CENTRE FOR APPLIED MACRO - AND PETROLEUM ECONOMICS (CAMP)

CAMP Working Paper Series
No 2/2012

What drives oil prices?
Emerging versus developed economies.

Knut Are Aastveit, Hilde C. Bjørnland
and Leif Anders Thorsrud

NORWEGIAN
BUSINESS SCHOOL

What drives oil prices? Emerging versus developed economies*

Knut Are Aastveit† Hilde C. Bjørnland‡ Leif Anders Thorsrud§

May 28, 2014

Abstract

We explore the role of demand from emerging and developed economies as drivers of the real price of oil. Using a FAVAR model that identifies shocks from different regions of the world, we find that demand from emerging economies (most notably from Asian countries) is more than twice as important as demand from developed countries in accounting for the fluctuations in the real oil price and in oil production. Furthermore, geographical regions respond differently to adverse oil market shocks that drive up oil prices, with Europe and North America being more negatively affected than countries in Asia and South America.

JEL-codes: C32, E32, F41

Keywords: Oil prices, emerging and developed countries, demand and supply shocks, factor augmented vector autoregressions

*We thank the editor, Fabio Canova, three anonymous referees, Christiane Baumeister, Sandra Eickmeier, James D. Hamilton, Ana Maria Herrera, Lutz Kilian, Francesco Ravazzolo and Benjamin Wong as well as seminar and conference participants at Deutsche Bundesbank, Ghent University, International Monetary Fund, University of Oslo, the American Economic Association 2013 Meeting in San Diego and Norges Bank's conference on 'Modeling and Forecasting Oil Prices' for valuable comments. This paper is part of the research activities at the Centre for Applied Macro and Petroleum economics (CAMP) at the Norwegian Business School. The usual disclaimer applies. The views expressed in this paper are those of the authors and do not necessarily reflect the views of Norges Bank.

†Norges Bank. Email: `Knut-Are.Aastveit@norges-bank.no`

‡BI Norwegian Business School. Email: `hilde.c.bjornland@bi.no`

§BI Norwegian Business School. Email: `leif.a.thorsrud@bi.no`

1 Introduction

Since the seminal work by Hamilton (1983), a large body of literature has suggested that there is a significant negative relationship between oil price increases and economic activity in a number of different countries (see, e.g., Burbidge and Harrison (1984), Gisser and Goodwin (1986), Bjørnland (2000) and Hamilton (1996, 2003, 2009) among many others). Higher energy prices typically lead to an increase in production costs and inflation, thereby reducing overall demand, output and trade in the economy.

Most of the research on oil price changes has focused on quantifying the effects on developed countries. During the last two decades, the global economic landscape has shifted dramatically. Emerging market economies have experienced rapid growth in economic activity and international trade, outperforming most developed countries across the world. At the same time, the real oil price has more than doubled, without any apparent severe negative effects on the global economy. This has led to an intense research effort to quantify the importance of global demand for changes in the real price of oil, see e.g., Kilian (2009).

This paper breaks ground in this area by explicitly examining the importance of demand from emerging and developed economies as drivers of the real price of oil. Our main focus is to determine whether the increased demand for oil originates from emerging economies, which are growing at a pace twice that of the developed economies, or from the developed world, which historically has been the main consumer of oil. Having established where demand originates, we analyze how different geographical regions respond to the oil market shocks that drive up oil prices.

To explore these questions, we estimate a factor augmented vector autoregressive (FAVAR) model that includes separate activity factors for emerging and developed economies in addition to observable measures of global oil production and the real price of oil. The advantage of this modeling strategy is that we can preserve the parsimonious data representation offered by factor modeling techniques, while also including a large cross-section of countries and variables in a single model. Traditionally, empirical studies investigating the interaction between oil prices and the macro economy have employed one or many small-scale vector autoregressions (VAR), typically only including a single country (e.g.

Hamilton (1983)) or one index of aggregate global activity (e.g., Kilian (2009)) in each model.

To differentiate between oil supply and demand shocks we build on Kilian (2009), identifying these shocks using recursive zero restrictions in a structural VAR model. Novel to our approach is the separation of demand shocks in emerging and developed countries, identified using a mixture of sign and zero restrictions.

Our paper contributes to the large and growing empirical literature investigating the role of demand for oil prices. Barsky and Kilian (2002) first suggested that one should allow for a "reverse causality" from macroeconomic variables to oil prices. Subsequently, Kilian (2009), Lippi and Nobili (2012), Peersman and Van Robays (2012) and Aastveit (2013), among others, have shown that if the increase in the oil price is driven by an increased demand for oil associated with fluctuations in global activity and not disruptions of supply capacity, economic activity may even be positively affected, at least in the short-run.

We extend the literature in three ways. First, this is, to the best of our knowledge, the first paper to explicitly analyze the contribution of demand from developed and emerging countries on the real price of oil. While it is commonly believed that growth in emerging markets (in Asia in particular) is the main driver of the increased demand for oil (see, e.g., the discussion in Kilian (2009), Baumeister and Peersman (2013b) and Hicks and Kilian (2013)), no studies have analyzed this question explicitly using a structural model. Therefore, very little is known about the effect that increased growth in emerging economies has had on the real price of oil and, equally important, how economic activity in different regions of the world is affected by the oil market. Second, the identification strategy adopted to isolate the various demand shocks is novel in the oil literature, as is the use of the FAVAR model for this purpose. Finally, given the large number of countries included in the analysis, this is also the most comprehensive analysis to date of the relationship between oil prices and macroeconomic activity.

We have three main findings, which are robust to numerous robustness checks. First, we show that demand shocks in emerging and developed economies together account for 50-60 percent of the fluctuations in the real price of oil over the last two decades.

Second, demand shocks in emerging markets, particularly in Asia, are more than twice as important as demand shocks in developed economies in explaining fluctuations in the real price of oil and global oil production. This is a new finding in the literature.

Third, countries respond differently to the adverse oil market shocks that drive up oil prices. In particular, while economic activity in Europe and the US declines substantially following oil market shocks, economic activity in emerging markets in Asia and South America declines by a substantially smaller amount and in some cases actually temporarily increases. We find that country characteristics, such as high investment as a share of GDP and a high degree of openness, are important in explaining the heterogeneity observed. Hence, we show that it is imperative not only to identify where demand is coming from, but also to allow countries to respond differently to various oil market shocks.

The remainder of the paper is structured as follows: Section 2 describes the model, the identification strategy and the estimation procedure. We report the results of the study and of our robustness checks in Section 3, while Section 4 concludes.

2 The Factor Augmented VAR model

We specify a model that includes separate measures for activity in developed and emerging economies, in addition to global oil production and the real price of oil. The activity measures are intended to capture the respective shifts in the demand for oil in developed and emerging markets and are constructed by applying factor modeling techniques. More precisely, our full model is a FAVAR that builds on the general setup of Bernanke et al. (2005) and Boivin et al. (2009).

It is instructive to represent the model in a state space form. The transition equation is specified as:

$$F_t = \beta(L)F_{t-1} + u_t, \tag{1}$$

where $F_t = \begin{bmatrix} \Delta prod_t, & devAct_t, & emeAct_t, & \Delta rpo_t \end{bmatrix}'$ is a (4×1) vector containing the first differences of the logarithm of global oil production, an unobserved developed economy activity factor, an unobserved emerging economy activity factor and the first difference of the logarithm of the real price of oil[1], respectively. $\beta(L)$ is a conformable lag polynomial of order p, and u_t is a 4×1 vector of reduced form residuals. The structural disturbances

[1]Since various stationarity tests indicate that the real price of oil is non-stationary in our sample, we include the first difference of the logarithm of the real price of oil in our model.

follow $u_t = \Omega^{1/2}\varepsilon_t$, with $\varepsilon \sim N(0,1)$ and $\Omega = A_0(A_0)'$, where Ω is the covariance of the reduced form residuals and A_0 is a (4×4) matrix describing how the structural shocks relate to the reduced form errors.

