) and Bredin et al. (2011).

uncertainty, and indirectly cause a decline in the responsiveness of oil production.[3] Subsequently, this option value to wait effect would lower the elasticity of *oil supply.* Other studies providing evidence on this channel are for example Kellogg (2010). Guiso and Parigi (1999) find the effect of demand uncertainty on the responsiveness of investment to be stronger if it is harder to reverse investment decisions and if the firm has more market power, which is characteristic to oil firms. Similarly, concerning the elasticity of *oil demand*, the responsiveness of oil demand under uncertainty could be lower as oil consumers prefer to wait with reducing their demand following an oil supply shock that pushes oil prices upwards. In addition, uncertainty could reduce the tendency of oil consumers to substitute oil for other energy products, or at least delay substitution until there is more certainty about the effect of the oil shock.

Second, *futures markets* might also play a role in explaining why oil demand and supply elasticities vary over time. Baumeister and Peersman (2012b) note that hedging against oil price movements could weaken the responsiveness of oil demand and supply. Accordingly, if higher macroeconomic uncertainty leads to an increased use of futures contracts, which is plausible given that futures markets exist to transfer risks, it could cause the oil price elasticity of demand and supply to decline.

Third, the oil supply elasticity could decline during uncertain periods because oil producers prefer to *leave oil reserves below the ground* when uncertainty rises. In a two period equilibrium model, Litzenberger and Rabinowitz (1995) show that uncertainty increases the value of oil reserves below the ground for any level of the extraction cost. As oil producers will not extract oil as long as the net value of oil below the ground is higher than that above the ground, an increase in uncertainty will lower the extraction of oil. Litzenberger and Rabinowitz (1995) also find empirical support for this.

Finally, uncertainty could also affect price setting in the oil spot and futures markets without the need for immediate oil demand and supply adjustments. Singleton (2012) argues that *heterogeneous beliefs* about public information concerning the future course of economic events can induce higher price volatility, price drifts and even booms and

[3] Higher macroeconomic volatility can lead to stronger oil price volatility for a fixed price elasticity of oil demand and supply, in case of stronger macroeconomic shocks for example. It is however important to stress that the focus of this paper is different; we argue that higher macroeconomic volatility, via its effect on oil price volatility, might lead to a change in the oil price elasticity. In other words, a single oil demand or supply shock, independent of its magnitude, will have a stronger impact on the oil price and a lower impact on oil production when it occurs in uncertain times, as the oil price elasticity is lower.

busts in prices. The release of new information about oil supply and demand can have a large effect on prices as investors learn about the economic environment. Although Singleton (2012) uses these arguments to explain the role of financial flows on oil prices, they could also help in understanding why in times of higher macroeconomic uncertainty, when investors' beliefs typically diverge more than in normal times, shocks to oil demand and supply have a larger impact response on prices.

3 Model and identification

In this section, we discuss the methodology in more detail. First we describe the properties of the threshold VAR model and the data. In the second part, we test for the presence of threshold effects and identify the different uncertainty regimes. After that, we discuss the sign restriction identification of the oil shocks, and finally, we explain how we circumvent the problem of endogeneity between uncertainty and oil shocks.

3.1 Threshold VAR model and data

In this paper, we want to analyze whether uncertainty changes the way in which shocks to oil demand and supply affect oil prices, and thereby explains why oil price volatility changes over time. To address this issue empirically, we need a model that allows for non-linear effects depending on the state of the economy and that in addition can directly test whether this non-linearity is due to changes in uncertainty. Therefore, we rely on a structural threshold vector autoregressive (TVAR) model. The threshold model is attributed to Tong (1978) and has been extensively used afterwards in different fields besides oil, see Hansen (2011) for an overview.

The TVAR model has several advantages compared to other empirical models in addressing our research question. First, it enables us to identify different regimes endogenously. Assuming strict exogeneity between macroeconomic uncertainty and oil shocks might not be realistic given the bidirectional link between oil prices and the macroeconomy. The regime determination occurs endogenously as the transition variable that determines the economic regime, the so-called threshold variable, is a variable included in the model. Second, the interpretation of the differences across regimes is relatively straightforward in a TVAR model. This is because the threshold variable can be chosen to be any of the variables in the model, as long as the model does not reject non-linearity with respect to it. Given our research question, we choose the threshold variable to be a function of macroeconomic uncertainty. Under the assumption that the shocks in the model can only trigger a regime change with a lag, the differences in the responses across regimes can then be interpreted as due to macroeconomic uncertainty.[4] In Markov-switching models, for example, interpretation is more difficult as the transition variable is typically not observed, or as the regimes are determined exogenously. Similar to Balke (2000), we estimate a two-regime TVAR model of the following form:

$$Y_t = \mu_1 + A^1 Y_t + B^1 (L) Y_{t-1} + \left(\mu_2 + A^2 Y_t + B^2 (L) Y_{t-1}\right) I_t(c_{t-d} \geq \gamma) + u_t$$

The vector of endogenous variables Y_t captures the global dynamics in the oil spot market, i.e. world oil production (Q_{oil}), the price of crude oil expressed in US dollars (P_{oil}), a measure of world economic activity (Y_w) and oil inventories (I_{oil}). To model different uncertainty regimes, we also add a measure of macroeconomic uncertainty denoted by U. The variable c_{t-d} is the threshold variable and $I_t(.)$ is an indicator function that takes value one when the d-lagged value of the threshold variable is higher or equal to the estimated threshold γ, and zero otherwise. This indicator function thus determines the regimes based on the value of c_{t-d} relative to γ. As the threshold variable c_{t-d} is a function of macroeconomic uncertainty and subsequently a function of an endogenous variable in the TVAR model, shocks to the oil market as well as to macroeconomic uncertainty are allowed to determine whether the economy is in a high or low uncertainty regime. As will be explained below, the lag on the threshold variable is needed to avoid an endogeneity bias in the results, and to make sure that the differences in the results across regimes are due to uncertainty. μ is a vector of constants, $B(L)$ is a matrix polynomial in the lag operator L and A is the contemporaneous impact matrix of the vector of orthogonalized error terms u_t. The TVAR model allows for non-linearity in the effects across regimes as each regime has different autoregressive matrices. If $I_t = 0$, the dynamics of the system are given by μ_1, A^1 and $B^1(L)$, and if $I_t = 1$, the relevant coefficients are $\mu_1 + \mu_2$, $A^1 + A^2$ and $B^1(L) + B^2(L)$. Note that the contemporaneous impact of the shocks is allowed to vary, which is crucial for our analysis of the price elasticities on impact.

The TVAR model is estimated using monthly data over the period 1986:01-2011:07. We choose 1986 as our starting point for two reasons. First, Baumeister and Peersman

[4] We discuss this assumption in detail later on.

(2012a,b) document a structural break in the oil price elasticities around the mid-1980s, after which both the oil demand and oil supply elasticity became substantially smaller. This decline is typically explained by a reduction in spare capacity which reduces the responsiveness of oil supply, and a more limited scope for substitution away from oil which reduces the responsiveness of oil demand.[5] Second, the Great Moderation in the mid-1980s caused a downward shift in the level of uncertainty as macroeconomic volatility declined, which implies a downward shift of the threshold in our model. Including these two events in our sample period could therefore significantly bias the identification of the regimes and the estimation results.[6]

The oil price is the nominal refiner acquisition cost of imported crude oil, which has extensively been used in the literature as the best proxy for the free market global price of imported crude oil. We use the nominal price of oil because theoretically, this should allow for a better identification of the different types of oil shocks.[7] We proxy global economic activity by the OECD measure of global industrial production, which covers the OECD countries and the six major non-OECD economies, including e.g. China and India. Following Kilian and Murphy (2013), we proxy global crude oil inventories as total US crude oil inventories scaled by the ratio of OECD petroleum stocks over US petroleum stocks. Global macroeconomic uncertainty is proxied by the volatility of world industrial production growth, which is modelled as a GARCH(1,1) process.[8] To ensure robustness of our findings, we construct two additional measures of uncertainty. Following Baum and Wan (2010), the first alternative measure is the conditional variance of US GDP production growth. We generate a monthly GDP series by interpolating quarterly GDP using industrial production based on the Chow-Lin procedure, after which we model the

[5] Note that, as Baumeister and Peersman (2012a,b) document a break in the oil price elastcities around the mid 1980s, we try to explain the variability in the price elasticities that still occured after that date.

[6] The fact that macroeconomic uncertainty decreased around the same time that the price elasticity of oil declined does not contradict our results, i.e. increased uncertainty lowers the price elasticity of oil. This is because the break in the oil price elastcicity around the mid-1980s is found to be *exogenous* to the macroeconomy, see Baumeister and Peersman (2012a,b).

[7] For example, when we would deflate the nominal price of oil by the US CPI, it could be that a domestic positive demand shock to the US could wrongly be identified as a negative oil supply shock because real oil prices fall, and oil production and economic activity do not decline (see Section 3.3 for more details on the shock identification).The results are however very similar then when using real oil prices.

[8] The GARCH(1,1) gives the best specification for modelling the conditional variance according to various information criteria. We estimated the conditional variance over the period 1985-2011 to avoid a possible bias due to the Great Moderation.

conditional variance as a GARCH(1,1) process. As a second alternative, we consider the Chicago Board of Exchange VXO stock market volatility measure. The VXO index is based on a hypothetical at the money S&P100 option, and is the measure of uncertainty used by Bloom (2009). We constructed a monthly series of the VXO index by taking monthly averages of the daily closing price. As noted by Baum and Wan (2010), these different measures capture different types of uncertainty. The measure based on GDP growth is designed to reflect the overall uncertainty of the macroeconomic environment, whereas the measure based on industrial production disregards uncertainty about the service sector. The VXO stock market volatility measure is more closely related to financial market uncertainty. Note that the first two measures of uncertainty are backward looking as these are based on GARCH models, whereas the measure based on the VXO index is essentially forward looking.[9] As data on world GDP growth is not available, we have a trade-off between modeling volatility on a global scale using industrial production (and hence excluding the service sector), or using the volatility of total economic activity but then at the level of the US. Given that oil prices are set at a global level, we choose the global industrial production measure as our preferred indicator of macroeconomic uncertainty. The results show that the conclusions hold for the other measures of uncertainty as well.

Based on the conventional lag length criteria, we include four lags of the endogenous variables. All variables are included in stationary form, by taking the first difference of the natural logarithm of all endogenous variables besides macroeconomic uncertainty which is a volatility measure. The results are robust to different specifications of the variables and the structural TVAR model, see Section 5 for a more detailed discussion.

3.2 Test for threshold effects and identification of regimes

Before testing whether the model is indeed non-linear depending on the level of uncertainty, we have to decide on the exact specification of the threshold variable, i.e. macroeconomic uncertainty. In the TVAR literature, two assumptions are typically made to prevent potential problems of endogeneity. First, the threshold variable is assumed to have a certain delay in determining the regimes. Under this assumption, oil shocks cannot trigger a change in the uncertainty regimes on impact. We discuss the restrictiveness of this

[9] More specifically, a GARCH(1,1) model specifies the variance as a function of a lagged squared error term and the lagged variance: $\sigma_t^2 = c + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2$.

assumption later on. In our case, however, as we model uncertainty as a GARCH process, shocks can, by construction, only affect uncertainty with a delay. Hence, we assume no additional delay in the TVAR model. Second, the threshold variable is typically modeled as a moving average process depending on the persistence of the series (Balke 2000). As the measures of uncertainty that we employ are highly volatile, we model the threshold variable as a moving average process of order three to allow for some persistence in the uncertainty regimes, which corresponds to the average volatility of the past quarter.

To test for the significance of threshold effects, we use the approach described in Balke (2000). If the threshold value γ was known, the test of linearity under the null hypothesis against the presence of non-linear threshold behavior would simply come down to testing whether $\mu_2 = A^2 = B^2(L) = 0$. As this is not the case, we have to rely on non-standard inference. A commonly used approach is to estimate the model for each possible value of the threshold variable using least squares. The range of possible thresholds is trimmed by a certain percentage to allow for sufficient observations in each regime. As suggested by Hansen (1999), we choose a trimming parameter of 10 percent. Conditional on each threshold, we calculate the Wald statistic that evaluates the hypothesis of equality between the regimes. Three different summary test statistics are generated: the maximum Wald statistic (sup-Wald) over all threshold values, the average Wald statistic (avg-Wald) and a statistic calculated as a function of the sum of the exponential Wald statistics for all possible thresholds (exp-Wald). For the reason that the distribution of these test statistics is non-standard, we rely on the bootstrap technique proposed by Hansen (1996) to simulate the unknown asymptotic distributions. This enables us to derive the p-values associated with the test statistics and hence to evaluate the significance of the non-linear threshold effects. The estimated threshold value is the one that maximizes the log determinant of the variance-covariance matrix of residuals.

Table 1 shows the threshold test results for the different measures of uncertainty and some summary statistics on the identified regimes. There is strong evidence for significant threshold effects for all measures of uncertainty according to all the Wald summary test statistics. The threshold based on the preferred measure of macroeconomic uncertainty using world industrial production growth is estimated to be 0.3512, which splits the sample into high and low uncertainty regimes that represent respectively 17 and 83 percent of all observations. To put this into perspective, Panel A of Figure 1 illustrates the threshold variable, the estimated threshold and the identified regimes for this measure of uncertainty.

The shaded areas correspond to the high uncertainty states, when the threshold variable surpasses the threshold. Using world industrial production growth volatility, the main periods of higher global uncertainty are identified to be the slowdown in GDP growth across most industrialized countries in 2001, the 9/11 Terrorist Attacks at the end of 2001, and the financial crisis that hit the global economy in 2008. Global uncertainty was already elevated before the financial crisis hit due to a recession in the US and a decline in economic growth in other major industrialized countries. More recently, concerns about the sovereign debt crisis in the euro area might explain why uncertainty is again higher. When comparing Panel A with Panel B and C in Figure 1, it is clear that the different measures of uncertainty correspond to somewhat different definitions of uncertainty. The US GDP volatility measure is more closely related to US economic downturns in addition to global uncertainty periods. In general, it succeeds well in capturing the periods that are typically regarded as uncertain, see e.g. Bloom (2009).[10] The periods identified to be highly uncertain, which are not captured by the global measure, are Black Monday at the end of 1987, the US recession in the early 1990s, the Russian financial crisis in 1998, and the US recession of the early 2000s. On the other hand, the VXO measure captures financial market uncertainty more closely.[11]

3.3 Identifying oil shocks using sign restrictions

3.3.1 Sign restriction identification

In the reduced-form VAR model, we face the problem that the contemporaneous errors could be correlated. In order to make the shocks orthogonal and thereby economically interpretable, we need to impose structure on the model to identify the different shocks. Given that we only want to evaluate whether uncertainty acts as a reinforcer of oil shocks, we are only interested in identifying the oil shocks.

The oil literature has increasingly recognized that different factors can drive oil price

[10] Bloom (2009) identifies 17 volatility shock events that substantially increased uncertainty, which he uses as 'arguably exogenous' shocks to empirically evaluate the effect of uncertainty shocks. Most of these shocks are caused by economic events, war or terrorism.

[11] Using the VXO index, high uncertainty is concentrated around the Black Monday event, the Russian and Long-Term Capital Management (LTCM) default, 9/11 Terrorist attack, the Enron and Worldcom accounting scandals, Gulf War II and the financial crisis. The working paper version of Bloom (2009) provides more details on these events.

movements, and that the economic effects of those shocks crucially depend on the underlying source of the oil price change (e.g. Barsky and Kilian 2004, Kilian 2009, Peersman and Van Robays 2009, 2012). Not accounting for the driving force behind the oil price increase could therefore significantly bias the results. It is also crucial to separate oil demand from supply shocks when evaluating the role of uncertainty, as uncertainty can affect the behavior of oil producers and consumers differently, which implies a different impact on the price elasticity of oil supply and demand. We will identify three different types of oil shocks using sign restrictions: oil supply shocks, oil demand shocks driven by global economic activity and oil-specific demand shocks. Sign restriction identification is particularly useful as we do not have to rely on zero impact restrictions to separate oil demand and supply shocks. Calculating the short-run oil demand elasticity is not possible if we assume that oil supply does not respond to oil demand shocks on impact, implying a zero price elasticity of supply, see the assumption made in Kilian (2009). Similar to Peersman and Van Robays (2009, 2012), Baumeister, Peersman and Van Robays (2010) and Baumeister and Peersman (2012b), we identify the oil shocks by relying on the following set of sign restrictions:

STRUCTURAL SHOCKS	Q_{oil}	P_{oil}	Y_w	I_{oil}
Oil supply	≤ 0	≥ 0	≤ 0	
Oil demand driven by economic activity	≥ 0	≥ 0	≥ 0	
Oil-specific demand	≥ 0	≥ 0	≤ 0	

The sign restrictions are derived from a simple supply-demand scheme of the oil market. An *oil supply shock* is an exogenous shift of the oil supply curve to the left and therefore moves oil prices and production in opposite directions. Production disruptions caused by military conflicts in the Middle-East are natural examples. As oil prices are higher, global industrial production will not increase following this supply shock. In contrast, shocks on the demand side of the oil market will result in a shift of oil production and oil prices in the same direction. On the one hand, demand for oil can endogenously increase because of changes in macroeconomic activity. A change in the demand for commodities from emerging economies like China or India for example, will shift world economic activity, oil prices and oil production in the same direction. We define such a shock as an *oil demand shock driven by economic activity*. On the other hand, oil demand can also vary for reasons not related to economic activity. We label these shocks as *oil-specific demand shocks*. Shocks to expected net oil demand in the future, which increases oil inventory demand as a precaution, and oil-gas substitution shocks are two examples. In contrast to demand shocks driven by economic activity, oil-specific demand shocks do not have a positive effect on global economic activity as oil prices are higher.

Concerning our identification, we differ from Kilian and Murphy (2013) in two aspects. First, in contrast to Kilian and Murphy (2013), we do not impose sign restrictions on oil inventories. This is because we do not need to impose these additional restrictions to disentangle the different types of oil shocks we are interested in. Also, by restricting inventories to increase following the oil-specific demand shock, as done in Kilian and Murphy (2013), the oil-specific demand shocks that do not lead to an *immediate* increase in inventories will end up in the residuals, although they might be relevant.[12] Instead, we prefer to have a complete representation of all possible oil shocks, which we define in three broad categories as just described: oil supply, oil demand driven by economic activity and oil-specific demand. Moreover, it is well-known that the available inventory data is a poor proxy of global inventory levels, in particular due to the poor data availability for non-OECD inventories (IEA 2009). Second, we differ from Kilian and Murphy (2012, 2013) by not imposing bounds on the estimated price elasticities to identify the shocks. This is for the reason that they assume the price elasticities of oil demand and supply to be fixed over time. Baumeister and Peersman (2012b) show that this clearly is not the case, which we confirm in this paper by finding that the price elasticity of oil demand and supply is actually dependent on the level of uncertainty.

3.3.2 Estimation procedure

The estimated threshold value splits the sample period into two subsamples, corresponding to high and low uncertainty states. Conditional upon these two subsamples, we generate two sets of impulse response functions, one estimating the effects in the high uncertainty state and the other in the low uncertainty state. In other words, we analyze the change in impact of oil shocks on the oil price elasticity by constructing *conditional* impulse response functions, i.e. conditional upon a specific uncertainty regime. By doing this, we assume

[12] Examples of these types of shocks are subtitution shocks between oil and other forms of energy or financial shocks in the oil futures market that do not lead to an immediate increase in oil inventories, but do affect the oil spot price. Lombardi and Van Robays (2012) show the relevance of these latter types of shocks, and in line with this, Sockin and Xiong (2013) argue that restricting inventories to identify a financial speculation shock is not desirable.

that the impact is linear within a regime, but that the size and persistence of the responses to similar oil shocks can differ across regimes. In the TVAR literature, the effects of shocks are often evaluated using so-called 'generalized impulse response functions', which allow shocks to cause a switch in regime over the duration of the response.[13] A possible caveat of not allowing shocks to cause switches in regimes during the response is that the impulse response functions might differ slightly beyond the impact response. In this paper, however, the main question is whether the *impact* price elasticities differ depending on the level of uncertainty. Therefore, this assumption does not make any difference for our main findings. Given that we estimate the effects of oil shocks conditional upon the uncertainty regime, we identify the structural oil shocks in a model that only includes oil prices, oil production, world economic activity and oil inventories to save some degrees of freedom.[14]

We estimate the impulse response functions by following the sign restriction procedure of Peersman (2005) and Uhlig (2005), which we apply to both subsample VAR models simultaneously. We use a Bayesian approach for estimation and inference. Our prior and posterior distributions of the reduced form VAR belong to the Normal-Wishart family. To draw the 'candidate truths' from the posterior, we take a joint draw from the unrestricted Normal-Wishart posterior for the VAR parameters as well as a random possible decomposition of the variance-covariance matrix, which allows us to construct impulse response functions. If the impulse response functions from a particular draw satisfy the imposed sign conditions, the draw is kept. Otherwise, the draw is rejected by giving it a zero prior weight. We simultaneously rotate the model conditional upon high and low uncertainty, and restrict both sets of impulse response functions to satisfy the sign restrictions of all three shocks. To improve identification of the shocks, we impose the sign conditions to hold for the first three months, see Paustian (2007). A total of 1000 'successful' draws from the posterior are then used to construct the 68 percent posterior range of possible impulse responses. For each simultaneous rotation of the two regime-dependent models, we

[13] See for example the working paper version of Calza and Sousa (2006) for more details, as they construct both the conditional and the generalized impulse response functions.

[14] Constructing generalized impulse response functions when using sign restrictions instead of recursive identification also proves to be quite difficult. In order to model the transition between the regimes following a structural shock, it is necessary that the shocks come from the same model, which is not the case when using sign restrictions. Up to our knowledge, only Candelon and Lieb (2011) have used TVAR models in combination with sign restrictions, and they make the same assumption as we do here.

also generate the *difference* in estimated impulse response functions across regimes, which allows us to also calculate the 68 percent posterior range of the difference in estimated effects across regimes.

3.4 The problem of endogeneity between uncertainty and oil shocks

As mentioned above, when identifying the uncertainty regimes, the TVAR model allows macroeconomic uncertainty to endogenously respond to oil shocks. It is well known that oil shocks can lower economic activity and cause recessions (e.g. Hamilton 1983, 2009; Bjørnland 2000; Peersman and Van Robays 2009, 2012). Assuming strict exogeneity between macroeconomic uncertainty and oil prices is therefore not realistic. Because of this link between oil prices and uncertainty, however, the results might be subject to an endogeneity bias. For example, suppose that only large oil shocks would cause the economy to switch from a low to a high uncertainty state, then differences in the effects of oil shocks across regimes might be merely driven by differences in shock size. Therefore, the issue of potential endogeneity deserves some discussion.

There are several reasons to believe that an endogeneity bias is negligible if not non-existent, and that the differences in elasticities we find are due to uncertainty. First and foremost, the threshold variable is assumed to only switch regimes with a delay of one period. This implies that we only evaluate the oil shocks that occur *within* a certain regime. The differences in the effects of oil shocks can, for example, not be driven by large oil shocks that trigger the economy to immediately go into a high uncertainty state, and thereby bias the strength of the oil prices response. Second, the threshold variable is defined as a moving average process of macroeconomic uncertainty which requires some persistence in the increase of macroeconomic uncertainty before shocks can trigger a regime switch. These two assumptions are exactly imposed to delete the endogeneity bias, and are commonly used in the TVAR literature, see Balke (2000) for example.

There are good reasons to believe these assumptions are not restrictive. When estimating the model over the total sample, the different types of structural oil shocks do not significantly affect uncertainty on impact, which is exactly the assumption we impose. The conditional variance decompositions indicate that the contribution of oil shocks in explaining variability in macroeconomic uncertainty is small, and even the simple correlation between oil price changes and macroeconomic uncertainty is actually negative.[15] More-

[15] The contribution of the different types of oil shocks to the contemporaneous median variance decom-

over, most of the high uncertainty events identified are not directly linked to oil shocks, and the results are robust to using financial uncertainty instead of macroeconomic uncertainty. Finally, coming back to the example given above, it is excluded that the decline in price elasticities that we find is merely due to the occurrence of larger oil shocks in uncertain times. This is because larger shocks would simultaneous have a stronger effect on both oil prices and oil production, thereby leaving the oil price elasticities unchanged. As we will show now, this is not the case.

4 Effects of oil shocks in different uncertainty regimes

Figure 2 shows the estimated effects of the variables in the TVAR model to different types of oil shocks in the two regimes. In order to make the effects comparable across regimes, we normalized the contemporaneous response of oil production to a one percent change. The conditional impulse responses are accumulated and shown in levels over the first two years after the shock. The shaded responses in the figure represent the 68 percent posterior range of the estimated effects in the high uncertainty regime and the dotted ones represent those conditional on low uncertainty.[16] Note that in Figure 2, as we are using sign restrictions, the posterior range represents the uncertainty concerning the *model specification*. An overlap between the estimated responses across regimes could thus partly be due to the fact that we are comparing different model specifications that each satisfy the sign restrictions imposed. In Figure 3, we therefore also show the significance of the difference in estimated responses across regime *per* model specification, i.e. for each joint draw from the posterior, we generate the difference of the responses in the high and the low uncertainty regime, and this for all 1000 'successful' model specifications, of which the 68 percent posterior range is shown in Figure 3.[17]

The first two rows of Figure 2 show the effects of the different types of oil shocks on oil prices and oil production, while the third row displays the estimated oil price elasticity

position of macroeconomic uncertainty is around 4%.

[16] Note that as we report the posterior *range* of possible outcomes, the results are not subject to the Fry and Pagan (2011) critique, which only applies when some kind of summary measure such as the median is used.

[17] The difference in effects across regimes is calculated as the response in the high uncertainty regime minus the low uncertainty regime. A positive difference for the (typically negative) oil demand elasticity and a negative difference for the (typically positive) oil supply elasticity thus implies that the estimated elasticities are significantly lower in the high uncertainty regime.

of oil demand (following the oil supply shock) and the estimated oil price elasticity of oil supply (following both oil demand shocks). These price elasticities are calculated as the ratio of the oil production response to the oil price response following a particular shock. Remarkably, it is clear that for all three oil shocks a similar impact change in oil production has a much stronger impact effect on the oil price in the high uncertainty regime. This indicates that oil shocks have larger effects on oil prices when macroeconomic conditions are highly uncertain, compared to more normal times.[18] In line with this, the elasticity of both oil demand and supply falls considerably when uncertainty is high. In other words, the oil demand and supply curve become steeper.

More specifically, following the *oil supply shock*, we estimate the oil demand elasticity to decrease from within a range of -0.52 to -0.15 in the low uncertainty regime, to a value within the range of -0.36 to -0.11 in the high uncertainty regime. There is quite some overlap in estimated elasticities across the regimes (which can be due to model uncertainty, cfr. above), but Figure 3 shows that the difference in estimated oil demand elasticities across regimes is statistically significant. Hence, higher uncertainty causes the oil price elasticity of demand to be significantly lower. Given that the oil price elasticity in the high uncertainty regime might be less than half its value of the low uncertainty regime, the effect is also economically very significant. Although we use a different methodology, the estimated oil demand elasticities are broadly in line with those estimated in the literature. Hamilton (2009), Dahl (1993) and Cooper (2003) report oil demand elasticities between -0.05 and -0.07, whereas more recent analyses by Baumeister and Peersman (2012b), Bodenstein and Guerrieri (2011) and Kilian and Murphy (2013) arrive at estimates ranging from -0.26 to -0.44, which is at the higher end of our estimation range. Interestingly, using a time-varying VAR model and thereby allowing for non-linearity, Baumeister and Peersman (2012b) estimate the median price elasticity of oil demand to fluctuate within a range of -0.05 to -0.25 since 1986, with 68 percent posterior credible sets reaching up to -0.60, which broadly covers our estimation range over the two regimes. Therefore, the variation of the oil demand elasticity within their sample might well be explained by varying levels of macroeconomic uncertainty.

For the reason that we have two types of oil demand shocks, we can estimate the cur-

[18] During some periods of high uncertainty, small changes in oil demand were associated with enormous variation in the oil price, which might explain why the estimation uncertainty surrounding the oil price response following the oil demand shock driven by economic activity is so high. For example, in the fourth quarter of 2008, oil demand felt by 0.6 percent whereas oil prices plummeted by more than 111 percent.

vature of the oil supply curve following shocks in oil demand driven by economic activity and following oil-specific demand shocks. Figure 2 shows that also the elasticity of oil supply, as proxied by both types of oil demand shocks, tends to be lower when uncertainty is higher. Indeed, Figure 3 clearly confirms that the price elasticity of oil supply is statistically significantly lower in the uncertain regime. Following the *oil demand shock driven by economic activity*, the estimated oil supply elasticity drops from a maximum value of 0.21 in the low uncertainty regime to a maximum of 0.15 when uncertainty is high. The minimum estimated elasticity of oil supply reduces from 0.03 to 0.02. Again, these estimates correspond well with the estimates in the literature. Baumeister and Peersman (2012b), for example, estimate the median oil supply elasticity to lie in between 0.02 and 0.25. When the oil supply elasticity is estimated following an *oil-specific demand shock*, the results also show a reduction in the oil supply elasticity conditional on high uncertainty, although the magnitudes differ slightly. These differences could be due to the fact that the oil-specific demand shock captures a broad set of shocks, i.e. all demand shocks that are not driven by global economic activity. Shocks to expected net oil demand and oil-gas substitution shocks are two examples, and also speculation shocks are thought to be part of it.[19] For the reason that these shocks could trigger diverging responses in oil demand and supply, the estimation of the oil price elasticities could be subject to noise. As noted by Baumeister and Peersman (2012b), the differences in the estimated elasticities could also be explained by a different reaction of oil supply to both shocks in oil demand.

Not only the oil price elasticity, but also the real economic effects of oil shocks appear to differ considerably when uncertainty is high. Row four of Figure 2 shows that economic activity appears to react more strongly following oil shocks in the high regime. The difference in real impact effects across regimes is statistically significant for all three shocks, see Figure 3. Again, the uncertainty effect is also economically relevant as the impact response in the high uncertainty regime might be twice as large than when uncertainty is low, which could be explained by increased sensitivity of the oil price. At first sight, there is no apparent difference between the reaction of oil inventories across regimes. Nevertheless, Figure 3 indicates that following the oil demand shocks on impact, the reaction of inventories is stronger when uncertainty is high, which corresponds well with increased precautionary inventory building motivated by increased uncertainty (Pirrong 2011).

[19] See for example Lombardi and Van Robays (2011) and Kilian and Murphy (2013).

A simple back-on-the-envelope calculation illustrates the economic relevance of the difference in estimated elasticities. In the aftermath of the financial crisis that hit the global economy in summer 2008, oil demand dropped considerably. Global oil demand declined with about two percent between 2008Q3-2009Q2 and oil prices decreased from about USD 112 to USD 58 per barrel which, under some assumptions, implies an oil supply elasticity of 0.04.[20] Based on our estimates, uncertainty concerning the macroeconomy was already high before the financial crisis hit (see Figure 1). If the economy would have been in a low uncertainty state instead, in which the elasticity of oil supply is estimated to be in between 0.01 and 0.06 higher in absolute terms (see the results in Figure 3), the percentage price decline would have been smaller by about 12 to at least 32 percent. The finding that the oil demand elasticity and the oil supply elasticity are smaller when uncertainty is higher is robust to using other measures of uncertainty, see Panel B and C of Figure 4 in comparison with Panel A.

5 Robustness of the results

The main results on the lower price elasticity of oil demand and supply in times of higher uncertainty, and the stronger real economic impact of oil shocks, hold for various specifications of the model used. First, our conclusions hold for using the real oil price, reasonable variation in the number of lags given our data sample (2, 3 and 5 lags), only imposing the sign restrictions on impact and for different measures of uncertainty as described in the main text. Second, if we identify regimes of negative growth instead of regimes of higher uncertainty, the overall results remain the same although the significance of the difference across regimes disappears. This indicates that our findings concerning the uncertainty effect can not be solely explained by a different effect of oil shocks on oil prices in recessions versus expansions. Third, the result of a lower price elasticity under high uncertainty cannot be ascribed to time-variation in OPEC‘s spare capacity, as in periods of high uncertainty, spare capacity typically increased instead of declined because of a reduction in demand, which is clearly seen in the data. Finally, as mentioned before, the finding that oil prices are more responsive in the high uncertainty regime cannot be solely

[20] For simplicity, we made the assumption that the two percent drop in global oil demand is entirely caused by an oil demand shock driven by economic activity, and that the drop in production is equal to the drop in demand. These assumptions could be restrictive and the results should therefore be interpreted with caution.

explained by larger shocks, as this would generate a stronger response in both oil prices and oil production, leaving the price elasticity unchanged. These results are available upon request.

6 Conclusions

This paper analyzes whether the impact of oil shocks differs in times of high and low macroeconomic uncertainty. It is well documented that uncertainty can affect the decision behavior of economic agents, and it could equally impact on the strength at which shocks to oil demand and supply affect oil prices, oil production and economic activity. Several important insights emerge from our analysis. First, a test for the significance of threshold effects indicates that our oil model is non-linear and behaves differently in regimes of high uncertainty, which are mostly associated with periods of slowing economic growth, recessions and financial crises. Second, higher macroeconomic uncertainty causes oil prices to respond more strongly given a certain change in oil production, implying that the price elasticity of oil demand and supply decreases when uncertainty is higher. The reduction in the oil price elasticity is both statistically and economically significant. A third, possibly related, finding is that the effect of all types of oil shocks on economic activity is more aggressive in times when macroeconomic volatility is already high. These findings are robust to variations in the specification of the model, identification of the shocks and the measure of uncertainty.

As far as we are aware, this is the first paper considering a role for macroeconomic uncertainty in explaining changes in the impact of oil shocks, and that can endogenously explain time-variation in the elasticity of oil demand and supply. Concerning the discussion on the driving factors behind the recent rollercoaster ride in oil prices, we provide an empirical explanation for why the oil price elasticity has lowered over the past decade and caused fundamental shocks in oil demand and supply to impact more strongly on the oil price, i.e. because of macroeconomic uncertainty.

References

[1] Arrow, K.J. (1968): Optimal capital policy with irreversible investment, *in Wolfe, N.J. (ed.), Value, Capital and Growth, Essays in Honour of Sir John Hicks, Edinburgh.*

[2] Balke, N. (2000): Credit and economic activity: Credit regimes and nonlinear propagation of shocks, *The Review of Economics and Statistics 82(2): 344-349.*

[3] Barsky, R.B. and L. Kilian (2004): Oil and the Macroeconomy since the 1970s, *Journal of Economic Perspectives, 18(4): 115-34.*

[4] Baum, C.F. and C. Wan (2010): Macroeconomic uncertainty and credit default swap spreads, *Applied Financial Economics 20(15): 1163-1171.*

[5] Baumeister, C. and G. Peersman (2012a): Time-varying effects of oil supply shocks on the US economy, *Bank of Canada Working Paper 2012-2, January 2012..*

[6] Baumeister, C. and G. Peersman (2012b): The role of time-varying price elasticities in accounting for volatility changes in the crude oil market, *Journal of Applied Econometrics, doi: 10.1002/jae.2283.*

[7] Baumeister, C., G. Peersman and I. Van Robays (2010): The economic consequences of oil shocks: differences across countries and time, *in Fry, R., C. Jones and C. Kent (eds.), Inflation in an Era of relative Price Shocks, Sydney: 91-137.*

[8] Bernanke, B.S. (1983): Irreversibility, uncertainty, and cyclical investment, *The Quarterly Journal of Economics, 98(1): 85-106.*

[9] Bjørnland, H.C. (2000): The dynamic effects of aggregate demand, supply and oil price shocks - A comparative study, *The Manchester School of Economic Studies, 68: 578-607.*

[10] Bloom, N. (2009): The impact of uncertainty shocks, *Econometrica, 77(3): 623-685.*

[11] Bloom, N., S. Bond and J. Van Reenen (2007): Uncertainty and investment dynamics, *Review of Economic Studies, 74: 391-415.*

[12] Bodenstein, M. and L. Guerrieri (2011): Oil efficiency, demand and prices: A tale of ups and downs, *International Finance Discussion Papers Federal Reserve Board 1031, October 2011.*

[13] Bredin D., J. Elder and S. Fountas (2011): Oil volatility and the option value of waiting: An analysis of the G-7, *The Journal of Futures Markets, 31(7): 679-702.*

[14] Brennan, M.J. and E.S. Schwartz (1985): Evaluating natural resource investments, *The Journal of Business, 58(2): 135-157.*

[15] Calza, A. and J. Sousa (2006): Output and inflation responses to credit shocks: Are there threshold effects in the euro area?, *Studies in Nonlinear Dynamics & Econometrics, 10(2), Article 3.*

[16] Candelon, B. and L. Lieb (2011): Fiscal policy in good and bad times, *Maastricht University Working paper RM/11/001.*

[17] Cooper, J.C.B. (2003): Price elasticity of demand for crude oil: Estimates for 23 countries, *OPEC Review, 27: 1-8.*

[18] Dahl, C.A. (1993): A survey of oil demand elasticities for developing countries, *OPEC Review, 17: 399-419.*

[19] Elder, J. and A. Serletis (2010): Oil price uncertainty, *Journal of Money, Credit and Banking, 42(6): 1137-1159.*

[20] Favero, C.A., M.H. Pesaran and S. Sharma (1994): A duration model of irreversable oil investment: Theory and empirical evidence, *Journal of Applied Econometrics, 9:95-112.*

[21] Ferderer, P.J. (1996): Oil price volatility and the macroeconomy, *Journal of Macroeconomics, 18(1): 1-26.*

[22] Fry, R., Pagan A., (2011). Sign restrictions in structural vector autoregressions: a critical review. *Journal of Economic Literature, 49(4): 938–960.*

[23] Guiso, L. and G. Parigi (1999): Investment and demand uncertainty, *The Quarterly Journal of Economics, 114(1): 185-227.*

[24] Hamilton, J.D. (1983): Oil and the Macroeconomy Since World War II, *Journal of Political Economy, 91(2): 228-248.*

[25] Hamilton, J.D. (2009): Causes and consequences of the oil shock of 2007-08, *Brookings Papers on Economic Activity, Spring 2009: 215-259.*

[26] Hansen, B.E. (1996): Inference when a nuisance parameter is not identified under the null hypothesis, *Econometrica, 64(2): 413-430.*

[27] Hansen, B.E. (1999): Testing for linearity, *Journal of Economic Surveys, 13(5): 551-576.*

[28] Hansen, B.E. (2011): Threshold Autoregression in Economics, *Statistics and Its Interface, 4: 123-127.*

[29] Henry, C. (1974a): Option values in the economics of irreplaceable assets, *The Review of Economic Studies, 41, Symposium on the Economics of Exhaustible Resources (1974): 89-104.*

[30] Henry, C. (1974b): Investment decisions under uncertainty: The "irreversibility effect", *The American Economic Review, 64(6): 1006-1012.*

[31] IEA (2009), International Energy Agency Oil Market Report, December 2009.

[32] Kellogg, R. (2010): The effect of uncertainty on investment: Evidence from Texas oil drilling, *NBER Working Paper 16541.*

[33] Khan, M.S. (2009): The 2008 oil price "bubble", *Peterson Institute for International Economics Policy Brief PB09-19.*

[34] Kilian, L. (2009): Not all oil price shocks are alike: disentangling demand and supply shocks in the crude oil market, *American Economic Review 99: 1053-1069.*

[35] Kilian, L. and D. Murphy (2012): Why agnostic sign restrictions are not enough: understanding the dynamics of oil market VAR models, *Journal of the European Economic Association, 10(5), October 2012.*

[36] Kilian, L. and D. Murphy (2013): The role of inventories and speculative trading in the global market for oil, *Journal of Applied Econometrics, forthcoming.*

[37] Lee, K., S. Ni and R.A. Ratti (1995): Oil shocks and the macroeconomy: The role of price variability, *The Energy Journal, 16(4): 39-56.*

[38] Litzenberger, R.H. and N. Rabinowitz (1995): Backwardation in oil futures markets: Theory and empirical evidence, *The Journal of Finance, 50(5): 1517-1545.*

[39] Lombardi, M. and I. Van Robays (2011): Do financial investors destabilize the oil price?, *Working Paper Series, European Central Bank, 1346.*

[40] Majd, S. and R.S. Pindyck (1987): The learning curve and optimal production under uncertainty, *MIT Working papers 1948-87.*

[41] Paustian, M. (2007): Assessing Sign Restrictions, *The B.E. Journal of Macroeconomics, 7(1), Article 23.*

[42] Peersman, G. (2005): What caused the early millennium slowdown? Evidence based on vector autoregressions, *Journal of Applied Econometrics, 20: 185-207.*

[43] Peersman, G. and I. Van Robays (2009): Oil and the euro area economy, *Economic Policy, 24(60): 603-651.*

[44] Peersman, G. and I. Van Robays (2012): Cross-country differences in the effects of oil shocks, *Energy Economics, 34(5), September 2012: 1532-1547.*

[45] Pindyck, R.S. (1980): Uncertainty and exhaustible resource markets, *The Journal of Political Economy, 88(6): 1203-1225.*

[46] Pindyck, R.S. (1991): Irreversibility, uncertainty and investment, *Journal of Economic Literature, 29(3): 1110-1148.*

[47] Pindyck, R.S. (2004): Volatility and commodity price dynamics, *The Journal of Futures Markets, 24(11): 1029-1047.*

[48] Pirrong, C. (2011): Stochastic fundamental volatility, speculation, and commodity storage, *in Commodity Price Dynamics. A Structural Approach, Cambridge University Press: 109-130.*

[49] Singleton, K. (2012): Investor flows and the 2008 boom/bust in oil prices, *Stanford University Working Paper.*

[50] Sockin, M. and W. Xiong (2013): Feedback effects of commodity futures prices, *Working Paper Princeton University.*

[51] Tang, K. and W. Xiong (2012): Index investing and the financialization of commodities, *Financial Analysts Journal 68 (6), 2012,: 54-74.*

[52] Tong, H. (1978): On a threshold model, *Pattern Recognition and Signal Processing, (eds.) C.H.Chen, Amsterdam: Sijhoff & Noordhoff.*

[53] Uhlig, H. (2005): What are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure, *Journal of Monetary Economics, 52(2): 381-419.*

PANEL A. Macroeconomic uncertainty based on conditional variance world industrial production growth

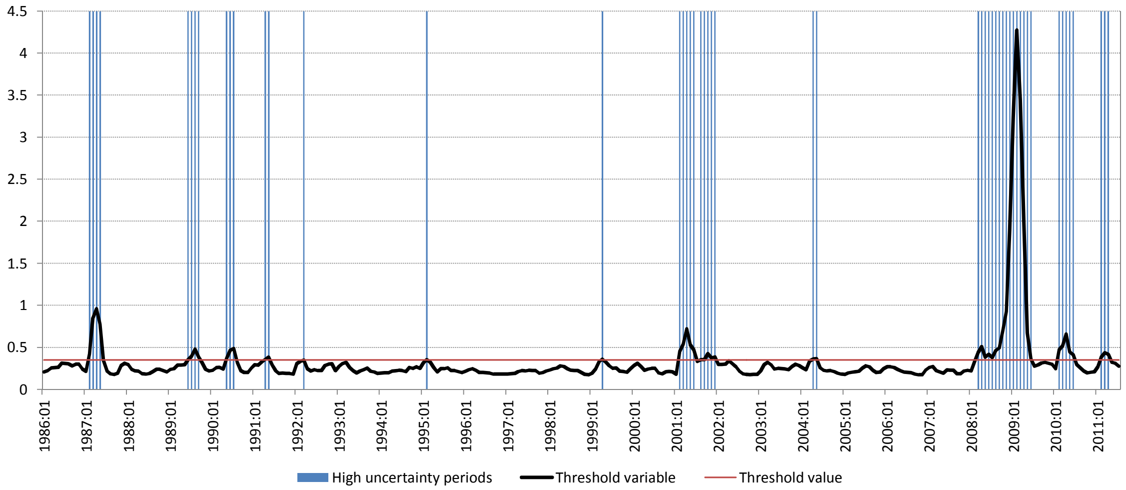

PANEL B. Macroeconomic uncertainty based on conditional variance US GDP growth

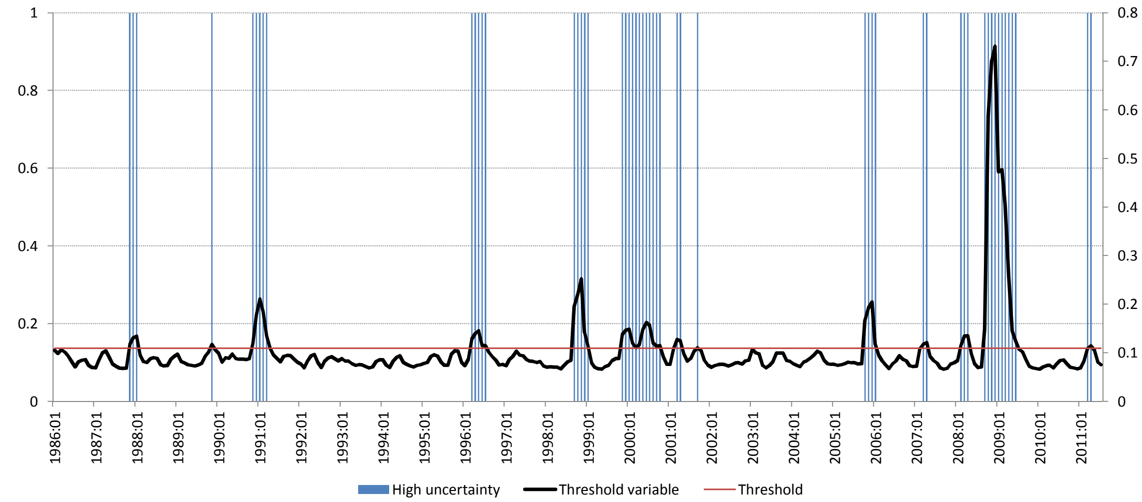

PANEL C. Macroeconomic uncertainty based on CBOE VXO stock market volatility

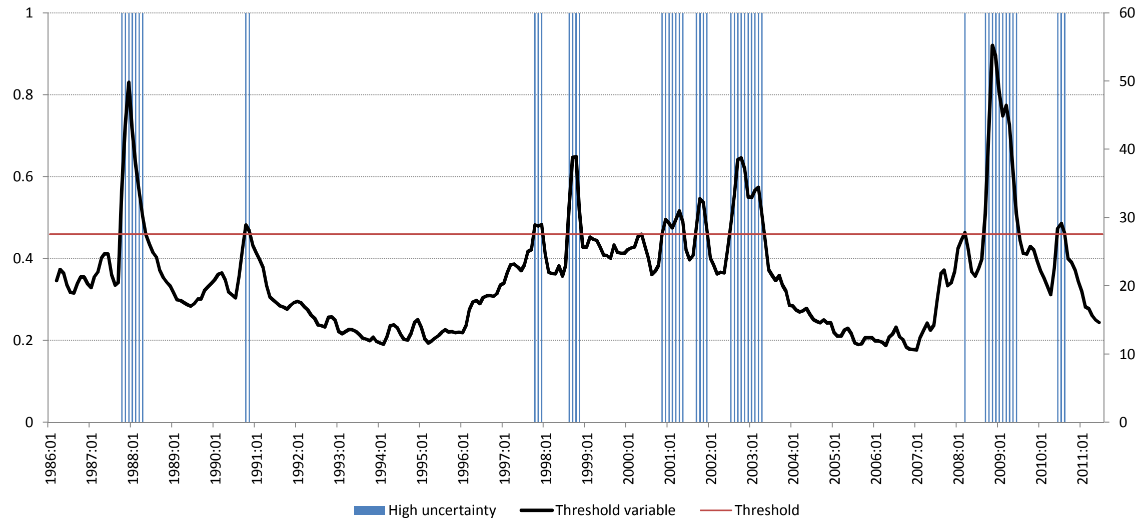

Figure 1. Threshold variable related to uncertainty, estimated threshold and identified periods of high uncertainty

Notes: the threshold variable is constructed as a three-period moving average of the respective measure of uncertainty.

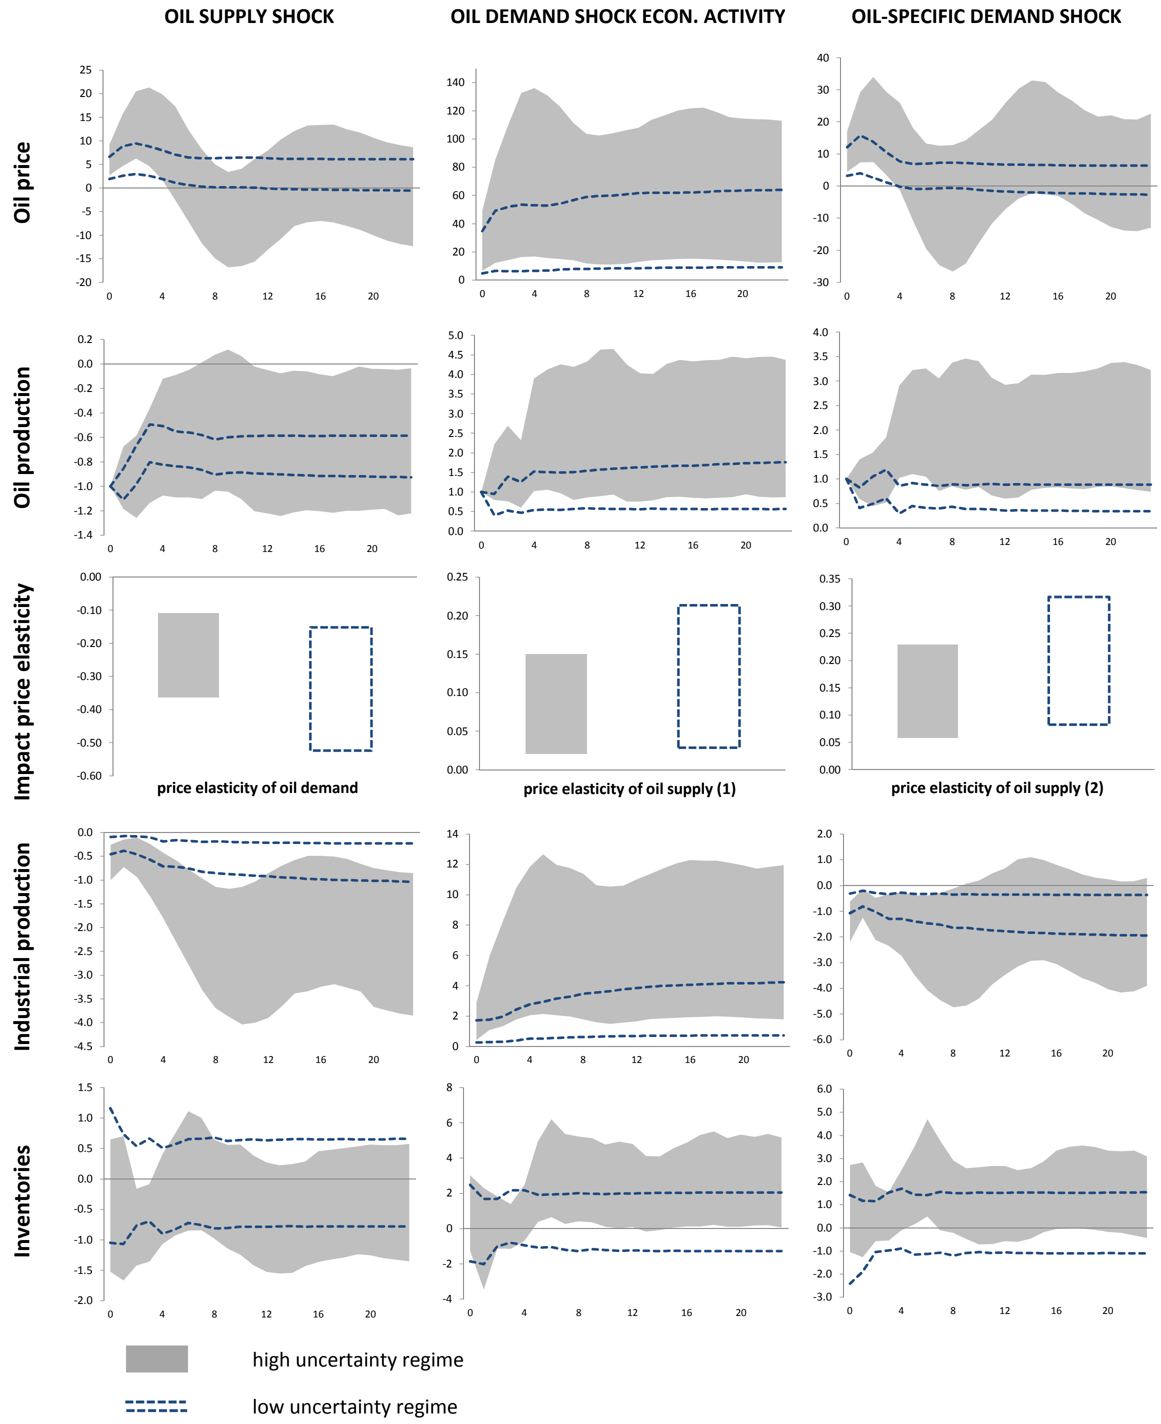

Figure 2. Impact of different types of oil shocks in different regimes of macroeconomic uncertainty

Notes: Figures are 68 percent posterior probability regions of the estimated conditional impulse response functions normalized on a 1 percent change in oil production, horizon is monthly and the measure of uncertainty is the conditional variance of world industrial production growth.

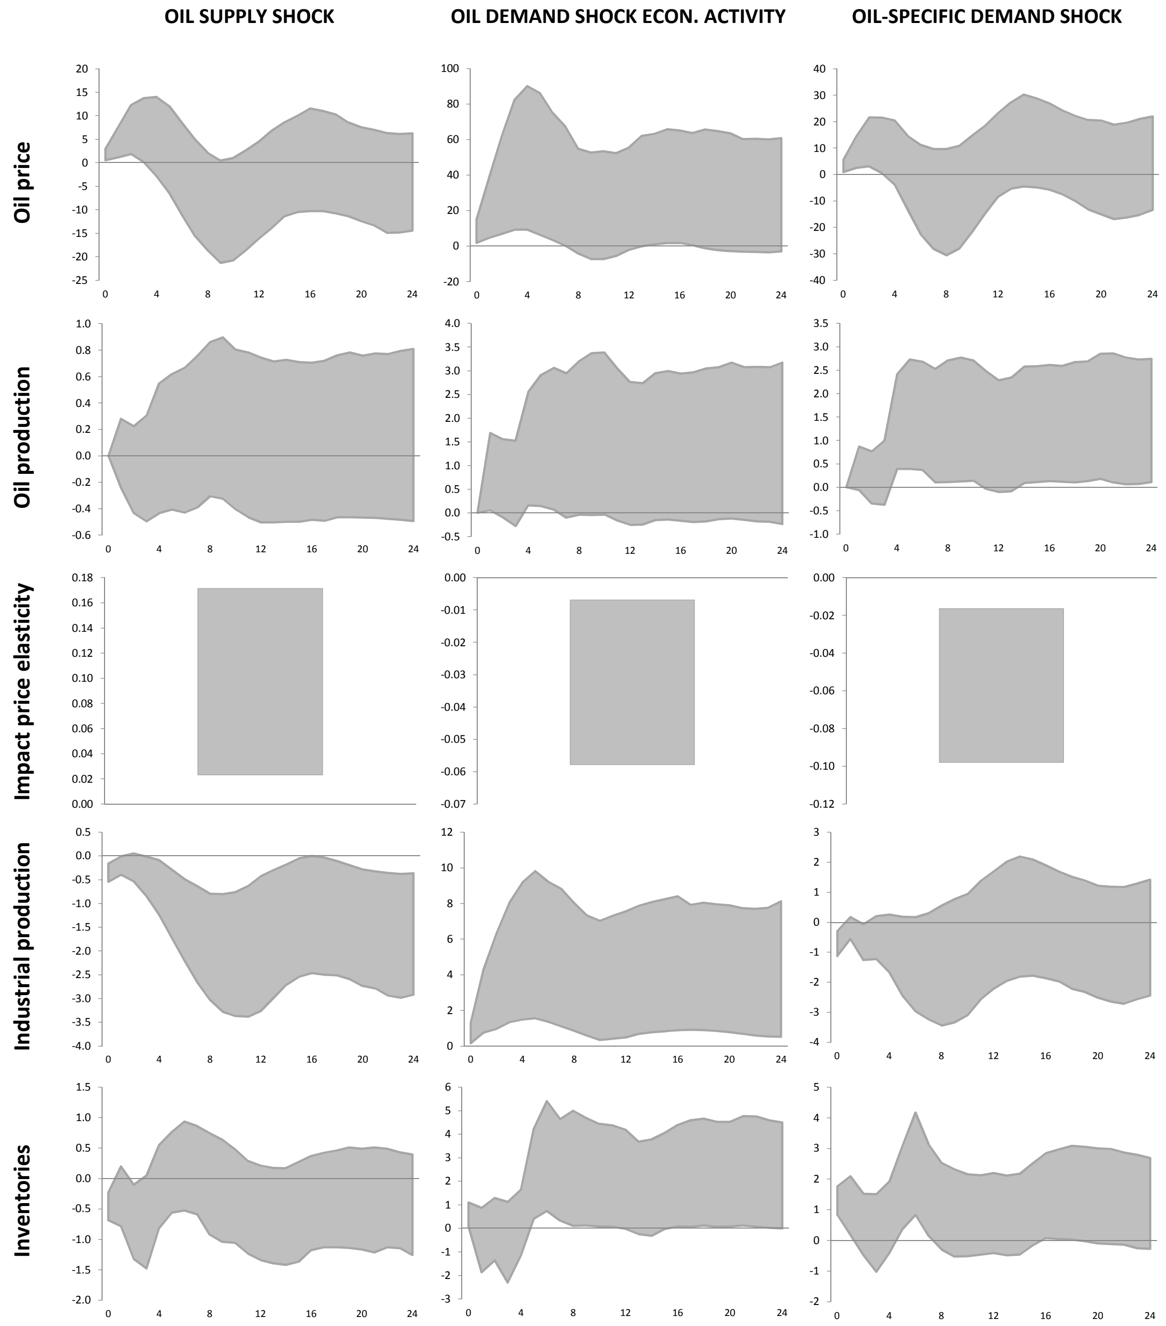

Figure 3. Signifcance of the difference in impact between high uncertainty and low uncertainty regime

Notes: Figures are 68 percent posterior probability regions of the difference in estimated conditional impulse response functions in the high uncertainty regime minus the low uncertainty regime. The impulse response functions normalized on a 1 percent change in oil production, horizon is monthly and the measure of uncertainty is the conditional variance of world industrial production growth.

PANEL A. Estimated impact elasticities using uncertainty proxied by conditional variance world industrial production growth

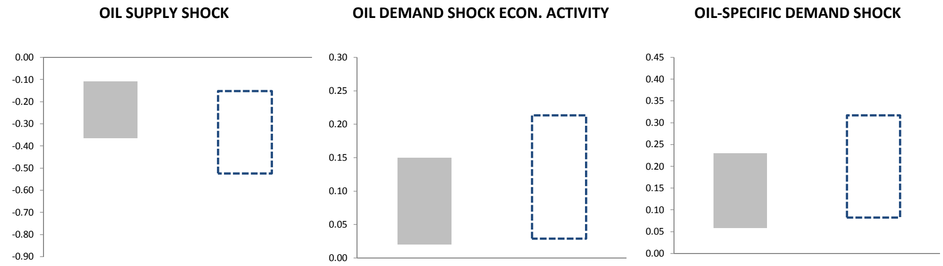

PANEL B. Estimated impact elasticities using uncertainty proxied by conditional variance US GDP growth

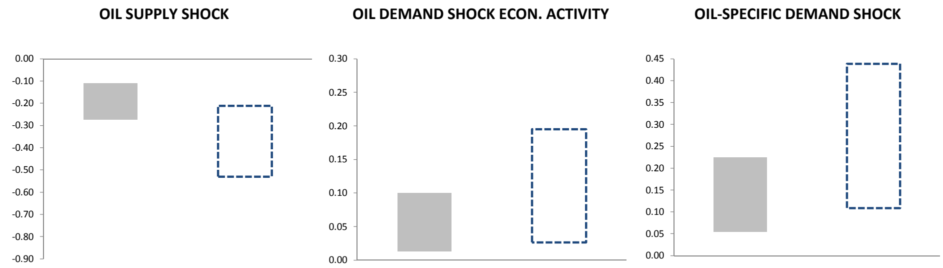

PANEL C. Estimated impact elasticities using uncertainty proxied by CBOE VXO stock market volatility

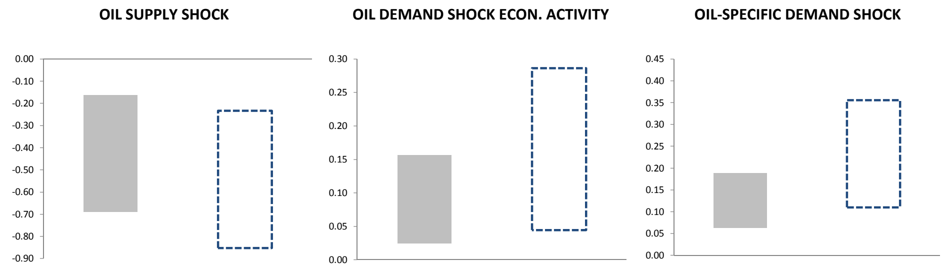

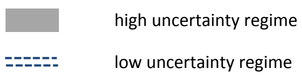

Figure 4. Robustness impact elasticities of oil demand and supply to various uncertainty measures

Notes: estimated impact elasticities estimated conditional on high or low uncertainty regimes identified based on threshold values shown in Table 1

Threshold Variable	Estimated threshold	Wald Statistics			% observations in high uncertainty regime	Duration of the high uncertainty regimes in months (min; max; mean)
		Sup-Wald	Avg-Wald	Exp-Wald		
World industrial production growth GARCH(1,1)	**0.3512**	431.45 (0.00)	211.88 (0.00)	210.43 (0.00)	17%	(1; 16; 8.5)
US GDP growth GARCH(1,1)	**0.1095**	435.40 (0.00)	179.62 (0.00)	212.86 (0.00)	18%	(1; 12; 6.5)
CBOE VXO monthly average of daily closing price	**27.5467**	389.15 (0.00)	176.60 (0.00)	189.35 (0.00)	17%	(1; 10; 5.5)

Table 1. Test for threshold effects

Notes: Tests are performed for the reduced form of the 5-variable TVAR model described in equation (1) with four lags of the endogenous variables, no delay parameter and three moving average terms for the threshold variable. The p-values based on the simulation technique of Hansen (1996) for 500 replications are in parenthesis. GDP and CBOE VXO stand respectively for gross domestic product and the Chicago Board of Option Exchange VXO US stock market volatility measure. The sample period is 1986:01-2011:07.