

Thiem, Christopher

Working Paper

Oil price uncertainty and the business cycle: Accounting for the influences of global supply and demand within a VAR GARCH-in-mean framework

Ruhr Economic Papers, No. 674

Provided in Cooperation with:

RWI – Leibniz-Institut für Wirtschaftsforschung, Essen

Suggested Citation: Thiem, Christopher (2017) : Oil price uncertainty and the business cycle: Accounting for the influences of global supply and demand within a VAR GARCH-in-mean framework, Ruhr Economic Papers, No. 674, ISBN 978-3-86788-782-3, RWI - Leibniz-Institut für Wirtschaftsforschung, Essen, <https://doi.org/10.4419/86788782>

This Version is available at:

<https://hdl.handle.net/10419/155286>

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.



RUHR

ECONOMIC PAPERS

Christopher Thiem

Oil Price Uncertainty and the Business Cycle: Accounting for the Influences of Global Supply and Demand Within a VAR GARCH-In-Mean Framework

Imprint

Ruhr Economic Papers

Published by

Ruhr-Universität Bochum (RUB), Department of Economics
Universitätsstr. 150, 44801 Bochum, Germany

Technische Universität Dortmund, Department of Economic and Social Sciences
Vogelpothsweg 87, 44227 Dortmund, Germany

Universität Duisburg-Essen, Department of Economics
Universitätsstr. 12, 45117 Essen, Germany

RWI Leibniz-Institut für Wirtschaftsforschung
Hohenzollernstr. 1-3, 45128 Essen, Germany

Editors

Prof. Dr. Thomas K. Bauer
RUB, Department of Economics, Empirical Economics
Phone: +49 (0) 234/3 22 83 41, e-mail: thomas.bauer@rub.de

Prof. Dr. Wolfgang Leininger
Technische Universität Dortmund, Department of Economic and Social Sciences
Economics – Microeconomics
Phone: +49 (0) 231/7 55-3297, e-mail: W.Leininger@tu-dortmund.de

Prof. Dr. Volker Clausen
University of Duisburg-Essen, Department of Economics
International Economics
Phone: +49 (0) 201/1 83-3655, e-mail: vclausen@vwl.uni-due.de

Prof. Dr. Roland Döhrn, Prof. Dr. Manuel Frondel, Prof. Dr. Jochen Kluve
RWI, Phone: +49 (0) 201/81 49-213, e-mail: presse@rwi-essen.de

Editorial Office

Sabine Weiler
RWI, Phone: +49 (0) 201/81 49-213, e-mail: sabine.weiler@rwi-essen.de

Ruhr Economic Papers #674

Responsible Editor: Volker Clausen

All rights reserved. Bochum, Dortmund, Duisburg, Essen, Germany, 2017

ISSN 1864-4872 (online) – ISBN 978-3-86788-782-3

The working papers published in the Series constitute work in progress circulated to stimulate discussion and critical comments. Views expressed represent exclusively the authors' own opinions and do not necessarily reflect those of the editors.

Ruhr Economic Papers #674

Christopher Thiem

**Oil Price Uncertainty and the Business
Cycle: Accounting for the Influences of
Global Supply and Demand Within a
VAR GARCH-In-Mean Framework**

Bibliografische Informationen der Deutschen Nationalbibliothek

Die Deutsche Bibliothek verzeichnet diese Publikation in der deutschen Nationalbibliografie; detaillierte bibliografische Daten sind im Internet über:
<http://dnb.d-nb.de> abrufbar.

Das RWI wird vom Bund und vom Land Nordrhein-Westfalen gefördert.

<http://dx.doi.org/10.4419/86788782>
ISSN 1864-4872 (online)
ISBN 978-3-86788-782-3

Christopher Thiem¹

Oil Price Uncertainty and the Business Cycle: Accounting for the Influences of Global Supply and Demand Within a VAR GARCH-In-Mean Framework

Abstract

This paper reinvestigates the influence of oil price uncertainty on real economic activity in the U.S. using a four-variable VAR, GARCH-in-mean, asymmetric BEKK model. In contrast to previous studies in this area, the analysis focuses on business cycle fluctuations and we control for global supply and demand factors that might affect the real price of oil, its volatility as well as the U.S. economy. We find that – even after accounting for these factors – oil price uncertainty still has a highly significant negative influence on the U.S. business cycle. Our computations show that the effect is economically important during several periods, mostly after a significant variance shift in the mid-1980s. We simultaneously estimate the effect on the global business cycle, but find that it is comparatively weak. A battery of robustness checks confirms these results. Finally, significant spillover effects in the GARCH model suggest that oil price volatility is a gauge and channel of transmission of more general macroeconomic shocks and uncertainty. These linkages are particularly strong in case of unexpected bad news.

JEL Classification: C32, E32, Q43

Keywords: Asymmetric BEKK model; crude oil; multivariate GARCH-in-mean; oil price volatility; real options; U.S. business cycle

February 2017

¹ Christopher Thiem, UDE. – All correspondence to: Christopher Thiem, University of Duisburg- Essen, Universitaetsstr. 12, 45117 Essen, Germany, e-mail: christopher.thiem@vwl.uni-due.de

1. Introduction

More than thirty years ago, Hamilton (1983) first pointed out that the majority of United States post-war recessions had been preceded by strong oil price increases. According to today's literature, there are indeed several direct and indirect channels of transmission through which oil price changes might affect economic activity.² The influence of oil price uncertainty, for example, suggests a non-linear and asymmetric relationship between the two. Measuring such effects, however, is still an important issue in modern-day macroeconometrics. This paper re-examines the importance of the uncertainty channel, while accounting for global supply and demand factors that might not only drive the real price of oil and its volatility, but also influence the U.S. business cycle.

Based on the theory of real options developed by Bernanke (1983), Pindyck (1991) and Bloom (2009), amongst others, it is expected that oil price uncertainty has a negative effect on firms' investment and production. Once a firm faces an investment decision where a significant part of a project's initial costs is irreversible, rising uncertainty about its returns makes it increasingly attractive to postpone the decision until new information arrives. Bernanke (1983) argues that when the uncertainty relates to important macroeconomic factors, such as energy prices, many firms might be subject to the uncertainty effect at the same time, leading to a reduction of aggregate demand for investment goods and a corresponding economic downturn. Once the uncertainty dissipates, however, the reactivation of unused capacities will cause a "boom" in aggregate investment. In this way, the uncertainty effect might explain business cycle fluctuations. Furthermore, this mechanism may not only play a role in investment but in all economic decisions with a certain degree of irreversibility, such as hiring or the purchase of consumer durables (see, e.g., Edelstein and Kilian 2009). Bloom (2009) shows that surges in macroeconomic uncertainty – including the ones related to oil price shocks – cause a sharp decline and rebound in output, employment and productivity within just six months after the initial shock.

The first attempts to measure the direct influence of oil price uncertainty on U.S. economic activity empirically were conducted by Ferderer (1996), and Guo and Kliesen (2005). Both studies find a significant negative effect on different macroeconomic variables. However, they both also suffer from a certain degree of inconsistency, as they implicitly treat oil prices and oil price volatility as exogenous to the U.S. economy (cf. Jo 2014, p. 1114). In an important contribution, Elder and Serletis (2010) were able to overcome this problem by using a structural VAR model with GARCH-in-mean errors. The two main advantages of this framework are i) that all parameters are estimated simultaneously and ii) that the proxy for oil price uncertainty is generated in a consistent manner using the model's internal information. In their baseline setup, the authors estimate a bivariate model with quarterly data from 1975Q2 to 2008Q1 and find that real oil price volatility has had an economically important adverse impact on U.S. GDP growth. Since then, using the VAR GARCH-in-mean framework in such analysis has become popular. Several other studies confirm that oil price uncertainty significantly depresses economic activity in the U.S. and the rest of the G7. Elder and Serletis (2011) and Bredin et al. (2011), for instance, use a four-variable model that also includes measures for inflation and interest rates. Rahman and Serletis (2011, 2012) not only control for the influence of economic growth uncertainty by relaxing questionable zero-restrictions maintained by earlier studies, but also show that there are significant asymmetries in the reactions of the uncertainty measures depending on whether the relevant mean innovation is associated with good news or bad news. Finally, Pinno and Serletis (2013) confirm the detrimental influence of oil price uncertainty at the sector level and also find that the effect was generally weaker in a 1980s sample than in the 1990s or 2000s.

² See Kilian (2014) for a well-written overview.

Another branch of the literature has focused on the idea that the time-varying effects of oil price changes can be explained by differences in the underlying shocks to the global oil market. Barsky and Kilian (2002, 2004) were the first to challenge the, up to this point, common view that oil prices are mostly driven by political events in the Middle East and, hence, can be treated as exogenous with respect to the economy. Instead, they argue, prices are determined by the interplay of global supply and demand for crude oil, which both also contain endogenous components. Estimating an empirical model of the global oil market, Kilian (2009) finds that global aggregate demand shocks are much more important than previously thought.³ He further documents that the effects of an oil price rise on the U.S. economy may vary substantially depending on whether the rise was caused by a positive demand or negative supply shock. One explanation might be that specifically global demand shocks influence the U.S. both directly via the demand for goods and services, for instance, and indirectly via their impact on the oil price. As a consequence, not accounting for global oil market forces might lead to imprecise and unstable conclusions when studying the effects of oil price changes on country-specific variables. By now, a number of follow-up studies seem to confirm this view.⁴

To our best knowledge, however, all studies using the VAR GARCH-in-mean methodology to analyse the effects of oil price uncertainty have so far treated the oil price and national economic activity as exogenous with respect to the global business cycle and the global supply of crude oil.⁵ Therefore, their results may be subject to an omitted variable bias, which leads to an overestimation of the uncertainty effect. If, for example, oil price uncertainty is on average higher during a global recession, not accounting for the influence of weakening global demand might attribute too large of a portion of the following decline in U.S. activity to oil price volatility.

Thus far, only a study by Jo (2014) sheds – at least in part – some light on this issue. The author builds on the work by Kilian (2009) and estimates a Bayesian VAR of the global oil market with quarterly data from 1958Q2 to 2008Q3. She models oil price uncertainty as a stochastic volatility process that is driven by random, exogenous volatility shocks uncorrelated to changes in the variables' means. Jo (2014) confirms a significant negative effect of oil price uncertainty on global economic growth. Substituting her global activity measure with U.S. GDP growth, though, she finds a much weaker influence on the U.S. economy compared to the previous results of Elder and Serletis (2010). However, the three-variable model does still not account for the influence of global demand on the U.S. economy, which – according to the relevant literature – is likely to be important. Furthermore, it does not control for other uncertainty influences, such as uncertainty about GDP itself, for instance, as Jo (2014) imposes zero-restrictions on all but one of the elements in the relevant coefficient matrix.

This paper builds on the existing empirical literature by utilising a VAR GARCH-in-mean model to investigate the influence of oil price uncertainty on U.S. and global economic activity. Like other studies in this area, we define the conditional standard deviation of the one-period-ahead forecast error of the change in the real oil price as our uncertainty proxy. However, our approach differs from previous work with respect to several important aspects:

³ Kilian (2009) uses a proxy of the global business cycle to capture shifts in global aggregate demand for crude oil. In this regard, our study goes in a similar direction. Hence, keep in mind that, whenever we are discussing global business cycle fluctuations in the remainder of this paper, we are also implicitly referring to global oil demand, and vice versa.

⁴ See, for example, Kilian and Park (2009), Fukunaga et al. (2011), Lippi and Nobili (2012), and Carstensen et al. (2013).

⁵ Caporale et al. (2014) estimate the influence of oil price uncertainty on sectoral stock returns in China and allow for the effect to depend on whether a certain time period is dominated by any of the three structural shocks identified by Kilian (2009). They implement this by including slope dummies in the mean equation of the stock returns.

Most notably, this study accounts for the influence of changes in global demand and supply of crude oil on the first and second moment of the real oil price and U.S. economic activity. Following Kilian (2009) and others, we estimate a four-variable model that not only includes the growth rate of the real oil price and a measure of U.S. economic activity, but also the change of global crude oil production and a proxy for the global business cycle. This allows us to re-examine the influence of oil price uncertainty on the U.S. economy with a model that is less likely to suffer from omitted variables bias. Also, conditioning on the additional information contained in the global oil market variables most likely helps us to obtain a more precise measure of oil price uncertainty – cf. Elder and Serletis (2011, p. 385).

Beyond that, we improve the model's specification in a number of other ways. Following Rahman and Serletis (2011, 2012), the conditional covariance process is governed by a BEKK model that allows for cross-influences of shocks, volatility spillovers, and asymmetric reactions to positive and negative innovations. Therefore, we utilise one of the most general multivariate GARCH models in the literature. Moreover, this gives us the opportunity to learn more about oil price uncertainty by analysing which of our model's variables drives oil price volatility and under what circumstances. Furthermore, influenced by the observations of Baumeister and Peersman (2013), the results of Pinno and Serletis (2013), and the advent of the "Great Moderation", our model accounts for a one-time structural change in the mid-1980s of the otherwise constant part of the covariance process. Likewise, we deviate from the overly restrictive assumption of a multivariate normal distribution of the error terms and allow for a fat-tailed distribution instead. Both novelties are not only appropriate on logical grounds, but also fit the data well and lead to a significant improvement of the Akaike information criterion.

Apart from the important conceptual differences regarding the measurement of oil price uncertainty, this study also differs from Jo's (2014) approach in that we control for uncertainty about economic activity, and in that it is feasible to estimate our model with monthly instead of quarterly data. The latter is especially important, given that the real options literature has mostly highlighted the short-term ramifications of uncertainty shocks and their potential role in business-cycle fluctuations. As a matter of fact, we are also able to look at these predictions more closely than other papers, because we do not measure economic activity simply as the growth rate of the relevant industrial production index. Instead, we derive more precise measures of the U.S. and global business cycles using the Hodrick-Prescott (HP) filter to calculate the percent deviation from the original series' long-run trend. The similar characteristics of the two series then allow us to compare the influence of oil price uncertainty on the U.S. directly with the influence on the global output gap.

We estimate our baseline model with new data ranging from August 1975 to March 2014. Our results show significant non-diagonalities, variance spillovers and asymmetries in the multivariate GARCH process. Oil price uncertainty is not only caused by unexpected shocks to the real oil price itself, but also by unexpected shocks to the global supply and demand of crude oil, as well as shocks to the U.S. business cycle. The linkages are stronger when the respective shocks are associated with bad news instead of good news. Our results imply that – at least partially – oil price volatility also reflects general macroeconomic uncertainty.

The main finding of this study, however, is that oil price uncertainty has a highly significant, negative effect on the U.S. business cycle – even after accounting for the influence of global oil market forces on both variables. Our computations show that the effect is also economically meaningful, especially during several episodes in the period after 1985. The influence of oil price uncertainty on the global business cycle, on the contrary, tends to be much weaker. The relevant coefficient often shows the expected sign, but it is typically small and not always statistically significant. A battery of robustness checks confirms our findings. These include the estimation of bi- and trivariate versions of the model,

using different proxies for the oil price and global economic activity, and changing the model's specification in various ways.

Besides providing strong evidence in favour of the theory of real options, our study is also relevant because it raises interesting questions for future research. One point will be to re-estimate our model with data for different countries and compare the results with our findings for the U.S. business cycle. Also, future studies should try to integrate others factors into our framework that potentially explain oil price uncertainty, as well as changes in economic activity.

The remainder of this paper is organised as follows: Section 2 presents the VAR GARCH-in-mean framework and elaborates on the details of the model's specification. Section 3 discusses the sources of our data series and provides a brief overview of their characteristics. In Section 4, we illustrate the appropriateness of our empirical approach and discuss the estimation results. Section 5 examines the robustness of our findings, while Section 6 concludes.

2. Empirical Methodology

In our study we use a VAR, GARCH-in-mean, asymmetric BEKK model, which goes back to work by Sims (1980), Engle and Kroner (1995), and Grier et al. (2004), and, since then, has been employed to the question of oil price uncertainty particularly by Rahman and Serletis (2011, 2012). In order to estimate the influence of oil price uncertainty on the U.S. economy while accounting for shocks in the global oil market, however, we augment the authors' model by two additional variables representing global aggregate supply and global aggregate demand for crude oil. Hence, our model is also related to previous oil market studies like Kilian (2009), Kilian and Park (2009) and Carstensen et al. (2013).

2.1 VAR GARCH-in-mean Framework

Our mean model can be written as follows:

$$\mathbf{y}_t = \boldsymbol{\mu} + \sum_{m=1}^6 \boldsymbol{\Gamma}_m \mathbf{y}_{t-m} + \boldsymbol{\Psi} \sqrt{\mathbf{h}_t} + \mathbf{e}_t \quad (1)$$

where $\mathbf{y}_t = (y_{1,t}, y_{2,t}, y_{3,t}, y_{4,t})'$ is a 4×1 vector of endogenous variables containing measures for the global supply of crude oil ($y_{1,t}$), global aggregate demand ($y_{2,t}$), the real price of oil ($y_{3,t}$), and U.S. economic activity ($y_{4,t}$). In addition, $\boldsymbol{\mu}$ is a 4×1 vector of constants and $\boldsymbol{\Gamma}_m$ for $m = 1, \dots, 6$ are 4×4 coefficient matrices.⁶ As we will see further below, the choice of including $p = 6$ lags of the model's four endogenous variables is sufficient to capture their dynamics at the time horizons relevant for our study. We further assume that, given the information set available at time t , Ω_{t-1} , the error terms in the 4×1 vector \mathbf{e}_t follow a multivariate t -distribution with a conditional mean of zero, time-varying conditional covariance matrix \mathbf{H}_t , and ν degrees of freedom:

$$\mathbf{e}_t | \Omega_{t-1} \sim t(\mathbf{0}, \mathbf{H}_t, \nu) \text{ with } \mathbf{H}_t = \begin{bmatrix} h_{11,t} & h_{12,t} & h_{13,t} & h_{14,t} \\ h_{12,t} & h_{22,t} & h_{23,t} & h_{24,t} \\ h_{13,t} & h_{23,t} & h_{33,t} & h_{34,t} \\ h_{14,t} & h_{24,t} & h_{34,t} & h_{44,t} \end{bmatrix},$$

where $h_{ij,t}, \forall i, j = 1, \dots, 4$, are the conditional variances and covariances of the innovations.

⁶ As it turns out to be the cause for a large outlier in the standardised residuals, we also include a dummy variable in the U.S. equation controlling for the impact of hurricane Katrina in September 2005. This, however, is not crucial for our conclusions and mainly serves computational purposes. Estimation results for the baseline model without the dummy are available upon request.

As we will also show later, assuming a multivariate t -distribution of the underlying shocks, instead of the regularly chosen multivariate normal, is warranted by the data and tends to improve the explanatory power of the model.

The third term in Equation (1) represents the GARCH-in-mean part of the model. In detail, it can be written as follows:

$$\Psi \sqrt{\mathbf{h}_t} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & \psi_{22} & \psi_{23} & 0 \\ 0 & 0 & \psi_{33} & 0 \\ 0 & 0 & \psi_{43} & \psi_{44} \end{bmatrix} \begin{bmatrix} \sqrt{h_{11,t}} \\ \sqrt{h_{22,t}} \\ \sqrt{h_{33,t}} \\ \sqrt{h_{44,t}} \end{bmatrix} \quad (2)$$

The elements in $\sqrt{\mathbf{h}_t}$ are the conditional standard deviations of the one-period-ahead forecast errors. We use them as our proxies for the uncertainty regarding each of the four variables. The coefficients in the 4×4 matrix Ψ thus measure the influence of these different kinds of uncertainty. In previous studies, the authors typically either choose to restrict the elements of Ψ so that only the influence of oil price uncertainty on one of the economic activity measures may be non-zero (ψ_{23} for the global business cycle, or ψ_{43} for U.S. activity), or decide to estimate a completely unrestricted version of the GARCH-in-mean part.⁷ In our case, we allow for the following parameters to be estimated freely:⁸

First, we allow both, ψ_{23} and ψ_{43} , to be non-zero. As these measure the influence of oil price uncertainty on the global and U.S. business cycle, respectively, they are the coefficients of particular interest. Second, it is important to control for the influence of business cycle uncertainty as the theory of real options predicts that it should also depress investment and, hence, economic activity. Thus, we estimate the effect of business cycle volatility on the respective variable's mean in the form of ψ_{22} and ψ_{44} . Third, we presume that uncertainty about the real price of oil has a positive influence on the variable's own mean. As Alquist and Kilian (2010), and Kilian and Murphy (2014) point out, speculative demand in the crude oil market might not only be driven by changes in expectations regarding future demand and supply of crude oil, but may also reflect precautionary motives. As uncertainty about future oil-supply shortfalls rises, firms' so-called convenience yield from having access to oil inventory increases because – by holding appropriate reserves – they might be able to prevent production breakdowns after unexpected shocks. Higher uncertainty, therefore, raises precautionary demand for crude oil and, hence, today's spot price. Consequently, we expect $\psi_{33} > 0$.⁹

⁷ As examples for the former case see Elder and Serletis (2010, 2011), Bredin et al. (2011) or Jo (2014). Rahman and Serletis (2011, 2012), on the contrary, do not impose any restrictions on Ψ in their bivariate model.

⁸ Note that, as the influence of the uncertainty variables depends on the dynamics in the VAR, as well as the GARCH model, each additional coefficient makes the estimation and, particularly, the interpretation of the results more complex, which is why we do not use an unrestricted version of Ψ . Apart from this, we also estimated the model with several other sets of possible restrictions and found that the one given in (2) fits the data specifically well. The results are available upon request.

⁹ Note that our perspective slightly differs from Alquist and Kilian's (2010) argument. In our view, higher uncertainty about expected oil-supply shortfalls is inextricably linked to oil price uncertainty, which then raises the real price of oil. It can be assumed, for instance, that the distribution of market participant's expectations widens in times of heightened uncertainty about fundamentals, and the same is then likely to follow for the distribution of oil price changes. Further, if oil price uncertainty would not be relevant, one would have no additional reason to worry about still being readily able to buy crude oil at the expected market price any time in the future. Thus, the economic rationale for an increase in precautionary demand for oil would be missing.

2.2 Asymmetric BEKK Model with Variance Shift

The dynamics of the conditional covariance matrix \mathbf{H}_t are assumed to be appropriately described by a multivariate GARCH(1,1) process. Specifically, we employ a modified version of the BEKK specification introduced by Engle and Kroner (1995) to model the time path of the conditional variances and covariances:¹⁰

$$\mathbf{H}_t = (\mathbf{C} + x_t \mathbf{E})(\mathbf{C} + x_t \mathbf{E})' + \mathbf{A}' \mathbf{e}_{t-1} \mathbf{e}_{t-1}' \mathbf{A} + \mathbf{B}' \mathbf{H}_{t-1} \mathbf{B} + \mathbf{D}' \mathbf{u}_{t-1} \mathbf{u}_{t-1}' \mathbf{D}, \quad (3)$$

where \mathbf{C} is a 4×4 lower triangular matrix of constants, and \mathbf{A} , \mathbf{B} and \mathbf{D} are 4×4 coefficient matrices.

One of our objectives is to improve the explanatory power of our model by taking into account recent findings in the empirical literature. Therefore, we decide to augment the standard BEKK model following two rationales:

First, Baumeister and Peersman (2013) point out that there has been a substantial and systematic increase in the variance of the real oil price since 1986. Around the same time, the variance of global crude oil production started to develop in the opposite direction. Also during the 1980s, the U.S. and other economies around the world entered the so-called “Great moderation”, a period characterised by low inflation and a reduced volatility of GDP growth and industrial production – see, for example, Stock and Watson (2003), or Summers (2005). Taken together, these facts make a strong case for a presumable structural change of the variances and covariances of all four of our variables during the mid-1980s. This potential trait of the data has so far not been taken into account by any of the previous studies using the VAR GARCH-in-mean approach. In order to change this, we allow for a one-time structural change in the constant components of the covariance process by including the products of the dummy variable x_t – which takes on the value 1 until June 1985 and 0 thereafter – and the constants in the 4×4 lower triangular matrix \mathbf{E} .

Secondly, the findings of Rahman and Serletis (2011, 2012) suggest that unexpected bad news about any of the variables might lead to a stronger increase of uncertainty than unexpected good news associated with a shock of similar magnitude. With the last term in (3), we incorporate this possibility into our model. In particular, the vector

$$\mathbf{u}_{t-1} = \begin{bmatrix} u_{1,t-1} \\ u_{2,t-1} \\ u_{3,t-1} \\ u_{4,t-1} \end{bmatrix} = \begin{bmatrix} e_{1,t-1} \cdot I_{e_{1,t-1} < 0} \\ e_{2,t-1} \cdot I_{e_{2,t-1} < 0} \\ e_{3,t-1} \cdot I_{e_{3,t-1} > 0} \\ e_{4,t-1} \cdot I_{e_{4,t-1} < 0} \end{bmatrix}$$

includes the product of the model’s residuals and an indicator variable, which takes on the value 1 if the respective condition for bad news is met, and zero otherwise. As can be seen, with respect to global oil production, and the global and U.S. business cycle variables we interpret negative shocks as bad news. Regarding the real price of oil, on the contrary, we follow Rahman and Serletis (2011, 2012), and assume that sudden price increases are viewed as unfavourable news by the majority of economic agents.

In summary, the effects of any shock on the covariance matrix depend on its size and sign. Via the inclusion of the GARCH-in-mean terms, we thus allow for a highly non-linear relationship between the causes and effects of oil price movements and of uncertainty.

¹⁰ The BEKK model offers a high degree of flexibility by allowing for volatility spillover effects and time-varying covariances between the variables. At the same time, it guarantees the positive-definiteness of \mathbf{H}_t .

2.3 Maximum Likelihood Estimation

We estimate the model's 175 parameters simultaneously using a maximum likelihood approach and robustified estimators for the standard errors of the coefficients. The calculations are conducted with version 9.00 of the WinRATS software package distributed by Estima.¹¹ As previous studies like Elder and Serletis (2010), Rahman and Serletis (2011), and Pinno and Serletis (2013) have already noted, the estimation of multivariate VAR GARCH-in-mean models involves a considerable number of complex calculations and is prone to computational difficulties. We have to acknowledge that, in this regard, our framework constitutes no exception. Therefore, we implement a self-developed estimation routine that involves running the main optimisation algorithm with a large number of randomly generated starting values that – to a certain degree – tend to improve with each step. While this renders the estimation very computationally intensive and time-consuming, it helps us to cope with these problems and to achieve convergence while covering a broad area of the likelihood surface. Further details are given in the appendix.¹²

3. Data & Descriptive Analysis

For our study of the global oil market and the U.S. economy, we use monthly data from 1975:01 to 2014:07. Following Kilian (2009), Baumeister and Peersman (2013), and others, global oil supply is measured by the growth rate of world crude oil production (first differences in logs multiplied by 100). The data is obtained from the U.S. Energy Information Administration (EIA). Regarding our series for U.S. and global economic activity, we construct two business cycle measures that have not been used in the relevant literature before and that are conceptually closely related to the output gap. As a basis for the global series, we use the industrial production index for the total of OECD countries, which we retrieve from the organisation's website. For the U.S. proxy, similar data is obtained from the IMF's International Financial Statistics (IFS) database. We then estimate the long-run trend of each series using a Hodrick-Prescott (HP) filter with $\lambda = 14400$, and, finally, calculate the deviation from this trend in percentage points.¹³ For y_3 , the real price of oil, we acquire the West Texas Intermediate (WTI) spot price series from the IFS. This price benchmark is often used in empirical investigations of the effects of oil price uncertainty.¹⁴ After deflating the series with the consumer price index (CPI) from the U.S. Bureau of Labor Statistics (excluding energy prices), it is transformed into log differences and multiplied by 100. To mitigate the end-point problem associated with the HP filter, we drop the last four of our observations. This leaves us with an effective sample from 1975:08 to 2014:03 for the estimation.

Table 1 shows some descriptive statistics for the three oil market variables and the U.S. business cycle series. The important thing to note here is that all four time series show a great amount of excess kurtosis. Their distributions are thus heavy-tailed. Not surprisingly, the resulting statistics of the Jarque-Bera (1980) test are highly significant with p-values of less than 1%, and, for each variable, the null hypothesis of normality has to be rejected. This clearly speaks in favour of the previously made

¹¹ Note that the scale parameter ν of the multivariate t -distribution is also calculated during the estimation process. For our regressions it typically lies in the range of 7 to 13.

¹² Note that, despite our extensive efforts, we were not always able to achieve convergence at the highest likelihood computed by our estimation routine. As a consequence, we had to discard some of the more complex robustness checks originally intended for the paper. For the remaining regressions, however, the differences are typically very small or zero, and, hence, should not lead to significantly different estimates. Table A1 in the appendix provides an overview of the calculated and converged values up to the 5th decimal.

¹³ We do conduct multiple robustness checks to ensure that our estimates are not driven by our measure for the global business cycle or choice of λ . The results are reported in Section 5.

¹⁴ See, for instance, Rahman and Serletis (2011, 2012), Elder and Serletis (2011) or Caporale et al. (2014).

assumption that the error terms are more appropriately described by a multivariate t -distribution instead of a multivariate normal.

Panel B of the table gives the results for several unit root- and stationarity tests we carry out to ensure the stationarity of our data. The test statistics of the augmented Dickey-Fuller (1981) (ADF) tests are presented first. Using the Akaike information criterion (AIC) to select the lag order for each test automatically, all four test statistics turn out to be highly significant. Likewise, the Kwiatkowski et al. (1992) (KPSS) tests do not indicate any kind of unit root or non-stationarity behaviour of the series. We therefore conclude that all variables which enter our empirical model are stationary.

Table 1: Summary Statistics

	<i>Oil Production Growth y_1</i>	<i>Global Business Cycle y_2</i>	<i>Oil Price Change y_3</i>	<i>U.S. Business Cycle y_4</i>
A. Descriptive Statistics				
Mean	0.082	-0.020	0.152	-0.032
Variance	2.492	3.301	57.669	3.394
Skewness	-1.659	-1.264	-0.361	-0.827
Excess kurtosis	10.368	6.528	4.452	2.602
Jarque-Bera normality test	2320.790 (0.000)	959.734 (0.000)	398.414 (0.000)	186.125 (0.000)
B. Unit Root & Stationarity Tests				
ADF unit root tests	-5.855 (0.000)	-7.661 (0.000)	-10.219 (0.000)	-7.152 (0.000)
KPSS stationarity tests	0.050 (> 0.100)	0.014 (> 0.100)	0.090 (> 0.100)	0.016 (> 0.100)

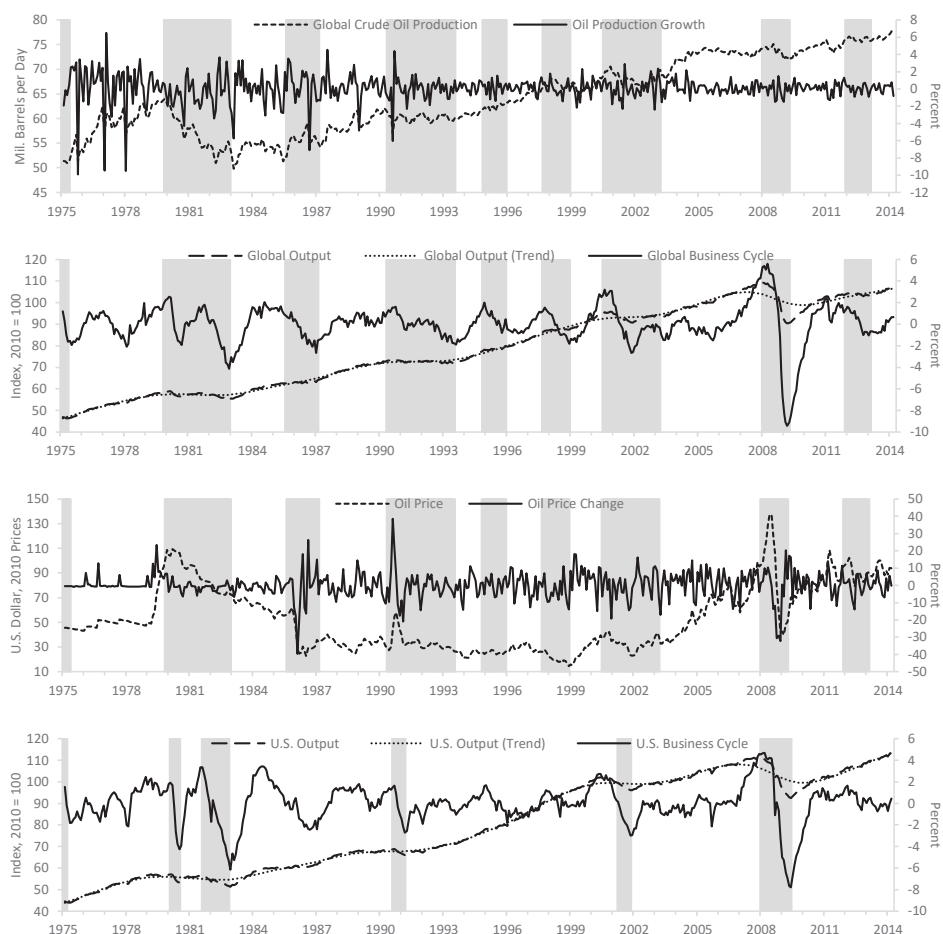
Notes: p-values in parentheses. Unit root and stationarity tests are specified with an intercept (no trend). ADF tests: Lag length determined by AIC. KPSS tests: Newey-West automatic bandwidth selection, Bartlett Kernel.

Figures 1 to 4 depict the variables used for the estimation, the respective series before the transformation of the data, and – in the cases of the business cycle measures – the estimated trend components. The shaded areas indicate OECD and U.S. recessions, respectively.

As can be seen in the second and fourth diagram, our global and U.S. activity indices seem to capture the state of the economies reasonably well. During each recession identified by the indicators from the Federal Reserve Bank of St. Louis database, we can see an upper-left to lower-right movement from the peak to the trough. Note further that in Figure 3 the previously discussed shift in the variance of the real oil price after 1985 becomes immediately evident. The same, though, is not true for the volatility of global oil production implicitly depicted in Figure 1. The findings of Baumeister and Peersman (2013) imply, however, that the bulk of the structural change in the global oil market took place during this period.¹⁵ Only due to the fact that during the late 1980s large oil supply disturbances were still prevalent in the data, the structural change in the variance is less obvious for the oil production series.

¹⁵ Specifically, see Figures 4-6 in Baumeister and Peersman (2013).

Figures 1-4: Original Data and Transformed Time Series



Note further that the strong decline in the real price of oil during the 1985/86 episode seems not to have led to a significant improvement in the U.S. business cycle. Previous studies, like Elder and Serletis (2010) or Rahman and Serletis (2011), attribute this to the rapidness of the price reduction. Their argument is that this unexpected shock led to a strong increase of oil price uncertainty, which, in turn, hampered economic output. A similarly abrupt decrease of the WTI can be observed during the recent financial crisis of 2008/09 – again without a substantial increase in the business cycle. Regarding the asymmetric effects of oil price shocks these two episodes are thus of particular interest. We will investigate in the next section whether oil price uncertainty has played a role during these events, or whether other factors such as shocks to the global business cycle might explain the apparent lack of a positive impulse to the U.S. economy.

4. Empirical Results

Before we actually discuss our baseline estimation results, it is important to analyse the empirical validity of our approach and also to test our model for signs of possible misspecification. Regarding the latter, we analyse the properties of the standardised residuals $\hat{\varepsilon}_t = \hat{H}_t^{-1/2} \hat{e}_t$, where $\hat{H}_t^{-1/2}$ is the inverse of the Cholesky factor of the conditional covariance matrix computed at each point in time. If our model is specified correctly, $\hat{\varepsilon}_t$ should be free from autocorrelation and ARCH effects.¹⁶

4.1 Model Specification

To test the two hypothesis, we apply the multivariate Q -test described by Hosking (1981) to the levels (serial correlation) and the squares of the standardised residuals (heteroscedasticity). In addition, we use the Akaike information criterion to distinguish between different model specifications regarding their explanatory power. Table 2 gives the corresponding results.

Table 2: Model Selection

Model	Components	AIC	$Q_{MV}(9)$	$Q_{MV}(18)$	$Q_{MV}^2(9)$	$Q_{MV}^2(18)$
(1)	Constants	8104	3008.875 (0.000)	3990.678 (0.000)	2249.718 (0.000)	2869.365 (0.000)
(2)	+ VAR(6)	6283	69.956 (0.999)	246.801 (0.962)	441.729 (0.000)	646.502 (0.000)
(3)	+ Diagonal BEKK	5897	106.319 (0.992)	280.429 (0.614)	233.520 (0.000)	349.750 (0.007)
(4)	+ t-distribution	5791	95.877 (0.999)	277.126 (0.667)	287.943 (0.000)	410.547 (0.000)
(5)	+ asymmetries, +variance shift	5715	103.434 (0.996)	279.693 (0.626)	195.015 (0.003)	325.316 (0.064)
(6)	+ spillover effects, + GARCH-in-mean terms (full model)	5691	113.531 (0.971)	279.939 (0.622)	116.911 (0.952)	263.704 (0.845)

Notes: *p*-values in parentheses. Test statistics significant at the 5% level are printed in bold.

From top to bottom, the table illustrates the changes in the test statistics and the AIC as we add more of the previously described components to our model. Starting only with constants in the mean and variance equations, we can see that all four test statistics are of a very large magnitude and highly significant with *p*-values smaller than 1%. Therefore, the null hypotheses of the tests have to be rejected. Multivariate autocorrelation and conditional heteroscedasticity are strongly present in the data, which confirms the appropriateness of our VAR GARCH approach.

Once we add the six lags of the vector of endogenous variables \mathbf{y}_t , the Q_{MV} -statistics testing for autocorrelation of orders up to $m = 9$ and 18, respectively, become insignificant.¹⁷ Additionally, the value of the information criterion reduces from 8104 to 6283. Thus, our model sufficiently captures the system's mean dynamics at a horizon of at least one and a half years. Next, we also allow for a dynamic evolution of \mathbf{H}_t by adding the GARCH terms $\mathbf{A}'\mathbf{e}_{t-1}\mathbf{e}_{t-1}'\mathbf{A}$ and $\mathbf{B}'\mathbf{H}_{t-1}\mathbf{B}$. In this step, however, the coefficient matrices \mathbf{A} and \mathbf{B} are restricted to be non-zero only on their main diagonals. This so-called diagonal BEKK specification does not allow for a transmission of shocks and volatility spillovers to the conditional second-moments of other variables. As a result, the AIC and the test

¹⁶ For a more detailed discussion of the desired properties of the standardised residuals and the challenges surrounding the diagnostic checking of multivariate GARCH models, see Ding and Engle (2001), and Bauwens et al. (2006), respectively.

¹⁷ Additional results for lag orders 12 and 24 can be found in Table A2 in the appendix.

statistics for conditional ARCH effects improve once more, but the latter still reject the null hypothesis of homoscedasticity at all conventional levels of significance. This finding also remains fundamentally unchanged when we move from assuming a multivariate normal to a multivariate t-distribution, and when we further include the mid-1980s structural shift and the asymmetry terms in the covariance model. Only when we finally arrive at our fully specified model by allowing for volatility spillovers and adding the GARCH-in-mean terms, all four test statistics become insignificant. On aggregate, the standardised residuals of the baseline model thus have the desired properties and do not show any significant signs of misspecification. In addition, the full model yields the lowest value of the information criterion, suggesting a high capability to explain the observed data.

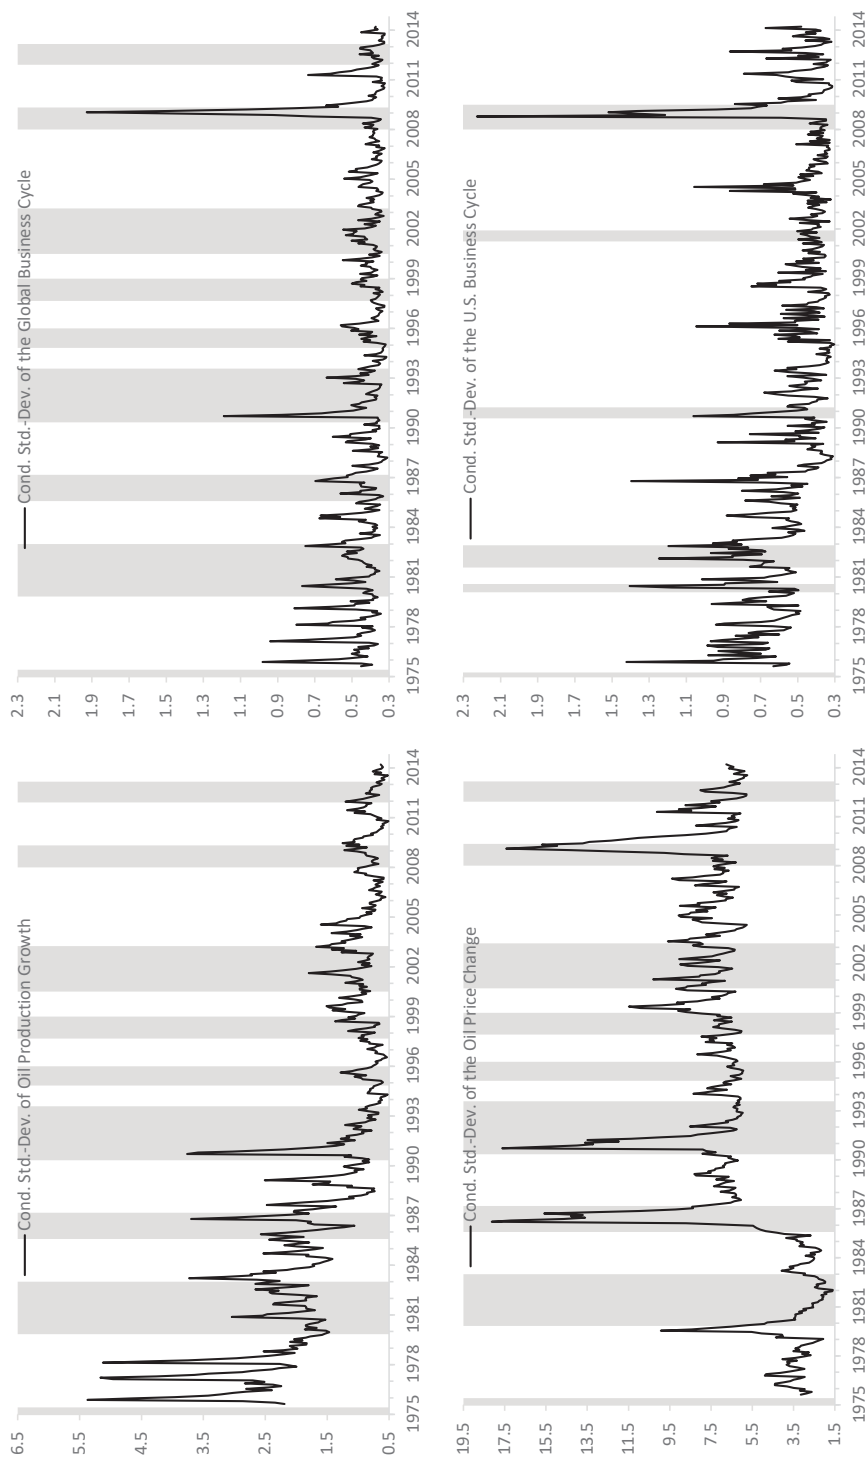
4.2 Evolution and Causes of Oil Price Uncertainty

We begin the discussion of the results by first looking at the uncertainty measures generated by our model. These are depicted in Figures 5-8 together with the relevant recession indicators. For three of the four volatility series the implications of the one-time structural change in the conditional variances become evident immediately. Consistent with the findings of Baumeister and Peersman (2013), we see a decline of the conditional standard deviation of oil supply and an increase in the volatility of the real price of oil since the mid-1980s (see Figures 5 and 7 in the upper and lower left, respectively). Especially the latter is abrupt and of a large magnitude, which confirms our view that ignoring this trait of the data might have led to rather imprecise measures of oil price uncertainty being used in previous studies. The observable decline in U.S. business cycle volatility mostly fits the descriptions of the great moderation. As we do not see a similar change in the uncertainty measure depicted in Figure 6, however, the narrative does not seem to be applicable to the global level. Nonetheless, the general importance of the changes can also be confirmed statistically: When we conduct a test for the joint significance of the coefficients in $x_t E$, the null hypothesis of the absence of a covariance shift is strongly rejected with a test statistic of $\chi^2(10) = 201.665$ and a corresponding p-value of 0.000.

Looking more closely at Figure 7, we are able to identify four distinct episodes of heightened oil price uncertainty. The first one, leading to a local maximum of the measure in September 1979, can be attributed to the Iranian revolution. While showing relatively high values of oil price uncertainty compared to the rest of the decade after 1975, however, the spike is tiny compared to the amounts of uncertainty that can be measured during the following episodes. This finding, first noted by Elder and Serletis (2010), becomes even more evident with our model. In addition, the height of uncertainty precedes the global and U.S. recessions of the early 1980s by several months, making it unlikely that oil price volatility somehow contributed to the economic downturn. The surge in uncertainty after 1985, on the contrary, does coincide with a global recession. Even though the volatility rise is triggered by a drop in the real price of oil due to Saudi Arabian overproduction, the abruptness of the unexpected price change causes oil price uncertainty to increase to its highest value in the sample. Moreover, it takes more than a year for it to return to a level close to its long-run average. The 1990 spike associated with the Iraqi invasion of Kuwait is even more interesting because we can observe a significant movement in all four uncertainty measures simultaneously. In fact, all conditional standard deviations reach their highest value since 1986 in September 1990 – immediately after the invasion. This suggests that the connectedness of the variables' means identified in the previous literature might also extend to their second moments.

During the financial crisis, the behaviour of the volatilities is similarly salient. As we do not see simultaneous, but successive peaks of the uncertainty proxies, however, this event's character is somewhat different from the 1990 episode. Within one month, from September to October 2008, uncertainty regarding U.S. economic activity almost triples and reaches its sample maximum. In the following three months, the conditional standard deviations of the global business cycle measure and the real price of oil rise by 164.3 % and 78.8 %, respectively, both reaching local peaks in January 2009.

Figures 5-8: Volatility Series



Afterwards, it takes at least six months for each variable to return in proximity to its sample average. The only exception is the volatility of global oil production, which hardly rises in absolute terms. This marks another difference to the 1990s volatility surge. Nevertheless, the overall level of uncertainty as measured by our model not only reaches unprecedented heights during the 2008/2009 crisis, but also stays relatively high for an extended amount of time.

Jo (2014) argues that a shortcoming of the GARCH-in-mean approach is that changes in the volatility process can only occur together with changes in the mean of the variables. While this is factually correct, one does not necessarily have to agree that this is a disadvantage. In contrast to the author's stochastic volatility approach, we do not need to use outside information like a realised volatility component to obtain a precise measure of oil price uncertainty.¹⁸ Moreover, in our framework, the uncertainty proxy is truly endogenous. As such, it can be explained by the evolution of the shocks to each variable and the dynamics of their conditional second moments. Therefore, our model allows us to get a better understanding of what actually causes oil price uncertainty by analysing which other variable's shocks and variance might also influence the series depicted in Figure 7. We can infer whether such relationships exist from the estimates of Eq. (3).¹⁹

$$\hat{A} = \begin{bmatrix} \mathbf{-0.542} & -0.037 & -0.175 & -0.019 \\ (0.000) & (0.067) & (0.088) & (0.488) \\ 0.085 & \mathbf{0.279} & -0.597 & \mathbf{0.198} \\ (0.380) & (0.000) & (0.336) & (0.001) \\ 0.000 & \mathbf{0.007} & \mathbf{0.403} & 0.003 \\ (0.966) & (0.022) & (0.000) & (0.496) \\ \mathbf{0.174} & \mathbf{-0.210} & \mathbf{1.029} & \mathbf{-0.160} \\ (0.009) & (0.000) & (0.001) & (0.006) \end{bmatrix}, \hat{D} = \begin{bmatrix} 0.059 & \mathbf{-0.089} & \mathbf{0.278} & \mathbf{-0.213} \\ (0.525) & (0.000) & (0.009) & (0.000) \\ \mathbf{0.369} & \mathbf{0.495} & 1.286 & -0.003 \\ (0.003) & (0.000) & (0.080) & (0.970) \\ \mathbf{0.045} & \mathbf{0.014} & -0.040 & -0.010 \\ (0.000) & (0.002) & (0.513) & (0.084) \\ -0.189 & 0.040 & \mathbf{-0.912} & \mathbf{0.774} \\ (0.205) & (0.672) & (0.022) & (0.000) \end{bmatrix},$$

$$\hat{B} = \begin{bmatrix} \mathbf{0.794} & -0.007 & 0.068 & \mathbf{0.094} \\ (0.000) & (0.608) & (0.442) & (0.007) \\ -0.302 & \mathbf{0.676} & 0.224 & \mathbf{0.927} \\ (0.083) & (0.000) & (0.708) & (0.000) \\ -0.006 & \mathbf{-0.009} & \mathbf{0.819} & -0.002 \\ (0.204) & (0.000) & (0.000) & (0.822) \\ 0.327 & 0.078 & \mathbf{1.327} & \mathbf{-0.648} \\ (0.210) & (0.483) & (0.043) & (0.000) \end{bmatrix}$$

As described before, the coefficients in \hat{A} and \hat{D} illustrate the relationship between the model's mean innovations e_t and the conditional covariances in H_t , while the elements of \hat{B} quantify the influence of H_{t-1} on H_t , including possible variance spillovers. Due to the form of the BEKK model, however, both linkages are highly nonlinear and thus difficult to interpret.²⁰ We will hence simplify the analysis by mostly looking at the coefficients' significance.

¹⁸ Admittedly, this might also have to do with the fact that we are using monthly instead of quarterly data.

¹⁹ Numbers in parenthesis are p-values and each coefficient significant at the 5%-level is printed in bold.

²⁰ Note that each coefficient ultimately enters the equations for the conditional variances in different squared and multiplicative forms. The former are typically found to measure the direct influence of any shock to one variable onto the conditional variance of itself or of another. $\hat{a}_{13}^2 e_{1,t-1}^2$, for instance, would measure the direct impact of any shock to global crude oil production on the conditional variance of the real oil price. The same shock, however, enters the equation for $h_{33,t}$ three more times with $2\hat{a}_{23}\hat{a}_{13}e_{2,t-1}e_{1,t-1}$, $2\hat{a}_{33}\hat{a}_{13}e_{3,t-1}e_{1,t-1}$, and $2\hat{a}_{43}\hat{a}_{13}e_{4,t-1}e_{1,t-1}$. As one can see, apart from the constant direct influence, the total impact of any given shock $e_{1,t-1}$ on the conditional variance of the oil price depends on the size and sign of several coefficients, and on the total composition of all residuals in the previous period – again with respect to their size and sign. The matter is complicated further by the inclusion of the asymmetry terms with their influence being represented by the coefficients in the \hat{D} matrix and the possibility of volatility spillovers as represented by the elements of \hat{B} . In addition, the coefficients are also not readily comparable to each other in terms of their magnitude, because the

The off-diagonal elements in $\hat{\mathbf{A}}$ contain information about whether there exists a link between the mean innovations of one variable and the uncertainty measured for another. The third column of the matrix, in particular, gives an indication of how unexpected shocks to other variables directly tend to raise oil price uncertainty. As both, $\hat{a}_{13} = -0.175$ and $\hat{a}_{43} = 1.029$, are statistically significant, we first see that this is the case for the shocks to global crude oil production and U.S. economic activity. The influence of sudden changes of global demand – at least at first sight – seems not to be important, as \hat{a}_{23} is insignificant with a p-value of 0.336.

The direct transmission to oil price uncertainty is significantly stronger, however, when the unexpected mean shock is associated with bad news rather than good news. This can be seen from the elements at the corresponding position in the $\hat{\mathbf{D}}$ matrix. If, for instance, there is an unexpected decline in the global business cycle, the relevant coefficient increases more than fivefold, from just $\hat{a}_{23}^2 = (-0.597)^2 = 0.356$ for a positive shock, to $\hat{a}_{23}^2 + \hat{d}_{23}^2 = (-0.597)^2 + (1.286)^2 = 2.010$ for a negative shock of similar magnitude. In this particular case, the influence also becomes statistically significant at the 10% level.

With respect to the other variables of the model, there are such relationships as well. E.g., the first row of $\hat{\mathbf{D}}$ shows that negative oil production shocks tend to significantly increase the uncertainty measures of all other three variables, while the last row of $\hat{\mathbf{A}}$ indicates the same for any particular shock to U.S. economic activity. Furthermore, as there are several highly significant off-diagonal elements in $\hat{\mathbf{B}}$, the conditional second moments themselves are also closely interrelated. There are significant variance spillovers from each variable to at least one of the others. Looking again at oil price volatility in particular, our results show that higher economic uncertainty in the U.S. causes a direct increase of h_{33} , since \hat{b}_{43} is significant at the 5% level.

In summary, the statistical significance of the majority of GARCH coefficients once more confirms the appropriateness of our modelling approach. Furthermore, our findings suggest that oil price uncertainty is not only caused by sudden changes in the real price of oil, but also by variance spillovers and unexpected shocks that simultaneously tend to raise the uncertainty measured for other macroeconomic variables. In fact, testing for the joint insignificance of the off-diagonal elements in $\hat{\mathbf{A}}$, $\hat{\mathbf{B}}$ and $\hat{\mathbf{D}}$ leads to a clear rejection of the null hypothesis [$\chi^2(36) = 1037.251$ (0.000)]. Significant asymmetries in the conditional second moments imply that these effects are stronger in case of bad news. Altogether, this leads to the important conclusion that – at least to some degree – oil price volatility can also be viewed as a measure of general macroeconomic uncertainty and distress.

4.3 Effects of Oil Price Uncertainty

After discussing the possible causes of oil price uncertainty, we will now turn to its potential effects. The main question our study tries to answer is whether we can still find a negative influence on the U.S. economy once we account for the endogeneity of the real oil price with respect to global demand and supply factors as proposed in the literature. The estimated influences of the uncertainty measures on the variables' means are given by $\hat{\Psi}$ as follows:²¹

latter depends on the average sample size of the shocks the coefficient is multiplied with and on the average magnitude of the variance whose equation it enters.

²¹ Again, numbers in parenthesis are p-values and each coefficient significant at the 5%-level is printed in bold.

$$\Psi = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & -\mathbf{0.628} & -0.007 & 0 \\ 0 & 0 & \mathbf{0.303} & 0 \\ 0 & 0 & -\mathbf{0.041} & -\mathbf{0.325} \end{bmatrix}$$

(0.000) (0.288) (0.003) (0.000)


Note first that all five of the estimated GARCH-in-mean coefficients show the a priori expected signs. Increased uncertainty about the movements in the global and U.S. business cycles depresses economic activity, as $\hat{\psi}_{22}$ and $\hat{\psi}_{44}$ are negative and highly significant. What's more, we find that the conditional second moment of the real price of oil has a positive influence on the variable's mean and that this influence is significantly different from zero ($\hat{\psi}_{33} = 0.303$). In combination with the substantial interdependences we found within the GARCH model, we interpret this as indirect evidence in favour of the hypothesis that increased uncertainty about potential future supply-shortfalls leads to higher oil price volatility, which, in turn, raises precautionary demand and hence the real oil price.

More importantly, $\hat{\psi}_{23}$ and $\hat{\psi}_{43}$ both have a negative sign as well. Therefore, as predicted by the theory of real options, oil price uncertainty tends to have a detrimental influence on economic activity. However, only for the U.S. business cycle the point estimate of -0.041 is found to be significant at all conventional thresholds (p-value of 0.000). The first coefficient, which measures the influence on global demand, is associated with a p-value of 0.288 and is thus not statistically meaningful. Since the series y_2 and y_4 have sample variances of similar magnitudes, in this case it is also possible to directly compare the two coefficients with respect to their sizes. As we can see, the direct influence of oil price uncertainty on the U.S. business cycle is almost six times as high as its direct influence on the global measure. To conclude, there is still a highly significant negative impact of oil price uncertainty even when taking into account the potential influences of global demand and supply shocks on the U.S. economy. Regarding this aspect, our study thus joins the ranks of Elder and Serletis (2010) and the numerous follow up studies which also utilized the GARCH-in-mean approach. On the other hand, though, with our baseline model we are not able to confirm Jo's (2014) finding of a significant effect on global economic activity.

Apart from looking at the statistical importance of these effects, we can also conduct some straightforward calculations to get an impression of their economic relevance. In relation to the original two business cycle series and the growth rate of the real oil price, Figures 9 to 11 show the continuous influence of oil price uncertainty shocks as they run through the equations of the estimated VAR model. Note that, as we simply define these uncertainty shocks as deviations from the sample mean, the overall informational content of these diagrams should not be overstressed.²² That being said, they still yield some interesting insights. We can see in Figure 9, for instance, that, in relation to the overall dynamics of the global business cycle, the influence of oil price uncertainty is rather small for most of the time. Only during the three episodes of extraordinary high volatility described earlier, a negative contribution to global economic activity becomes clearly visible.

²² Basically, this is a dynamic version of the impact calculations conducted by previous studies like Pinno and Serletis (2013), for example. The calculated uncertainty influence always represents the weighted-average of the composition of underlying mean innovations, which determine the changes in oil price volatility at any given point in time. Hence, we do not distinguish between the individual contributions of these shocks to oil price uncertainty (directly or indirectly via variance spillovers). Also, we only consider their effects through the oil price uncertainty channel and do not compare them with any simultaneous effects through the mean. As an upside, however, similarly to Jo's (2014) approach, the calculations do not require an identification scheme for structural shocks, which, given our setup, might not be trivial to come up with.

Figures 9-11: Dynamic Effects of Oil Price Uncertainty²³



Similarly, in the next figure, the precautionary demand effect originating from oil price uncertainty is strongest in 1986, 1990/91 and 2008/09. Throughout a twelve-month window beginning with each peak, the average value of the volatility influence amounts to about 15 to 23 percent of the standard deviation of the WTI growth rate during the respective time frame. A small, but non-negligible effect.²⁴ In 1979/1980, on the contrary, our estimates imply that oil price uncertainty just explains a small fraction of the protracted price increase. The influence increases relatively little and is above zero for a mere three months. This suggests that the significant increase in speculative demand identified by Kilian and Murphy (2014) for that time might have been largely driven by shifts in expectations and not precautionary motives.

Finally, the time-varying influence of oil price uncertainty on the U.S. business cycle is depicted in Figure 11. The estimates indicate that a large proportion of the recessionary movements in the business cycle during the mid-1980s and early-1990s might be attributed to heightened oil price volatility. Put in numbers, from its local peak in mid-1984 to the trough in September 1986, U.S.

²³ The influence on the growth rate of oil production is not depicted because – even though it is technically not zero – it is so small that it would not be visible in the diagram.

²⁴ The specific time windows for the calculation were 1986:03-1987:02, 1990:09-1991:08 and 2009:01-2009:12.

industrial production falls from 3.43 percentage points above its long-run trend to 2.45 percentage points below. During the same period, the negative influence of oil price volatility rises by roughly 3.24 percentage points, potentially explaining about 55% of the decline.²⁵ Further, from September 1990 to March 1991, the business cycle measure falls by 4.37 points, while the calculated influence of oil price uncertainty increases by 1.15 in absolute terms, again potentially explaining about 26% of the recessionary development. After 1991 the effect somewhat retracts and does not seem to contribute largely to any of the short-term up- or downward movements of U.S. economic activity for about eighteen years. It then rises again significantly during the financial crisis, explaining approximately 16% of the unprecedented drop in economic activity between early 2008 and mid-2009.

Three more interesting points should be noted. First, in each case, the subsequent economic recovery also coincides with the dissipation of uncertainty. Circa 21.5% of the sharp rebound between March and September 1991, for example, may be attributed to subsiding oil price volatility. Hence, during these three episodes, the changes in the influence of oil price uncertainty fit the timing of the business cycle fluctuations surprisingly well. Second, our empirical analysis suggests that during the decade between 1975 and 1985 lower-than-average oil price uncertainty actually fostered economic activity in the U.S. As our model also suggests, however, this effect does not become visible in the data because it was mitigated by the above-average volatility of the U.S. business cycle before the advent of the Great Moderation.²⁶ It would be interesting to see, whether future research might be able to find any connection between the reduction of production volatility in the U.S., and the structural changes in the global oil market after the mid-80s identified by Baumeister and Peersman (2013). Third, we cannot find any evidence that oil price uncertainty somehow contributed to the 1980 and 1982 recessions. We do agree with Elder and Serletis (2010) that oil price uncertainty, even though small given the overall picture, was relatively high shortly before the 1980 downturn. However, our calculations of the dynamic influence clearly show that the timing of the volatility surge fits the small retraction in 1979, but is strongly off for the 1980 and 1982 episodes. Therefore, we remain sceptical towards this narrative, which was originally brought up by Pindyck (1991).

To summarise, our empirical results confirm the a priori expectations of a negative effect of oil price uncertainty on economic activity and of a positive influence on the real oil price itself through a precautionary demand effect. Hence, our qualitative findings are largely consistent with the predictions made by the theory of real options and also with the empirical results of previous studies employing the GARCH-in-mean methodology – even after controlling for the influences of global oil supply and demand on the U.S. economy. Quantitatively, though, given the importance of energy prices in typical investment projects, the influence of oil price uncertainty we measure is surprisingly strong – specifically in the U.S. As a matter of fact, previous studies of oil price uncertainty have already been criticised on similar grounds.²⁷ However, we find significant volatility spillovers and influences of other variables' shocks in our multivariate GARCH model, which suggests that oil price uncertainty, to a certain degree, also reflects and is a channel of transmission of other kinds of macroeconomic uncertainty and shocks. Given that Kilian (2009), Alquist and Kilian (2009), and Kilian and Murphy (2014) have argued that oil prices in part reflect expectations and uncertainty regarding future economic conditions, this seems not to be too far-fetched. This result also underpins the plausibility of measuring a large influence of oil price volatility because macroeconomic uncertainty might negatively affect all kinds of future-related economic and political decisions – not just investment.

²⁵ When we start the calculation in 1986:01, this number even rises to 86%.

²⁶ Note that, even though the measured positive effect of low oil price volatility comes out smaller, it is not necessary to include the variance shift component in the model to arrive at this conclusion.

²⁷ See Kilian (2014).

5. Sensitivity Analysis

To ensure that the core insights discussed in the previous section are not driven by the specificities of our empirical approach, we conduct a battery of robustness checks. These can roughly be divided into three different categories: Modifications of the model's specification, variations in the underlying data, and the estimation of models with a reduced number of endogenous variables. The results are shown in Table 3.

Changes in Model Specification

First, we vary the number of lags in the VAR model. When we reduce m to 3 lags, for instance, we obtain similar results to the baseline estimation regarding the effects of oil price uncertainty. The influence on the global business cycle remains negative, but highly insignificant with a p-value of 0.609. Likewise, it is once more much smaller than the effect on the U.S. business cycle, as $|\hat{\psi}_{23}| = |-0.006| < |\hat{\psi}_{43}| = |-0.044|$. The latter again turns out to be significant at all common thresholds. In this case, both estimates are also very close in magnitude to the original results from the baseline model, $\hat{\psi}_{23} = -0.007$ and $\hat{\psi}_{43} = -0.041$.²⁸ With $m = 9$, however, we do see a few changes with respect to both coefficients. In particular, $\hat{\psi}_{23}$ becomes highly significant and increases relative to $\hat{\psi}_{43}$. Nonetheless, the latter still remains almost double in size.

Next, we examine whether the results are sensitive to the specific date chosen for the shift in the variance-covariance structure. Corresponding to the remarks in the literature regarding the advent of the great moderation, we first choose 1984:01 as an alternative shift date.²⁹ Afterwards, we conduct a second estimation where we move the shift date forward to 1993:01. As can be seen from the table, in both cases our results are once more surprisingly robust.³⁰

Other robustness checks we conduct involve leaving out one or more of the standard and non-standard components of the model. With the exception that $\hat{\psi}_{23}$ is sometimes statistically significant, the qualitative results from the baseline model remain largely unchanged and, hence, we decide not to report all of these results in detail.³¹ However, the most rigorous way to ensure that the earlier conclusions are not driven by our specific estimation approach or econometric model, is to abandon these elements altogether.

To be more precise, for the last robustness check in this category, we discard the multivariate GARCH model and our simultaneous estimation approach, and, instead, estimate Eq. (1) with regular OLS estimators. Thereby, we add the uncertainty series generated by our baseline model as exogenous regressors.³² Even though the coefficients and standard errors might then be affected by the generated regressors problem discussed in Pagan (1984), we find that the parameters retain their expected signs. In addition, while the influence of oil price uncertainty on the U.S. business cycle is to some extent weaker and the influence on the global business cycle somewhat stronger than before, there are only minor changes with respect to the two coefficient's p-values.

²⁸ Note, however, that in general the coefficient sizes are only roughly comparable across estimations as the properties of the relevant volatility series tend to change with the specification of the model and the data. Looking at the sign and significance of the estimated parameters is typically more informative.

²⁹ See, for instance, Stock and Watson (2003) or Summers (2005).

³⁰ Nevertheless, changing the shift date leads to a significant deterioration of the AIC and of the uncertainty proxies' general properties. Also, several previously significant shift coefficients become insignificant.

³¹ The results are available upon request. Among others, we re-estimated the baseline model without the GARCH terms, $B'H_{t-1}B$, the variance shift, variance spillovers or asymmetry terms.

³² This idea is based on the additional robustness checks in Bredin et al. (2011).

Table 3: Robustness Checks

		Global Business Cycle y_2		Oil Price Change y_3	U.S. Business Cycle y_4	
	Sample	$\hat{\psi}_{22}$	$\hat{\psi}_{23}$	$\hat{\psi}_{33}$	$\hat{\psi}_{43}$	$\hat{\psi}_{44}$
<i>A. Model Specification</i>						
Baseline model ($m = 3$)	1975:05- 2014:03	-1.280 (0.010)	-0.006 (0.609)	0.250 (0.000)	-0.044 (0.002)	-0.690 (0.001)
Baseline model ($m = 9$)	1975:11- 2014:03	-0.362 (0.003)	-0.010 (0.000)	0.241 (0.000)	-0.018 (0.000)	0.090 (0.488)
Variance shift in 1984:01	1975:08- 2014:03	-0.637 (0.000)	-0.004 (0.272)	0.287 (0.000)	-0.042 (0.000)	-0.332 (0.000)
Variance shift in 1993:01	1975:08- 2014:03	-2.141 (0.019)	-0.013 (0.011)	0.356 (0.000)	-0.045 (0.000)	-0.841 (0.040)
Eq. (1) estimated with OLS	1975:08- 2014:03	-0.933 (0.000)	-0.011 (0.201)	0.122 (0.338)	-0.025 (0.016)	-0.068 (0.668)
<i>B. Data Modifications</i>						
Reduced sample excl. the financial crisis	1975:08- 2008:01	-0.357 (0.000)	-0.011 (0.004)	0.182 (0.000)	-0.036 (0.000)	-0.345 (0.000)
y_3 : Nominal WTI	1975:08- 2014:03	-0.643 (0.478)	-0.006 (0.364)	0.249 (0.000)	-0.040 (0.016)	-0.327 (0.071)
y_3 : U.S. RAC	1975:08- 2014:03	-0.326 (0.007)	-0.011 (0.006)	0.099 (0.001)	-0.023 (0.000)	-0.273 (0.010)
y_2 : proxy based on Lippi and Nobili's RoW output ($m = 3$)	1975:05- 2009:02	0.621 (0.112)	0.054 (0.138)	0.160 (0.004)	-0.035 (0.000)	-0.632 (0.001)
y_2 : Kilian's GAI ($m = 3$)	1975:05- 2014:03	0.282 (0.358)	0.078 (0.155)	0.233 (0.003)	-0.044 (0.000)	-0.963 (0.047)
y_2, y_4 : HP-filtered with $\lambda =$ 129600	1975:08- 2014:03	-0.576 (0.000)	-0.006 (0.087)	0.278 (0.000)	-0.033 (0.000)	-0.161 (0.109)
<i>C. Reduced Models</i>						
Trivariate model global $\{y_1, y_2, y_3\}$	1975:08- 2014:03	-0.738 (0.000)	-0.009 (0.020)	0.276 (0.001)	-	-
Trivariate model U.S. $\{y_1, y_3, y_4\}$	1975:08- 2014:03	-	-	0.267 (0.000)	-0.029 (0.000)	-0.562 (0.000)
Bivariate model global $\{y_2, y_3\}$	1975:08- 2014:03	-0.761 (0.008)	-0.007 (0.043)	0.239 (0.005)	-	-
Bivariate model U.S. $\{y_3, y_4\}$	1975:08- 2014:03	-	-	0.253 (0.000)	-0.023 (0.018)	-0.428 (0.137)

Notes: *p*-values in parentheses. Test statistics significant at the 5% level are printed in bold.

Data modifications

For the next category, we leave the specification of the baseline model unchanged but modify the underlying data. The reader might wonder, for instance, whether our results are somehow driven by the extraordinary levels of volatility in y_2 , y_3 and y_4 that we observe during the financial crisis of 2008/09. In order to make sure that they are not, we cut the sample at 2008:01. As a result, $\hat{\psi}_{23}$ is again larger in magnitude and becomes significant with a *p*-value of less than 1%. The effect for the U.S. slightly weakens to -0.036 , but its high significance remains unaltered.

Results very similar to the ones from the baseline model can also be found when we do not deflate the WTI and use the nominal oil price series instead. With this setup, only the p-value of $\hat{\psi}_{43}$ increases slightly to 0.016 so that significance at the 1% level is not maintained. Regarding the WTI itself, it is sometimes argued that it may not be an optimal benchmark for the global price of crude oil because over the sample period it was influenced by many U.S.-specific factors.³³ This could be an explanation for why we consistently find that the influence of oil price uncertainty is more relevant for the U.S. business cycle than for global economic activity. To examine this issue, we follow Elder and Serletis (2010) and substitute our original price series with the composite refiner's acquisition costs (RAC) provided by the EIA. Indeed, the relevant parameter raises in magnitude to $\hat{\psi}_{23} = -0.011$ with a p-value of 0.006. Nonetheless, with $\hat{\psi}_{43} = -0.023$, the depressing effect on the U.S. business cycle is still considerably larger.

Next, we substitute our measure of the global business cycle with two alternative proxies. The first one is Kilian's (2009) global activity index, for which we obtain the data from the author's homepage.³⁴ Second, we construct another proxy for the global output gap by detrending Lippi and Nobili's (2012) Rest of the World (RoW) output measure, which allows for a clear differentiation between demand shocks originating in the U.S. and elsewhere.³⁵ Interestingly, in both cases, the negative signs of the coefficients for business cycle and oil price uncertainty, $\hat{\psi}_{22}$ and $\hat{\psi}_{23}$, disappear.³⁶ However, the estimates are not significant at any of the common levels and the remaining free parameters in Ψ are only slightly affected by the changes of the series for y_2 .³⁷ The same largely applies to the coefficients, when we follow Ravn and Uhlig (2002), and use a HP filter with $\lambda = 129600$ in the derivation of the business cycle proxies from the original industrial production data.

Tri- and Bivariate Models

For the last category, we replicate the lower-dimensional models of previous studies with our particular specification and data. Following Jo (2014), we first estimate two different trivariate models by removing the variable for global and U.S. economic activity, respectively. As a consequence, the direct influence of oil price uncertainty on the remaining business cycle variable is slightly higher when the U.S. series is excluded from the VAR, and lower when the global series is omitted. Nonetheless, we can see from the table that the influence on the U.S. economy is again negative, highly significant and much stronger than the effect on the global aggregate. Finally, to mimic the bivariate models employed by Elder and Serletis (2010), Rahman and Serletis (2011, 2012) and others, we further reduce the dimension of each model by excluding the global supply series. Specifically for the U.S., the estimated influence of oil price uncertainty is once again weaker than before. We attribute this to a loss in precision of the generated uncertainty proxy when not accounting for the influence of the global supply and demand variables. Other than that, however, the conclusions derived from the baseline model remain valid.

³³ Until the end of price regulation in 1981, for instance, the U.S. went through four different phases of price controls. Also, in 1975 the Energy Policy and Conservation Act formally introduced an export ban on crude oil which still remains active today. For further details see Bordoff and Houser (2015).

³⁴ <http://www-personal.umich.edu/~lkilian/>

³⁵ The data for the RoW measure was obtained from the online supplement of the paper available on the homepage of the Journal of the European Economic Association (<http://onlinelibrary.wiley.com/doi/10.1111/j.1542-4774.2012.01079.x/supinfo>).

³⁶ Note that, due to computational difficulties still being prevalent in both cases, we had to reduce the number of lags in the model to $m = 3$.

³⁷ When we conduct calculations similar to the ones that we used to produce Figures 10 and 11, we find that the same is true with respect to the dynamic effects of the uncertainty proxies.

In summary, having conducted a battery of robustness checks that involved several modifications of our model's specification and underlying data, we are able to confirm the results of the previous section. In all cases, we find a negative influence of oil price volatility on the U.S. business cycle that is statistically significant at the 5% level. With respect to the influence on the global business cycle, the results are more ambiguous. While the relevant coefficient is negative, but insignificant in the baseline estimation, it turns out to be significant for a non-negligible number of the subsequent robustness checks. Overall, we view our empirical evidence in the way that we cautiously confirm Jo's (2014) finding of a negative impact of oil price uncertainty on global economic activity – specifically because $\hat{\psi}_{23}$ is typically negative. Under all circumstances examined, however, this influence remains considerably weaker than the one measured for the U.S., suggesting that the latter is more susceptible to oil price uncertainty than the global average. Likewise, in this section we obtained strong additional evidence in favour of our hypothesis that oil price volatility raises precautionary demand and the spot price. Looking at the relevant column in Table 3, it becomes immediately evident that the parameter $\hat{\psi}_{33}$ is highly significant and positive in all but one of the estimations. Finally, note also that the estimated influences of the business cycle volatilities, $\hat{\psi}_{22}$ and $\hat{\psi}_{44}$, remain consistent with the theory of real options throughout most of our robustness checks.

6. Conclusion & Outlook

This study utilises a VAR, GARCH-in-mean, asymmetric BEKK model to re-examine the influence of oil price uncertainty on real economic activity in the U.S. and globally, while accounting for changes in world demand and supply of crude oil. Hence, in contrast to other papers in this area, we employ a four-variable model that includes the rate of change of the real oil price, a proxy for U.S. activity, as well as measures of global oil production and global real economic activity. We add a number of non-standard components to increase the model's explanatory power and the precision of our uncertainty proxy. The model is estimated with monthly data from 1975:08 to 2014:03 and economic activity is measured in terms of business cycle fluctuations, instead of economic growth as in previous studies.

Our results show that oil price volatility is endogenous with respect to the other variables in the model. In addition to its inherent dynamics, oil price uncertainty is caused by unexpected shocks and volatility in the global supply and demand series, as well as in the U.S. business cycle. Significant asymmetries in the variance-covariance process suggest that these linkages are stronger when the respective shocks are regarded as bad news. Based on this evidence, we conclude that oil price volatility, to a certain extent, is also a gauge and channel of transmission of more general macroeconomic uncertainty and shocks.

This might also be one explanation as to why we still find a highly significant negative influence of oil price uncertainty on the U.S. business cycle – even when accounting for the influence of changes in global supply and demand for crude oil, and for any direct effects of business cycle volatility. Our main result is confirmed by an extensive number of robustness checks. Furthermore, the estimates imply that the effect is economically important as well, particularly during the mid-1980s, early-1990s and during the crisis of 2008/09. Before 1985, on the other hand, oil price volatility is systematically lower than during the rest of the sample. Consequently, we find no evidence that oil price uncertainty might have contributed to the 1980 or 1982 recessions. Even though there is a volatility surge in 1979, our calculations document clearly that neither the timing, nor the magnitude of this increase would fit such a narrative. Similarly, the regressions consistently show that the impact on the global business cycle is much less pronounced. In our baseline model, for instance, the relevant coefficient is insignificant and less than a fifth in magnitude than its counterpart in the U.S. equation. Finally, we present evidence that oil price uncertainty has a highly significant impact on the real oil price itself, most likely reflecting higher demand in times of uncertainty due to precautionary motives.

To conclude, our study provides strong evidence in favour of the theory of real options and the role of uncertainty in business cycle fluctuations. The results are also largely consistent with that of earlier studies on the relationship between oil price uncertainty and economic growth.

An interesting question for further research will be to find explanations for why the effect on the U.S. is comparably strong. A good starting point would be to estimate our model for countries with different economic structures. Also, given a proper identification scheme for structural shocks, our framework could be used to analyse the relative importance of first- and second-moment shock transmissions in more detail. Finally, future research might try to integrate other relevant factors into our model to get further insight about what oil price uncertainty actually represents and to control for additional sources of uncertainty. Some promising ground for this has already been laid out by Nazlioglu et al. (2015), and Antonakakis et al. (2014), for instance, who find significant relationships between oil price uncertainty and financial uncertainty, and oil price shocks and policy uncertainty, respectively.

Appendix I: Estimation Approach

As in Fiorentini et al. (2003), our empirical approach is based on the assumption that the vector of endogenous variables \mathbf{y}_t is generated by the following stochastic process:

$$\mathbf{y}_t = \mathbf{M}_t(\boldsymbol{\theta}_0) + \mathbf{H}_t^{1/2}(\boldsymbol{\theta}_0)\boldsymbol{\varepsilon}_t, \quad (\text{A1})$$

where the conditional mean vector $\mathbf{M}_t(\boldsymbol{\theta}_0)$ and the conditional covariance matrix $\mathbf{H}_t(\boldsymbol{\theta}_0)$ are dependent on the vector of true parameter values $\boldsymbol{\theta}_0$ and the information set given at time $t - 1$, Ω_{t-1} . We further assume that the innovations $\boldsymbol{\varepsilon}_t$ are identically, independently distributed and follow a standardised multivariate t -distribution, i.e., $\boldsymbol{\varepsilon}_t | \Omega_{t-1} \sim t(\mathbf{0}, \mathbf{I}_N, \nu_0)$.

In order to estimate the true parameters $\boldsymbol{\theta}_0$ and ν_0 , the log-likelihood function for the sample of size T , $L_T(\boldsymbol{\Phi})$, is maximised with respect to the parameter set $\boldsymbol{\Phi} = (\boldsymbol{\theta}', \nu)'$. More specifically, $L_T(\boldsymbol{\Phi})$ can be written as

$$L_T(\boldsymbol{\Phi}) = \sum_{t=1}^T \log f(\mathbf{y}_t | \boldsymbol{\Phi}, \Omega_{t-1}) \quad (\text{A2})$$

with

$$f(\mathbf{y}_t | \boldsymbol{\Phi}, \Omega_{t-1}) = \frac{\Gamma\left(\frac{\nu + N}{2}\right)}{\Gamma\left(\frac{\nu}{2}\right) [\pi(\nu - 2)]^{\frac{N}{2}} |\mathbf{H}_t(\boldsymbol{\theta})|^{\frac{1}{2}}} \left[1 + \frac{\boldsymbol{\varepsilon}_t'(\boldsymbol{\theta}) \boldsymbol{\varepsilon}_t(\boldsymbol{\theta})}{\nu - 2} \right]^{-\frac{\nu + N}{2}}, \quad (\text{A3})$$

where $\Gamma(\cdot)$ is Euler's Gamma function.³⁸ For simplicity, (A2) ignores some initial conditions for the mean and variance of the model. Due to our model's structure, however, in practice $L_T(\boldsymbol{\Phi})$ is only defined recursively and we have to choose some values of \mathbf{M}_t and \mathbf{H}_t to initialise the likelihood function. For the mean model we condition on the presample values of \mathbf{y}_t . Regarding the initial value of \mathbf{H}_t , we decided to use the residual covariance matrix from the VAR model, which we calculated for the sample period until the presumed date of the variance shift.³⁹

To solve the numerical optimisation problem, we make use of the algorithms implemented in the WinRATS 9.0 software package distributed by Estima. Specifically, we employ the method of Broyden, Fletcher, Goldfarb and Shanno (BFGS) as the main optimisation algorithm. Since BFGS is rather sensitive to the initial guess values for the parameter vector $\boldsymbol{\Phi}$, however, we use the simplex algorithm as a preliminary, derivative-free method to improve the starting values before they are handed over to BFGS.⁴⁰

As previous studies like Elder and Serletis (2010), Rahman and Serletis (2011), and Pinno and Serletis (2013) have noted, the estimation of multivariate VAR GARCH-in-mean models involves a large number of complex calculations and, hence, is prone to computational difficulties. This is not only because of the significant number of parameters that have to be estimated simultaneously, but also due to the

³⁸ Cf. Fiorentini et al. (2003) and Bauwens et al. (2006).

³⁹ The common choice would be to use an unconditional expectation of the matrix derived from an earlier estimation – see, for example, Elder and Serletis (2010, 2011), and Bredin et al. (2011). In our case, however, due to the asymmetries and the suspected structural shift in the covariance structure, this is not feasible. We did, however, test a few alternative values with the baseline model and found that – except for a slight worsening of the statistical properties of the volatility series – our estimation results remained largely unchanged.

⁴⁰ Cf. Estima (2014), p. 120.

high degree of non-linearity within these models.⁴¹ We have to acknowledge that, in this regard, our framework constitutes no exception. To cope with the aforementioned problems, we, therefore, decided to estimate our model with a large number of different starting values for Φ . In general, the guess values are generated randomly, but to a certain degree they also improve systematically. To be more precise, we proceed as follows:

Step 1. First, we estimate the model with the default starting values determined by RATS. We improve on these guess values with the simplex algorithm, before continuing the optimisation with the BFGS. The number of iterations for the simplex ρ is chosen randomly between 1 and 15.⁴²

Step 2. We draw new initial guess values Φ^I from a multivariate t-distribution with ten degrees of freedom, whereby the expected value of the distribution Φ^E and its covariance matrix are derived from the results in *Step 1*.⁴³ The estimation is then started again with Φ^I and a newly randomized ρ .

Step 3. *Step 2* is repeated $i = 15$ times. Whenever the random starting values lead to a higher log-likelihood, Φ^E is updated to the corresponding parameter set.⁴⁴ Also, when the algorithm achieves convergence to a maximum, we save Φ_{max}^I and ρ_{max} for future replication. Both variables are overwritten in case the BFGS later indicates convergence at an even higher likelihood.

Step 4. *Steps 1-3* are repeated $j = 50$ times. We thus run the optimization algorithm for a total of 800 times. After these repetitions, Φ_{max}^I and ρ_{max} are used to obtain the final result.

⁴¹ This concerns not only the link between the mean and variance equations, which is created by the inclusion of the GARCH-in-mean terms, but also other parts of the model. The BEKK coefficients, for instance, are not uniquely identified with respect to their signs because they enter the respective covariance equations only in squared or multiplicative form – see Engle and Kroner (1995).

⁴² This ensures that the starting values for the main algorithm are randomized from the beginning and, hence, that the subsequent steps are not always conditioned on the same set of initial results. The preiterations interval is based on the recommendations by Estima (2014, p. 120).

⁴³ Note that we do not impose any restrictions on the model's coefficients or their initial guess values. We do verify, however, that no starting value for v is below four. In case any v^I does not match this criterion, it is reset to the RATS default value of 20. This ensures that the first four moments of the distribution exist and, hence, that the likelihood function is defined – cf. Harvey et al. (1992). Note further that utilizing the coefficient covariance matrix from *Step 1* is actually incorrect from a technical point of view, if the optimization algorithm does not achieve convergence. In practice, however, we find that in most of these cases the matrix is still relatively close to its counterparts in case of convergence and, hence, sufficiently accurate to be used for the random draws.

⁴⁴ In this way, the initial values may improve in course of the i repetitions, as the distribution's centre is shifted towards a higher likelihood. However, to avoid that the estimation gets stuck in a certain direction or drifts off into extreme or even non-computable territory indefinitely, it is necessary to reset the procedure after a certain number of repetitions.

Table A1: Log-likelihoods for the results discussed in the paper

	Highest log-likelihood converged	Highest log-likelihood computed	Lowest log-likelihood computed
<i>A. Main Analysis</i>			
Baseline model ($m = 6$)	-2670.52819	-2670.19556	-2692.67479
<i>B. Robustness Checks</i>			
Baseline model ($m = 3$)	-2732.56822	-2732.56822	-2745.28394
Baseline model ($m = 9$)	-2620.25289	-2619.30136	-2630.40378
Variance shift in 1984:01	-2676.11412	-2676.10200	-2692.67408
Variance shift in 1993:01	-2676.60321	-2676.59947	-2696.90708
Reduced sample excl. the financial crisis	-2228.00594	-2228.00420	-2251.15956
y_3 : Nominal WTI	-2670.93469	-2670.93451	-2692.34105
y_3 : U.S. RAC	-2530.28141	-2530.27661	-2545.17745
y_2 : proxy based on Lippi and Nobili's RoW output ($m = 3$)	-3095.11432	-3094.77970	-3112.68134
y_2 : Kilian's GAI ($m = 3$)	-3941.14407	-3941.02052	-3950.27121
y_2, y_4 : HP-filtered with $\lambda = 129600$	-2694.90155	-2694.90105	-2716.84155
Trivariate model global $\{y_1, y_2, y_3\}$	-2456.75789	-2456.75789	-2457.26571
Trivariate model U.S. $\{y_1, y_3, y_4\}$	-2508.16255	-2508.16255	-2514.49665
Bivariate model global $\{y_2, y_3\}$	-1765.17273	-1765.17273	-1765.40179
Bivariate model U.S. $\{y_3, y_4\}$	-1821.47881	-1821.47881	-1825.79338

Appendix II: Additional Results

Table A2: Model selection with alternative lag orders

Model	Components	AIC	$Q_{MV}(12)$	$Q_{MV}(24)$	$Q_{MV}^2(12)$	$Q_{MV}^2(24)$
(1)	Constants	8104	3166.592 (0.000)	4856.786 (0.000)	2460.835 (0.000)	3032.769 (0.000)
(2)	+ VAR(6)	6283	141.047 (0.998)	375.518 (0.612)	521.391 (0.000)	722.931 (0.000)
(3)	+ Diagonal BEKK	5897	173.256 (0.830)	424.163 (0.077)	279.439 (0.000)	415.989 (0.126)
(4)	+ t-distribution	5791	173.801 (0.823)	416.128 (0.125)	331.895 (0.000)	471.796 (0.002)
(5)	+ asymmetries, +variance shift	5715	173.150 (0.832)	419.343 (0.104)	231.199 (0.028)	413.097 (0.147)
(6)	+ spillover effects, + GARCH-in-mean terms (full model)	5691	179.767 (0.727)	412.117 (0.155)	153.636 (0.981)	360.615 (0.799)

Notes: *p*-values in parentheses. Test statistics significant at the 5% level are printed in bold.

References

- Alquist, R. and Kilian, L. (2010). What Do We Learn from the Price of Crude Oil Futures?. *Journal of Applied Econometrics*, Vol. 25, pp. 539-573.
- Antonakakis, N., Chatziantoniou, I. and Filis, G. (2014). Dynamic spillovers of oil price shocks and economic policy uncertainty. *Energy Economics*, Vol. 44, pp. 433-447.
- Barsky, R. and Kilian, L. (2002). Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative. *NBER Macroeconomics Annual 2001*, Vol. 16, pp. 137-183.
- Barsky, R. and Kilian, L. (2004). Oil and the Macroeconomy Since the 1970s. *Journal of Economic Perspectives*, Vol. 18, No. 4, pp. 115-134.
- Baumeister, C. and Peersman, G. (2013). The Role of Time-Varying Price Elasticities in Accounting for Volatility Changes in the Crude Oil Market. *Journal of Applied Econometrics*, Vol. 28, pp. 1087-1109.
- Bauwens, L., Laurent, S. and Rombouts, J. (2006). Multivariate GARCH Models: A Survey. *Journal of Applied Econometrics*, Vol. 21, pp. 79-109.
- Bernanke, B. (1983). Irreversibility, Uncertainty, and Cyclical Investment. *Quarterly Journal of Economics*, Vol. 98, pp. 85-106.
- Bloom, N. (2009). The Impact of Uncertainty Shocks. *Econometrica*, Vol. 77, No. 3, pp. 623-685.
- Bordoff, J. and Houser, T. (2015). Navigating the U.S. Oil Export Debate. *Columbia SIPA, Center on Global Energy Policy*.
- Bredin, D., Elder, J. and Fountas, S. (2011). Oil Volatility and the Option Value of Waiting: An Analysis of the G7. *The Journal of Futures Markets*, Vol. 31, No. 7, pp. 679-702.
- Carstensen, K., Elstner, S. and Paula, G. (2013). How Much Did Oil Market Developments Contribute to the 2009 Recession in Germany?. *The Scandinavian Journal of Economics*, Vol. 115, No. 3, pp. 695-721.
- Caporale, G., Menla Ali, F. and Spagnolo, N. (2014). Oil price uncertainty and sectoral stock returns in China: A time-varying approach. *China Economic Review* (in press).
- Dickey, D. and Fuller, W. (1981). Likelihood ratio statistics for autoregressive time series with a unit root. *Econometrica*, Vol. 49, No. 4, pp. 1057-1072.
- Ding, Z. and Engle, R. (2001). Large Scale Conditional Covariance Matrix Modeling, Estimation and Testing. *Academia Economic Papers*, Vol. 29, No. 2, pp. 157-184.
- Edelstein, P. and Kilian, L. (2009). How Sensitive Are Consumer Expenditures to Retail Energy Prices?. *Journal of Monetary Economics*, Vol. 56, pp. 766-779.
- Elder, J. and Serletis, A. (2010). Oil Price Uncertainty. *Journal of Money, Credit and Banking*, Vol. 42, pp. 1137-1159.
- Elder, J. and Serletis, A. (2011). Volatility in Oil Prices and Manufacturing Activity: An Investigation of Real Options. *Macroeconomic Dynamics*, Vol. 15, pp. 379-395.
- Engle, R. and Kroner, K. (1995). Multivariate Simultaneous Generalized ARCH. *Econometric Theory*, Vol. 11, pp. 122-150.

Estima (2014). RATS Version 9.0 User's Guide. Evanston, IL, US.

Ferderer, J. (1996). Oil price Volatility and the Macroeconomy. *Journal of Macroeconomics*, Vol. 18, pp. 1-26.

Fiorentini, G., Sentana, E. and Calzolari, G. (2003). Maximum Likelihood Estimation with Inference in Multivariate Conditionally Heteroscedastic Dynamic Regression Models with Student *t* Innovations. *Journal of Business & Economic Statistics*, Vol. 21, No. 4, pp. 532-546.

Fukunaga, I., Hirakata, N. and Sudo, N. (2011). The Effects of Oil Price Changes on the Industry-Level Production and Prices in the United States and Japan. *Commodity Prices and Markets, East Asia Seminar on Economics*, Vol. 20, pp. 195-231.

Grier, K., Henry, O., Olekalns, N. and Shields, K. (2004). The Asymmetric Effects of Uncertainty on Inflation and Output Growth. *Journal of Applied Econometrics*, Vol. 19, pp. 551-565.

Guo, H. and Kliesen, K. (2005). Oil Price Volatility and U.S. Macroeconomic Activity. *Federal Reserve Bank of St. Louis Review*, Vol. 86, No. 6, pp. 669-683.

Hamilton, J. (1983). Oil and the Macroeconomy since World War II. *Journal of Political Economy*, Vol. 91, No. 2, pp. 228-248.

Harvey, A., Ruiz, E. and Sentana, E. (1992). Unobserved component time series models with ARCH disturbances. *Journal of Econometrics*, Vol. 52, pp. 129-157.

Hosking, J. (1981). Equivalent Forms of the Multivariate Portmanteau Statistic. *Journal of the Royal Statistical Society, Series B (Methodological)*, Vol. 43, No. 2, pp. 261-262.

Jarque, C. and Bera, A. (1980). Efficient tests for normality, homoscedasticity, and serial independence of regression residuals. *Economic Letters*, Vol. 6, pp. 255-259.

Jo, S. (2014). The Effects of Oil Price Uncertainty on Global Real Economic Activity. *Journal of Money, Credit and Banking*, Vol. 46, No. 6, pp. 1113-1135.

Kilian, L. (2009). Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market. *American Economic Review*, Vol. 99, No. 3, pp. 1053-1069.

Kilian, L. (2014). Oil Price Shocks: Causes and Consequences. *Annual Review of Resource Economics*, Vol. 6, pp. 133-154.

Kilian, L. and Murphy, D. (2014). The Role of Inventories and Speculative Trading in the Global Market for Crude Oil. *Journal of Applied Econometrics*, Vol. 29, pp. 454-478.

Kilian, L. and Park, C. (2009). The Impact of Oil Price Shocks on the U.S. Stock Market. *International Economic Review*, Vol. 50, No. 4, pp. 1267-1287.

Kwiatkowski, D., Philips, P., Schmidt, P. and Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root. How sure are we that economic time series have a unit root?. *Journal of Econometrics*, Vol. 54, pp. 159-178.

Lippi, F. and Nobili, A. (2012). Oil and the Macroeconomy: A Quantitative Structural Analysis. *Journal of the European Economic Association*, Vol. 10, No. 5, pp. 1059-1083.

Nazlioglu, S., Soytas, U. and Gupta, R. (2015). Oil Prices and financial stress: A volatility spillover analysis. *Energy Policy* (in press).

Pagan, A. (1984). Econometric Issues in the Analysis of Regressions with Generated Regressors. *International Economic Review*, Vol. 25, No. 1, pp. 221-247.

Pindyck, R. (1991). Irreversibility, Uncertainty, and Investment. *Journal of Economic Literature*, Vol. 29, No. 3, pp. 1110-1148.

Pinno, K. and Serletis, A. (2013). Oil Price Uncertainty and Industrial Production. *The Energy Journal*, Vol. 34, No. 3, pp. 191-216.

Rahman, S. and Serletis, A. (2011). The Asymmetric Effects of Oil Price Shocks. *Macroeconomic Dynamics*, Vol. 15, Suppl. 3, pp. 437-471.

Rahman, S. and Serletis, A. (2012). Oil Price Uncertainty and the Canadian Economy: Evidence from a VARMA, GARCH-in-Mean, asymmetric BEKK model. *Energy Economics*, Vol. 34, pp. 603-610.

Ravn, M. and Uhlig, H. (2002). On Adjusting the Hodrick-Prescott Filter for the Frequency of Observations. *The Review of Economics and Statistics*, Vol. 84, No. 2, pp. 371-380.

Sims, C. (1980). Macroeconomics and Reality. *Econometrica*, Vol. 48, No. 1, pp. 1-48.

Stock, J. and Watson, M. (2003). Has the Business Cycle Changed and Why?. *NBER Macroeconomics Annual 2002*, Vol. 17, Gertler, M. and Rogoff, K. (eds.), MIT Press.

Summers, P. (2005). What Caused The Great Moderation? Some Cross-Country Evidence. *Kansas City FED Economic Review*, Third Quarter 2005, pp. 5-32.