The observation equation of the system is:

$$X_t = \Lambda F_t + e_t, \tag{2}$$

where X_t is a $N \times 1$ vector of observable variables, Λ is a $N \times 4$ matrix of factor loadings, and e_t is a $N \times 1$ vector of idiosyncratic, zero mean, disturbances. Below we discuss details on data, estimation and identification of the FAVAR model.

2.1 Data

We use real GDP growth and industrial production growth as measures of economic activity for 33 different countries. In total, our sample countries account for approximately 80 percent of world GDP, measured by purchasing power parity[2]. All GDP series are at constant prices. The industrial production series are volume indexes, and refer, with few exceptions, to the manufacturing industry.

We determine a priori which countries should be considered developed and emerging economies. The 18 countries that are members of the OECD at the beginning of our sample are considered developed economies. The remaining 15 countries are considered emerging economies, see the on-line appendix for details on the data set.

To measure oil production and the real price of oil, we use world crude oil production, in millions of barrels per day, and US real refiner' acquisition cost for imported crude oil, respectively. The nominal oil price has been deflated using the US consumer price index. These are the same variables used in Kilian (2009) and many other papers.

2.2 Estimation and model specification

We estimate the reduced form model (equation (1) and (2)) in a two-step procedure: First, the unobserved activity factors for developed and emerging economies are estimated and identified using the principal components method and normalizing restrictions. Prior to

[2]Authors calculations based on 2009 estimates from the IMF.

estimating the factors, all variables are transformed to induce stationarity, by using the first difference of the logarithm of the respective variable, and standardized. The identified factors are then used as observed variables in a standard VAR framework. The lag length is set to 4, and the VAR residuals pass standard diagnostic tests.[3] In our baseline model, $N = 66$, and we estimate the model over a sample period from 1992:Q1 to 2009:Q4, yielding $T = 72$ observations.

To construct distributions for the impulse response functions, and accurately account for the problem of generated regressors in the second estimation step, we employ a residual bootstrap procedure for the entire system with 5000 replications, see Goncalves and Perron (2013).[4]

Identification in this model is affected by two issues: First, we need to estimate the unobserved factors such that the developed economy and emerging economy factors are identified, and, second, we need to identify the structural shocks.

2.3 Identifying the factors

For the exposition it is instructive to re-write equation (2) as:

$$X = \Lambda F' + e \tag{3}$$

where the dimensions of X, Λ, and F are $N \times T$, $N \times 4$, and $T \times 4$, respectively. As described above, two of the factors in our system are observable, $\Delta prod_t$ and Δrpo_t. Thus, we only need to estimate and identify the two unobserved activity factors, $devAct_t$ and $emeAct_t$. Accordingly, identification of the factors and loadings proceeds in two steps.

We start by partitioning X into a part consisting only of the activity variables: $X^{act} = \left[X^{US'}, \quad X^{China'}, \quad X^{dev'}, \quad X^{eme'} \right]'$, which is a $(N-2) \times T$ matrix where X^{dev} and X^{eme} are $(N^{dev}-1) \times T$ and $(N^{eme}-1) \times T$ matrices of developed and emerging activity variables, respectively, and X^{US} and X^{China} are real GDP growth in the US and China. The reason

[3] As shown in, e.g., Hamilton and Herrera (2004), an overly restrictive lag length can produce misleading results regarding the effects of oil market shocks on the macro economy, while increasing the lag length to over one year has negligible effects.

[4] The confidence bands for the impulse response functions are bias-adjusted using Hall's percentile intervals, see Hall (1992). In the on-line appendix we also show that our results are robust to estimating the model using likelihood-based Gibbs sampling techniques. However, the two-step procedure is simpler and less computationally intensive.

for this exact partition will become clear below. Then, we estimate the two first factors from X^{act} employing the principal components method, with the standard normalizing assumption that the factors are orthogonal.[5] We denote the estimates of the two factors as F^{act1} and F^{act2} and define the $T \times 4$ matrix $F^* = \begin{bmatrix} \Delta prod', & F^{act1}, & F^{act2}, & \Delta rpo' \end{bmatrix}$. We then regress F^* on X^{act} to get an estimate of the initial factor loadings, denoted Λ^{act} (which is a $(N-2) \times 4$ matrix).

In the second identification step, we rotate the space of the common components such that they can be given an economic interpretation. This can easily be done based on the above estimates. To comply with the ordering suggested for F_t (in equation (1)) we first order the variables in X such that $X = \begin{bmatrix} \Delta prod', & X^{US'}, & X^{China'}, & \Delta rpo', & X^{dev'}, & X^{eme'} \end{bmatrix}'$, and construct Λ^*:

$$\Lambda^* = \begin{bmatrix} 1 & 0 & 0 & 0 \\ \Lambda^{act}_{11} & \Lambda^{act}_{12} & \Lambda^{act}_{13} & \Lambda^{act}_{14} \\ \Lambda^{act}_{21} & \Lambda^{act}_{22} & \Lambda^{act}_{23} & \Lambda^{act}_{24} \\ 0 & 0 & 0 & 1 \\ \Lambda^{act}_{i1} & \Lambda^{act}_{i2} & \Lambda^{act}_{i3} & \Lambda^{act}_{i4} \\ \vdots & \vdots & \vdots & \vdots \\ \Lambda^{act}_{M1} & \Lambda^{act}_{M2} & \Lambda^{act}_{M3} & \Lambda^{act}_{M4} \end{bmatrix}$$

where $i = 3, \ldots, M$, $M = N - 2$ and Λ^{act}_{1j} and Λ^{act}_{2j} for $j = 1, \ldots, 4$ are factor loadings for US and Chinese GDP growth, respectively. Λ^* highlights that oil production and the real price of oil are observable factors that naturally have factor loadings equal to one on themselves. Finally, we define Λ_r as the 4×4 upper component of Λ^*, and, following the method proposed by Bai and Ng (2013), compute the identified factor loadings and factors as:

$$\Lambda = \Lambda^* \Lambda_r^{-1} \text{ and } F = F^* \Lambda_r' \tag{4}$$

Note, that the observable factors ($\Delta prod$ and Δrpo) are not rotated due to the zero restrictions in Λ^*. Furthermore, US and Chinese GDP growth have a loading equal to one on the $devAct_t$ and $emeAct_t$ factors, respectively and zero everywhere else, i.e., the upper 4×4 component of Λ in equation (4) is an identity matrix.

[5] The choice of two activity factors is motivated by our economic question. However, the different information criteria discussed in Bai and Ng (2002) also suggest that two factors are appropriate for our data set.

Importantly, this identification strategy imposes no restrictions on the correlations between the factors, but still yields $r^2 = 16$ restrictions, thus ensuring unique identification of the factors and loadings (see Bai and Ng (2013) for details). The factors are displayed in an on-line appendix, where we also show that our results are robust to alternative choices regarding how we estimate the factors.

The normalizing restrictions we apply do not imply that the identified activity factors are equal to US and Chinese GDP growth, nor that these variables are forced to be the main drivers of the two activity factors.[6] However, the US and China have by far the highest daily oil consumption rates among the countries in the developed and emerging data blocks, respectively. Thus, they naturally serve as normalizing variables in our global oil market model.

2.4 Identifying the shocks

To identify the structural shocks in the FAVAR model, we build on Kilian (2009), differentiating between oil supply and demand shocks in a structural VAR model. The novelty of our study is the identification of separate demand shocks in emerging and developed economies.

Accordingly, to identify the structural innovations in the model as oil supply shocks, developed-country oil demand shocks, emerging-country oil demand shocks and other oil-specific demand shocks, we employ a mixture of sign and zero restrictions. In particular, we restrict A_0, defined in section 2, as:

$$\begin{bmatrix} u^{prod} \\ u^{devAct} \\ u^{emeAct} \\ u^{rpo} \end{bmatrix} = \begin{bmatrix} x & 0 & 0 & 0 \\ x & + & + & 0 \\ x & + & + & 0 \\ x & x & x & x \end{bmatrix} \begin{bmatrix} \varepsilon^{\text{oil supply}} \\ \varepsilon^{\text{developed demand}} \\ \varepsilon^{\text{emerging demand}} \\ \varepsilon^{\text{oil-specific demand}} \end{bmatrix} \tag{5}$$

where $+$ indicates that the effect of the shock must be positive, x leaves the effect unrestricted, and, finally, zero imposes contemporaneous exclusion restrictions.

The identification strategy imposes the following restrictions. First, crude oil supply

[6] In the on-line appendix we show, by reporting the correlation between the factor estimates and the individual activity measures, that the identified factors are influenced by all the variables in the data set.

shocks ($\varepsilon_t^{oilsupply}$) are defined as unpredictable innovations to global oil production. The supply shocks are allowed to affect oil production, all activity measures and the price of oil within the quarter, while oil production itself responds to all shocks except the oil supply shock with a lag. This implies a vertical short-run supply curve for crude oil. As adjusting oil production is costly, and the state of the crude oil market is uncertain, these are plausible restrictions, (see e.g. discussion in Hamilton (2009) and Kilian (2009)). Baumeister and Peersman (2013a) also estimate the price elasticity of oil supply to be very small in our estimation period, consistent with the view that the short-run supply curve is nearly vertical.

Second, innovations to the activity factors for developed and emerging economies (that cannot be explained by global oil supply shocks) are referred to as, respectively, developed-country oil demand shocks ($\varepsilon^{\text{developed demand}}$ for short) and emerging-country oil demand shocks ($\varepsilon^{\text{emerging demand}}$ for short). The real price of oil, as well as the developed and the emerging activity factors, can be affected on impact by these demand shocks. Thus, we allow for a simultaneous reaction to demand shocks in emerging and developed countries. This is plausible given the relative sizes of the economies (or block of countries) and the potential interaction due to trade and financial integration.[7]

To fully identify the structural shocks, we impose two additional sign restrictions on the short-run impulse responses: To identify $\varepsilon^{\text{developed demand}}$, we impose that the response of $devAct - emeAct > 0$, and to identify $\varepsilon^{\text{emerging demand}}$ we impose that the response of $emeAct - devAct > 0$. The restriction implies that after an emerging-country demand shock (that increases activity in emerging countries), activity in developed countries is also restricted to increase, but by less than in the emerging countries (and vice versa for a developed-country demand shock).[8] This type of restriction allows us to identify demand from different groups of countries (regions), while also allowing for interaction between them. The sign restrictions are set to hold for 2 quarters only. Results are robust to increasing or decreasing the horizon by one quarter.[9]

[7] At the end of 2009, the emerging and developed economies in our sample accounted for approximately 32 and 47 percent of world GDP based on purchasing power parity respectively.

[8] Restrictions on one variable relative to another have been applied previously by, among others, Farrant and Peersman (2006) and Eickmeier and Ng (2011), but in a very different context.

[9] Note that the restriction does not prevent a shock in developed countries to have a substantial effect on emerging countries (and vice versa) over time. This is consistent with findings in Canova (2005) and Maćkowiak (2007), analyzing the impact of US monetary policy shocks on some emerging countries.

Finally, innovations to the real price of oil that cannot be explained by the three aforementioned shocks are referred to as oil-specific demand shocks ($\varepsilon^{\text{oil-specific demand}}$). Although these shocks may encompass other oil market specific shocks not explained in the model, Kilian (2009) argues that such shocks primarily capture precautionary demand for oil driven by the uncertain availability (scarcity) of future oil supply.

With minor modifications, the sign restrictions are implemented following the procedure outlined in Rubio-Ramirez et al. (2010) and Mumtaz and Surico (2009) and is explained in detail in the on-line appendix. As is now well known in the literature, the sign restrictions will not yield unique identification (see Fry and Pagan (2011)). Instead a large number of models are consistent with the restrictions. To circumvent this problem, we follow Fry and Pagan (2011) and adopt the following procedure: For each set of reduced form parameters, we draw 1000 accepted candidate impulse responses (based on A_0 above), and compute the mean impulse response function among these accepted draws. We then compute the mean squared error between all candidate functions and the mean impulse response function. The impulse response function with the lowest score is stored. Thus, for each set of parameter estimates, the identified structural shocks are orthogonal.

3 Results

In the following we investigate what drives the real price of oil and oil production, before we examine how the different regions and countries are affected by the adverse oil market shocks.

3.1 What drives oil prices?

Figure 1 reports the impulse responses of the model. Each row contains the responses of the level of a specific variable to the four different shocks. To compare the developed-country and emerging-country demand shocks, we normalize both shocks to increase activity in their respective regions by one percent. The oil supply shock is normalized to decrease oil production by one percent, while the oil-specific demand shock is normalized to increase

However, as also emphasized in Canova (2005), US real demand and supply shocks do not have substantial spillover to emerging countries, suggesting a delayed reaction, in line with what we have imposed here. We check explicitly the role of US monetary policy in the robustness section below.

Figure 1: **Impulse responses**

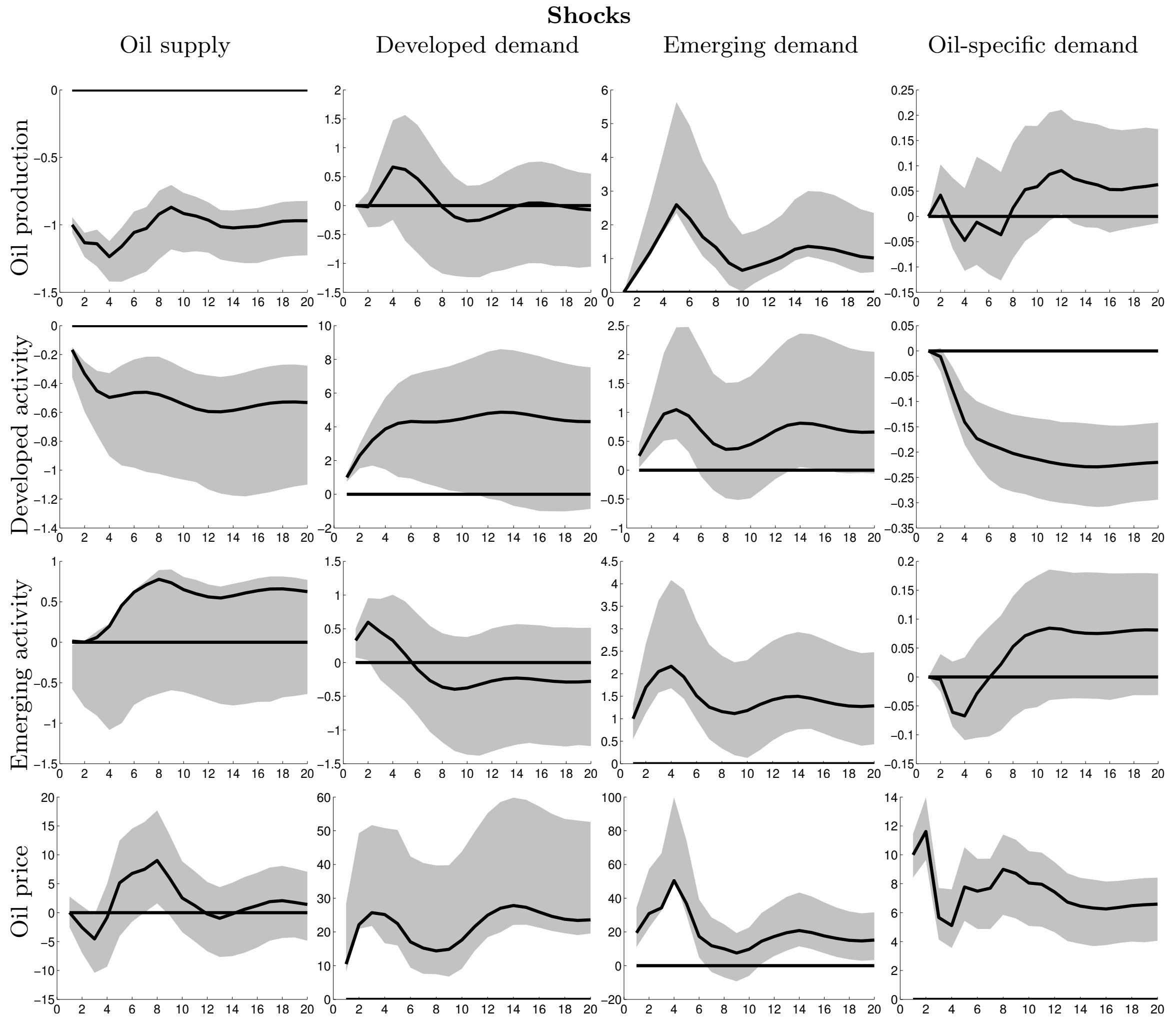

Note: The responses are displayed in levels of the variables. The developed-country and emerging-country demand shocks are normalized to increase activity in developed and emerging countries by one percent, respectively. To facilitate comparison with earlier studies, the oil supply shock is normalized to decrease oil production by one percent, while the oil-specific demand shock is normalized to increase the real oil price by 10 percent. The gray shaded areas represent 68 percent confidence bands (bootstrapped), while the black solid lines are the point estimates.

the real oil price by an initial 10 percent. While the normalization of the two demand shocks allows us to compare the contributions of developed and emerging countries, the normalization of the two “oil market” shocks is selected to facilitate comparison with previous studies.[10]

[10]The FAVAR is estimated on standardized data. To make the normalized shocks interpretable in terms of the scale of the observables (prior to standardization), the shocks are scaled with the standard deviation of the variables used to identify the factors. For our sample these standard deviations are 0.9, 16, 0.7, and 1.3 for oil production, the real price of oil, and GDP growth for the US and China respectively.

Beginning with our focal question, **what drives oil prices?**, we examine the bottom row in Figure 1. Demand shocks in both the developed and emerging countries increase the real oil price significantly for 1-2 years. However, the effect of a normalized emerging-country demand shock has a stronger impact on the real price of oil than a developed-country demand shock, increasing the oil price by an initial 20 percent versus 10 percent. Interestingly, a shock to the emerging activity factor also has the strongest effect on oil production (upper row), increasing the total supply of oil significantly for a year. The effect of a developed-country demand shock on oil production is insignificant.

The impulse responses in the lower left and right corners of Figure 1 report the oil price responses after an oil supply and oil-specific demand shock, respectively. A one percent disruption in oil production due to an oil supply shock eventually increases the oil price by 5-10 percent. The delayed response may reflect the fact that oil deliveries are based on future contracts. Hence, it might take some time before supply disruptions affect prices. The last shock, interpreted in Kilian (2009) as a precautionary oil-specific demand shock (ultimately driven by expectations about future oil supply shortfalls), triggers an immediate and sharp increase in the real price of oil (normalized to increase by 10 percent). Such expectations can change almost immediately in response to, e.g., exogenous political events, and therefore tend to increase prices without any subsequent effects on oil production.

As the effect on the real price of oil is stronger following an emerging-country demand shock compared to a developed-country demand shock, at least in the short-run, our results suggest that the emerging countries have higher income elasticity than the developed countries. Typically, as a country becomes more developed (richer), growth in petroleum use declines (as the country produces less manufacturing goods and more services), and hence income elasticity also declines. Consistent with this, Hamilton (2009) has suggested that while the income elasticity of oil consumption in the US has declined over time (to below 0.5), income elasticity in newly industrialized countries may be closer to unity.

However, the divergence in income elasticities between emerging and developed countries may still not be sufficient to explain the large effect from emerging-country demand shocks on oil prices. In particular, over the last two decades, oil consumption in developed countries has on average been twice as large as in emerging countries. For this to trigger a price effect such as the one documented here would require an even larger spread in income elasticity between the two groups of countries. However, as shown below, emerging-country demand shocks have been particularly important for driving up oil prices during the latter part of our sample. The discrepancy between the observed oil consumption rates, assumed income elasticities, and oil price effects can be explained by using the oil consumption shares for this latter period.[11] Moreover, heavy subsidizing of energy prices in many emerging countries makes the case for an even larger spread in income elasticities between emerging and developed economies. Finally, Figure 1 shows that the spillovers from an emerging-country demand shock to the developed-country activity factor are large and positive for a prolonged period. The opposite effects are small and insignificant after two quarters. Thus, emerging-economy demand shocks have contributed to increasing the real price of oil by also stimulating demand in the rest of the world.[12]

Turning to **the reverse causality**, we ask, what are the effects of the two oil market shocks on the macro economy? First, an adverse supply shock, normalized to decrease oil production by one percent, eventually increases the real oil price by 5-10 percent and reduces activity in developed economies by approximately 0.5 percent. The effect of an oil supply shock on emerging countries is positive, although the standard error bands are wide, suggesting an insignificant response in the emerging group as a whole.[13] The response reported for the developed country group resembles that found for aggregate activity in Kilian and Murphy (2013) and Kilian (2009) (using a global index as their measure of activity), although the impact is somewhat smaller, most likely as the responses in the emerging countries (which are less negative) are also contained in the aggregate index.

Second, a shock to oil-specific demand, normalized to initially increase the real price of oil by 10 percent, reduces GDP in the developed countries (0.2 percent), while in the emerging countries, the effect is not significantly different from zero. The responses docu-

[11] Rolling sub-sample regressions also indicate that the impact of emerging demand shocks relative to developed demand shocks on the real price of oil is stronger towards the second half of the sample. Results can be obtained on request.

[12] This is consistent with findings in Thorsrud (2013), who shows that during last two decades, Asia is the region with the strongest spillover effects to the rest of the world. The results presented in Kose and Prasad (2011), and the references therein, also document the growing importance of emerging countries to business cycles in the world.

[13] Note that following both an oil supply and an oil-specific demand shock, the uncertainty bands around the responses in the emerging activity factor are particularly large. This likely reflects that the emerging market economies are less homogeneous than the developed economies, as observed in the correlation numbers and the individual country impulse responses, reported in the on-line appendix.

mented here are different from those in Kilian (2009), where global real economic activity actually increases after an oil-specific demand shock. This result could reflect the increasing role of emerging countries in Kilian's global index, since, as we have documented here, emerging countries do not contract significantly following this shock. Thus, it is imperative not only to identify where demand is coming from, but also to allow regions/countries to respond differently to various oil market shocks.

Variance decompositions for the real price of oil and oil production, displayed in Panels (a) and (b) of Figure 2, respectively, allow us to compare the relative contributions of all shocks. The figure confirms the results found above. Shocks to the emerging activity factor are more important than the shocks to the developed activity factor in explaining fluctuations in the oil price and oil production. In fact, for 1-2 years, approximately 40 percent of the variation in the oil price is explained by emerging demand shocks, while developed demand shocks explain approximately 15 percent, (see Panel a). Turning to oil production, 40-50 percent of the variation is explained by emerging demand shocks, while less than 10 percent is due to developed demand shocks (see Panel b). We conclude that demand shocks from emerging countries has been more than twice as important as demand shocks from developed economies in explaining the variance in the oil price and up to five times more important in explaining the variance in oil production.

Turning to the two oil market shocks, Figure 2 shows that oil supply shocks explain a small share of the variation in the real price of oil, while 40-50 percent of the variance in oil production is explained by these shocks. This is consistent with findings in Baumeister and Peersman (2013b), who use a time-varying SVAR approach to demonstrate that oil supply shocks have become a less important source of oil price movements in recent years. Finally, the oil-specific demand shocks explain the remaining 40-50 percent of the variation in the real price of oil after 1-2 years, but a negligible share of oil production at all horizons.

The results in Figure 2 reflect the average contributions of the various shocks over the last two decades. To examine the different periods in greater detail, Figure 3 plots the accumulated contribution of each structural shock to the real price of oil based on a historical decomposition of the data. Panel (a) shows for each quarter from 1992 to 2009 the real price of oil in levels (solid line) and the real price minus the contribution from the emerging-country shocks, i.e., what the real price of oil would have been in the absence of

Figure 2: **Variance decomposition**

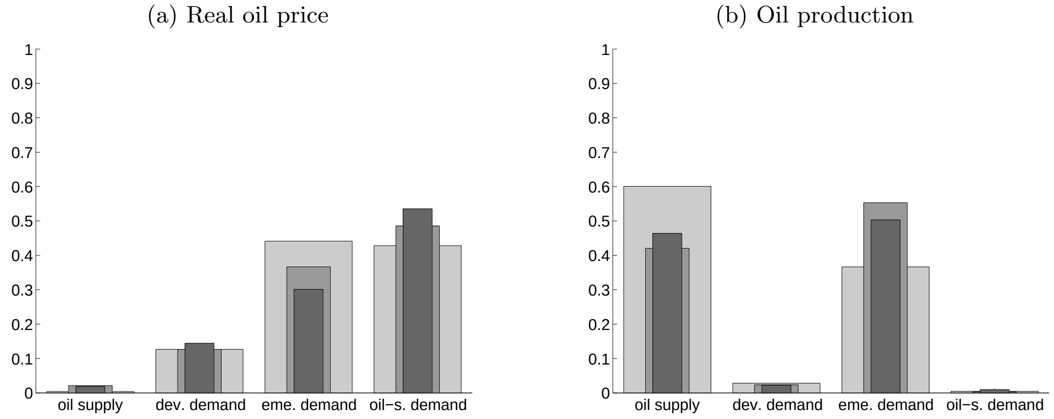

Standard errors

	Horizon	Oil supply	dev. demand	eme. demand	oil-s. demand
Real oil price	4	0.03	0.13	0.23	0.19
	12	0.08	0.16	0.22	0.22
Oil production	4	0.17	0.11	0.17	0.04
	12	0.19	0.12	0.22	0.10

Note: The bars display the variance decomposition with respect to the shocks for horizons 4, 8 and 12 quarters. The widest bars correspond to the shorter horizon. The standard errors of the variance decompositions are reported in the table above (for horizons 4 and 12)

emerging-country shocks (dotted line). In Panels (b), (c) and (d), the dotted line displays the real price of oil minus the respective contributions of demand from developed countries, oil supply shocks and oil-specific demand shocks.[14]

The figure emphasizes the importance of emerging markets as drivers of the real oil price. This was especially pronounced in the middle of the 1990s and from 2002/2003 and onwards (see Panel a). In fact, demand in emerging markets added approximately 20 dollars to the peak oil price (approximately 55 dollars in real terms) in 2008. Thus, in our sample, the strong positive contribution from emerging economies has been steadily increasing and was only interrupted by the East Asian crisis (1997/1998) and the broader global economic downturn around 2001.

However, demand in developed countries contributed negatively to the real oil price throughout much of the sample period (Panel (b)). Only during the period leading up to the onset of the global financial crisis did developed economies contribute to driving up the oil price. This emphasizes the importance of distinguishing between the various

[14] We scale the initial values such that the total variance explained by each structural shock (the bars in Figure 3) is in accordance with the variance decompositions in Table 2.

Figure 3: **Historical shock decompositions**

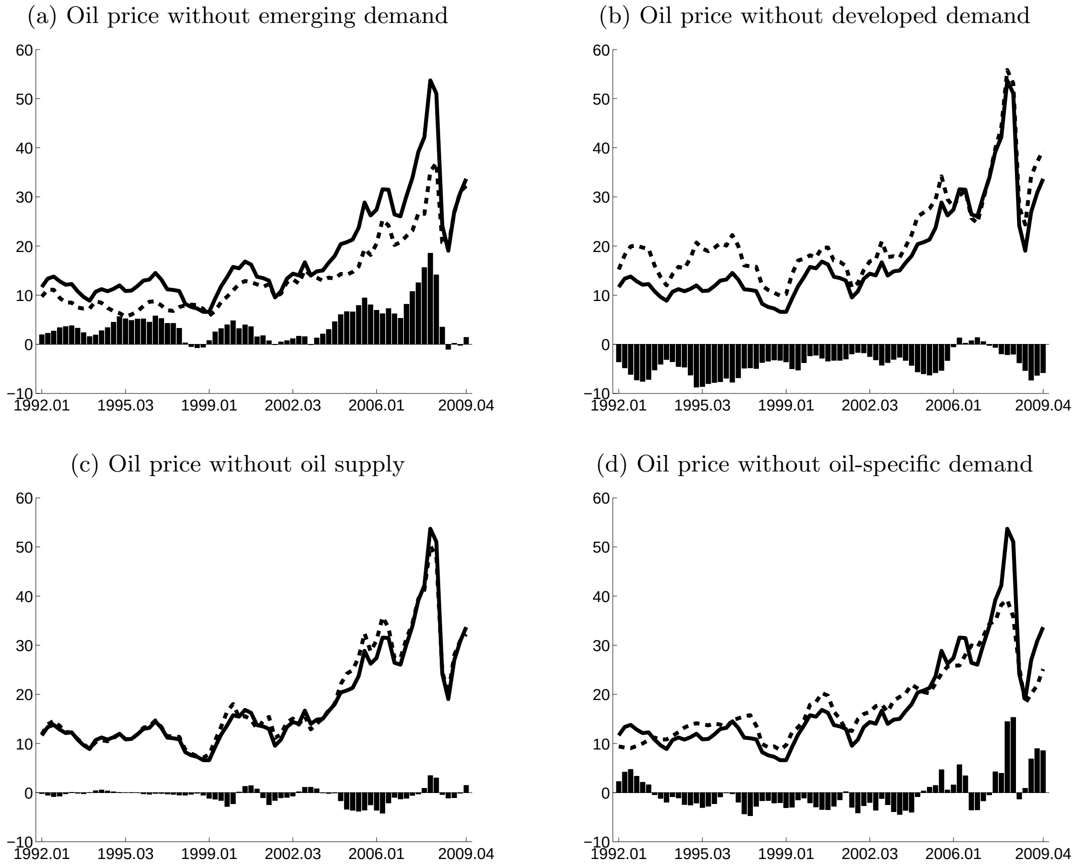

Note: The solid lines display the actual real oil price. The dotted lines display what the oil price would have been if we exclude one of the structural shocks. The bars show the difference between the solid and dotted lines. A positive value indicates that the structural shock contributed to increasing the price of oil.

sources behind the demand. While the negative contribution from the developed countries resembles the results for the aggregate global demand shocks identified in Kilian (2009), especially during the 1990s, we show that from 2005 onwards, the positive contribution is primarily due to increased demand from emerging countries.

The contrasting results for emerging and developed economies as drivers of the oil price have not been previously documented in the literature, but are well in line with the changes in global oil consumption patterns, where, e.g., the share of total world oil consumption attributed to non-OECD countries has grown by almost 40 percent since the beginning of the 1990s.

The results in Panel (c) reiterate the discussion above; oil supply shocks have contributed little in explaining oil price fluctuations over the last two decades, a finding supported by a number of recent studies. However, as described in, e.g., Hamilton (2013), the only geopolitical events that have potentially affected world oil production since 1992 were the Venezuelan unrest and the second Persian Gulf War, which both occurred around 2003. Although the results in Panel (c) suggest that oil supply shocks increased the price of oil during this period, the effects are small.

Finally, Panel (d), which graphs the historical contribution of oil-specific demand shocks, reveals a more erratic pattern than any of the other shocks, in particular during the financial crisis (2007-2009). This could be consistent with the interpretation that the shock primarily captures precautionary demand for oil, driven by the uncertainty of the future oil supply, as described in Kilian (2009). However, an alternative view, advocated by among others Lombardi and Van Robays (2011), is that speculation may have exacerbated oil price fluctuations, particularly prior to and during the financial crisis. Therefore, we cannot rule out the possibility that speculation accounts for part of our identified oil-specific demand shock. However, the results presented in Panel (d) do not indicate that oil-specific demand shocks have been the main factor driving up oil prices during the 2003-2008 period.

3.2 Region and country details

An advantage of our FAVAR methodology is that we can calculate the responses of various shocks across a large panel of countries and separately investigate the individual country effects. Thus, we add a dimension to the traditional literature concentrating on either one or a few developed (OECD) countries.

For ease of interpretation, we calculate the point estimates of the GDP response in levels for each individual country in the FAVAR model, but report in Figure 4 the median response at the two-year horizon across countries within the same geographical region: Asia, Europe, North America (NA) and South America (SA), respectively.[15]

A demand shock (normalized to increase activity in either developed or emerging countries by one percent initially) has a positive effect on GDP across all geographical regions. The emerging-country demand shock has by far the strongest effect on Asia followed by South America, while the developed-country demand shock has the strongest effect on Eu-

[15] All the individual country impulse responses are reported in the on-line appendix.

Figure 4: **Effect of macroeconomic and oil market shocks on the level of GDP in different regions**

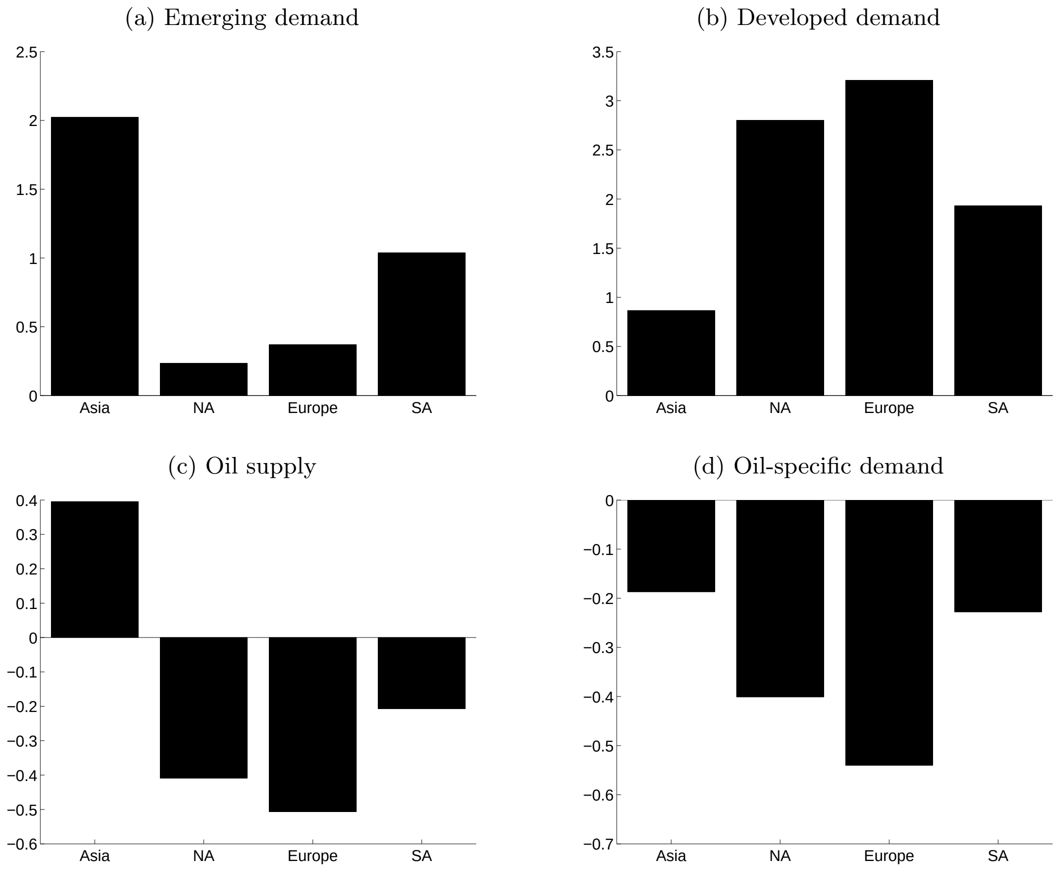

Note: All shocks are normalized to increase the real price of oil. The y-axis reports the median response at the two year horizon across countries within the same geographical region: Asia, Europe, North America (NA) and South America (SA). Each individual country's response is based on the point estimate for that country and horizon, see the on-line appendix and Figure 1 for further details.

rope, followed closely by North America. Interestingly, South America is more positively affected by the demand shock in developed countries than Asia. This is consistent with observed trade patterns, where a larger share of South America's trade is with developed countries than is the case for Asia.

The responses to the oil supply and the oil-specific demand shocks are more varied. Panel (c) of Figure 4 shows that while economic activity in North America and Europe is negatively affected following an adverse oil supply shock (that eventually increases the real oil price), activity in emerging countries in Asia and South America is much less affected; activity in Asia even increases.[16] The same divergence between the regions can

[16]Although the median response in Asia is substantial, the response is particularly strong in Indonesia and

also be found following oil-specific demand shocks, see Panel (d). While all countries are now affected negatively following such an oil shock, the effect is again less severe for the emerging Asian and South American countries than for Europe and the US.

In sum, our results highlight a large degree of heterogeneity in terms of spillovers from the oil market to the macro economy, with emerging countries in Asia and South America being more important drivers of the real oil price, yet responding less severely to adverse oil market shocks.

3.3 The Asian puzzle - further examination

The heterogeneity in responses across countries and regions to disturbances in the oil market requires further examination. First, as previously mentioned, some of the countries in the sample are commodity exporters, where the terms of trade increase with higher oil prices (oil prices and most commodity prices are highly correlated). However, this cannot explain the positive response in the majority of Asian countries that are not commodity exporters. Nor can it explain why South America does not respond more positively, given the large share of commodity exporters in this region.

Second, the structure of a country matters. According to Hamilton (2009), a key parameter in determining the consequences of an oil price increase is the share of energy purchases in total expenditures. In particular, a low expenditure share combined with a low price elasticity of demand will imply very small negative effects (if any) of an oil price increase. While the oil consumption share in the US and other industrial economies has generally been flat since the 1980s, it has risen sharply in emerging countries such as China. However, as China began from a much lower level, per capita oil consumption in the US is still 10 times larger than in China (cf. IMF WEO 2011). This may suggest that the (negative) price elasticity in the US is larger (in absolute terms) than in China, and thus that the US and other industrial countries respond more negatively to adverse oil market shocks than emerging countries.

Finally, as pointed out by Edelstein and Kilian (2009) and Hamilton (2009), a key factor transmitting energy price shocks to the domestic economy has been the automobile

Malaysia, two energy-rich countries. In South America, only Brazil and Peru respond positively, explaining why the overall response in South America is negative, see the on-line appendix for more details.

Table 1: **Characteristics and correlations**

(a) Correlations between country structure and IRF levels

	Consumption	**Investment**	**Open**
Oil supply shocks	-0.44	0.59	0.46
	(0.01)	(0.00)	(0.01)
Oil demand shocks	-0.28	0.48	0.22
	(0.11)	(0.00)	(0.22)

(b) Regional structure

	Consumption	**Investment**	**Open**
Emerging Asia	55	33	145
Developed countries	66	23	60
Emerging South America	70	22	50

Note: Panel (a) reports the correlations between the impulse responses from the two oil market shocks (over the two-year horizon) with the means of the three indicators; the consumption share, investment share and degree of openness (export + import), all as a share of PPP converted GDP per capita at 2005 constant prices. (The mean of the indicators are calculated over the sample period 1992 - 2009). The second row for each shock are p-values. Panel (b) graphs the consumption share, investment share and degree of openness in each region.

sector. In particular, higher energy prices have typically implied an increase in the demand for energy-efficient small cars at the expense of energy-inefficient large cars (SUVs). Consistent with this, Hamilton (2009) has shown that US domestic light truck sales were down almost 30 percent between July 2007 and July 2008 (the year when oil prices peaked), while US import car sales were up 14 percent in the same period. The increased demand in the US for import of energy-efficient cars may have benefitted emerging countries in Asia in particular.

In Table 1, we examine these issues by analyzing informally to what extent the composition of output is relevant. In particular, as investments are less reversible (due to long-term plans) than overall consumption, we examine whether countries with a low consumption share and high investment share are less negatively affected by higher oil prices. We also examine if the degree of openness is relevant. As discussed above in relation to the mix of industries, countries that have tied most of their capital to the export sector may be able to export some of the increase in oil prices to the importers, whose price elasticity may be small.

Panel (a) in Table 1 reports the correlations between the impulse responses from the two oil market shocks (oil supply and oil-specific demand shocks), with the means of the three relevant indicators, the consumption share, investment share and degree of openness (all as share of GDP) over the sample period (1992 - 2009). Panel (a) suggests that countries with a high consumption share have negative correlations with the two oil shocks, while countries with high investment shares and a high degree of openness have positive correlations with the oil shocks (oil supply shocks in particular).

Panel (b) shows that emerging Asia matches these features, having a low consumption share, high investment share and a high degree of openness. For South American countries, however, the picture is reversed (high consumption shares, low investment shares and low degrees of openness), potentially explaining why, despite being positively affected by increased terms of trade, they do not respond positively overall.

On a final comment, subsidies may also play a role. Price controls prevent the full cost of a higher imported oil price from being passed through to the end user, thereby dampening the responsiveness of consumption to increases in prices. Studies in the IMF's World Economic Outlook reports from 1999 and 2009 show that pervasive under-pricing of energy resources occurs in several non-OECD countries, including China, India, Indonesia and South Africa. This may also help to explain the small (and occasionally positive) effects of oil supply and oil demand shocks in these countries. We leave this issue to be explored further in another study.

3.4 Robustness checks

To examine the robustness of our findings, we perform a series of tests. The main results are summarized briefly here. For details we refer to the on-line appendix. First, we analyze the separate role of Asia versus South America as drivers of demand for oil. Results suggest that Asia is the primary driver of demand for oil. Second, monetary conditions are often cited as a driver of commodity prices, see, e.g., Anzuini et al. (2013). By including the federal funds rate explicitly into the baseline model, we find that shocks to the federal funds rate can explain a considerable fraction of the fluctuations in the emerging activity factor at longer horizons. Yet, the effect on the fluctuations in the real price of oil is marginal and our main results remain unchanged.

Third, to avoid a direct dependence on the factor loading structure imposed by the

FAVAR, we regress the estimated structural oil supply and oil-specific demand shocks from the FAVAR model on the individual countries' GDP growth using standard OLS. The findings confirm the baseline results reported in Figure 4 and the accompanying discussion. Finally, our results are robust to alternative model specifications and identification restrictions, estimating the model using Bayesian techniques, as well as using alternative measures of the oil price. We also document that the uncertainty bands reported in Figure 1 are only partially affected by the fact that some of the shocks in the model are set identified, i.e., sign identified.

4 Conclusion

In this paper we estimate a FAVAR model with separate activity factors for emerging and developed economies in addition to global oil production and the real price of oil and study two main questions: 1) How demand shocks in emerging and developed economies affect the real price of oil and global oil production, and 2) what are the effects of oil supply and oil-specific demand shocks on emerging and developed economies. We have three main findings that are robust across numerous robustness checks.

First, we demonstrate that demand shocks to emerging and developed economies have accounted for 50-60 percent of the fluctuations in the real price of oil in the last two decades. Second, demand shocks to emerging markets, in Asia in particular, are far more important than demand shocks in developed economies in explaining fluctuations in the real price of oil and in global oil production. Third, we find that various regions respond differently to adverse oil market shocks, with Europe and North America being more negatively affected than countries in Asia and South America.

Our results suggest that it is imperative not only to identify where demand for energy is coming from, but also to allow countries to respond differently to various oil market shocks. For policy makers, knowledge of these heterogeneities is important.

References

Aastveit, K. A. (2013). Oil price shocks and monetary policy in a data-rich environment. Working Paper 2013/10, Norges Bank.

Anzuini, A., M. J. Lombardi, and P. Paganoa (2013). The impact of monetary policy shocks on commodity prices. *International Journal of Central Banking 9*, 119–144.

Bai, J. and S. Ng (2002). Determining the number of factors in approximate factor models. *Econometrica 70*(1), 191–221.

Bai, J. and S. Ng (2013). Principal components estimation and identifiaction of static factors. *Journal of Econometrics 176*(1), 18–29.

Barsky, R. B. and L. Kilian (2002). Do we really know that oil caused the great stagflation? a monetary alternative. In B. S. Bernanke and K. Rogoff (Eds.), *NBER Macroeconomics Annual 2001*, pp. 137–183. MIT Press: Cambridge, MA.

Baumeister, C. and G. Peersman (2013a). The role of time-varying price elasticities in accounting for volatility changes in the crude oil market. *Journal of Applied Econometrics 28*(7), 1087–1109.

Baumeister, C. and G. Peersman (2013b). Time-varying effects of oil supply shocks on the us economy. *American Economic Journal: Macroeconomics 5*(4), 1–28.

Bernanke, B., J. Boivin, and P. S. Eliasz (2005). Measuring the effects of monetary policy: A factor-augmented vector autoregressive (FAVAR) approach. *Quarterly Journal of Economics 120*(1), 387–422.

Bjørnland, H. C. (2000). The dynamic effects of aggregate demand, supply and oil price shocks - A comparative study. *Manchester School of Economic Studies 68*(5), 578–607.

Boivin, J., M. P. Giannoni, and I. Mihov (2009). Sticky prices and monetary policy: Evidence from disaggregated US data. *American Economic Review 99*(1), 350–84.

Burbidge, J. and A. Harrison (1984). Testing for the effects of oil-price rises using vector autoregressions. *International Economic Review 25*(2), 459–84.

Canova, F. (2005). The transmission of US shocks to Latin America. *Journal of Applied Econometrics 20*(2), 229–251.

Edelstein, P. and L. Kilian (2009). How sensitive are consumer expenditures to retail energy prices? *Journal of Monetary Economics 56*(6), 766–779.

Eickmeier, S. and T. Ng (2011). How do credit supply shocks propagate internationally? A GVAR approach. CEPR Discussion Papers 8720.

Farrant, K. and G. Peersman (2006). Is the exchange rate a shock absorber or a source of shocks? New empirical evidence. *Journal of Money, Credit and Banking 38*(4), 939–961.

Fry, R. and A. Pagan (2011). Sign restrictions in structural vector autoregressions: A critical review. *Journal of Economic Literature 49*(4), 938–60.

Gisser, M. and T. H. Goodwin (1986). Crude oil and the macroeconomy: Tests of some popular notions: A note. *Journal of Money, Credit and Banking 18*(1), 95–103.

Goncalves, S. and B. Perron (2013). Bootstrapping factor-augmented regression models. *Journal of Econometrics*, Forthcoming.

Hall, P. (1992). *The bootstrap and Edgeworth expansion.* Springer: New York, NY.

Hamilton, J. D. (1983). Oil and the macroeconomy since world war II. *Journal of Political Economy 91*(2), 228–248.

Hamilton, J. D. (1996). This is what happened to the oil price-macroeconomy relationship. *Journal of Monetary Economics 38*(2), 215–220.

Hamilton, J. D. (2003). What is an oil shock? *Journal of Econometrics 113*(2), 363–398.

Hamilton, J. D. (2009). Causes and consequences of the oil shock of 2007-08. *Brookings Papers on Economic Activity 40*(1), 215–283.

Hamilton, J. D. (2013). Historical oil shocks. In R. E. Parker and R. M. Whaples (Eds.), *Routledge Handbook of Major Events in Economic History*, pp. 239–265. New York: Routledge Taylor and Francis Group.

Hamilton, J. D. and A. M. Herrera (2004). Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy: Comment. *Journal of Money, Credit and Banking 36*(2), 265–86.

Hicks, B. and L. Kilian (2013). Did unexpectedly strong economic growth cause the oil price shock of 2003-2008? *Journal of Forecasting 32*(5), 385–394.

Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review 99*(3), 1053–1069.

Kilian, L. and D. Murphy (2013). The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics*, Forthcoming.

Kose, M. A. and E. S. Prasad (2011). *Emerging Markets: Resilience and Growth Amid Global Turmoil*. Washington D.C.: Brookings Institution Press.

Lippi, F. and A. Nobili (2012). Oil and the macroeconomy: A quantitative structural analysis. *Journal of the European Economic Association 10*(5), 1059–1083.

Lombardi, M. J. and I. Van Robays (2011). Do financial investors destabilize the oil price? ECB Working Paper 1346, European Central Bank.

Maćkowiak, B. (2007). External shocks, US monetary policy and macroeconomic fluctuations in emerging markets. *Journal of Monetary Economics 54*(8), 2512–2520.

Mumtaz, H. and P. Surico (2009). The transmission of international shocks: A factor-augmented VAR approach. *Journal of Money, Credit and Banking 41*(s1), 71–100.

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

Rubio-Ramirez, J. F., D. F. Waggoner, and T. Zha (2010). Structural vector autoregressions: Theory of identification and algorithms for inference. *Review of Economic Studies 77*(2), 665–696.

Thorsrud, L. A. (2013). Global and regional business cycles. shocks and propagations. Working Paper 2013/08, Norges Bank.

ONLINE APPENDIX TO

What drives oil prices?

Emerging versus developed economies

Knut Are Aastveit* Hilde C. Bjørnland† Leif Anders Thorsrud‡

May 28, 2014

*Norges Bank. Email: Knut-Are.Aastveit@norges-bank.no
†BI Norwegian Business School.Email: hilde.c.bjornland@bi.no
‡BI Norwegian Business School. Email: leif.a.thorsrud@bi.no

1 Introduction

In this on-line appendix we provide details on the data, sources and diagnostics of the model, including a discussion of the estimated factors and the implementation of sign restriction. Thereafter we report the estimated country responses in greater detail, and discuss various robustness tests in terms of the chosen estimation, specification and identification strategy.

2 Data, sources and oil market exposure

Our data set includes variables from 33 different countries, where we use real GDP growth and industrial production growth as measures of economic activity for each country. In total, our sample countries account for approximately 80 percent of world GDP, measured by purchasing-power-parity (authors calculations based on 2009 estimates from the IMF).

We determine a priori which countries should be considered developed and emerging economies. Countries that are members of the OECD at the beginning of our sample are considered developed economies. The remaining countries are considered emerging economies. Accordingly, the following 18 countries are considered developed economies: Australia, Belgium, Canada, Denmark, Finland, France, Germany, Italy, Japan, Luxembourg, Netherlands, New Zealand, Norway, Spain, Sweden, Switzerland, the UK and the US. The following 15 countries are considered emerging economies: Argentina, Brazil, Chile, China, Hong Kong, India, Indonesia, Korea, Malaysia, Mexico, Peru, Singapore, South Africa, Taiwan and Thailand. Of these countries, four developed countries (Canada, Denmark, Norway and the UK) and four emerging countries (Argentina, Indonesia, Malaysia and Mexico) are net oil exporters over the period (1991-2009). However, many other countries are commodity producers (e.g., Australia, New Zealand and Peru), where export prices may have been highly correlated with oil prices over the period (see Table 1 below). To these categorizations it is worth noting that Chile, Korea and Mexico are now members of the OECD, and although Brazil is not a net oil exporter over the entire period, in recent years Brazil has been a major producer and a net exporter.

Most of the data series were collected from Thomson Reuters Ecowin. Other series were collected from the following sources: Gross Domestic Product (GDP) in China and

Indonesia were found in the GVAR data set constructed by Pesaran et al. (2009). Industrial production (IP) in Argentina, Indonesia, Mexico and the Netherlands were collected from Datastream. Industrial production in Denmark and Portugal were taken from OECD, while industrial production in Norway was collected from Statistics Norway.

All GDP series are at constant prices. The industrial production series are volume indexes, and refer, with few exceptions, to the manufacturing industry. For Argentina, China, Indonesia, Italy, