

Hahn, Elke; Mestre, Ricardo

Working Paper

The role of oil prices in the euro area economy since the 1970s

ECB Working Paper, No. 1356

Provided in Cooperation with:

European Central Bank (ECB)

Suggested Citation: Hahn, Elke; Mestre, Ricardo (2011) : The role of oil prices in the euro area economy since the 1970s, ECB Working Paper, No. 1356, European Central Bank (ECB), Frankfurt a. M.

This Version is available at:

<https://hdl.handle.net/10419/153790>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



EUROPEAN CENTRAL BANK

EUROSYSTEM

WORKING PAPER SERIES

NO 1356 / JUNE 2011

**THE ROLE OF OIL
PRICES IN THE EURO
AREA ECONOMY
SINCE THE 1970s**

by Elke Hahn
and Ricardo Mestre



EUROPEAN CENTRAL BANK

EUROSYSTEM



WORKING PAPER SERIES

NO 1356 / JUNE 2011

THE ROLE OF OIL PRICES IN THE EURO AREA ECONOMY SINCE THE 1970s

by Elke Hahn
and Ricardo Mestre¹



In 2011 all ECB publications feature a motif taken from the €100 banknote.



NOTE: This Working Paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB.

This paper can be downloaded without charge from <http://www.ecb.europa.eu> or from the Social Science Research Network electronic library at http://ssrn.com/abstract_id=1858615.

¹ Both authors: European Central Bank, Kaiserstrasse 29, D-60311 Frankfurt am Main, Germany; e-mails: elke.hahn@ecb.europa.eu and ricardo.mestre@ecb.europa.eu

© European Central Bank, 2011

Address

Kaiserstrasse 29
60311 Frankfurt am Main, Germany

Postal address

Postfach 16 03 19
60066 Frankfurt am Main, Germany

Telephone

+49 69 1344 0

Internet

<http://www.ecb.europa.eu>

Fax

+49 69 1344 6000

All rights reserved.

Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

Information on all of the papers published in the ECB Working Paper Series can be found on the ECB's website, <http://www.ecb.europa.eu/pub/scientific/wps/date/html/index.en.html>

ISSN 1725-2806 (online)

CONTENTS

Abstract	4
Non-technical summary	5
1 Introduction	7
2 Methodology	11
2.1 The VAR framework with time varying parameters and stochastic volatility	11
2.2 Identification of the structural shocks	17
3 Empirical results	19
3.1 Evolution of the volatilities of the variables	20
3.2 Contribution of the identified shocks to the reduced form residuals	20
3.3 Counterfactual histories	23
3.4 Impulse responses	28
3.5 Variance decompositions	30
4 Conclusions	31
References	33
Appendices	36
Figures	49

Abstract

This paper explores the role of oil prices in the euro area economy since the 1970s by applying a VAR framework with time varying parameters and stochastic volatility in which oil supply and global demand shocks are identified. Our results show that both types of shock contributed substantially to the oil price surges during historical oil crises and likewise to those over the past decade. Counterfactual histories of the price and activity variables, moreover, reveal much larger adverse contributions of both shocks to HICP inflation and GDP in the first half of the sample than in the second, which suggests that changes related to these shocks have contributed to the Great Moderation. Impulse responses, moreover, show that a decline in the pass-through of the two shocks has added to the moderating contribution over time, while variance decompositions indicate no change in the relative importance of the two shocks over time.

JEL classification: E3

Keywords: Oil prices, Great Moderation, time-varying parameter VAR model, stochastic volatility, euro area.

Non-technical summary

The empirical literature on oil prices has for a long time provided evidence for substantial instability in the relationship between oil and the economy over time. A stable linear relationship between oil price changes and US GDP found in very early empirical studies appears to have broken down in the mid-1980s. In order to re-establish a stable relationship in the late 1980s researchers started to use non-linear transformations of the oil price variable (an “asymmetric specification”, the “scaled oil price specification” and the “net oil price increase”) which helped to restore or at least improve the stability of the relationship.

Yet, doubts remain concerning all of these non-linear oil price transformations as well as the proceeding of creating specific nonlinear oil price variables in general. Though intuitively plausible, all of these transformations are very specific and ad hoc in nature and the focus on one concrete kind of non-linearity appears far too narrowly considered when trying to understand the general relationship between oil prices and the macroeconomy. Many different transformations of the oil price variable may be considered in trying to capture the same underlying mechanisms as the above specifications and many of them may be equally admissible. Moreover, besides the above described non-linearities in the oil price - GDP relationship extensively explored in the literature, many other reasons for non-linearities exist which would deserve to be considered as well.

The presence of a multitude of factors prone to creating non-linearities and changes in the relationship between oil prices and the macroeconomy calls for a much more general approach than the very specific and narrow focus on one concrete aspect of non-linearity in order to understand how oil prices have impacted on the economy at different points in time over the past decades. Time varying parameter models offer such a more general approach. A particularly well suited approach in this context is the recently proposed time varying structural vector autoregressive (VAR) model that allows for time variation in both the variance covariance matrix of the shocks and the coefficients of the model (see e.g. Primiceri (2005), Benati and Mumtaz (2007) and also Cogley and Sargent (2005)). This model has been intensively used to analyse the causes of the “Great Moderation” but although being obvious candidates for having contributed to macroeconomic instability in industrialised economies, the role of oil prices has not been explored much in this discussion.

In this paper we apply a VAR framework with time varying parameters and stochastic volatility to explore the role of oil prices for the macroeconomic developments of the euro

area economy since the 1970s. Given numerous issues of interest, the analysis is conducted in the framework of two closely related complementary models and besides the classical oil supply shocks, in line with the recent literature, also the impact of global demand shocks is examined.

Our empirical results show that both oil supply and global demand shocks contributed substantially to the oil price surges during historical oil crises and both shocks likewise contributed to driving up oil prices over the past decade, albeit concerning the latter period oil supply shocks to a larger extent than global demand shocks. The results also reveal much larger adverse contributions of both types of shocks to HICP inflation and euro area/global GDP in the earlier part of the sample, i.e. from the 1970s up to the mid-1980s, than thereafter, an exception being the impact of demand shocks on GDP during the latest recession. This suggest that changes related to both oil supply and global demand shocks have contributed to the Great Moderation, but also taken together the two shocks do not seem to account fully for the changes observed with the Great Moderation but other factors appear to have been at play as well. Given the findings of Benati and Surico (2009) we refrain from trying to disentangle the relative importance of the two traditionally explored opposite explanations of the Great Moderation. It is also not possible to analyse changes in the volatilities of our structural shocks over time in this framework, but the increase in the volatility of oil prices over time raises doubts about the possibility of a decline in the volatility of important shocks to oil prices over time. We have, however, examined whether the pass-through of the considered shocks might have changed over time. The results show that both oil supply and global demand shocks are transmitted to a smaller degree to prices and activity today than in the past, which has added to the moderating contribution of these shocks to the economy over time. At the same time, variance decompositions, show fluctuations but neither a trend-like in- nor decrease in the relative importance of either of the two shocks for the variables over time. Overall, the evidence suggests that there have been substantial changes related to oil prices over time and these changes have played a role for the Great Moderation. Changes in the structure of the economy or policy, which affect the transmission of the considered shocks to oil prices have likely contributed to the Great Moderation. Their relative importance for the Great Moderation compared with potential changes in the volatility of the shocks over time, however, remains unclear and a topic for future research.

1. Introduction

The empirical literature on oil prices has for a long time provided evidence for substantial instability in the relationship between oil and the economy over time. Very early empirical studies were able to find a stable linear relationship between oil price changes and GDP growth in the US, but in the mid-1980s this relationship appears to have broken down (see e.g. Hamilton (1983), Hooker (1996a) and Hamilton (2008)). While highly significant effects of oil price changes on GDP growth were found in earlier samples, this significance declined dramatically in extended samples and the estimated magnitude of the effects was substantially smaller in these samples as well. This led to the perception that the link had become non-linear or was more complex than what could be captured by a linear relationship.

In order to re-establish a stable relationship in the late 1980s researchers started to use non-linear transformations of the oil price variable, which helped to restore or at least improve the stability of the relationship. Three different kinds of non-linear transformations were suggested in the literature. Tatom (1988) and Mork (1989) were the first to explore an asymmetric specification of oil price changes which allows for differences in the effects of oil price increases and decreases.¹ While Tatom (1988) did not find evidence for asymmetry in the effects in a sample extending to 1986Q3,² Mork (1989) showed for a slightly longer sample up to mid-1988 that the asymmetric specification helps to restore the stability of the oil price - GDP relationship and provides evidence for large negative effects of oil price increases and only much smaller positive effects of oil price declines. The two further suggested non-linear transformations of the oil price variable of Lee, Ni and Ratti (1995) and Hamilton (1996) are refinements of this asymmetric specification. In Lee, Ni and Ratti's (1995) so-called scaled oil price specification the oil price variable is defined as the unanticipated oil price increase scaled by the time varying conditional variance of oil prices.

¹ Potential explanations for the emergence of asymmetries in the GDP effects of oil price increases and declines include the response of monetary policy, adjustment costs, coordination problems, uncertainty and financial stress, and asymmetries in the petroleum product markets. For a summary of this literature see e.g. Brown, Yücel and Thompson (2003).

² This sample period includes the smaller oil price declines from 1980 onwards but the first large oil price decline of 1986 only at the very end of the sample.

This reflects the idea that the effect of an oil price increase on real activity may be smaller if it occurs in an environment of larger oil price volatility as in that case the oil price increase is more likely to be regarded as transitory. The reasoning behind Hamilton's (1996) "net oil price increase" variable has similarities to that of Lee, Ni and Ratti (1995). Hamilton argues that most of the oil price increases since 1986 were just reversals of previous oil price declines while for the effects of oil price increases on activity it matters more how unusual they are, i.e. economic agents will respond to oil price increases only if the oil price becomes larger than what they are accustomed to. For the suggested net oil price increase therefore the price of oil in a period is compared to its maximum level over the previous year rather than the level in the previous quarter and if it exceeds this level the percentage change in oil prices is recorded for that period and a zero value otherwise.

Yet, doubts remain concerning all of the above discussed non-linear oil price transformations as well as the proceeding of creating specific nonlinear oil price variables in general. And this is not only on account of Hooker (1996a, 1996b) demonstrating that both Mork's as well as Hamilton's specifications did not solve the problems in the later part of the samples and suggesting that the same is likely to also apply to Lee, Ni and Ratti's (1995) transformation.³ But it is rather that, though intuitively plausible, all of these transformations are very specific and ad hoc in nature and the focus on one concrete kind of non-linearity appears far too narrowly considered when trying to understand the general relationship between oil prices and the macroeconomy.⁴ Many different transformations of the oil price variable may be considered in trying to capture the same underlying mechanisms as the above specifications and many of them may be equally admissible. Moreover, many other reasons for non-linearities exist which would deserve to be considered as well. Apart from the direction of the oil price change, the duration of the movement as well as the general volatility of oil prices discussed so far, for instance, the level of oil prices may matter (i.e. there may be threshold effects), the absolute level change of oil prices might be important (and not only

³ He shows that significant Granger causality tests for the full sample based on these specifications resulted from improved significance in the first part of the sample, while the results for the latter sub-sample remained insignificant.

⁴ Although the literature on nonlinearities focuses more on the relationship between oil prices and GDP, such non-linearities or changes over time are also prevalent for the relationship with prices (see e.g. Hooker (2002), De Gregorio, Landerrechte and Neilsson (2007) and Baumeister and Peersman (2008)).

the percentage change) and also the general conjunctural environment in which the oil price change occurs might have a bearing on the impact (see e.g. ECB (2004)). In addition, the cause of the oil price change is likely to play a role as well (i.e. whether it is supply or demand driven) and, therefore, also the composition of the shocks underlying an oil price change, as recently suggested by Kilian (2009). Moreover, structural changes in the oil market, as documented by Baumeister and Peersman (2008, 2009), suggest that the underlying characteristics even of the same kind of shock and, hence, its impact on the economy may be different over time. In addition, changes in the effects of oil prices on the macroeconomy over time may also result from structural changes in other parts of the economy, such as a decline in the oil intensity or increased labour market flexibility, as well as changes in policy such as increased monetary policy credibility leading to smaller effects on both inflation and GDP contemporaneously (see Blanchard and Gali (2007)). Empirical evidence for such a declining impact over time of changes in oil prices has been found in Hooker (2002), Herrera and Pesavento (2009), De Gregorio, Landerrechte and Neilsson (2007), Edelstein and Kilian (2007), Blanchard and Gali (2007) and Baumeister and Peersman (2008). Finally, it is worth stressing that several of the above mentioned factors might apply at the same time, such that the effects might interfere and the combinations of contemporaneously occurring factors are likely to change over time as well.

The presence of such a multitude of factors prone to creating non-linearities and changes in the relationship between oil prices and the macroeconomy calls for a much more general approach than the very specific and narrow focus on one concrete aspect of non-linearity in order to understand how oil prices have impacted on the economy at different points in time over the past decades. Time varying parameter models offer such a more general approach. A particularly well suited approach in this context is the recently proposed time varying structural vector autoregressive (VAR) model that allows for time variation in both the variance covariance matrix of the shocks and the coefficients of the model (see e.g. Primiceri (2005), Benati and Mumtaz (2007) and also Cogley and Sargent (2005)). The time variation in the variance covariance matrix allows for heteroscedasticity of the shocks and time variation and non-linearities in the contemporaneous effects of the shocks, while the time varying coefficients should capture possible non-linearities and time variation in the lag structure of

the model. This model has been intensively used to analyse the causes of the “Great Moderation,”⁵ i.e. the increased macroeconomic stability that has been observed in important industrialised economies since the mid-1980s, with the aim to assess whether the Great Moderation was caused by declines in the volatilities of the shocks hitting the economy (i.e. “good luck”) or changes in the transmission mechanism related to structural changes or policy, the latter often termed “good policy”.

Surprisingly, the role of oil prices for macroeconomic instability in industrialised economies has not been explored much in the previous literature. Particularly high macroeconomic instability was found during the periods of the two big oil price shocks of the 1970s. Large changes in oil prices have been, however, also observed occasionally in later decades but these oil price movements were accompanied by much more muted price and output developments, at least when excluding the period of the latest recession. It is therefore insightful to see whether changes related to oil prices may have contributed to the increased macroeconomic stability over the past few decades. So far only Del Negro (2003) and Baumeister and Peersman (2008) have included commodity prices or oil prices in time-varying parameter, stochastic-volatility VAR models of this kind, in their case applied to the US economy.

In this paper we explore the role of oil prices for the macroeconomic developments of the euro area economy since the 1970s, i.e., among others, during the sharp historical and more recent strong oil price movements as well as the latest crisis in 2008/2009. Against the background of the above discussion, we refrain from imposing concrete non-linear transformations on the oil price variable but apply the time varying parameter VAR model with stochastic volatility to gain insights into the role of oil prices at each point in time over our sample period.⁶ In this way we remain largely agnostic about the kind of non-linearity prevalent in the oil price-macroeconomic relationship, but try to provide insights into whether changes related to oil prices might have contributed to the Great Moderation. The literature on the Great Moderation has stressed the evolution of volatility in the system

⁵ See for instance McConnell and Perez-Quiros (2000), Blanchard and Simon (2001), Stock and Watson (2002), Ahmed, Levine and Wilson (2004), Boivin and Giannoni (2006) and Smets and Wouters (2007).

⁶ Canova, Gambetti and Pappa (2007) also estimate a time varying parameter VAR model for the euro area. Their analysis, however, is not related to oil prices and their model also differs from ours in the way the stochastic volatilities are modelled.

(measured by the volatilities of the reduced form shocks) against the evolution of the transmission mechanism (measured by the impulse response function) as a gauge for “good luck” versus “good policy”, but Benati and Surico (2009) hint that the two are in fact indistinguishable with structural VAR models. We, therefore, refrain from attempting to gain insights into the relative importance of the two explanations, but nevertheless explore the evolution of impulse response functions to assess whether the pass-through of shocks to oil prices might have changed over time and also conduct variance decompositions to assess the relative importance of different shocks to oil prices at different points in time. In line with the recent literature, the identification of the different kinds of shocks to oil prices is achieved by applying sign restrictions.

The rest of the paper is structured as follows: Section 2 explains the applied methodology. In this section the VAR framework with time varying parameters and stochastic volatility is introduced and the strategy concerning the identification of the structural shocks is elucidated. Section 3 presents the empirical results. This starts by elucidating the evolution of the volatilities of the variables over time, which highlights the known main features of the Greta Moderation. The overall importance of the identified shocks concerning the developments in oil prices and the other macroeconomic variables in different periods since the 1970s is examined by looking, first, at their contributions to the reduced form shocks and exploring their impact on the historical developments of the variables. Thereafter, changes in the pass-through of the shocks over time are analysed by assessing the time pattern of impulse responses and insights into the relative importance of the identified shocks over time are gained by exploring the evolution of forecast error variance decompositions. Section 4 concludes.

2. Methodology

2.1. The VAR framework with time varying parameters and stochastic volatility

The empirical analysis on the role of oil prices in the euro area economy over the past decades is conducted in the framework of the VAR model with time varying parameters and stochastic volatility suggested by Primiceri (2005). In order to analyse all aspects of the role



of oil prices for the euro area economy we are interested in, the model would necessitate more than the 3 or 4 variables usually found in the previous literature for these kind of models, on account of their heavy parameterisation. We, therefore, decided to split the analysis and consider two related complementary models. The first model considered is one with a strong domestic focus, therefore henceforth denoted “domestic model”. It includes as variables the real crude oil price denominated in Euro, together with a set of basic domestic macroeconomic variables of the euro area economy: the Harmonised Index of Consumer Prices (HICP), GDP and a short-term interest rate.⁷ That is, this model is designed to explicitly consider the dynamic interrelationships of the domestic euro area variables, including that of interest rates, when analysing the role of shocks to oil prices - represented in this model by oil supply shocks, as in most previous literature - for oil prices, HICP inflation and GDP in the euro area.

There are, however, two interesting aspects which cannot be tackled in the setup of the domestic model. The first relates to the identification of the oil supply shocks. While the details of the identification of the shocks will be discussed in the next section, it is worth stressing at this stage that in some of the recent literature these shocks are identified not only by considering the price of oil (besides other macroeconomic variables) but both the price and the output side of the oil market. That is, this kind of identification schemes would necessitate the inclusion of oil production into the set of endogenous variables (see Kilian (2009), Baumeister and Peersman (2008, 2009) and Peersman and Van Robays (2009)). The second aspect which cannot be dealt with in our domestic model relates to the role of global demand shocks for oil prices, which are often considered to have gained in importance in recent years on account of rapidly growing global demand related to globalisation and growing emerging markets such as China and India (see e.g. Kilian (2009) and Hamilton (2008)). The analysis of this kind of shocks would necessitate the inclusion of a variable for global demand.

Given the above mentioned constraints concerning the dimension of the model, the inclusion of these two additionally interesting variables, however, needs to come at the cost of dropping two of the other variables. In our second model, with a stronger global focus,

⁷ A detailed description of the data is provided in Appendix A.

hence, denoted “global model”, we therefore decided to take out the euro area interest rate and GDP variables and include instead the two new variables, which allows us to analyse the role of both oil supply shocks and global demand shocks in line with the recent literature. The set of variables in the global model, hence, comprises oil production, oil prices (included in US Dollar terms), world GDP to represent global demand and the HICP, as the focus remains to analyse the role of oil prices on the euro area.⁸ The list of variables, however, highlights that the advantages offered by this model entail the drawback that, first, no further analysis of the effects on euro area GDP are possible and, second, the interrelationships with interest rates are omitted.⁹ The comparison of the results from both of our models will, however, provide us with some insights into the robustness of the results based on the different identification schemes of oil supply shocks as well as the sensitivity of the results to the omission of the two variables.

The data for both models are quarterly and included as log differences, except for the interest rate, which is used in levels. The dynamics of the system are captured by two lags of the endogenous variables.¹⁰ The sample period for estimation spans 1970 to 2009 for the domestic model and 1970 to 2008 for the global model (given the more delayed availability of world GDP data). That is, this period covers the 1970s, which is naturally of utmost interest for an analysis on oil prices, and includes the economically interesting period of the latest crisis, which represented one of the sharpest and globally widest-spread recessions since industrialised economies started to collect economic data.

The formal representation of the time varying parameter VAR model with stochastic volatility of Primiceri (2005) is given in equation (1) where the $r \times 1$ vector y_t includes the endogenous variables for period t , c_t is a $r \times 1$ vector referring to the time varying coefficients of the constants, the lags of the endogenous variables are multiplied by the $r \times r$ matrices $B_{t,l}$ of time varying coefficients for lag l and u_t is a $r \times 1$ vector of heteroscedastic

⁸ The data on the two additional variables are also described in more detail in Appendix A.

⁹ A version of the global model where euro area GDP was included instead of the HICP displayed problems with stability.

¹⁰ The choice of a lag length of two is based on the results from information criteria for the corresponding constant parameter VAR model and is also in line with the lag length usually chosen in VAR models with time varying parameters in the literature.

white noise processes distributed as i.i.d. Normal random variables with variance covariance matrix Ω_t .

$$y_t = c_t + \sum_{l=1}^L B_{t,l} y_{t-l} + u_t, \quad u_t \sim niid(0, \Omega_t), \quad t = 1, \dots, T \quad (1)$$

The time varying covariance matrix Ω_t can be decomposed, for the purpose of estimating the model, as indicated in equation (2), where A_t is a lower triangular matrix that captures the contemporaneous interactions of the endogenous variables (see equation (3)) and H_t is a diagonal matrix of the stochastic volatilities (see equation (4)).

$$\Omega_t = A_t^{-1} H_t H_t' (A_t^{-1})' \quad (2)$$

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

$$H_t = \begin{bmatrix} \sigma_{1,t} & 0 & 0 & 0 \\ 0 & \sigma_{2,t} & 0 & 0 \\ 0 & 0 & \sigma_{3,t} & 0 \\ 0 & 0 & 0 & \sigma_{4,t} \end{bmatrix} \quad (4)$$

Taking the decomposition of the variance covariance matrix into account and stacking the coefficients of the constant and B-matrices into a vector B_t , the model can be cast into state space form with observation equation (5)

$$y_t = X_t' B_t + A_t^{-1} H_t \varepsilon_t \quad (5)$$

whereby

$$X_t' \equiv I_r \otimes [1, y_{t-1}', \dots, y_{t-L}'] \quad (6)$$

$$Var(\varepsilon_t) = I_r. \quad (7)$$

The laws of motion of the time varying parameters of the model are specified in the transition equations (8) to (10). All parameters are assumed to follow random walk processes. This has the advantage of allowing for permanent parameter shifts and helps

reducing the number of parameters to estimate in this anyway heavily parameterised model. The elements of the vector B_t and the vector α_t , which includes the non-zero and non-one elements $\alpha_{ij,t}$ of the matrix A_t (stacked by rows), are postulated to evolve as driftless random walks. To properly model the positive values of the residual standard errors captured by the diagonal elements $\sigma_{ii,t}$ of the matrix H_t , these parameters are assumed to develop as geometric random walks. Allowing for time variation in the B_t , A_t and σ_t is aimed at capturing possible non-linearities and changes over time in the dynamics of the model and the contemporaneous relationships between the endogenous variables as well as possible heteroscedasticity in the variances of the shocks. Note that, as pointed out by Cogley and Sargent (2005), both stochastic volatility and time variation are needed, since lack of any of the two features would bias results in case both are present in the data.

$$B_t = B_{t-1} + \nu_t \quad (8)$$

$$\alpha_t = \alpha_{t-1} + \zeta_t \quad (9)$$

$$\log \sigma_t = \log \sigma_{t-1} + \eta_t \quad (10)$$

The innovations to the model are assumed to be normally distributed with variance covariance matrix V , which includes the hyperparameters of the model (see equation (11)). The innovations to the observation equation and the transition equations are assumed to be uncorrelated with each other. This is both a prudent and reasonable assumption as it substantially reduces the number of hyperparameters to be estimated and is necessary for the estimation technique used. Furthermore, it allows for a structural decomposition of the reduced-form VAR innovations following standard sign-restriction techniques. The matrices Q , S and W in equation (11) are positive definite. While Q and W are otherwise not constrained, for the matrix S the block diagonal structure indicated in equation (12) is assumed, where the blocks refer to the parameters belonging to the same equation. That is, the coefficients of the contemporaneous relations of the endogenous variables are assumed to be correlated only within the equations but uncorrelated across equations. This greatly simplifies the estimation of the model (see Primiceri (2005)).

$$V = \text{Var} \begin{pmatrix} \varepsilon_t \\ v_t \\ \zeta_t \\ \eta_t \end{pmatrix} = \begin{bmatrix} I_r & 0 & 0 & 0 \\ 0 & Q & 0 & 0 \\ 0 & 0 & S & 0 \\ 0 & 0 & 0 & W \end{bmatrix} \quad (11)$$

$$S_t = \begin{bmatrix} S_{1,t} & [0]_{1 \times 2} & [0]_{1 \times 3} \\ [0]_{2 \times 1} & S_{2,t} & [0]_{2 \times 3} \\ [0]_{3 \times 1} & [0]_{3 \times 2} & S_{3,t} \end{bmatrix} \equiv \begin{bmatrix} s_{11,t}^1 & 0 & 0 & 0 & 0 & 0 \\ 0 & s_{11,t}^2 & s_{12,t}^2 & 0 & 0 & 0 \\ 0 & s_{12,t}^2 & s_{22,t}^2 & 0 & 0 & 0 \\ 0 & 0 & 0 & s_{11,t}^3 & s_{12,t}^3 & s_{13,t}^3 \\ 0 & 0 & 0 & s_{12,t}^3 & s_{22,t}^3 & s_{23,t}^3 \\ 0 & 0 & 0 & s_{13,t}^3 & s_{23,t}^3 & s_{33,t}^3 \end{bmatrix} \quad (12)$$

Following Primiceri (2005), the system was estimated with Bayesian methods.¹¹ The details of the applied estimation procedure, i.e. the prior specifications and the Markov-Chain Monte Carlo algorithm as well as other methodological aspects are explained in Appendix B. One specific point related to the applied prior shall, however, be discussed at this stage. In time varying parameter VAR models, the prior is usually calibrated by estimating the corresponding constant parameter VAR model over a training sample which covers a certain time period before the start of the sample period used for estimation. For an analysis of oil prices it is obviously of utmost interest to include in the estimation the period of the 1970s with the two large oil price shocks. Since the euro area data included in the models start only in 1970, this period would be lost if a training sample was set up based on this data. We therefore decided to calibrate the prior based on US data, for which much longer backdata are available.¹² The period used for the prior calibration of the domestic model is 1947 to 1970 and that for the global model 1960 to 1970. The reasoning behind the use of the US prior is the assumption that the structures of the two economies are sufficiently close to get a reasonable starting point for the estimation of the euro area models based on these data. The

¹¹ We thank Giorgio Primiceri for providing us with the code of his Gibbs sampler.

¹² Canova and Ciccarelli (2009) use a different approach to tackle the issue of short euro area time series: they calibrate the prior using in-sample data. In our case, the much lower volatility of oil prices previous to 1970 up to the early 1970s relative to the later parts of the sample period advised calibrating the prior based on data corresponding to that period.

variances of the resulting prior were artificially widened to reduce their impact on the posterior, as an insurance against potential US-EA structural differences.

2.2. Identification of the structural shocks

Traditionally, in the literature oil price shocks were equalised with oil supply shocks and in the applied identification strategies reduced-form oil price shocks were often considered to represent the structural oil supply shocks. In VAR models, accordingly, these oil supply shocks were usually identified by applying a contemporaneous recursive identification scheme and ordering the oil price variable first in the model (see e.g. Hahn (2003) and Jiménez-Rodríguez and Sanchez (2005), Blanchard and Gali (2007)). This setting reflects the idea that the oil price variable is the most exogenous of the variables, such that the shocks to this equation (the oil supply shocks) are allowed to have also contemporaneous effects on the other variables of the model, while the shocks allocated to the other equations do not have a contemporaneous effect on oil prices. Alternatively, also sign restrictions (see e.g. Faust (1998), Canova and De Nicrolo (2002) and Uhlig (2005)) have been applied in the literature to identify oil supply shocks, imposing that these shocks lift oil prices as well as the overall consumer price index and lead to an increase in interest rates while reducing overall economic activity (see Peersmann (2005)).¹³

The exclusive focus on oil supply shocks caused by physical disruptions of the oil supply was, however, questioned in the recent literature and it was proposed that demand conditions play an important role for oil price movements as well (see e.g. Barsky and Kilian (2004), Hamilton (2008) and Hamilton (2009)). The idea hereby was that sudden oil price changes (oil price shocks) may not necessarily reflect oil supply disruptions (oil supply shocks) but might likewise be caused by other factors, i.e. demand shocks.

Kilian (2009) introduced the idea of different kinds of shocks to oil prices to the structural VAR literature. He proposed a recursive structural VAR model of the global crude oil market, in which, based on the variables global crude oil production, real economic activity and the real oil price, three different kinds of shocks driving oil prices can be identified.

¹³ A further somewhat different alternative strategy followed in the literature to measure exogenous oil supply shocks was to try to directly compile time series that capture the exogenous movements of oil prices (see e.g. Kilian (2008)).

Apart from the traditional oil supply shock, two demand shocks, an aggregate demand shock and an oil-specific demand shock, are identified. The aggregate demand shock is aimed at representing changes in overall demand, while the oil specific demand shock represents changes in the demand for oil on account of, for instance, precautionary demand for oil due to uncertainty about future oil supply shortfalls.

The same idea was pursued using sign restriction identification schemes by Peersman and Van Robays (2009). They distinguish oil supply shocks and the above mentioned two kinds of demand shocks by setting restrictions on the signs of the responses of global oil production, oil prices and overall economic activity. Oil supply shocks are assumed to be characterised by opposite movements in oil production and oil prices, e.g. a decline in oil production and an increase in oil prices, on account of which overall economic activity is assumed to fall as well. The two demand shocks, by contrast, are assumed to move oil prices and oil production in the same direction. They are distinguished by the sign of the movement in the third variable, global activity. In the case of the global demand shock the increase in oil prices and oil production is induced by an increase in global demand, i.e. there is a co-movement of all three variables, while in case of the oil-specific demand shocks the increased demand for oil (which drives up oil prices) is assumed to induce a decline in overall economic activity. This oil-market-based sign restriction identification scheme was applied to time varying parameter VAR models by Baumeister and Peersman (2008) for the identification of oil supply shocks and by Baumeister and Peersman (2009) in a three-variate model focusing solely on the oil market for all three kinds of shocks to oil prices.

We apply sign restrictions identification schemes in our two models.¹⁴ The identification of the shocks is delimited to the shocks of interest driving oil prices, i.e. in line with e.g. Baumeister and Peersman (2008) we only partially identify the models. The identification strategy followed in our domestic model for the oil supply shock is similar to the identification of this shock in Peersman (2005) whose model includes the same set of variables. The oil supply shock is assumed to be characterised by increases in oil prices and the HICP (which takes place mechanically via the energy component within the same quarter) and a decline in GDP, while concerning the response of interest rates we prefer to be

¹⁴ The sign restrictions are imposed for the impact period, which turned out to be sufficient. For further technical details on the implementation of the sign restrictions see Appendix B.

agnostic, firstly, on account of the synthetic euro area data before 1999 but also because of the structural break at the start of EMU. The imposed price and output responses are in line with standard textbook assumptions concerning the effects of an oil supply shock and are also generally confirmed empirically when the oil supply shock is identified just via restrictions on other variables (see e.g. Peersman and Van Robays (2009)). In our global model we extend the identification scheme beyond that of oil supply shocks and identify also global demand shocks. Given the set of variables in this model, however, oil supply shocks are identified in a different manner, involving in this case the interaction between oil price and oil production which is in line with the approach taken in the most recent literature. That is, the oil supply shock is now identified as a shock where oil prices rise and oil production, as well as overall economic activity, decline. Additionally, as for the domestic model, the HICP is assumed to increase on impact, as a rise in oil prices is mechanically transmitted to the energy component of the HICP within a quarter. This further restriction therefore helps to properly identify the oil supply shocks. For the global demand shocks, as in the recent literature, we assume that the increase in global demand corresponds with increases in both oil prices and production and again, in addition, we consider the channel that via the energy component the rise in oil prices is also within the same quarter reflected in an increase in the overall price level of the HICP.

3. Empirical results

The evidence on the role of oil prices for the macroeconomic developments since the 1970s will be presented in this section according to the following statistics obtained from the two models: the evolution of the volatilities of the variables over time; the contributions of identified shocks to the estimated reduced-form shocks; the historical reconstruction of variables in the absence of the identified shocks, respectively, in a counterfactual exercise; impulse response functions of the system as a measure of the mechanisms by which the economy absorbed the shocks and forecast error variance decomposition to assess the evolution of the relative importance of the shocks over time.

3.1. Evolution of the volatilities of the variables

We start the analysis on the role of oil prices in the euro area economy over time and, more specifically, its role in the Great Moderation, by looking at the evolution of the time varying standard deviation of the variables over the period since the 1970s in the two models (see Figures 1 and 15). The figures highlight the substantial decline in the volatilities of inflation and output growth over time, the key feature of the Great Moderation. The same pattern over time is also found for the volatility of short-term euro area interest rates as well as that of oil production. Interestingly, the picture, however, looks very different for the volatility of oil prices. The volatility of oil prices has not come down over time. There are large spikes in the volatilities of oil prices during certain periods of the first part of the sample in-between which the volatilities, however, fell back significantly. In the second half of the 1990 the volatilities increased, albeit to levels somewhat below the most extreme previous peaks, but have in contrast to the previous period remained persistently at this higher level with a large spike around 2009. This increase in the volatilities of oil prices over time is consistent with the findings in Baumeister and Peersman (2009).

3.2. Contribution of the identified shocks to the reduced form residuals

To get a first idea of the importance of the identified shocks in different periods we explore the contribution of the identified oil supply shocks in the domestic model and of the oil supply and global demand shocks in the global model to the reduced-form shocks of the equations since the 1970s (see Figures 2 to 5 and 16 to 19) and turn to their contribution to the historical developments of the variables in the next section.

A general observation is that the reduced form residuals of the equations of the variables which are included in both models are very similar. Concerning the residuals of the oil price equations, this is the case despite oil prices being denominated in euro in the domestic model and in US dollar in the global model (see Figures 2 and 16). The reduced form shocks to the oil price equations in both models show large spikes at the times of the first and second oil price shock of the 1970s (in 1973 related to the Yom Kippur war and between 1978 and 1980 due to the outbreak of the Iranian revolution and the Iran Iraq war), in the mid-1980 (due to the collapse of the OPEC cartel in 1985 and its attempts to reunite in 1987), around 1990/91

(on account of Iraq's invasion in Kuwait) and in the second half of the 1990s (related to the sharp movements of oil prices during the Asian crisis and following recovery) with a general increase in their magnitude since then and a particularly large spike again during the latest recession, when oil prices dropped sharply following their continuous surge during the previous years. The charts for both models highlight that although oil supply shocks can clearly not account for the entirety of these residuals during the periods of the first and second oil price shocks of the 1970s, they were quite important contributors. They contributed significantly to the large positive residuals in 1974/75 and provided continuous positive contributions, albeit of somewhat smaller magnitude, over several quarters also around 1979/80. Non-negligible upward effects are, however, also found in both periods to have come from the global demand shocks identified in the global model. The two shocks together explain large parts of the unexplained movements in oil prices during the two periods but there were also other forces not identified which have contributed to these oil price increases. Oil supply shocks, furthermore, added to the negative shocks to oil prices in 1985/86, when oil prices dropped very strongly for the first time, but the contribution from global demand shocks is estimated to have been much larger during that period. For the large shock to the oil price equation at the time of the Gulf war in 1990/91 oil supply shocks, however, provided a positive contribution and are found to have been more important than global demand shocks. Over the past years, since about 2000, the reduced form residuals of the oil price equation were mostly positive and oil supply shocks are estimated to have added mostly positively to the reduced form shocks and to explain a non-negligible share of these shocks, while global demand shocks are found to have provided both upward and downward contributions over that period. With regard to the oil price residuals during the latest recession (up to 2009Q4 in the domestic model and to 2008Q4 in the global model) it is evident that both oil supply and global demand shocks added positively to the residuals at the beginning of the recession and negatively to the large residual during the sharp part of the crisis with a much larger effect from demand shocks at that time but other factors appear to have been important as well.

The decomposition of the residuals of the HICP equations of the two models shows that oil supply shocks provided very large positive contributions to the shocks to this equation around the times of the first two oil price shocks, but the results from the global model show that also demand shocks provided large and mostly upward contributions such that the two

shocks together explain during these periods almost all of the reduced form HICP residuals (see Figures 3 and 19). The relative contribution of the two shocks to the reduced form HICP residuals appears, moreover, generally larger than that to the oil price residuals. In the years of the strong drops in oil prices during the mid-1980s both oil supply and global demand shocks provided large negative contributions. At the same time, the residual of the HICP equation at the time of the Gulf war of 1990/91 is rather small and the effect appears to have come more from oil supply shocks than demand shocks. Over the past decade both shocks together explain very large fractions of the unexplained movements in the HICP, but while oil supply shocks added mostly positively, the demand shocks contributed in both directions. The contributions from both shocks were inflationary at the onset of the latest recession, but dampening during the sharp part of the recession with a larger negative contribution coming from the demand shocks which were the largest contributors at that time also in absolute terms.

The decomposition of the shocks to the euro area GDP equation in the domestic model highlights that oil supply shocks contributed very strongly to the negative GDP residuals at the time of the first two large oil supply shocks, suggesting that they were important factors behind the two recessions that occurred at these times (see Figure 4). Some adverse effects of oil supply shocks are also measured for the 1992/93 recession while, by contrast, significant positive contributions are found around 1985/86. Over the past decade the oil supply shocks provided mainly negative contributions and they contributed also markedly to the adverse shocks to the GDP equation at the onset of the 2008/09 recession, but the sharp acceleration of the crises in autumn 2008 was clearly not related to oil supply shocks. Oil supply shocks, by contrast, in fact provided a significant positive contribution to the reduced form shocks to GDP in the very sharp part of the recession when oil prices fell strongly.

The picture concerning the contribution of oil supply shocks to the residuals of the global GDP equation from the global model is very similar to that for euro area GDP (see Figure 17). With regard to the global demand shocks, during the two large oil price shocks of the 1970s demand shocks contributed first positively to the residuals but followed in each period by significant negative contributions and also around mid-1980 a large negative contribution is identified. Over the past decade, in contrast to the mostly adverse contributions from oil supply shocks, global demand shocks provided a more balanced picture of contributions in

both directions but during the latest recession, global demand shocks provided by far the largest contributions to the huge negative residuals of the GDP equation.

With regard to the other variables included in the two models, Figure 5 shows that oil supply shocks also added strongly to the unexplained parts of the interest rate movements around the times of the two oil shocks of the 1970s. Their contributions to the shocks of the interest rate equation, however, declined strongly thereafter and remained small also in the last decade. Finally, looking at the residuals of the world oil production equation, persistent negative contributions from the oil supply shocks are visible during the periods of the two oil price shocks of the 1970s and around 1990/91, and positive contributions around the mid-1980s (see Figure 18). Also over the past decade there were periods of relatively persistent negative contributions from oil supply shocks. With regard to demand shocks, the effects are generally rather volatile, with more persistent effects mainly several times during the 1980s but also occasionally over the past decade, most notably during the recent recession.

3.3. Counterfactual histories

After having gained first insights into the role of the two identified shocks for the shocks that hit the economy in different periods, we now turn to explore directly how the oil supply and global demand shocks that occurred during our sample period have contributed to the variables and, in particular, to that of oil prices, HICP inflation and GDP growth. To that aim we compute counterfactual histories of the development of the series since the 1970s. More precisely, we run counterfactual simulations of the historical time series of the variables where we reconstruct the development of the variables assuming that no oil supply or global demand shocks occurred, respectively, and compare these developments with their actual time series (see Figures 6 to 9 and 20 to 23). The gap between the respective two series depicts the contribution of the considered identified shock to the variables in each period. For convenience, this gap is depicted again in the Figures 10, 24 and 25 normalised in such a way that a positive number reflects a positive contribution in that quarter of the considered shock to the respective variables and vice versa. Figures 11, 26 and 27, moreover, show the cumulated effects of the respective shocks on the development of the variables which are included in the models in terms of log differences, i.e. they display how the respective shocks have affected the levels of these variables over time.

The results from both models indicate that oil supply shocks had a substantial upward impact on oil prices at the time of the first large oil price shock of the 1970s and also during the second oil price shock at the end of the 1970s, though at that time of somewhat smaller magnitude in individual periods but of a more persistent nature, such that in cumulative terms the magnitude is similar to that during the first oil price shock. Oil supply shocks, moreover, contributed visibly to the decline in oil prices in the mid-1980s but only relatively moderately to their increase in 1990/91 according to the domestic model. The latter effect is, however, somewhat larger and more protracted for US Dollar oil prices in the global model. Over the past decade oil supply shocks added almost continually to the increases in oil prices. Global demand shocks are estimated to have contributed to the increases in oil prices during the two oil shocks of the 1970s to a similar degree as the oil supply shocks and they have also contributed markedly to the fall in oil prices in the mid-1980s. They appear, however, to have contributed less to the rise in oil prices around 1990/91. For the period since 2000, both upward and downward contributions are found for the global demand shocks with the overall inflationary contributions to oil prices much less pronounced than those found for oil supply shocks. Finally, a much larger contribution to the fall in oil prices during the recent recession is found for global demand shocks than for oil supply shocks. Overall, the two shocks together seem to explain large fractions of the movements in oil prices but, despite their importance, non-negligible parts of the occasionally large movements in oil prices are driven by other types of shocks.

Comparing our results with those of the literature, the finding of an important role of both oil supply and demand shocks for the oil price developments during most of the important historical episodes contrasts with those of Kilian (2009) who finds that oil supply shocks have historically in all periods made a comparatively small contribution to developments in oil prices and that oil price developments were mainly driven by (global and oil specific) demand shocks. Our results also differ from Hamilton's (2009) opposite finding that past oil price shocks were primarily caused by significant oil supply disruptions, but they are more consistent to Baumeister and Peersman's (2008) finding of an important role of oil supply shocks at least during certain periods of historical oil price shocks. By contrast, our findings concerning the latest period, i.e. that both oil supply and global demand shocks have contributed to driving up oil prices over the past decade prior to the recent recession, albeit the former to a larger extent than the latter, appear broadly consistent with Hamilton's (2009)

assessment that the main reasons for the latest oil price surge have been a failure of oil production to increase together with strong demand. The oil supply shocks captured by our models over the past decade probably mostly reflect such capacity constraints at times when oil production would have been increased if it would have been feasible. Our results are, however, also for this period not compatible with Kilian's (2009) finding of global demand shocks being the main driver of oil prices over the past years – except for the period of the sharp drop in oil prices during the recent recession (which is however not captured in Kilian's sample period). Strong demand might, however, have played a more important role in driving up oil prices over the past years than is captured by our global demand shocks as there might have been other demand shocks not identified in our model (e.g. oil specific demand shocks) which could have contributed to the increase in oil prices.

Very pronounced and persistent positive contributions of oil supply shocks are found for HICP inflation in the first part of the sample. Oil supply shocks provided very large positive contributions to inflation at the times of the first two oil price shocks: according to the domestic model, up to 0.7 percentage point to the quarterly inflation rate during the first oil price shock, where quarterly inflation reached a peak of 3.4% in 1974Q4, and up to 0.5 percentage points in the early 1980s, where the quarterly inflation rate remained below but close its 1974Q4 peak. Very similar results are found in the global model. These persistent positive contributions lasted basically over the whole period from 1974 to 1985. As the figures indicate, while oil supply shocks added substantially to inflation during these periods, a substantial part of the high inflation at that time, however, resulted also from other factors; the same applies to the moderation of the HICP inflation rates during the first half of the 1980s. Over the 15 years from the mid-1980s to 2000, oil supply shocks had mostly a moderating effect on HICP inflation, though of much smaller magnitude than the upward effect over the previous period, the general HICP inflation profile being relatively similar over this period in the absence of oil supply shocks. The situation is reverted after 2000, with almost continuous positive contributions of oil supply shock to HICP inflation, though again of relatively moderate and clearly much smaller magnitude than following the two oil price shocks of the 1970s. Demand shocks are also found to have had a substantial and persistent inflationary effect on the HICP over most of the 1970s and up to the early 1980s, though of somewhat smaller magnitude in cumulated terms than oil supply shocks. Around the mid-1980s, the contribution of the demand shocks to the HICP is estimated to have become and

remained negative over most of the subsequent sample, with particularly large negative effects on the HICP in the second half of the 1980s and a negative contribution during the recent recession. Taken together, the two identified shocks seem to explain a very substantial part of the movements of the HICP. Overall, these results for the euro area are in line with those of Blanchard and Gali (2007) for the US, who find a substantial but non-exclusive role of oil supply shocks for inflation during the periods of the first two large oil price shocks, but contrast with those of Baumeister and Peersman (2008) who come to the opposite conclusion that oil supply shocks explain little of the Great Inflation in the US.

With regard to euro area GDP, the figures highlight that oil supply shocks added strongly to the euro area recession in the mid-1970s, after the first oil price shock (with contributions to quarterly GDP growth up to -1.0 percentage point at that time) and also to the recession at the beginning of the 1980s, when smaller but more persistent shocks (up to -0.5 percentage point) led to an even larger cumulated effect. In both cases, however, important other factors were also at play depressing growth. At the same time, by contributing to the sharp drop in oil prices around 1985, oil supply shocks had a significant favourable effect on GDP growth at that time. Concerning the latest two euro area recessions, in 1992/93 and 2008/09, oil supply shocks are also found to have contributed to the declines in growth during these periods, but their adverse contributions were smaller than during the previous two recessions. And again for these two recessions and, in particular, the latest one, other factors than oil supply shocks were important.

A very similar pattern of the effects of oil supply shocks as on euro area GDP is found in the global model for global GDP. The global demand shocks additionally identified in this model are estimated to have also contributed to the fall in global GDP during the two recessions mid of the 1970s and at the beginning of the 1980s, though much less strongly than oil supply shocks. In the years after the mid-1980s, global demand shocks had large negative effects on global GDP and milder negative ones in the early 1990s, which were also smaller than the negative contributions of oil supply shocks during that period. Larger adverse effects of global demand shocks on global GDP are also found in the years before and after 2000. The largest negative contribution on global GDP over the whole sample period is recorded for the period of the recent recession at the end of 2008. This was also much larger than that of oil supply shocks at that time. In sum, it seems that the two

identified shocks have been very important drivers of the movements in global GDP. Overall, the findings for euro area and global GDP appear in line with those of Blanchard and Gali (2007), Baumeister and Peersman (2008) and Hamilton (2009) who point to important but non-exclusive adverse effects of oil supply shocks on GDP during historical US recessions as well as recessions in other countries.

The results concerning interest rates from the domestic model indicate that short-term interest rates in the euro area appear to have been persistently higher on account of the oil supply shocks from the late 1970s onwards up to the mid-1990s (by up to about two percentage points). By contrast, the impact of oil supply shocks on the level of interest rates thereafter, in special since the start of EMU, is found to be smaller.

Finally, looking at the contributions of oil supply and global demand shocks to world oil production, oil supply shocks are found to have contributed negatively to oil production in particular during the two oil shocks of the 1970s, around 1990/91 and around 2000 with basically continuous negative contributions thereafter. Global demand shocks, by contrast, are found to have had larger effects on world oil production in both directions up to the late 1980s, but thereafter they seem to have had much less capacity of moving world oil production. Overall, the two identified shocks together seem to explain over the whole period large parts of the movements of the series.

All in all, the evidence from the historical decompositions suggests that both oil supply and global demand shocks contributed significantly to the large oil price movements during historical oil crises and both shocks also contributed to driving up oil prices over the past decade, albeit over that period oil supply shocks to a larger extent than demand shocks. At the same time, the results show that both shocks had more harmful effects on the economy in the earlier part of the sample, i.e. from the 1970s up to the mid-1980s, than thereafter. Both the adverse contributions to HICP inflation and euro area/global GDP are estimated to have been much larger during that period than thereafter, an exception being the impact of demand shocks on GDP during the latest recession. This suggests that changes related to both oil supply shocks and global demand shocks have contributed to the Great Moderation, but also taken together the two shocks do not seem to account fully for the changes observed with the Great Moderation but further factors appear to have been at play as well.

3.4. Impulse responses

We now turn to examining whether the pass-through of the two identified shocks to the economy has changed since the 1970s. Figures 12, 28 and 30 display the median impulse responses of the variables to a one percent oil supply or global demand shock for a horizon of up to 20 quarters over the period 1970 to 2008/2009, respectively. The results point to substantial time-variation in the transmission of the considered shocks during that period but it is also important to note that the uncertainty of the estimates is rather large, which is quite common in time varying parameter models (see Figures 13, 29 and 31).

The results, starting with those of the oil supply shocks, can be summarised as follows: first, the response of oil prices denominated in euro to an oil supply shock seems to have decreased somewhat over time, but this effect is less evident when looking at oil prices in US dollar terms, such that it is not clear whether the potential change concerning oil prices in euro might perhaps just be due to changes in the exchange rate pass-through.

Second, concerning the HICP, both models show that an oil supply shock of the same magnitude has a much smaller inflationary effect today than in the 1970. There are some differences in the magnitudes of the estimated effects between the two models, probably owing to the differences in the currency denomination of the oil prices, but the general patterns are very similar. Both the contemporaneous as well as the long-term effect are estimated to have decreased since the 1970. The effects were very high throughout the 1970s but there was substantial time-variation also within that decade. Focusing in more detail on the quantitative estimates from the domestic model, at its peak during the 1970s, an oil price shock of one percent is estimated to have increased the HICP by around 0.5 percentage point after 5 years; the average long-run effect over that period amounted to about 0.3 percentage points, i.e. a pass-through of around 30 percent. Between around 1981 and mid-1990 the effects were clearly lower than in the 1970s, with the peak in the long-run effect now in the order of magnitude of less than 0.2 percentage points and the average magnitude of the long-run effects of about 0.1 percent. Since the mid-1990s the effects on the HICP have come down somewhat further. Now the peak in the long-run effect amounted to around 0.1 percentage points and the average effect since 1995 stands at about 0.07, i.e. a pass-through of 7%. For the global model the effects are even somewhat larger in the earlier part of the sample but lower over the recent period.

When compared with the results of models in the literature, the magnitude of the effects of the domestic model is somewhat larger than those identified in constant parameter VAR models for different sample periods (see ECB (2010)). This applies, in particular, to the results for the earlier part of the sample and is partly on account of the longer considered horizons of the impulse response functions (5 years compared with 3 years in the other analysis). Once this is accounted for, they are broadly comparable.

Third, the impulse responses for euro area GDP show large changes in the magnitude of the effects of an oil supply shocks of the same size over time as well. The adverse effects of an oil supply shock on GDP seem to have been much more muted in recent years than in earlier decades. But the particularly large effects in the first part of the sample appear to have been more concentrated in specific periods than was the case for the HICP effects, although also between these periods the adverse GDP effects were on average larger in the first part of the sample period than in the second. The estimates suggest that an oil supply shock of a certain magnitude would have had particularly large effects on GDP at the beginning of the 1970s, the periods of the two large oil price shocks of the 1970s, in 1985/86 and around 1988 when oil prices changed strongly and during the 1992/93 recession. A one percent increase in oil prices due to an oil supply shock is estimated to have reduced GDP by about 0.4 percentage points at the beginning of the 1970s and by around 0.2 to 0.25 percentage points during the other peaks. On average, the effect amounted to about 0.1 percentage point from 1970 to 1995 and about 0.04 percentage point in the period thereafter. The latter effect is close to what has been found in other models (see ECB (2010)). For global GDP a similar pattern of decline over time is observed as for euro area GDP.

As regards the other variables, there appear to have been also some changes in the response of interest rates to an oil supply shock of the same size over time. Interest rates are estimated to have been lowered following an oil supply shock in the early 1970s, but increased after an oil supply shock towards the end of the 1970s. The sign of the interest rate responses varied further in later years, but also for the interest rate response the absolute magnitude of the effect seems to have decreased over time. A declining effect over time is also found for the response of world oil production, which is estimated to be much less responsive to an oil supply shock today than in the first half of the sample period.

Turning to the impulse responses to the global demand shock, the inflationary response of the HICP has declined very strongly over time to that shock as well. Very large spikes in these impulse responses are found during the 1970s, while much more muted effects are estimated over the past decade. Also, the positive impact of a demand shock on global GDP has changed over time. This effect was relatively low in the 1970s where demand shocks appear to have been mostly transmitted to prices, but increased at the start of the 1980s and remained elevated until about 2000 (with a large spike around 1989) followed by a strong decline over the past decade. Moreover, as for oil supply shocks, oil production has become much less responsive to global demand shocks over time. By contrast, the response of oil prices to the demand shock increased somewhat since the 1970s.

Overall, the evidence from the impulse response functions suggests that the pass-through of both oil supply and global demand shocks to both HICP inflation and GDP has declined over the past decades, which has contributed to the moderating contribution of these shocks to the economy over time. The results of a declining magnitude over time in the transmission of oil supply shocks to prices and GDP in the euro area are in line with the available evidence on oil supply shocks in the literature for a larger number of other countries (see in particular Blanchard and Gali (2007), but also the other papers referred to above as well as the evidence provided by Baumeister and Peersman (2008), at least once their shocks are normalised to the same magnitude over time).

3.5. Variance decompositions

Finally, we look at variance decompositions to gain insights into whether the relative importance of the identified structural shocks might have changed over time. Figures 14 and 32 show the posterior mean of the time-varying forecast error variance decompositions of the variables in the two models 12 quarters after the shock. Concerning oil prices, the results highlight that oil supply shocks, while clearly not accounting for all of the forecast error variance of oil prices, do in fact explain a substantial share of it, which furthermore is relatively stable over time. It fluctuates mostly between 20 and 30 percent but, importantly, it does not show any systematic increase or decrease over time. Generally, global demand shocks are found to explain an only slightly smaller share of the oil price forecast error variance than oil supply shocks. Moreover, the share of these shocks varies somewhat over time, but again no systematic change is visible for them as well. Our results contrast with

Baumeister and Peersman's (2008) evidence of a decline in the share of the variance explained by oil supply shocks over the past years for USD denominated oil prices and their assessment that oil supply shocks appear to have become less important in explaining changes in oil prices over time as well as explicitly also with their conclusion for oil prices in USD of an increased importance of demand shocks, at least in the form of global demand shocks.

With regard to the other variables, the general picture from the forecast error variance decompositions is similar to that for oil prices with some variations over time in the relative importance of both oil supply and global demand shocks observed for all, but no systematic variation. For euro area GDP, for instance, there are some temporary increases in the share explained by oil supply shocks, in particular, in the mid-1970s and at the beginning of the 1990s (i.e. the importance of oil supply shocks seems to have increased temporarily during these recession periods) but overall the share has not changed much since the 1970s. Also for world GDP the share explained by oil supply shocks was relatively stable with some temporary fluctuations; the most notable change took place during the latest recession, where the share of oil supply shocks declined to its lowest over the whole period and that of demand shocks increased strongly, which suggests, in line with previous evidence, that demand shocks were the main driver of that recession. Finally, there is also some time variation in the share explained by the two shocks for the HICP, but again without any systematic increase or decrease.

4. Conclusions

In this paper we apply a VAR framework with time varying parameters and stochastic volatility to explore the role of oil prices for the macroeconomic developments of the euro area economy since the 1970s. Given numerous issues of interest, the analysis is conducted in the framework of two closely related complementary models and, in line with the recent oil price literature, besides the classical oil supply shocks also the impact of global demand shocks is examined.

Our empirical results show that both oil supply and global demand shocks contributed substantially to the oil price changes during historical oil crises and both shocks likewise contributed to driving up oil prices over the past decade, albeit concerning the latter period

oil supply shocks to a larger extent than global demand shocks. The results also reveal much larger adverse contributions of both types of shocks to HICP inflation and euro area/global GDP in the earlier part of the sample, i.e. from the 1970s up to the mid-1980s, than thereafter. An exception is the large negative contribution of the demand shocks on activity during the latest recession. These results suggest that changes related to both oil supply and global demand shocks have contributed to the Great Moderation, but they do not seem to account fully for the changes observed in the Great Moderation; other factors appear to have been at play as well. Given the findings of Benati and Surico (2009), we refrain from trying to disentangle the relative importance of the two traditionally explored opposite explanations for the Great Moderation. It is also not possible to analyse changes in the volatilities of our structural shocks over time in this framework, but the increase in the volatility of oil prices over time raises doubts about the possibility of a decline in the volatility of important shocks to oil prices over time. We have, however, examined whether the pass-through of the considered shocks might have changed over time. The results show that both oil supply and global demand shocks are transmitted to a smaller degree to prices and activity today than in the past and this change in the transmission mechanism has added to the moderating contribution of these shocks to the economy over time. At the same time, variance decompositions show fluctuations but no systematic increase or decrease in the relative importance of either oil supply or global demand shocks for the variables over time.

Overall, the evidence suggests that there have been substantial changes related to oil prices over time and these changes have played a role for the Great Moderation. Changes in the structure of the economy or policy, which affect the transmission of the considered shocks to oil prices, have likely contributed to the Great Moderation. Their relative importance for the Great Moderation compared with potential changes in the volatility of the shocks over time, however, is unclear and a topic for future research.

References

- Ahmed, S., Levin, A. and B. Wilson (2004), "Recent U.S. Macroeconomic Stability: Good Policies, Good Practices, or Good Luck?", *The Review of Economics and Statistics*, Vol. 86, No. 3, pp. 824-832.
- Barsky, R. and L. Kilian (2004), "Oil and the Macroeconomy Since the 1970s", *Journal of Economic Perspectives*, Vol. 18, No. 4, pp. 115-134.
- Baumeister, C. and G. Peersman (2008), "Time-Varying Effects of Oil Supply Shocks on the US Economy", mimeo, University of Ghent.
- Baumeister, C. and G. Peersman (2009), "Sources of the Volatility Puzzle in the Crude Oil Market", mimeo, University of Ghent.
- Benati, L. and H. Mumtaz (2007), "U.S. Evolving Macroeconomic Dynamics: A Structural Investigation", ECB Working Paper, No. 746 (April 2007).
- Benati, L. and P. Surico (2009), "VAR Analysis and the Great Moderation," *American Economic Review*, vol. 99(4), pp. 1636-52.
- Blanchard, O. and J. Simon (2001), "The Long and Large Decline in U.S. Output Volatility", *Brookings Papers on Economic Activity*, No. 1, pp. 135-174.
- Blanchard, O. and J. Gali (2007), "The Macroeconomic Effects of Oil Shocks: Why are the 2000s so different from the 1970s?", NBER Working Paper, No. 13368.
- Boivin, J. and M. Giannoni (2006), "Has Monetary Policy Become More Effective?", *The Review of Economics and Statistics*, Vol. 88, No. 3, pp. 445-462.
- Brown, Yücel and Thompson (2003), "Business Cycles: The Role of Energy Prices", Working Paper 0304, Federal Reserve Bank of Dallas.
- Canova, F. and G. de Nicolo (2002), "Monetary Disturbances Matter for Business Cycle Fluctuations in G-7", *Journal of Monetary Economics*, Vol. 49, No. 6, pp. 1131-1159.
- Canova, F., L. Gambetti and E. Pappa (2007), "The Structural Dynamics of Output Growth and Inflation: Some International Evidence", *Economic Journal*, Vol. 117, pp. 167-191.
- Canova, F. and M. Ciccarelli (2009), "Estimating Multicountry VAR Models", *International Economic Review*, Vol. 50, No. 3, August.
- Cogley, T. and T. J. Sargent (2005), "Drifts and Volatilities: Monetary Policies and Outcomes on the Post WWII U.S.", *Review of Economic Dynamics*, Vol. *, No. 2, pp. 262-302.
- De Gregorio, J., Landerrechte, O. and C. Neilson (2007), "Another Pass through Bites the Dust? Oil Prices and Inflation", Central Bank of Chile Working Paper, No. 417.
- Del Negro, M. (2003), "Discussion of Cogley and Sargent's "Drifts and Volatilities: Monetary Policy and Outcomes in the Post WWII U.S.", Federal Reserve Bank of Atlanta Working Paper, No. 2003-26, October 2003.
- ECB (2004), "Oil Prices and the Euro Area Economy", *Monthly Bulletin*, November 2004, European Central Bank.

- ECB (2010), "Energy Markets and the Euro Area Macroeconomy", Occasional Paper, No. 113, European Central Bank.
- Edelstein, P. and L. Kilian (2007), "The Response of Business Fixed Investment to Changes in Energy Prices: A Test of Some Hypotheses About the Transmission of Energy Price Shocks", *The B.E. Journal of Macroeconomics*, Vol. 7, No. 1, Contributions.
- Faust, J. (1998), "The Robustness of Identified VAR Conclusions about Money", *Carnegie Rochester Series on Public Policy*, Vol. 49, pp. 207-244.
- Hahn, E. (2003), "Pass-Through of External Shocks to Euro Area Inflation", ECB Working Paper, No. 243.
- Hamilton, J. D. (1983), "Oil and the Macroeconomy since World War II", *Journal of Political Economy*, Vol. 91, No. 2, pp. 228 – 248.
- Hamilton, J. D. (1996), "This is what happened to the Oil Price-Macroeconomy Relationship", *Journal of Monetary Economics* 38, pp. 215 – 220.
- Hamilton, J. D. (2008), "Oil and the Macroeconomy", In: *The New Palgrave Dictionary of Economics*, Second Edition, Ed. S. N. Durlauf and L. E. Blume, Houndmills, UK and New York, Palgrave Macmillan.
- Hamilton, J. D. (2009), "Causes and Consequences of the Oil Shock of 2007-08", *Brookings Papers on Economic Activity*, Spring 2009.
- Herrera, A. M. and E. Pesavento (2009), "Oil Price Shocks, Systematic Monetary Policy, and the "Great Moderation", *Macroeconomic Dynamics*, 13, pp. 107 – 137.
- Hooker, M. A. (1996a), "What Happened to the Oil Price – Macroeconomy Relationship?", *Journal of Monetary Economics* 38, pp. 195 – 213.
- Hooker, M. A. (1996b), "This is What Happened to the Oil Price – Macroeconomy Relationship: Reply", *Journal of Monetary Economics* 38, pp. 221 – 222.
- Hooker, M. A. (2002), "Are Oil Shocks Inflationary? Asymmetric and Nonlinear Specifications versus Changes in Regime?", *Journal of Money, Credit and Banking*, Vol. 34, No. 2 (May 2002), pp. 540 – 561.
- Jiménez-Rodríguez, R. and M. Sanchez (2004), "Oil Price Shocks and Real GDP Growth: Empirical Evidence for some OECD Countries", *Applied Economics*, Vol. 37, No. 2, pp. 201-228.
- Kilian, L. (2008), "Exogenous Oil Supply Shocks: How Big are They and How Much do They Matter for the U.S. Economy", *The Review of Economics and Statistics* 90, No. 2, pp. 216-240.
- Kilian, L. (2009), "Not All Oil Price Shocks are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market", *American Economic Review* 99, No. 3, pp. 1053-1069.
- Lee, K., Ni, S. and R. Ratti (1995), "Oil Shocks and the Macroeconomy: The Role of Price Variability", *The Energy Journal* 16, pp. 39-56.

- McConnell, M. and G. Perez-Quiros (2000), "Output Fluctuations in the United States: What Has Changed Since the Early 1980's?", *The American Economic Review*, Vol. 90, No. 5 (December 2000), pp. 1464-1476.
- Mork, K. A. (1989), "Oil and the Macroeconomy When Prices Go Up and Down: An Extension of Hamilton's Results", *The Journal of Political Economy*, Vol. 97, No. 3 (June 1989), pp. 740-744.
- Primiceri, G. (2005), "Time Varying Structural Vector Autoregressions and Monetary Policy", *The Review of Economic Studies*, Vol. 72 (July 2005), pp. 821-852.
- Peersman, G. (2005), "What Caused the Early Millennium Slowdown? Evidence Based on Vector Autoregressions", *Journal of Applied Econometrics*, Vol. 20, pp. 185-207.
- Peersmann, G. and I. Van Robays (2009), "Oil and the Euro Area Economy", *Economic Policy*, Vol. 24, pp. 603 – 651.
- Smets, F. and R. Wouters (2007), "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach", *The American Economic Review*, Vol. 97, No. 3, pp. 586-606.
- Stock, J. and M. Watson (2002), "Has the Business Cycle Changed and Why?", *NBER Macroeconomics Annual*, Vol. 17, pp. 159-218.
- Tatom, J. A. (1988), "Are the Macroeconomic Effects of Oil-Price Changes Symmetric?", *Carnegie Rochester Conference Series on Public Policy* 28, pp. 325 – 368.
- Uhlig, H. (2005), "What are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure", *Journal of Monetary Economics*, Vol. 52, pp. 381-419.

Appendix A: Data Sources

(1) Real oil prices in USD and Euro:

1947/1960 – 1969: Nominal oil price represented by West Texas intermediate (source: Dow Jones & Company, Wall Street Journal) deflated by US Consumer price index for all urban consumers: All items (source: US Department of Labour, Bureau of Labour Statistics).

1970 to 2009: Nominal world crude oil price in USD (source: Economic Outlook, OECD). For oil prices in Euro converted to Euro with the bilateral Euro/USD exchange rate (source: 1970: weighted average of euro area country exchange rates vis-à-vis the USD; 1971 to 1998: AWM database; 1999 to 2009: ECB). Deflation to real oil prices with the GDP deflator (source: 1970 to 1994: AWM database; 1995 to 2009: Eurostat).

(2) HICP

1947/1960 – 1969: US Consumer price index for all urban consumers: All items (source: US Department of Labour, Bureau of Labour Statistics).

1970 to 2009: HICP (source: 1970 to 1979: AWM database; 1980 to 2009: ECB and own calculations).

(3) GDP

1947 – 1969: US GDP (source: US Department of Commerce, Bureau of Economic Analysis).

1970 to 2009: Euro area GDP (source: 1970 to 1994: AWM database; 1995 to 2009: Eurostat)

(4) Interest rate

1947 – 1969: Three month Treasury bill rate, secondary market rate (source: Board of Governors of the Federal Reserve System).

1970 to 2009: 3-month interest rate/Euribor (source: 1970 to 1993: AWM database; 1994 to 2009: ECB).

(5) World GDP

1960 to 1969: World GDP approximated by industrial production data of the US (source: FRED database), euro area (source: Eurostat and own computations), Japan (source: annual industrial production data from Ameco database interpolated with Chow Lin methodology

using US industrial production as reference series) and UK (source: annual data industrial production data from Ameco database interpolated with Chow Lin methodology using US industrial production as reference series). Country data weighted together with relative GDP weights of 1970 in PPP from Ameco database

1970 to 2008: Sum of annual GDP from a number of countries, changing composition depending on availability (included countries are Australia, Austria, Belgium, Canada, Switzerland, Chile, China, Czech Republic, German, Denmark, Spain, Estonia, Finland, France, UK, Greece, Hungary, India, Ireland, Iceland, Israel, Italy, Japan, Korea, Luxembourg, Mexico, Netherlands, Norway, New Zealand, Poland, Portugal, Slovak Republic, Slovenia, Sweden, Turkey, USA, all GDP volume at the price levels and PPP of 2000 (million USD)) (source: Economic Outlook, OECD). A quarterly GDP series was derived by interpolating the annual series with the Chow Lin methodology. The quarterly series used as reference series for interpolation with the Chow Lin methodology was the quarterly GDP of a number of countries with changing composition. The quarterly series has a lower country coverage than the annual series (included countries are Australia, Austria, Belgium, Canada, Switzerland, Czech Republic, Germany, Denmark, Spain, Estonia, Finland, France, UK, India, Ireland, Iceland, Israel, Italy, Japan, Korea, Luxembourg, Mexico, Netherlands, Norway, New Zealand, Poland, Portugal, Slovak Republic, Sweden, US, all GDP volume at the price levels and PPP of 2000 (million USD)) (source: Economic Outlook, OECD).

(6) World oil production

1960 to 1972: Annual data of world oil production (source: Earth and Policy Institute, BP, IEA and Worldwatch) interpolated to quarterly data with Boot Freibes Lisman interpolation method.

1973 to 2008: World crude oil production data (source: US Department of Energy).

Appendix B: Bayesian Estimation

1. The model

Following Primiceri (2005), the model is in the form of a VAR with time-varying parameters and stochastic volatility, as in (1). The model is represented with 2 lags, but the extension to a different number of lags is straightforward.

$$\begin{aligned} y_t &= B_{0,t} + B_{1,t}y_{t-1} + B_{2,t}y_{t-2} + u_t \\ u_t &= A_t^{-1}H_t\varepsilon_t, \quad \varepsilon_t \sim \text{iid}(0, I_n) \end{aligned} \quad (1)$$

Data are in the form of an $n \times 1$ column vector y_t for period t . Dynamic parameters include a deterministic component, $B_{0,t}$, in our case an n -dimensional column vector with constants, and $n \times n$ matrices $B_{1,t}$ and $B_{2,t}$ for each of the 2 lags used in the paper. H_t is diagonal, with vector h_t on the diagonal. A_t is a lower-triangular $n \times n$ matrix with 1s (ones) on the diagonal; elements below the diagonal are gathered in a vector α_t .¹⁵ The error terms ε_t are uncorrelated and with unit variance.

The structure of the H_t matrix is shown in (2).

$$H_t \equiv \begin{bmatrix} h_{1,t} & 0 & \dots & 0 \\ 0 & h_{2,t} & \dots & 0 \\ \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & h_{n,t} \end{bmatrix} \quad (2)$$

The structure of the A_t matrix is shown in (3). Note that the notation of the elements in the α_t vector differentiate across equations, i.e. element $\alpha_{ij,t}$ is, correspondingly, the i^{th} contemporaneous parameter appearing in the j^{th} equation of the VAR, for period t . This distinction is important for the distributional assumptions made later.

$$A_t \equiv \begin{bmatrix} 1 & 0 & \dots & 0 & 0 \\ \alpha_{21,t} & 1 & \dots & 0 & 0 \\ \alpha_{31,t} & \alpha_{32,t} & \dots & 0 & 0 \\ \dots & \dots & \dots & \dots & \dots \\ \alpha_{n1,t} & \alpha_{n2,t} & \dots & \alpha_{nn,t} & 1 \end{bmatrix} \quad (3)$$

The model can be rewritten in vectorised form:

$$y_t = \tilde{y}_t B_t + u_t \quad (4)$$

$$\tilde{y}_t \equiv [1 \ y'_{t-1} \ y'_{t-2}] \otimes I_n, \quad B_t \equiv \text{vec}([B_{0,t} \ B_{1,t} \ B_{2,t}]) \quad (5)$$

¹⁵ Note that this imposes a Cholesky structure to the system. The estimated system is nevertheless a reduced-form one, the Cholesky orthogonalisation being merely a convenient device to pin down the posterior distribution.

The model is completed by assumed laws of motion for the time-varying parameters:

$$B_t = B_{t-1} + \eta_t, \quad \eta_t \sim \text{niid}(0, Q_B) \quad (6)$$

$$\alpha_t = \alpha_{t-1} + \mu_t, \quad \mu_t \sim \text{niid}(0, Q_A) \quad (7)$$

$$\ln h_t = \ln h_{t-1} + \nu_t, \quad \nu_t \sim \text{niid}(0, Q_H) \quad (8)$$

Note that the only parameters to estimate in (6), (7) and (8) are the elements entering the variance-covariance matrices Q_B , Q_A and Q_H . These matrices will be referred to below as *hyper-parameters*, while the term *VAR parameters* will be used for the parameters in the VAR system,(1). These matrices are assumed full (with an exception mentioned below), i.e. the time-variation in the parameters is assumed correlated across them.

Note also that the hyper-parameters are themselves not time-varying. Previous studies have adopted this convention, see e.g. Benati and Mumtaz (2008), in settings where hyper-parameters were diagonal matrices. The models estimated in the current paper stick instead to the more general approach followed by Primiceri (2005), in which parameters are allowed to be correlated among them. This approach cannot be easily extended to time-varying variance-covariance matrices due to the already large number of parameters estimated.

2. Estimation

Given a sample spanning periods 1 to T , the parameters $B_1^T = [B_1, B_2, \dots, B_T]$, $\alpha_1^T = [\alpha_1, \alpha_2, \dots, \alpha_T]$, $h_1^T = [h_1, h_2, \dots, h_T]$, Q_B , Q_A and Q_H are estimated by Bayesian methods. A prior is postulated for the parameters and the joint posterior is calculated according to (where $y_1^T \equiv [y_1, y_2, \dots, y_T]$):

$$p(B_1^T, \alpha_1^T, h_1^T, Q_B, Q_A, Q_H | y_1^T) \propto I(B_1^T) p(B_0, \alpha_0, h_0, Q_B^0, Q_A^0, Q_H^0) p(y_1^T | B_1^T, \alpha_1^T, h_1^T, Q_B, Q_A, Q_H) \quad (9)$$

In the expression, $I(B_1^T)$ is an index whose value is zero if any of the elements in B_1^T is dynamically unstable, i.e. all unstable draws are assigned probability zero and rejected.¹⁶

In order to ease the computations, a battery of Gibbs samplers is used. Priors are also defined. For each model a chain of 50000 draws is estimated. The first 20000 draws are used as burn in period and in addition thinning steps of 30 are used.

2.1 Priors

The VAR parameters (B , A , H) are all assumed to have a normal prior with mean and variance obtained from a prior sample, see (10) to (12). Priors for the B parameters are obtained by fitting a constant-parameter VAR system similar to (1) over the prior sample, using standard OLS; priors for the α parameters are obtained from the (Cholesky) orthogonalisation of reduced-form residuals of the

¹⁶ Rejecting unstable draws is standard in the literature but not necessary for most results reported in the paper.

OLS estimation, exploiting the recursive structure of A ; priors for $\ln h$ are obtained from the variances of the derived orthogonal error terms. The variances of the B and A parameters are increased by a factor of 10 to reduce the impact of the prior sample on the posterior estimates.

$$B_0 \sim N\left(\hat{B}_{OLS}, 10 \text{var}\left(\hat{B}_{OLS}\right)\right) \quad (10)$$

$$\alpha_0 \sim N\left(\hat{\alpha}_{OLS}, 10 \text{var}\left(\hat{\alpha}_{OLS}\right)\right) \quad (11)$$

$$\ln h_0 \sim N\left(\ln \hat{h}_{OLS} + \ln 10, I_n\right) \quad (12)$$

The hyper-parameters are all assumed to be inverse Wisharts with the minimum number of degrees of freedom (d.o.f.) permissible, see (13) to (15). For Q_B , the d.o.f. are the size of the prior sample. For Q_A , the d.o.f. are given by the size of the blocks conforming the matrix, i.e. the number of α parameters entering each of the equations in the system $u_t = A_t^{-1} H_t \varepsilon_t$, plus one: 2 d.o.f. in the 2nd equation, 3 d.o.f. in the 3rd one, and so on. Finally, for Q_H the d.o.f. are given by the number of variables in the VAR system plus one.

The prior-distribution parameter $\text{var}\left(\hat{\alpha}_{OLS}\right)$ is not a full matrix. It will be assumed that elements of $\hat{\alpha}_{OLS}$ pertaining to different contemporaneous equations are uncorrelated. The rationale for this assumption, together with a more detailed description of the resulting structure, is left for the discussion on the posterior draws.

$$Q_B^0 \sim IW\left(k_B^2 T_0 \text{var}\left(\hat{B}_{OLS}\right), T_0\right) \quad (13)$$

$$Q_A^0 \sim IW\left(k_A^2 (n_i + 1) \text{var}\left(\hat{\alpha}_{OLS}\right), (n_i + 1)\right) \quad (14)$$

$$Q_H^0 \sim IW\left(k_H^2 (n + 1) I_n, n + 1\right) \quad (15)$$

Constants k_B , k_A and k_H scale down the OLS variances for the B , A and H parameters, as done by Primiceri (2005). Their values are standard in the literature: $k_B=0.001$, $k_A=0.01$, $k_H=0.001$.

2.2 Estimation of B_t

Assuming $\{\alpha_1^T, h_1^T, Q_B, Q_A, Q_H\} \equiv \theta_B$ known, the parameter B_1^T can be computed recursively by drawing from a smoother for the state-space representation in (6) and (4), using Carter and Kohn (1994) formulas.

$$p\left(B_T^T \mid \theta_B, y_1^T\right) = p\left(B_T \mid \theta_B, y_1^T\right) \prod_{t=1}^{T-1} I(B_t) p\left(B_t \mid B_{t+1}, \theta_B, y_1^T\right) \quad (16)$$

Where $p\left(B_T \mid \theta_B, y_1^T\right)$ is obtained from the distribution of the standard Kalman filter recursions starting from B_0 and $I(B_t)$ indicates whether the B_t drawn is dynamically stable.

2.3 Estimation of α_t

Expressing the system (1) and (7) as a standard state space entails re-expressing the measurement equation. For this, we need to exploit the specific structure of A_t .

Assuming $\{B_1^T, h_1^T, Q_B, Q_A, Q_H\} \equiv \theta_\alpha$ known, u_1^T can be computed. (Note that u_t is observable if B_t is known). The system (1) can be summarised as $A_t u_t = H_t \varepsilon_t$, which is a set of disjoint systems of equations because of the recursive structure of A_t . The system can be combined with corresponding elements in (7) to build a battery of state-space systems.

For instance, the second equation in $A_t u_t = H_t \varepsilon_t$, together with the law of motion for $\alpha_{1,t}$, can be written:

$$\begin{aligned}\alpha_{21,t} &= \alpha_{21,t-1} + \mu_{2,1,t} \\ u_{2,t} &= -u_{1,t} \alpha_{21,t} + h_{2,t} \varepsilon_{2,t}\end{aligned}\tag{17}$$

The third equation:

$$\begin{aligned}\begin{bmatrix} \alpha_{31,t} \\ \alpha_{32,t} \end{bmatrix} &= \begin{bmatrix} \alpha_{31,t-1} \\ \alpha_{32,t-1} \end{bmatrix} + \begin{bmatrix} \mu_{3,1,t} \\ \mu_{3,2,t} \end{bmatrix} \\ u_{3,t} &= -\begin{bmatrix} u_{1,t} & u_{2,t} \end{bmatrix} \begin{bmatrix} \alpha_{31,t} \\ \alpha_{32,t} \end{bmatrix} + h_{3,t} \varepsilon_{3,t}\end{aligned}\tag{18}$$

The fourth equation:

$$\begin{aligned}\begin{bmatrix} \alpha_{41,t} \\ \alpha_{42,t} \\ \alpha_{43,t} \end{bmatrix} &= \begin{bmatrix} \alpha_{41,t-1} \\ \alpha_{42,t-1} \\ \alpha_{43,t-1} \end{bmatrix} + \begin{bmatrix} \mu_{4,1,t} \\ \mu_{4,2,t} \\ \mu_{4,3,t} \end{bmatrix} \\ u_{4,t} &= -\begin{bmatrix} u_{1,t} & u_{2,t} & u_{3,t} \end{bmatrix} \begin{bmatrix} \alpha_{41,t} \\ \alpha_{42,t} \\ \alpha_{43,t} \end{bmatrix} + h_{4,t} \varepsilon_{4,t}\end{aligned}\tag{19}$$

And so on if further equations are needed in the system, i.e. in case the number of variables in the VAR is higher than 4.

In order to be able to run the battery of Kalman systems in isolation, it is further assumed that the parameters are independent across equations, as done by Primiceri (2005). In practice, this means that Q_A has a block-diagonal structure, with zero in those elements of the variance-covariance matrix corresponding to α_t parameters in different equations in the VAR. This implies that the contemporaneous parameters in the VAR are correlated across them if they belong to the same equation, uncorrelated otherwise. As shown by Primiceri, this assumption can only be relaxed at the cost of forsaking the possibility of using Gibbs draws, which increases the numerical complexity of the estimation.

As for the B_1^T VAR parameters, each of the Kalman systems is run using the Carter and Kohn (1994) formulas:

$$p(\alpha_1^T | \theta_\alpha, y_1^T) = p(\alpha_T | \theta_\alpha, y_1^T) \prod_{t=1}^{T-1} p(\alpha_t | \alpha_{t+1}, \theta_\alpha, y_1^T)\tag{20}$$

Where $p(\alpha_T | \theta_\alpha, y_1^T)$ is obtained from the distribution of the standard Kalman filter recursions starting from α_0 .

2.4 Estimation of h_t

The diagonal elements of H_t must obviously be positive. Following Primiceri (2005), they are modelled in natural logarithm.

Assuming $\{B_1^T, \alpha_1^T, Q_B, Q_A, Q_H\} \equiv \theta_h$ known, $\tilde{\varepsilon}_t^T$ can be computed, where $\tilde{\varepsilon}_t \equiv A_t u_t = H_t \varepsilon_t$. Each of the elements of $\tilde{\varepsilon}_t$ can then be squared and its log mapped to the corresponding h_t . The i -th equation of the mapping would be:

$$\ln(\tilde{\varepsilon}_{i,t}^2 + \bar{c}) = 2 \ln h_{i,t} + \ln \varepsilon_{i,t}^2 \quad (21)$$

The constant \bar{c} is an offset to avoid numerical problems when $\tilde{\varepsilon}_{i,t}$ is very small; as in Primiceri (2005), it is set as $\bar{c} = 0.001$.

Stacking all the elements sets the system in vector form, as in (22), where $\tilde{\varepsilon}_t^*$ and ε_t^* are the column vectors collecting all the $\tilde{\varepsilon}_{i,t}^2$ and $\varepsilon_{i,t}^2$, respectively.

$$\ln \tilde{\varepsilon}_t^* = 2 \ln h_t + \ln \varepsilon_t^* \quad (22)$$

Equation (22), together with (8) forms a Kalman filter with a non-Gaussian error term in the measurement equation, each of which element is distributed as a $\ln \chi^2(1)$.¹⁷ In order to ease computations, a normal approximation is made using the approach of Kim, Shepard and Chib (1998). A mixture of 7 normal distributions is made to approximate each of the $\ln \varepsilon_{i,t}^2$, whose means, variances and mixture weights are calculated to approximate moments of a $\ln \chi^2(1)$ distribution.

The procedure entails drawing from the 7 normal distributions and giving each draw a probability according to its mixture weight. The specific normal used in each draw is selected with an indicator variable which is proportional to the mixture weight and the likelihood of $\ln(\varepsilon_{i,t}^2 + \bar{c})$ being a draw from the corresponding normal distribution, see (23).

$$\ln \varepsilon_{i,t}^2 \sim N(m_j - 1.2704, v_j^2) \text{ with probability } \Pr(s_{i,t} = j), j = 1, \dots, 7 \quad (23)$$

With the distribution of the $\ln \varepsilon_{i,t}^2$ errors so defined, equation (22) becomes the measurement equation of a Kalman system whose transition equation is given by (8). Draws are obtained, as before, using the Carter and Kohn (1994) formulas:

$$p(\ln h_1^T | \theta_h, y_1^T) = p(\ln h_T | \theta_h, y_1^T) \prod_{t=1}^{T-1} p(\ln h_t | \ln h_{t+1}, \theta_h, y_1^T) \quad (24)$$

The specific value drawn is selected from (23) using an indicator variable $s_{i,t}$ ($0 < s_{i,t} < 1$), itself drawn from (25), where $\Phi(x)$ is the CDF of a normal $N(0,1)$ for value x .

$$\Pr(s_{i,t} = j) \propto q_j \Phi_N \left(\frac{\ln \varepsilon_{i,t}^2 - m_j - 1.2704}{v_j} \right) \quad (25)$$

¹⁷ Note that the $\tilde{\varepsilon}_t^*$ processes are, by construction, uncorrelated among them. This implies that the ε_t^* processes are also uncorrelated, justifying the assumption that their distribution is a battery of $\ln \chi^2(1)$. But also note that Q_H is not constrained in the estimation.

The procedure entails drawing from the 7 normal distributions, calculating the probabilities in (25) and re-scaling them to add up to one. The actual draw is extracted by taking the draw from the 7 normal distributions whose probability is lower or equal than an index draw from the uniform distribution. This is performed for each period in the estimation sample.

The normal distributions are defined in Table 2-1.

Table 0-1. Mixing Normal Distributions

j	q _j	m _j	v _j
1	0.00730	-10.12999	5.79596
2	0.10556	-3.97281	2.61369
3	0.00002	-8.56686	5.17950
4	0.04395	2.77786	0.16735
5	0.34001	0.61942	0.64009
6	0.24566	1.79518	0.34023
7	0.25750	-1.08819	1.26261

2.5 Estimation of Q_B , Q_A and Q_h

Assuming B_1^T, α_1^T, h_1^T known, the error processes η_1^T, μ_1^T and ν_1^T of the law of motions of the parameters become observable. An estimate of their variance-covariances can be obtained by the usual formulas, which is then used to draw from an inverse Wishart distribution. This distribution is mixed to the prior distribution to result in the posterior draw.

2.6 Summary

The previous steps describe one single iteration of the Gibbs sampler. Each iteration (say, the i^{th} iteration) can be collected in the following sequence:

1. The B_1^T VAR parameters of the iteration are drawn using (16), assuming all other parameters known—in practice, they will be taken from the previous Gibbs iteration.
2. The α_1^T contemporaneous VAR parameters are drawn using (17), (18) or (19), as it may correspond, also assuming other parameters known.
3. The h_1^T VAR variances are drawn using (22) and (8), again assuming other parameters known.
4. The Q_B , Q_A and Q_H hyper-parameters are drawn using the formulas for the conjugate posterior distributions that apply:

$$Q_B^{(i)} = IW\left(T_0 Q_B^0 + T \text{var}\left(B_t^{(i-1)} - B_{t-1}^{(i-1)}\right), T_0 + T\right) \quad (26)$$

$$Q_A^{(i)}(k) = IW\left(k Q_A^0(k) + T \text{var}\left(\alpha_t^{(i-1)}(k) - \alpha_{t-1}^{(i-1)}(k)\right), k + T\right), k = 2, \dots, 4 \quad (27)$$

$$Q_h^{(i)} = IW\left(5 Q_h^0 + T \text{var}\left(h_t^{(i-1)} - h_{t-1}^{(i-1)}\right), 5 + T\right) \quad (28)$$

Where (27) applies to equation k of the VAR, since the contemporaneous parameters of the VAR are correlated only when they belong to the same equation.

5. The draw is rejected if any of the periods is dynamically unstable.

3. Computation of impulse responses

Koop et al. (1996) (KPP, henceforth) have proposed the basic procedure followed today to derive the impulse response functions (IRF, henceforth) for a non-linear model. Their procedure entails calculating the difference between one (stochastic) simulation with the model with all shocks according to their distribution and another simulation using the same set of shocks except one instance (i.e., a value for a given shock and a given period). This instance is replaced by a deterministic value. The proposal by KPP is to integrate out all the non-deterministic shocks.

KPP's procedure entails drawing repeatedly from the all the shocks in the model (i.e., ε_t , η_t , μ_t and v_t) and performing two simulations of (1): one with the complete draw and the other one with the same draw but one value for one shock (say, the p^{th} shock) replaced by a deterministic value δ selected by the analyst. The IRF for this draw is the difference between the two. The procedure is repeated for as many times as wished, and the resulting distribution is summarised in an appropriate manner.

Two caveats: firstly, as is done in the literature, both for VAR models in general and for VAR models of the kind used in this paper, we will concentrate on orthogonal shocks, the ε_t processes in our case; second, past history will not be concentrated out.¹⁸ Furthermore, we will not necessarily report the expectation of the resulting distribution as done by KPP but will refer below to the entire distribution.¹⁹

We will depart from KPP by following Gourieroux and Jasiak (2005) (GJ henceforth), in that we will replace KPP's deterministic value δ by a stochastic value corresponding to the drawn value for the shock *plus* a deterministic value, $\varepsilon_t + \delta e_p$, where e_p is a selection vector with one in the p^{th} position and zeroes elsewhere, and δ is a scalar. Note that this definition maps into a traditional impulse response function if the model considered is a linear VAR, as is the case with KPP; furthermore, the resulting distribution is identical for our model from what it would be using KPP's definition of the shock. We do the change because it slightly simplifies our calculations and (incidentally) because we prefer GJ's measure on intuitive grounds.

In a model of this kind, the conditioning information is crucial. The distribution of the parameters will change depending on what is assumed to be known, an occurrence not to be found in linear VARs. The assumption made in traditional VARs is that the entire sample is known and parameters are extracted from the OLS sample distribution. Applying this principle to our model, parameters should come from the posterior *smoothed* estimates of the parameters. If, instead, the conditioning only goes until the period from which simulations begin, let's call it t , the parameters should be drawn from posterior *filtered* estimates. If we note the entire sample by $y_1^T \equiv [y_1, y_2, \dots, y_T]$ and the partial sample by $y_1^t \equiv [y_1, y_2, \dots, y_t]$ ($t < T$), then a draw of period t $\{B_t, \alpha_t, h_t \mid y_1^t\}$ can be obtained from

¹⁸ That is to say, we will condition all our draws in the history previous to the initial period simulated. The KPP procedure allows in principle to also concentrate out past history.

¹⁹ KPP report the IRF as the expectation of the difference. We will tackle the entire distribution, which will then be reported using any statistic chosen by the analyst.

the posterior estimates of the parameters, since we use a smoother in the estimation. If instead $\{B_t, \alpha_t, h_t | y_1^t\}$ must be drawn, we would need a draw from the estimated posterior assuming we had used Kalman filters in the estimation. Furthermore, draws for subsequent periods would also be different: $\{B_{t+s}, \alpha_{t+s}, h_{t+s} | y_1^t\}$ ($s > 0$) can also be drawn from the posterior smoother, provided $t+s < T$, while $\{B_{t+s}, \alpha_{t+s}, h_{t+s} | y_1^t\}$ would need a forecast of the parameters using the filter.

Our choice has obviously been to derive the IRF using the full-sample conditioning, since this allows us to use directly the Gibbs draws of the estimation. We have nevertheless also calculated the partial-sample IRF, but using the smoothed posterior estimates for period t (for $t < T$) for simplicity.²⁰ In this case, we simulate the periods subsequent to the one with the shock, using the corresponding transition equations. For completeness, we have also calculated IRF in the traditional way, i.e. keeping parameters fixed at their value in t . This latter procedure is technically wrong but has been used in the past for this kind of models and is thus a good yardstick.

Last but not least, the specific form of the model allows for one simplification in the calculations. As said, the two simulations involved in the IRF have identical shock draws, except for one specific value which affects one orthogonal shock for one period. As the model is linear in the orthogonal shocks, it can be easily proved that the only shocks in the calculation that need to be projected past period t are the innovation errors to the dynamic VAR parameters.

This can be easily shown. Assuming that the initial shock takes place in period t , and that we want to simulate the IRF until period $t+S$, the calculations would be as follows:

The baseline simulations for periods t to $t+S$:

$$\begin{aligned}\bar{y}_t &= B_{0,t} + B_{1,t}y_{t-1} + B_{2,t}y_{t-1} + A_t^{-1}H_t\varepsilon_t \\ \bar{y}_{t+1} &= B_{0,t+1} + B_{1,t+1}\bar{y}_t + B_{2,t+1}y_{t-1} + A_{t+1}^{-1}H_{t+1}\varepsilon_{t+1} \\ \bar{y}_{t+s} &= B_{0,t+s} + B_{1,t+s}\bar{y}_{t+s-1} + B_{2,t+s}\bar{y}_{t+s-1} + A_{t+s}^{-1}H_{t+s}\varepsilon_{t+s}, \quad s = 2, \dots, S\end{aligned}$$

The shocked simulations for periods t to $t+S$:

$$\begin{aligned}\tilde{y}_t &= B_{0,t} + B_{1,t}y_{t-1} + B_{2,t}y_{t-1} + A_t^{-1}H_t(\varepsilon_t + e_p \delta) \\ \tilde{y}_{t+1} &= B_{0,t+1} + B_{1,t+1}\tilde{y}_t + B_{2,t+1}y_{t-1} + A_{t+1}^{-1}H_{t+1}\varepsilon_{t+1} \\ \tilde{y}_{t+s} &= B_{0,t+s} + B_{1,t+s}\tilde{y}_{t+s-1} + B_{2,t+s}\tilde{y}_{t+s-1} + A_{t+s}^{-1}H_{t+s}\varepsilon_{t+s}, \quad s = 2, \dots, S\end{aligned}$$

The IRF would be (for each period) the difference between the two.

For period t :

$$IRF_t \equiv \tilde{y}_t - \bar{y}_t = A_t^{-1}H_t e_p \delta$$

For period $t+1$:

$$IRF_{t+1} \equiv \tilde{y}_{t+1} - \bar{y}_{t+1} = B_{1,t+1}(\tilde{y}_t - \bar{y}_t) = B_{1,t+1} IRF_t$$

²⁰ This implies that we are under-estimating parameter variability for most of the sample, increasingly as we move backward in time.

For period $t+s$:

$$\begin{aligned} IRF_{t+s} &\equiv \tilde{y}_{t+s} - \bar{y}_{t+s} \\ &= B_{1,t+s} (\tilde{y}_{t+s-1} - \bar{y}_{t+s-1}) + B_{2,t+s} (\tilde{y}_{t+s-2} - \bar{y}_{t+s-2}) \\ &= B_{1,t+s} IRF_{t+s-1} + B_{2,t+s} IRF_{t+s-2} \end{aligned}$$

The only stochastic information needed, accordingly, is the value in period t of A_t and H_t and the values in periods $t+1$ to $t+S$ of $B_{1,t+s}$ and $B_{2,t+s}$.

3.1 Summary of IRF computations

We shall begin with the filtered version of the computations.

For as many times as requested, repeat the following steps:

1. Draw B_t, A_t, H_t from the posterior distribution.
2. Draw from $\eta_{t+1}^{t+S} \equiv [\eta_{t+1} \dots \eta_{t+S}]$.
3. Compute sequence of parameters for the $S-1$ periods to simulate (the first period is given by the Gibbs draw) of the B_{t+s} parameters:

$$B_{t+s} = B_{t+s-1} + \eta_{t+s}, \quad s = 1, \dots, S$$

4. Compute sequence of IRF:

$$IRF_t \equiv A_t^{-1} H_t e_p \delta$$

$$IRF_{t+1} \equiv B_{1,t+1} IRF_t$$

$$IRF_{t+s} \equiv B_{1,t+s} IRF_{t+s-1} + B_{2,t+s} IRF_{t+s-2}, \quad s = 2, \dots, S$$

The smoothed version of the IRF is identical, except that steps 2 and 3 are replaced by an extraction of $B_{1,t+s}$ and $B_{2,t+s}$ from the smoothed posterior, provided that $t+s \leq T$. If $t+s > T$, a forecast must be produced. So steps 2 and 3 are replaced by:

2. If $t+S > T$, draw from $\eta_{T+1}^{t+S} \equiv [\eta_{T+1} \dots \eta_{t+S}]$.
3. Retrieve B_{t+s} from the posterior draws until $t+s=T$. Then compute sequence of parameters for the remaining $t+S-T$ periods by drawing from:

$$B_{t+s} = B_{t+s-1} + \eta_{t+s}, \quad s = T+1, \dots, S$$

The traditional IRF computation skips steps 2 and 3 and uses $B_{1,t}$ and $B_{2,t}$ (i.e., the value of the parameters in period t) for all subsequent periods.

As said, all three methods were used (with the caveat that the filtered IRF used the smoothed estimates of A_t, B_t and H_t), even though the smoothed estimates were the preferred ones. Results nevertheless were very similar across methods, due to the relatively more important volatility of the impact matrix (i.e., $A_t^{-1} H_t$) than the dynamic VAR parameters.²¹

²¹ This by itself could be construed as evidence that shock volatility was more important than parameter volatility. In fact, structural change will very likely result in slow parameter time variation,

4. Variance decomposition

Variance decomposition has been done using traditional formulas. This implies that only contemporaneous parameters were used in the calculation of the variance decomposition in each period, i.e. decomposition in period t only used parameters dated at t . It is obviously possible to provide smoothed and filtered estimates, but this was not done for comparability with previous studies (which, to the best of our knowledge, have not used the alternative methods) and for simplicity.

Note that variance decomposition is reported in the main text through its expectations, as only those ensure an exact decomposition of the reported expected reduced-form variance. Reporting any other quantile of the distribution would have resulted in components that would not have added up exactly to the corresponding quantile of the reduced-form variance.

5. Shock and Historical decomposition

Historical decomposition was performed using the standard formulas.

The shown historical decomposition was done using the smoothed approach, i.e. the VAR parameters for each period (B_t , A_t , H_t) were drawn from the smoothed posterior.

6. Orthogonalisation

Reduced-form shocks are orthogonalised using the technique developed by Rubio, Waggoner and Zha (2010). The procedure entails computing a random matrix whose elements have independent unit-variance normal distributions and calculating its QR decomposition and conveniently normalising. The orthonormal Q matrix so obtained follows the uniform distribution and can be used to decompose the estimated $A_t^{-1} H_t$ matrix, which already corresponds to a Cholesky decomposition of the reduced-form variance-covariance matrix. The number of orthogonal residuals retained depends on the number of structural shocks under analysis. If this number is 1 (i.e., just one structural error is sought after), the method is equivalent to Uhlig (2005).

important changes in the transmission mechanism may coexist with little *period-to-period* changes in the parameters.

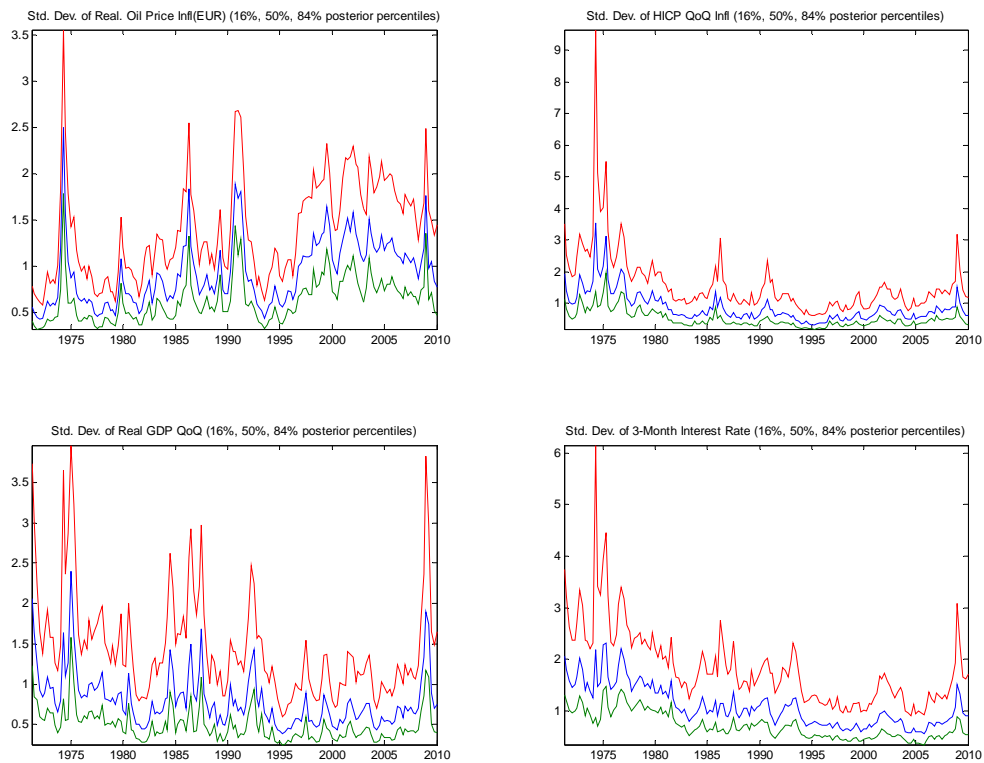
References

- Carter, C. K. and R. Kohn (1994), "On Gibbs Sampling for State Space Models," *Biometrika*, 81, 541-553.
- Gourieroux, C. and J. Jasiak (2005), "Nonlinear Innovations and Impulse Responses with Application to VaR Sensitivity," *Annales d'Economie et de Statistique*, ADRES, 78, 1-31.
- Kim, S., N. Shepard and S. Chib (1998), "Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models," *Review of Economic Studies*, 65, 361-393.
- Koop, G., M. H. Pesaran and S. M. Potter (1996). "Impulse response analysis in nonlinear multivariate models," *Journal of Econometrics*, Elsevier, 74(1), 119-147.
- Primiceri, G. (2005), "Time Varying Structural Vector Autoregressions and Monetary Policy," *Time Varying Structural Vector Autoregressions and Monetary Policy*, *Review of Economic Studies*, 72, 821-852
- Rubio-Ramírez, J. F., D. F. Waggoner and T. Zha (2010), "Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference," *Review of Economic Studies*, 77(2), 665-696.
- Uhlig, H. (2005), "What are the effects of monetary policy on output? Results from an agnostic identification procedure," *Journal of Monetary Economics*, vol. 52(2), 381-419.

Appendix C: Figures

1. DOMESIC MODEL

Figure 1: Time varying standard deviations of variables



Note: Posterior median, 16th and 84th percentiles of the time varying standard deviations of oil prices, the HICP, GDP and the short-term interest rate.

Figure 2: Contributions of oil supply shocks to the reduced form shocks of the oil price equation

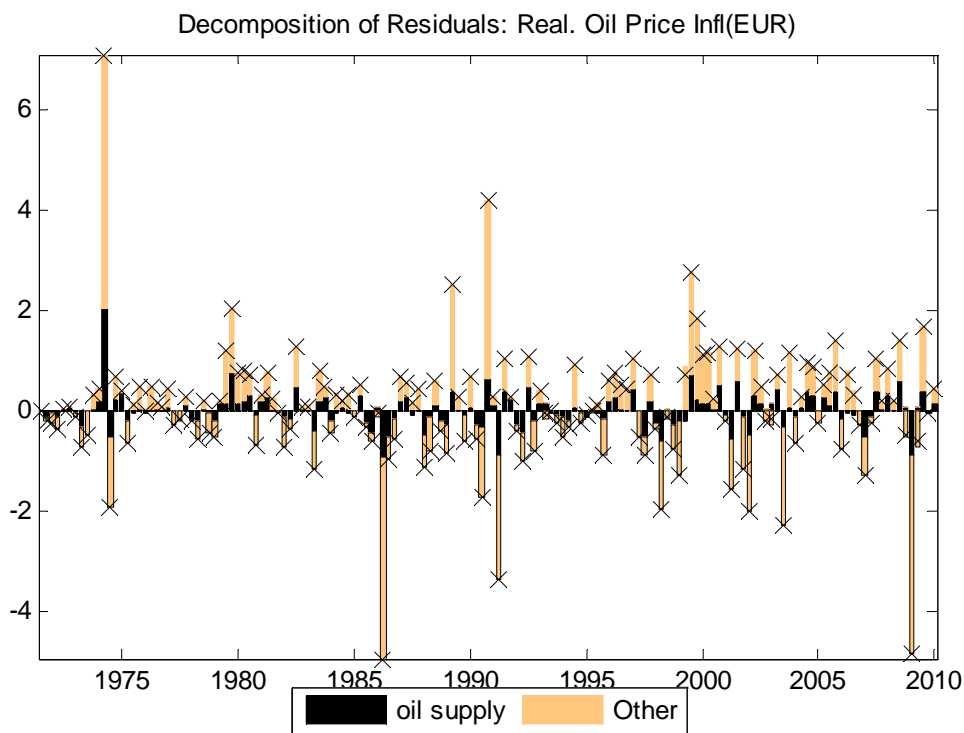


Figure 3: Contributions of oil supply shocks to the reduced form shocks of the HICP equation




Figure 4: Contributions of oil supply shocks to the reduced form shocks of the GDP equation

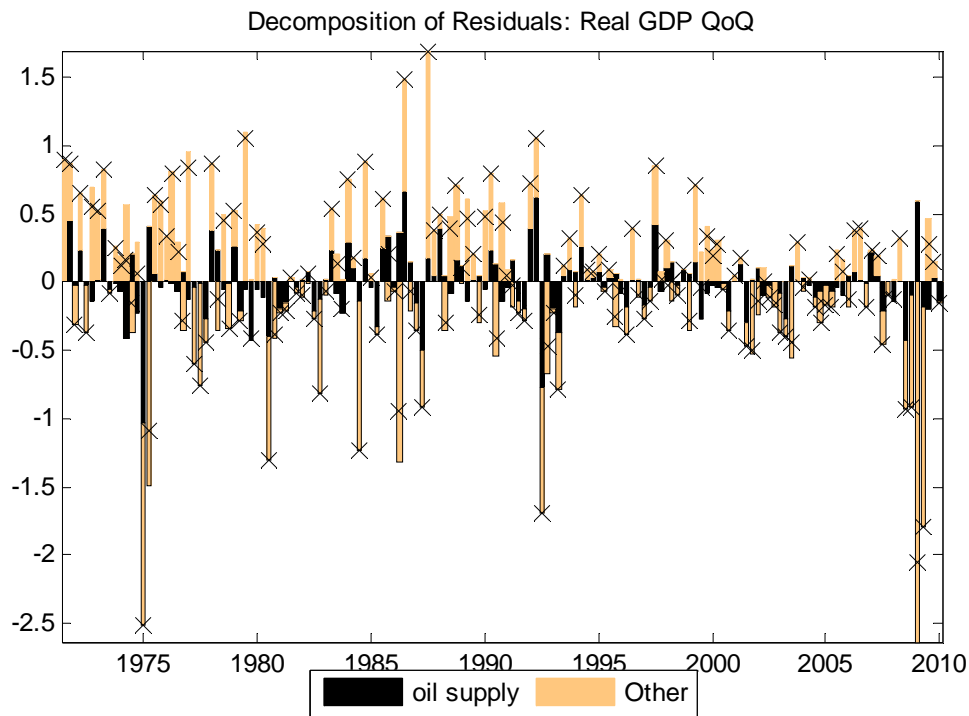


Figure 5: Contributions of oil supply shocks to the reduced form shocks of the interest rate equation

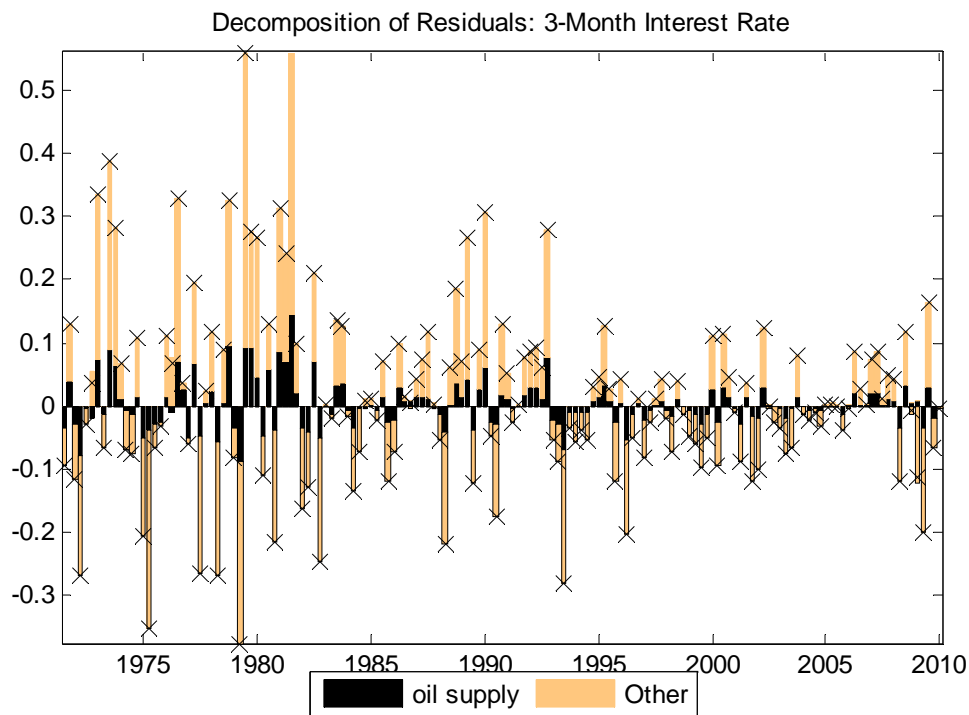
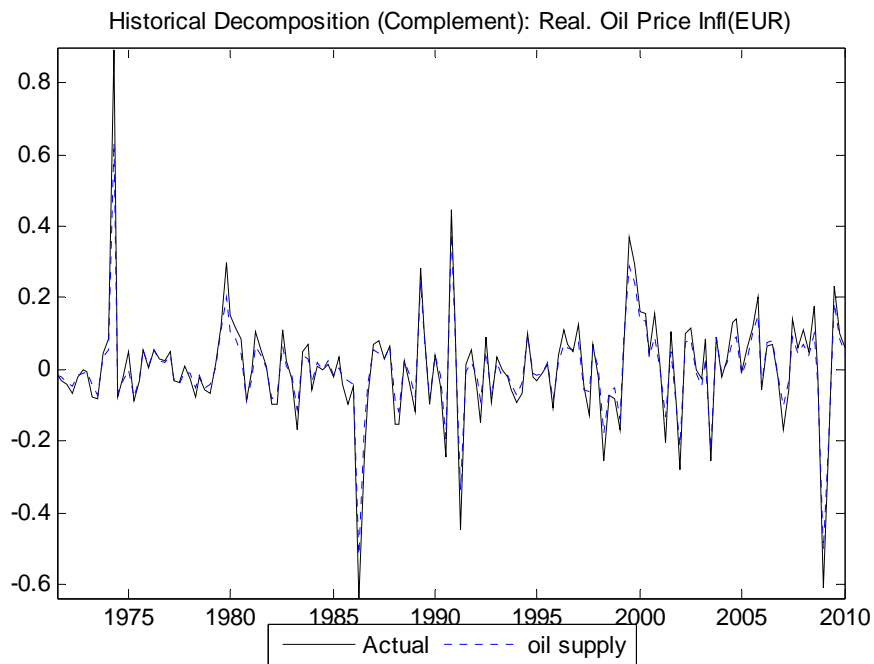
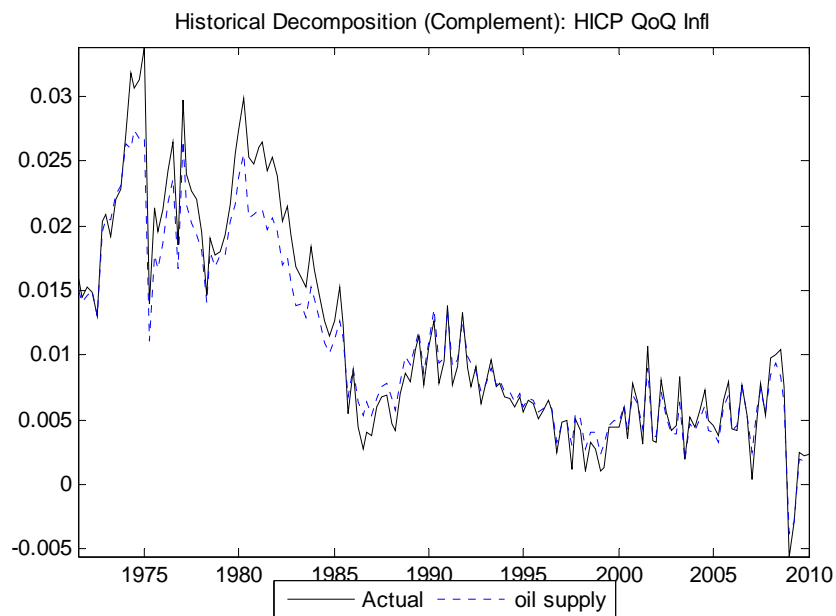


Figure 6: Quarterly change in oil prices and counterfactual development of the series without oil supply shocks



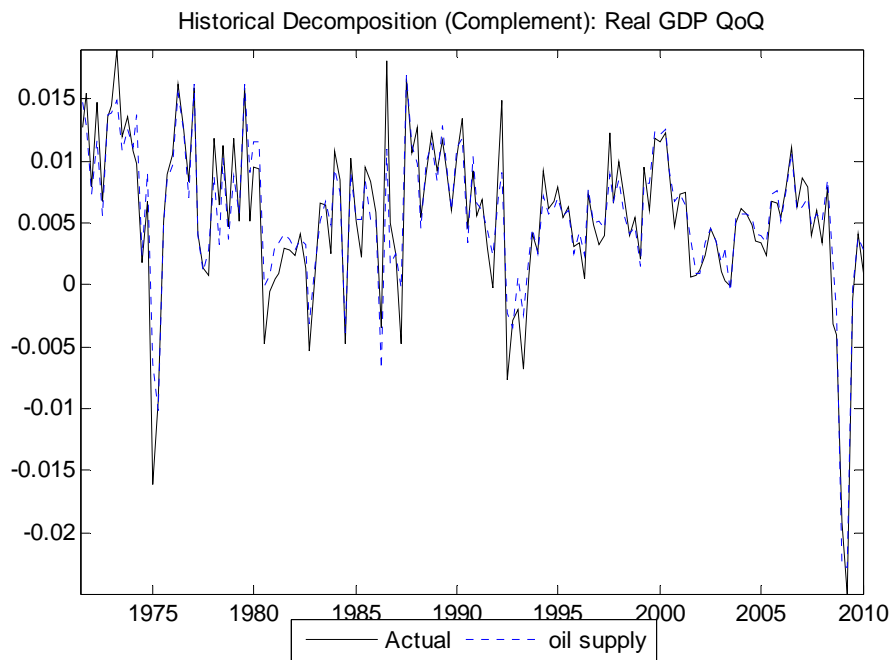
Note: Time series of the data and development of the series excluding the contribution of oil supply shocks.

Figure 7: Quarterly HICP inflation and counterfactual development of the series without oil supply shocks



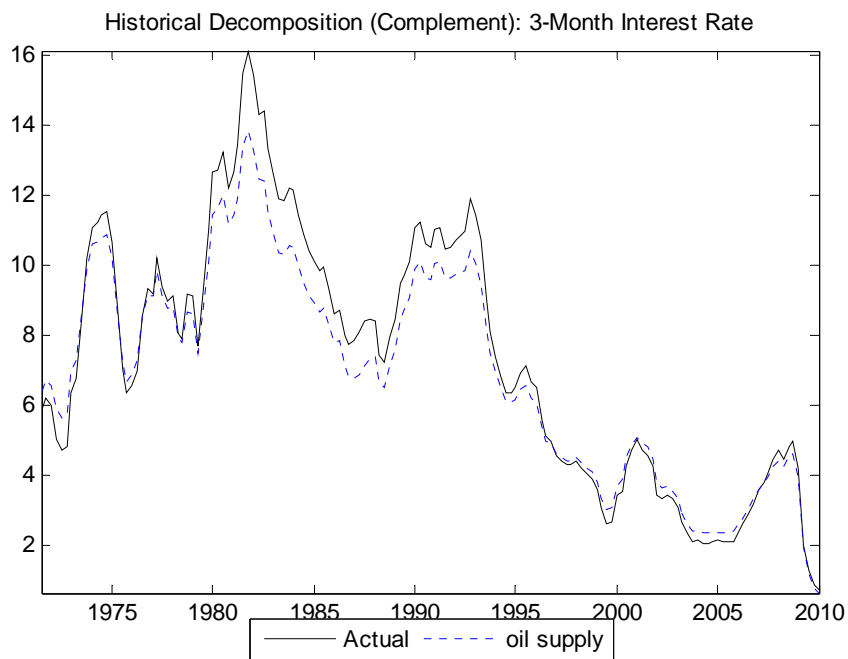
Note: Time series of the data and development of the series excluding the contribution of oil supply shocks.

Figure 8: Quarterly GDP growth and counterfactual development of the series without oil supply shocks



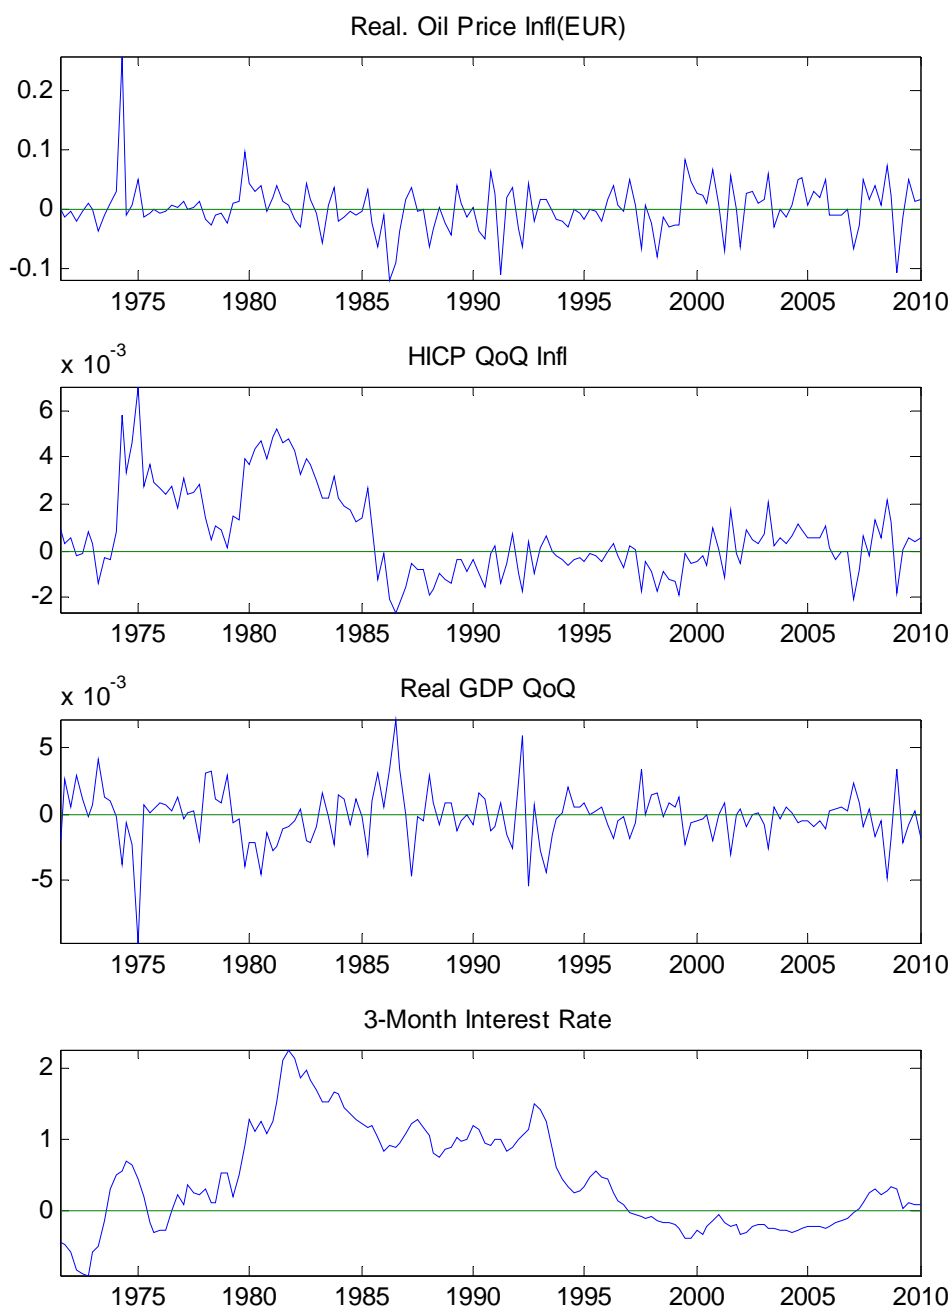
Note: Time series of the data and development of the series excluding the contribution of oil supply shocks.

Figure 9: Interest rate and counterfactual development of the series without oil supply shocks



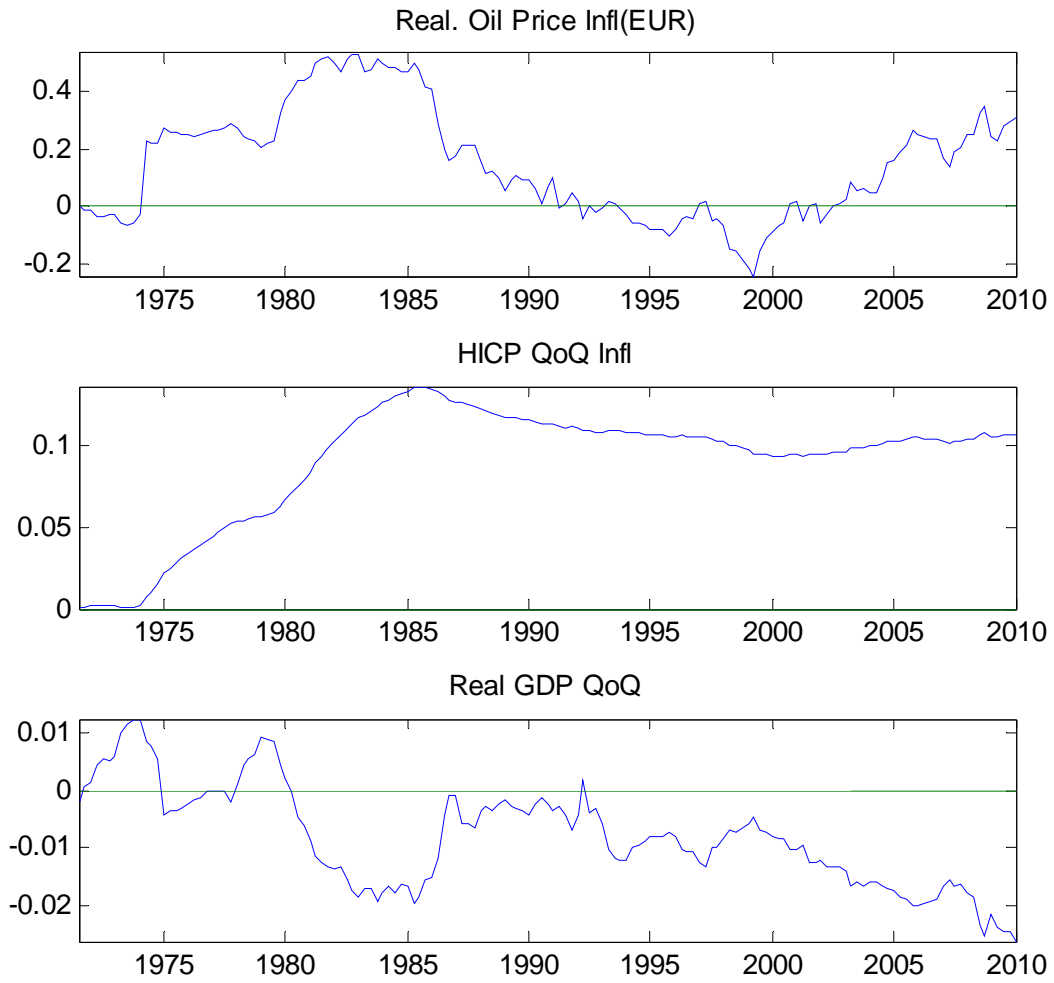
Note: Time series of the data and development of the series excluding the contribution of oil supply shocks.

Figure 10: Historical contribution of oil supply shocks



Note: Difference between actual time series and counterfactual history of the series by setting to zero the oil supply shocks, i.e. a positive figure indicates an upward effect of the shock on the series and vice versa.

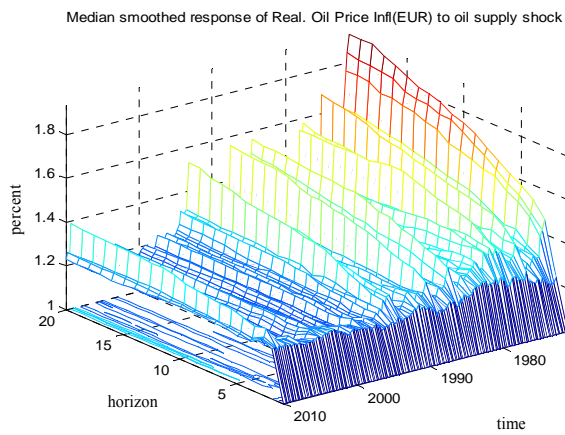
Figure 11: Historical cumulated contribution of oil supply shocks



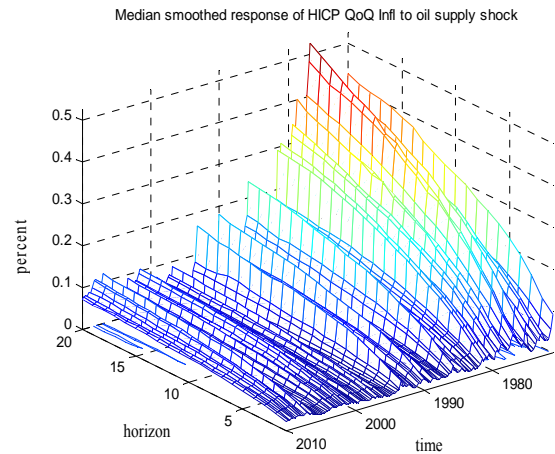
Note: Cumulated difference between actual time series and counterfactual history of the series by setting to zero the oil supply shocks.

Figure 12: Time varying impulse responses to a one percent oil supply shock

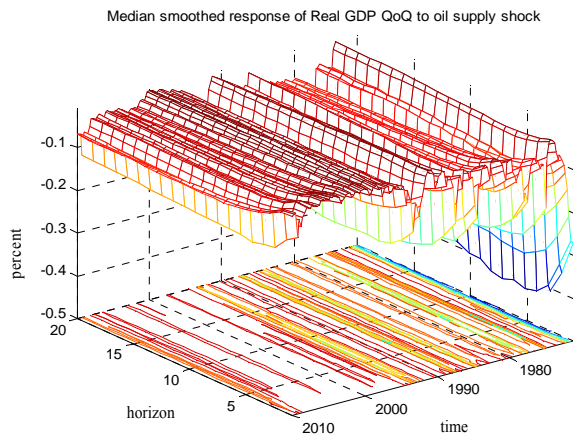
(a) Response of oil prices



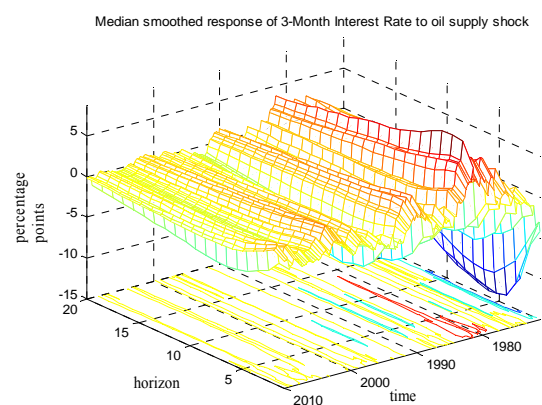
(b) Response of the HICP



(c) Response of GDP

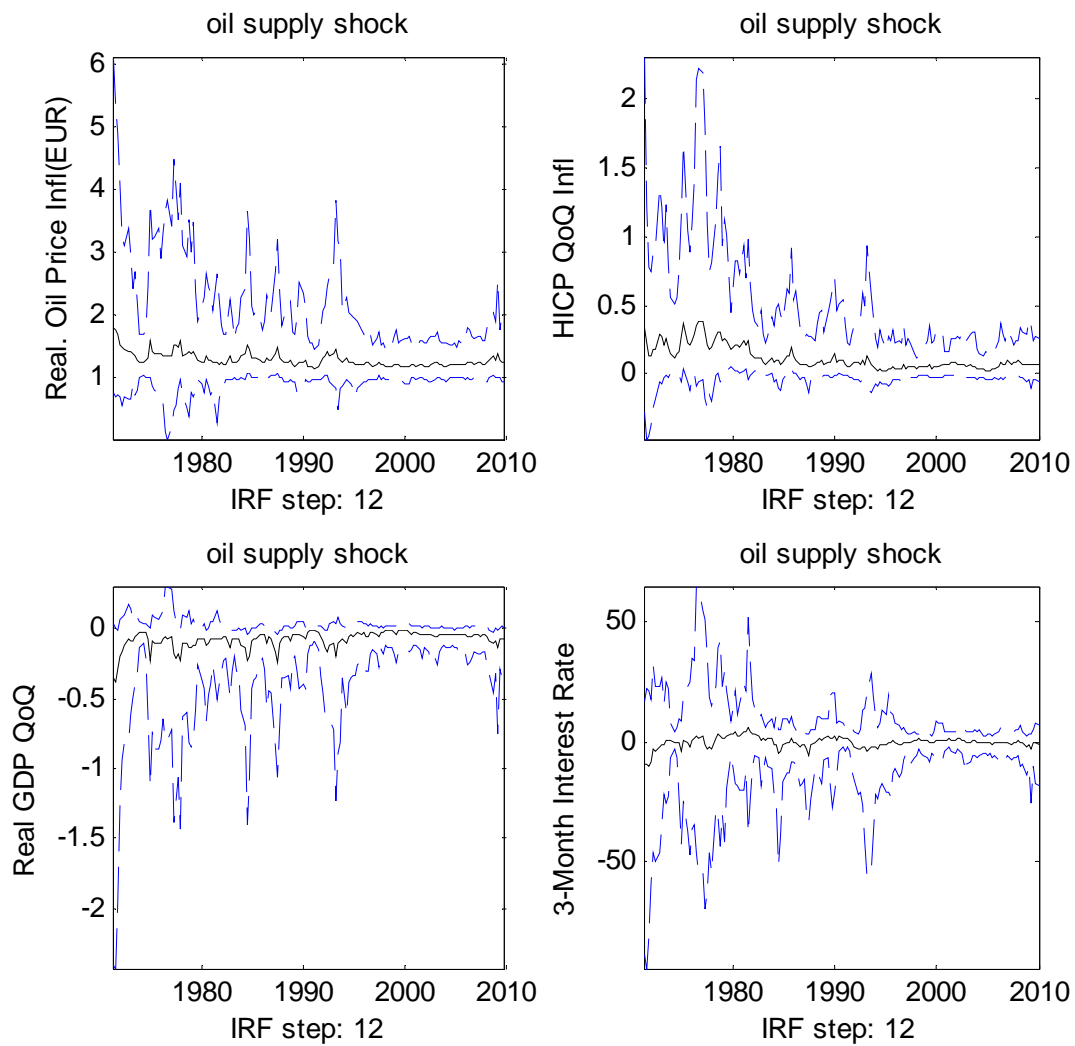


(d) Response of interest rates



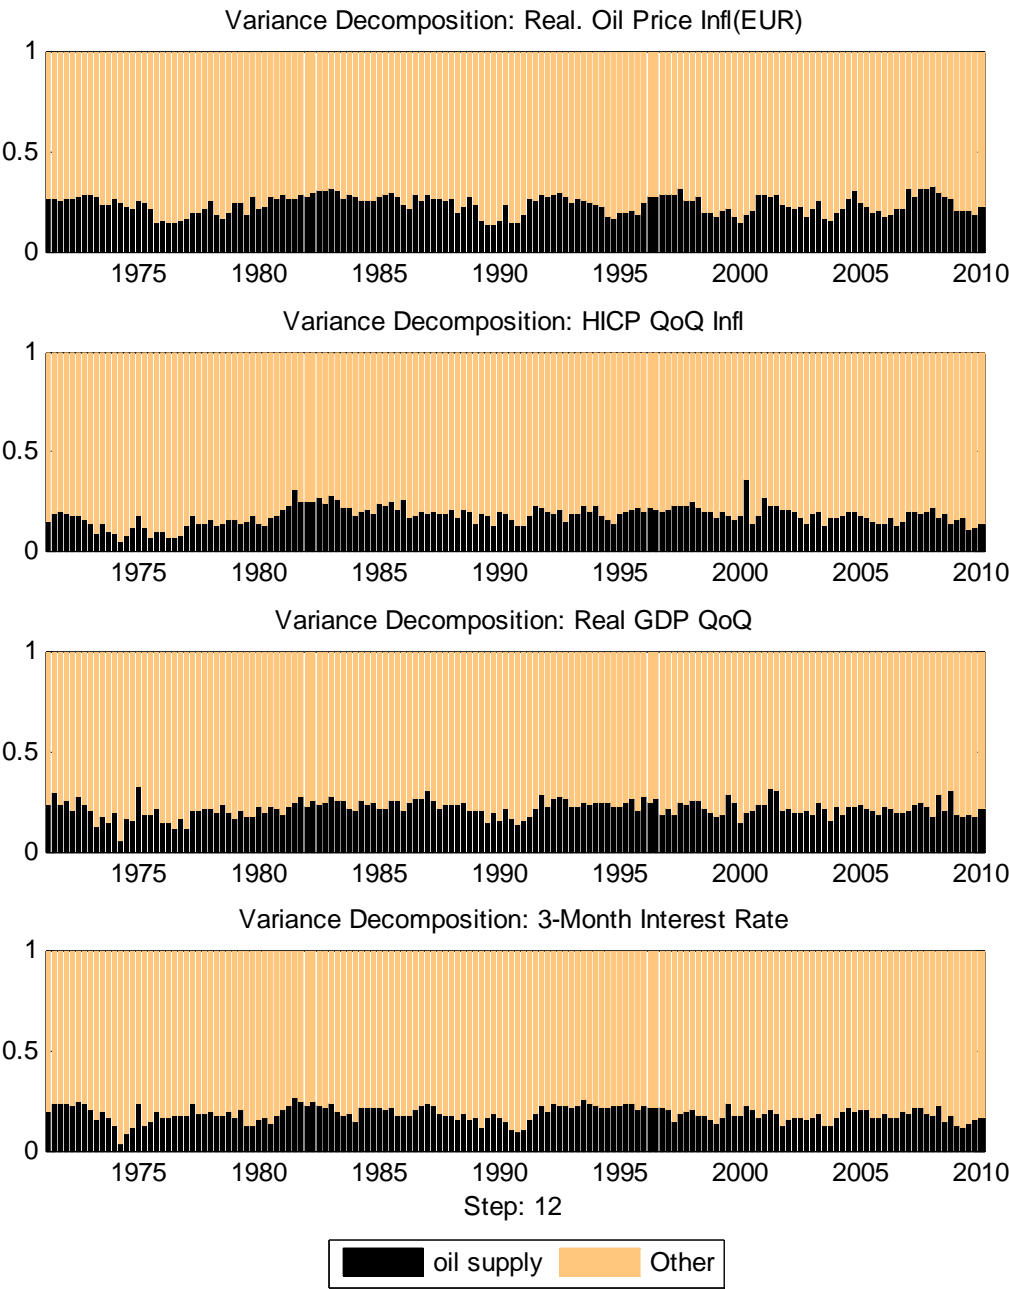
Note: Posterior median of the accumulated impulse response functions for oil prices, the HICP and GDP and the non-accumulated impulse responses for the interest rate.

Figure 13: Time varying impulse responses 12 quarters after a one percent oil supply shock



Note: Posterior median, 16th and 84th percentiles of the accumulated impulse responses for oil prices, the HICP and GDP and of the non-accumulated impulse responses for the interest rate.

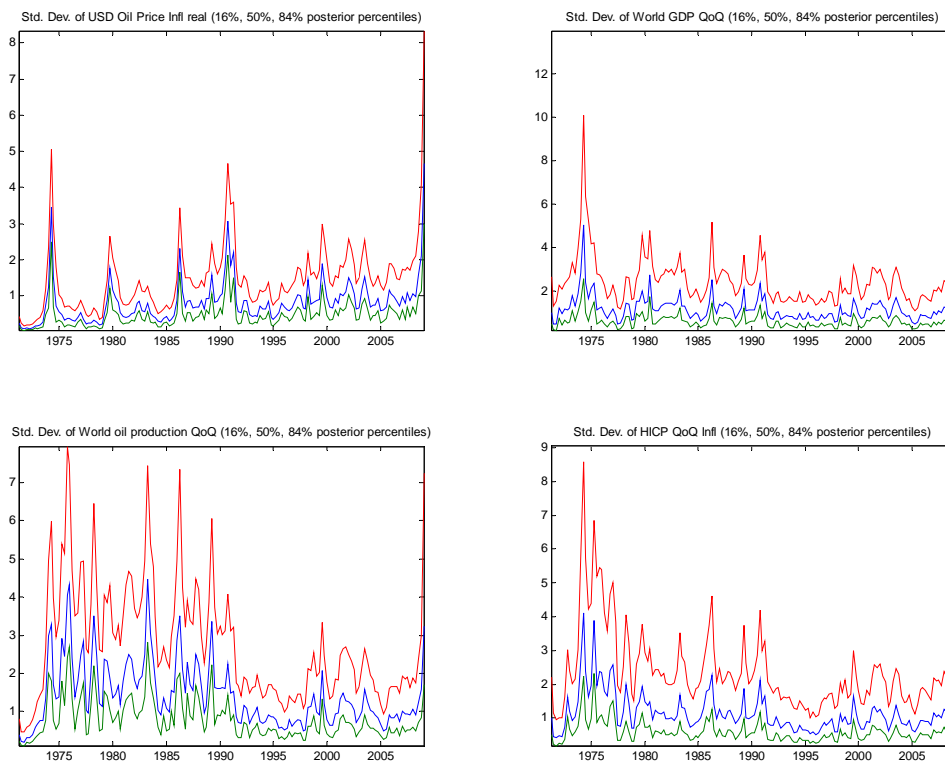
Figure 14: Time varying variance decompositions after 12 quarters



Note: Posterior mean of the variance decomposition.

2. Global model

Figure 15: Time varying standard deviations of variables



Note: Posterior median, 16th and 84th percentiles of the time varying standard deviations of oil prices, the HICP, GDP and the short-term interest rate.

Figure 16: Contributions of oil supply and demand shocks to the reduced form shocks of the oil price equation

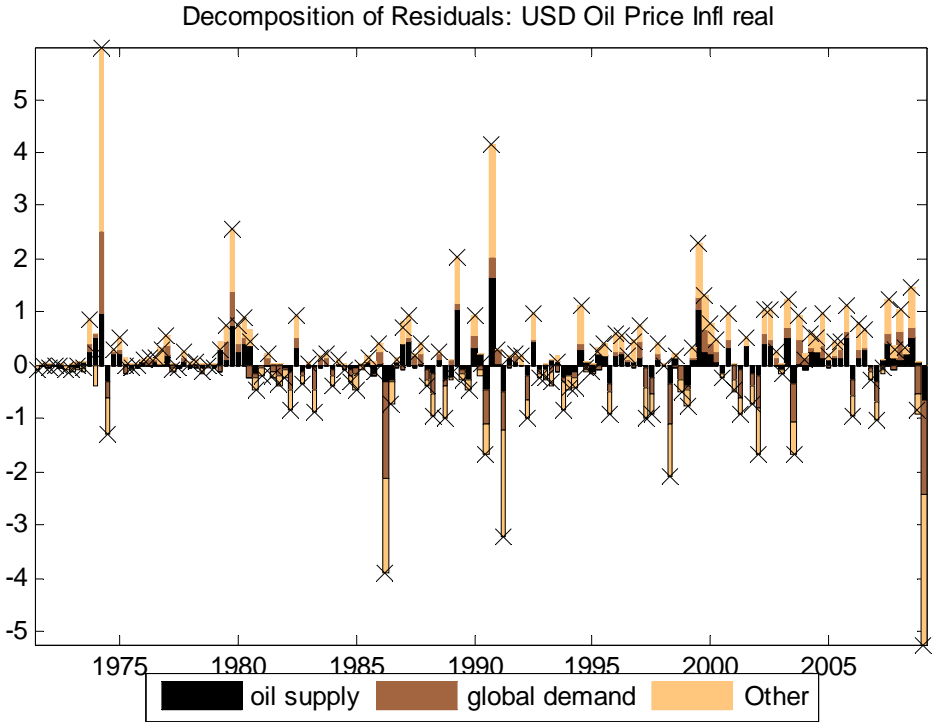


Figure 17: Contributions of oil supply and demand shocks to the reduced form shocks of the world GDP equation

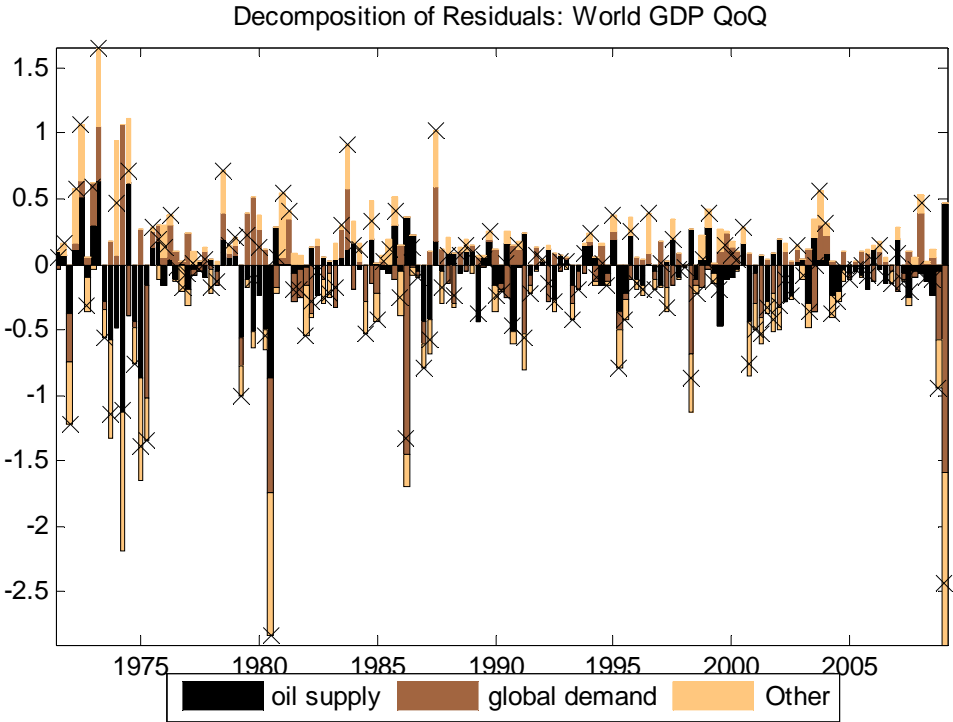


Figure 18: Contributions of oil supply and demand shocks the reduced form shocks of the world oil production equation

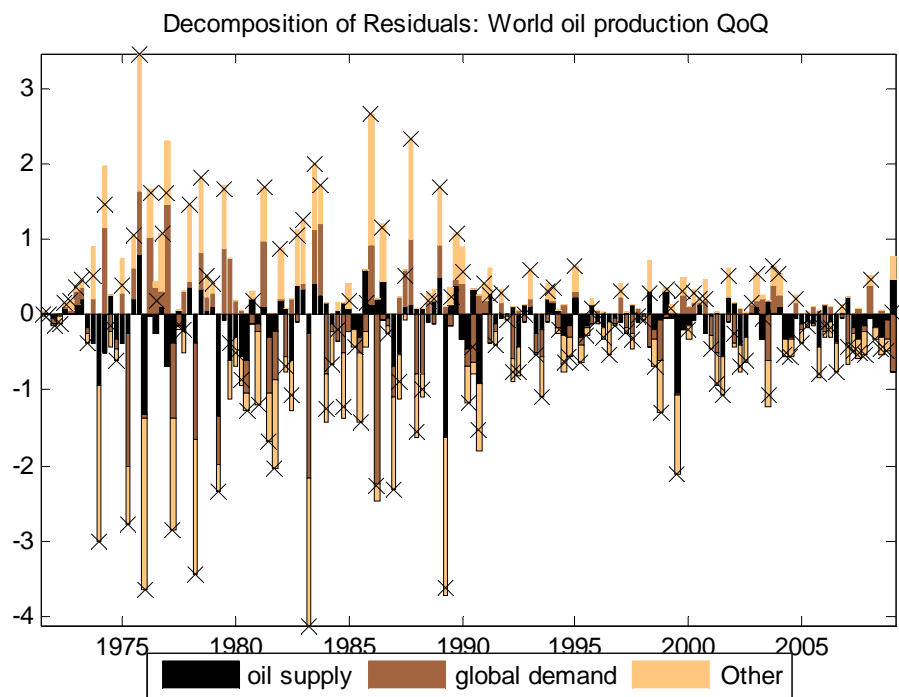


Figure 19: Contributions of oil supply and demand shocks to the reduced form shocks of the HICP equation

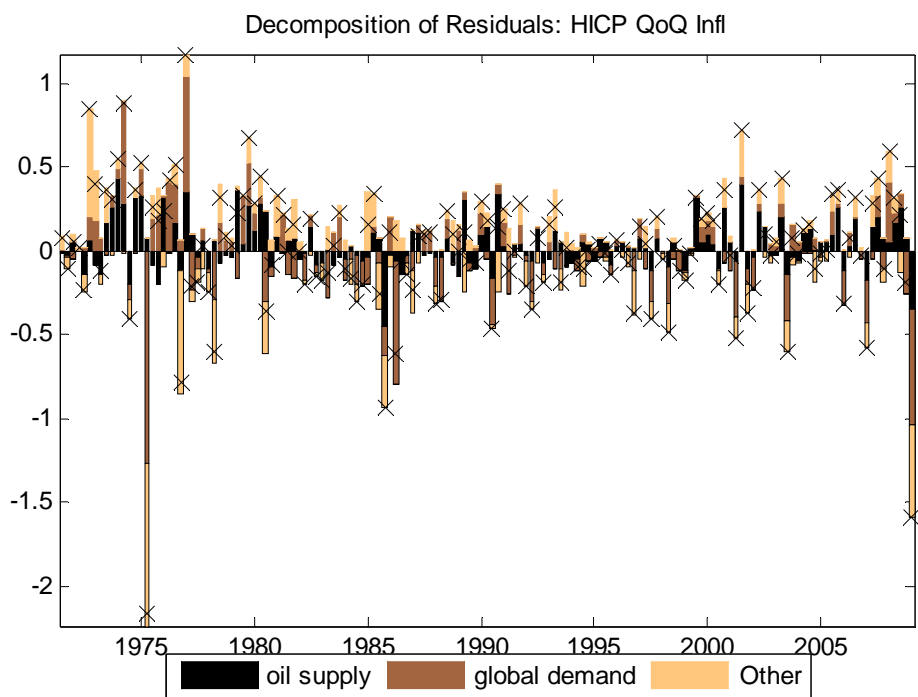
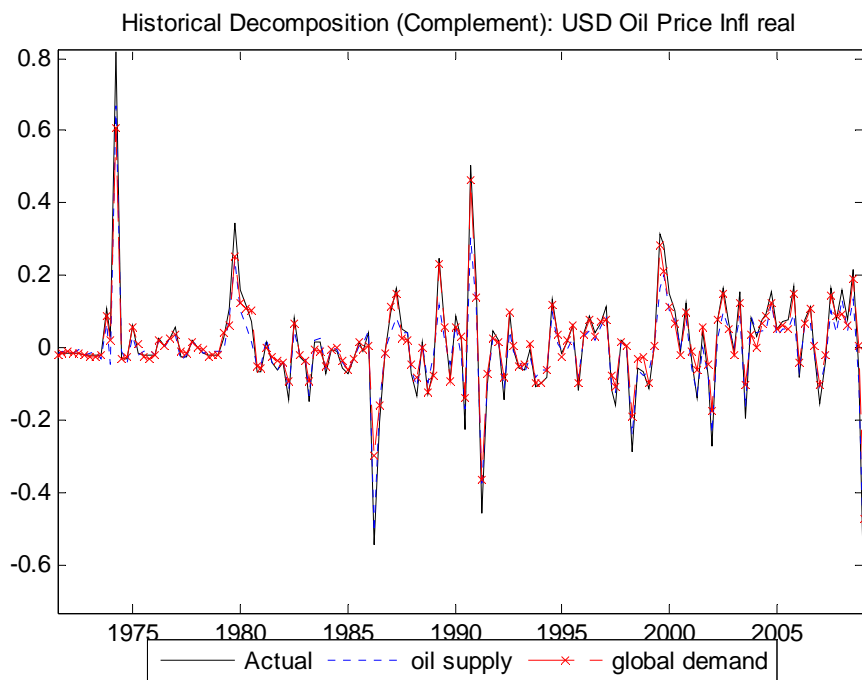
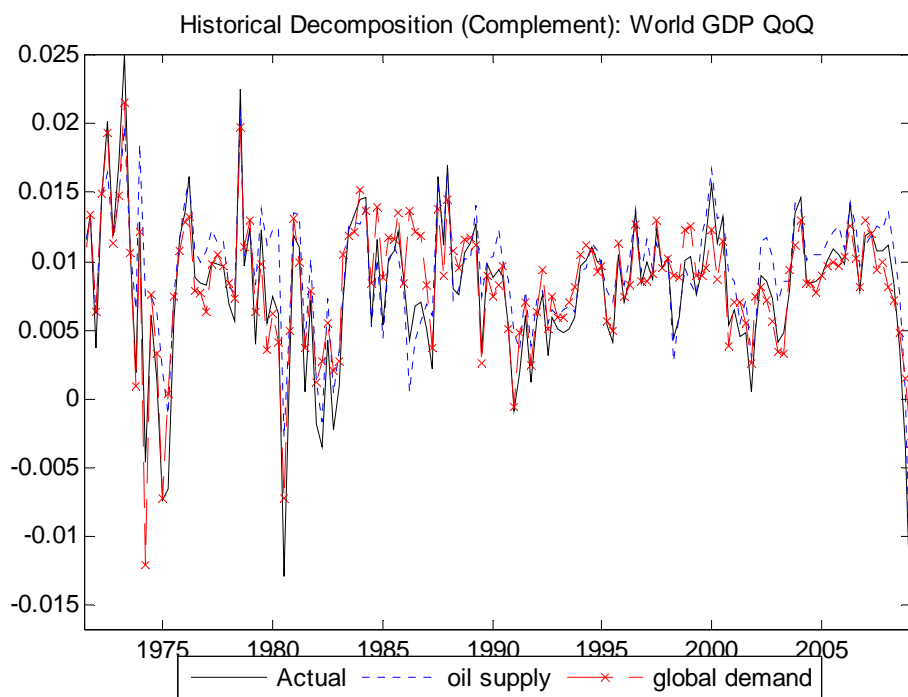


Figure 20: Quarterly changes in oil prices and counterfactual development of the series without oil supply and demand shocks respectively



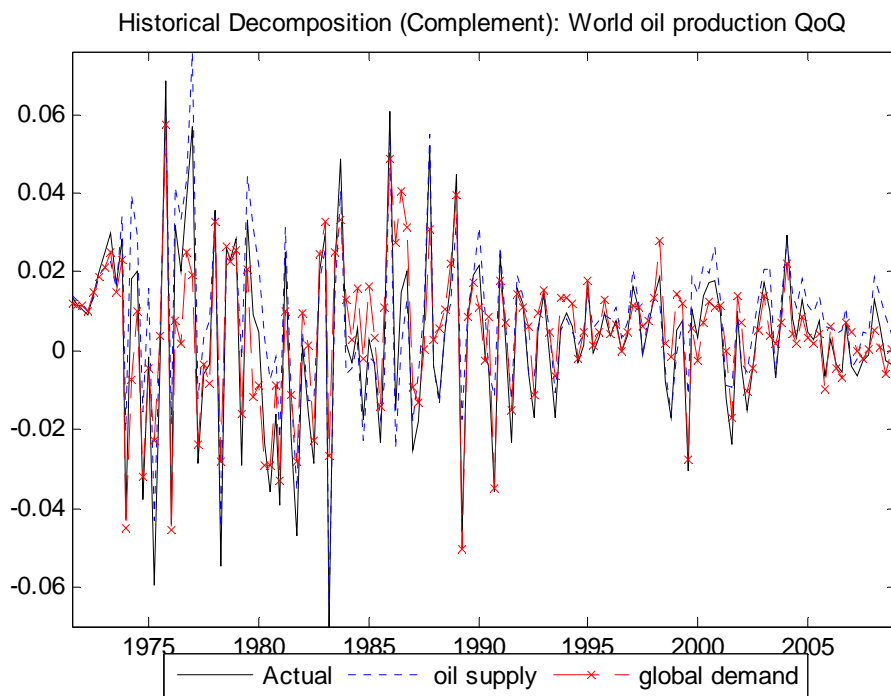
Note: Time series of the data and development of the series excluding the contribution of oil supply shocks and oil demand shocks respectively.

Figure 21: Quarterly world GDP growth and counterfactual development of the series without oil supply and demand shocks respectively



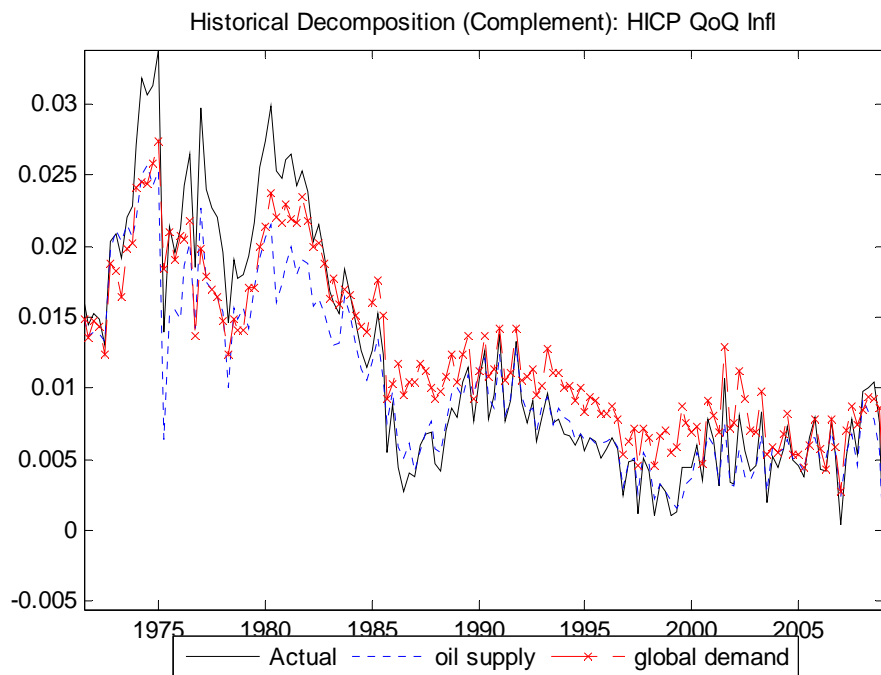
Note: Time series of the data and development of the series excluding the contribution of oil supply shocks and oil demand shocks respectively.

Figure 22: Quarterly world oil production growth and counterfactual development of the series without oil supply and demand shocks respectively




Note: Time series of the data and development of the series excluding the contribution of oil supply shocks and oil demand shocks respectively.

Figure 23: Quarterly HICP inflation and counterfactual development of the series without oil supply and demand shocks respectively



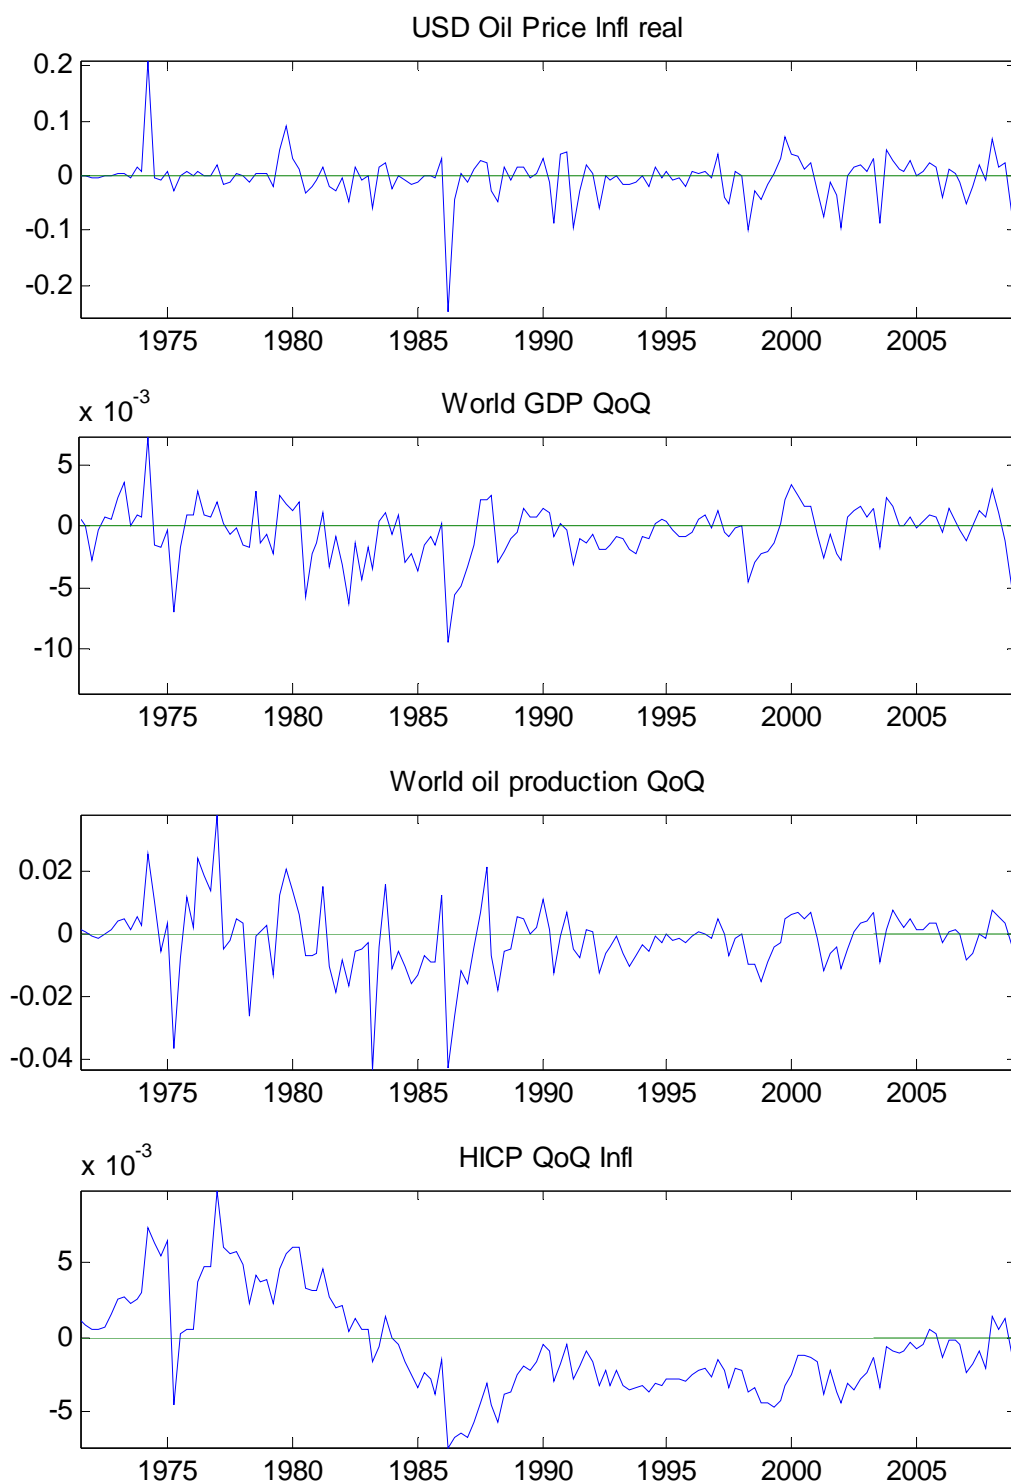
Note: Time series of the data and development of the series excluding the contribution of oil supply shocks and oil demand shocks respectively.

Figure 24: Historical contribution of oil supply shocks



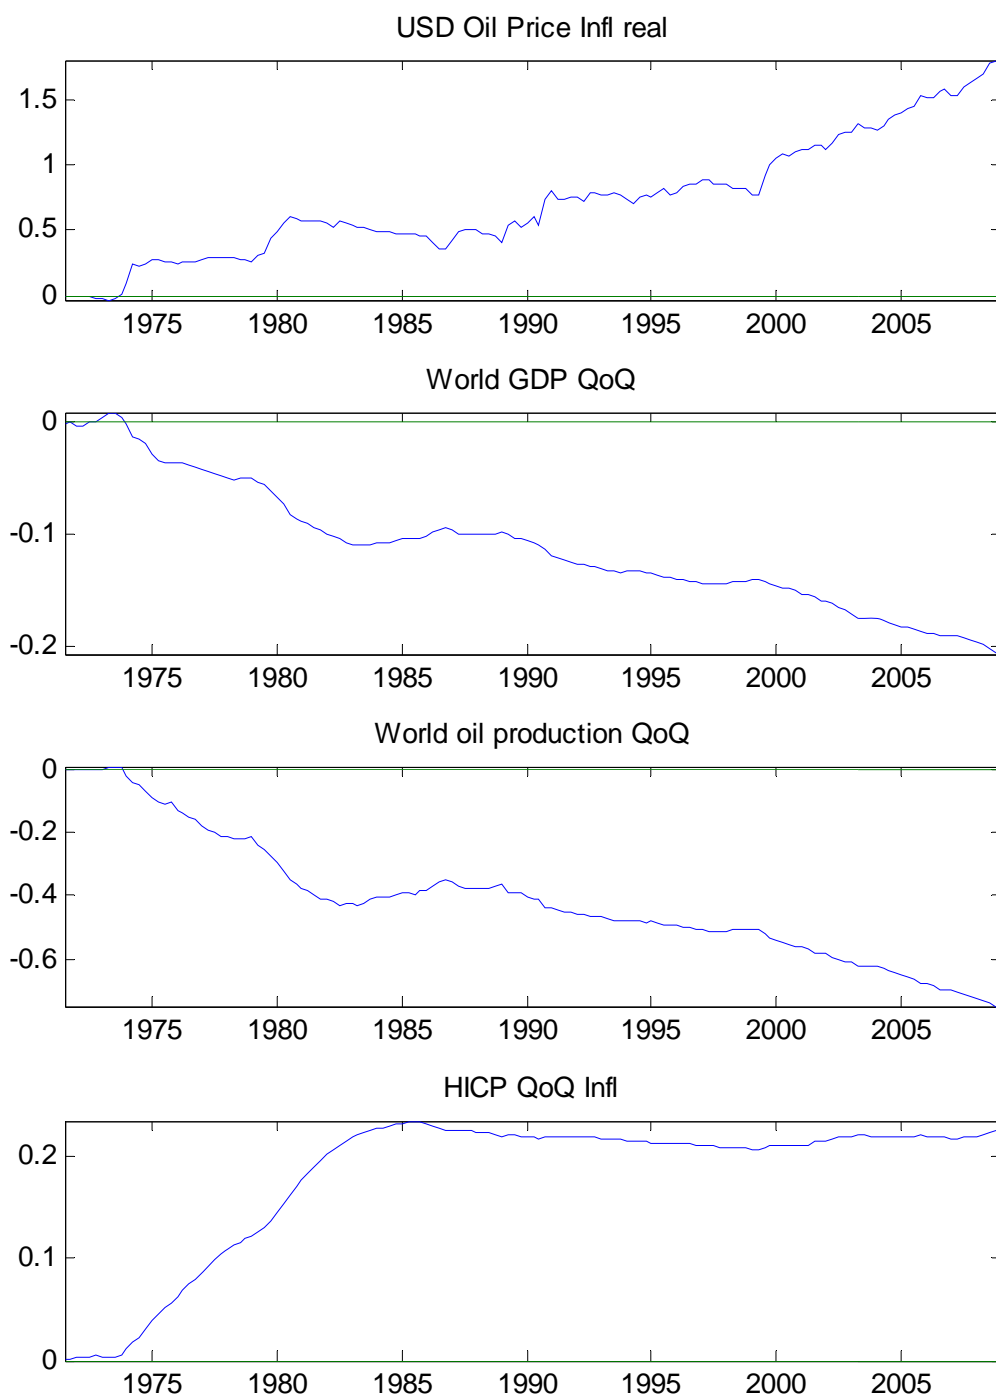
Note: Difference between actual time series and counterfactual history of the series by setting to zero the oil supply shocks, i.e. a positive figure indicates an upward effect of the shock on the series and vice versa.

Figure 25: Historical contribution of demand shocks



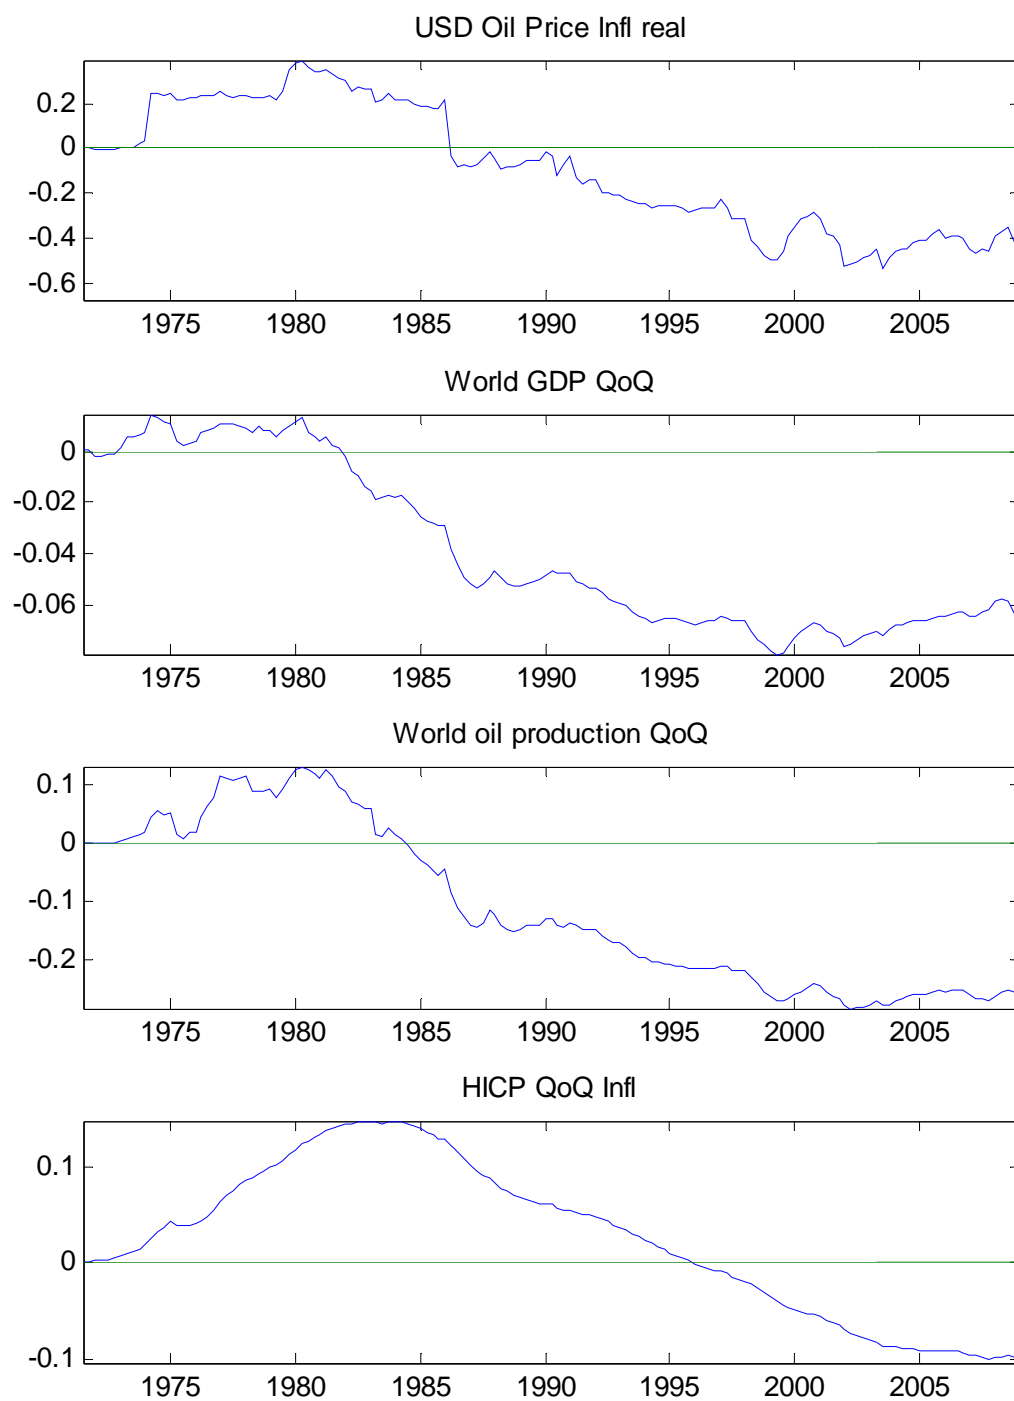
Note: Difference between actual time series and counterfactual history of the series by setting to zero the demand shocks, i.e. a positive figure indicates an upward effect of the shock on the series and vice versa.

Figure 26: Historical cumulated contribution of oil supply shocks



Note: Cumulated difference between actual time series and counterfactual history of the series by setting to zero the oil supply shocks.

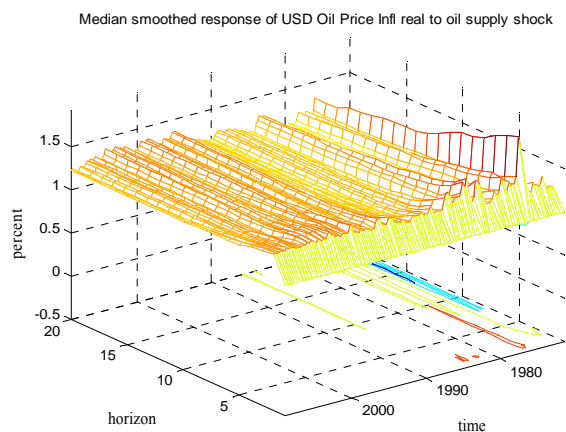
Figure 27: Historical cumulated contribution of demand shocks



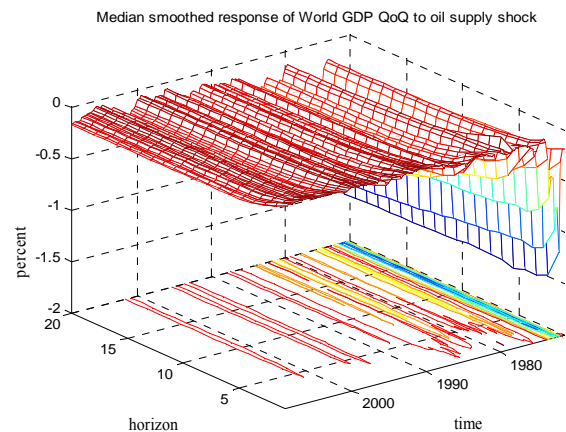
Note: Cumulated difference between actual time series and counterfactual history of the series by setting to zero the demand shocks.

Figure 28: Time varying impulse responses to a one percent oil supply shock

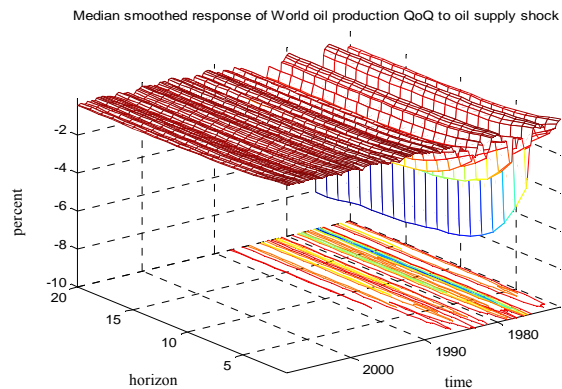
(a) Response of oil prices



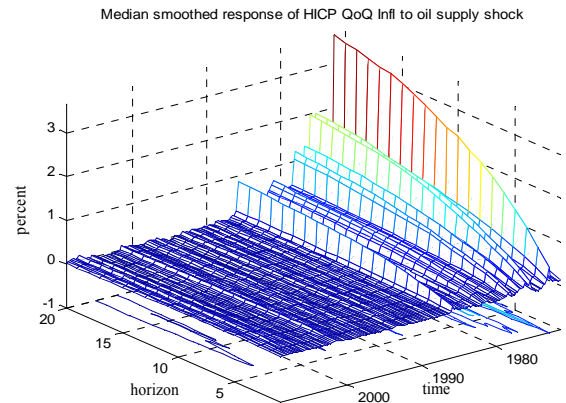
(b) Response of world GDP



(c) Response of world oil production

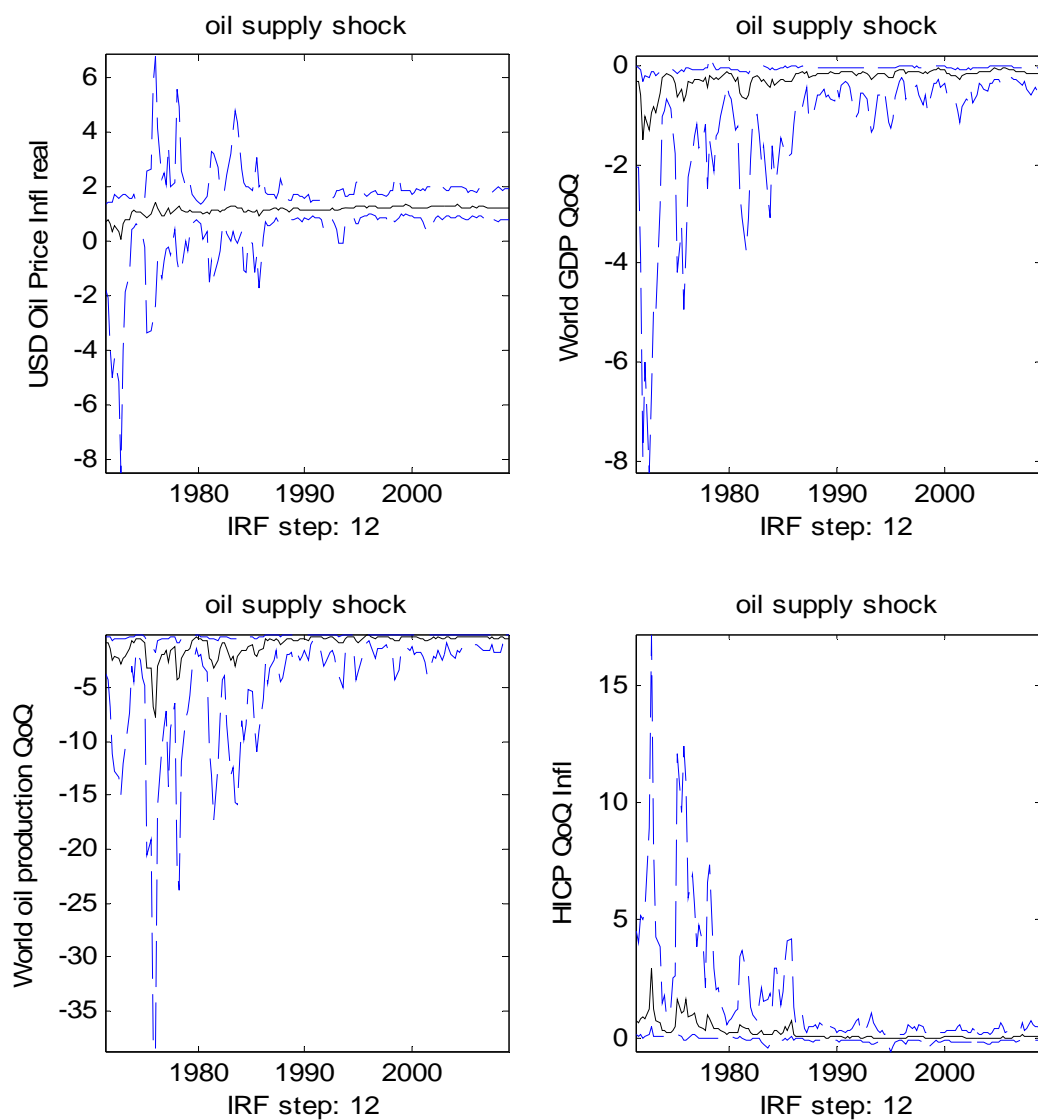


(d) Response of the HICP



Note: Posterior median of the accumulated impulse response functions for oil prices, world GDP, world oil production and the HICP.

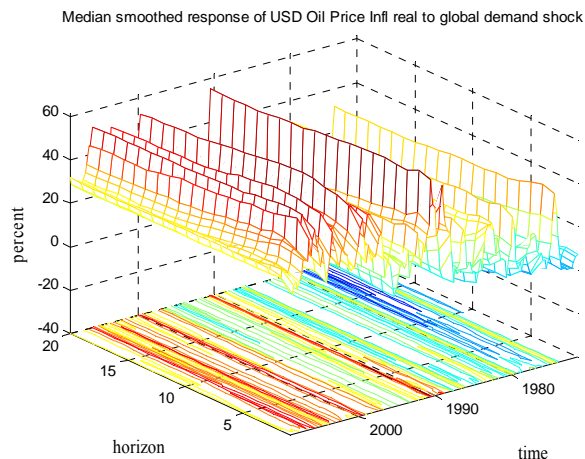
Figure 29: Time varying impulse responses 12 quarters after a one percent oil supply shock



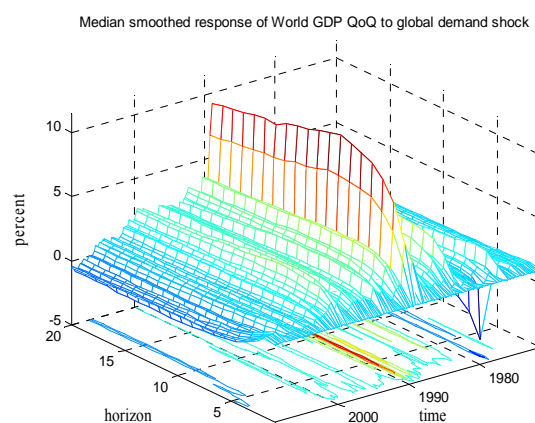
Note: Posterior median, 16th and 84th percentiles of the accumulated impulse responses for oil prices, world GDP, world oil production and the HICP.

Figure 30: Time varying impulse responses to a one percent demand shock

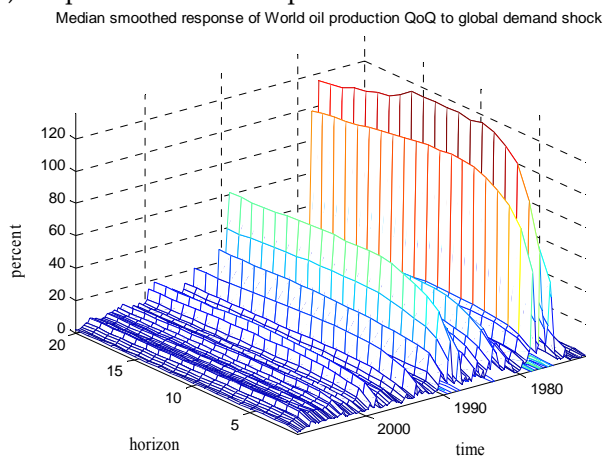
(a) Response of oil prices



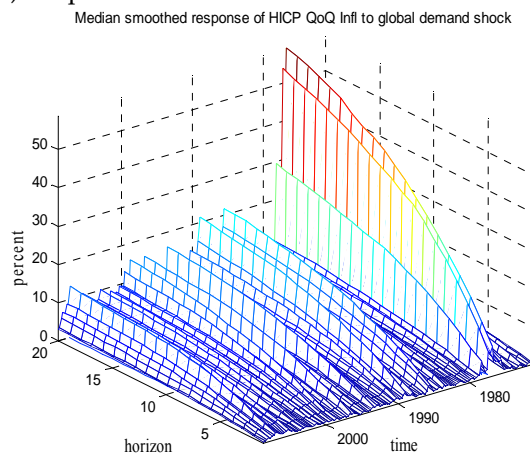
(b) Response of world GDP



(c) Response of world oil production

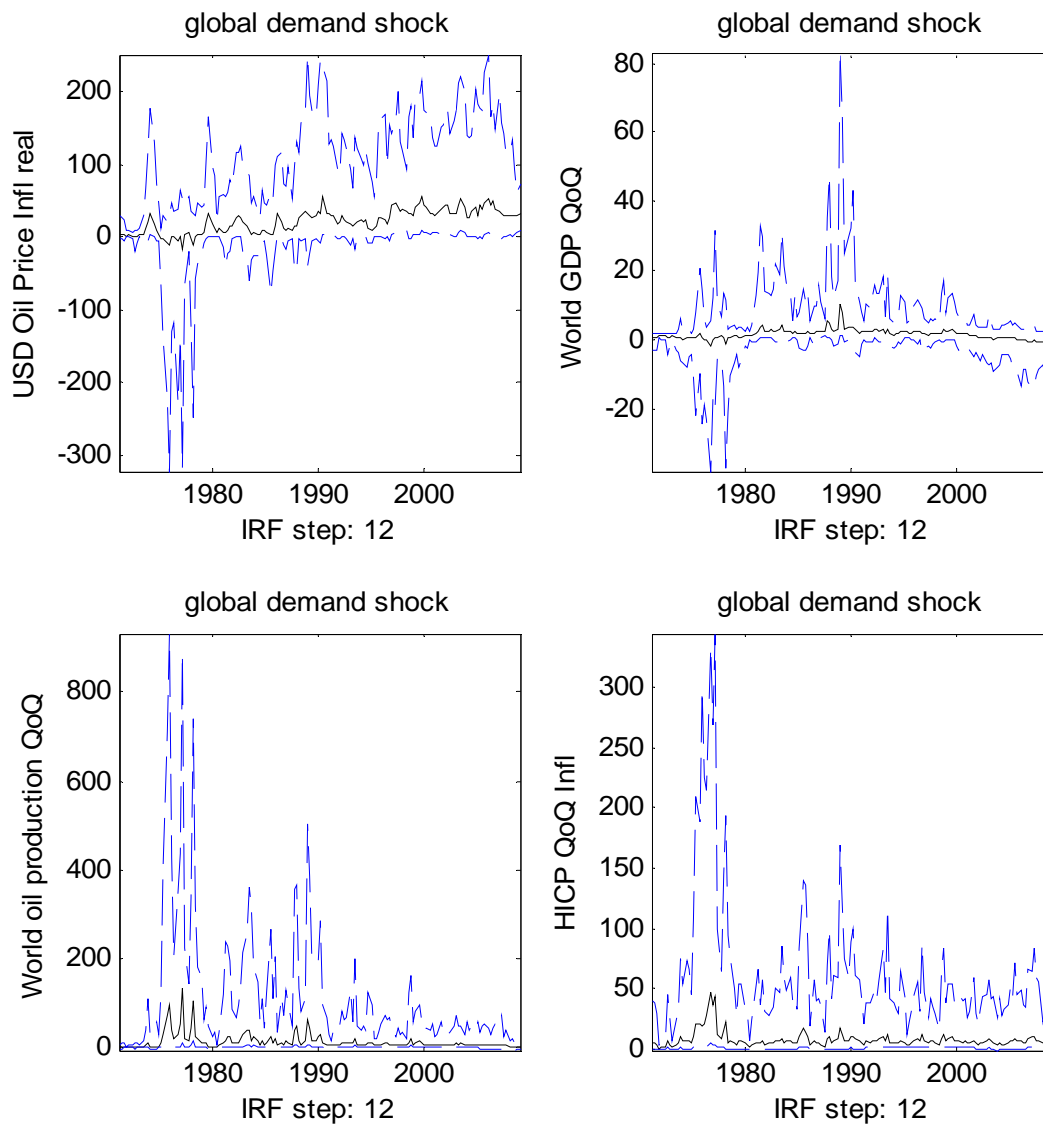


(d) Response of HICP



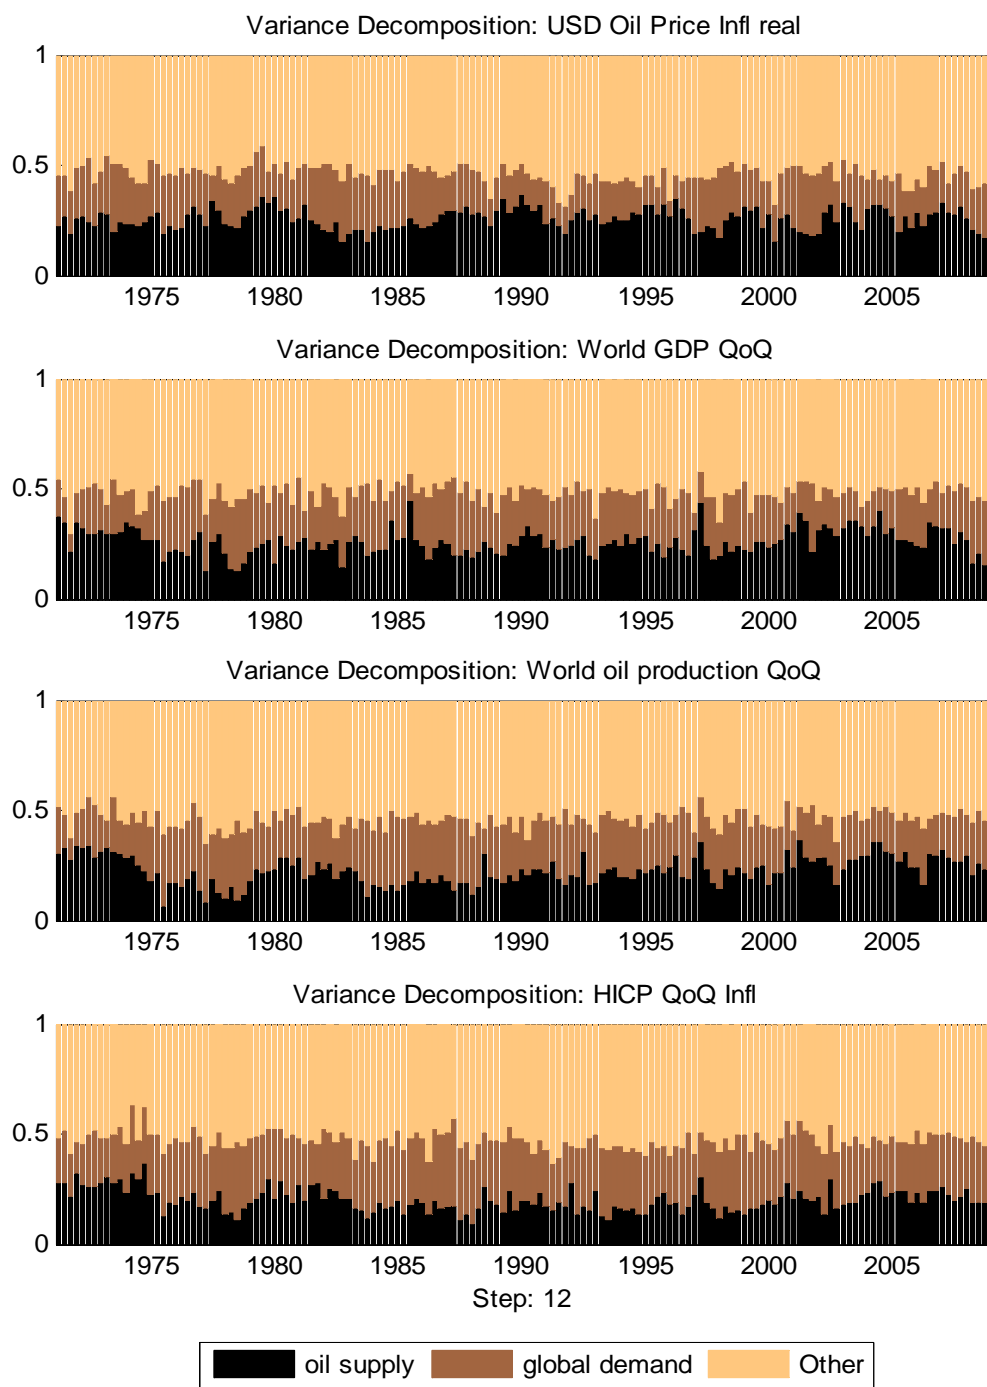
Note: Posterior median of the accumulated impulse response functions for oil prices, world GDP, world oil production and the HICP.

Figure 31: Time varying impulse responses 12 quarters after a one percent demand shock



Note: Posterior median, 16th and 84th percentiles of the accumulated impulse responses for oil prices, world GDP, world oil production and the HICP.

Figure 32: Time varying variance decompositions after 12 quarters



Note: Posterior mean of the variance decomposition.

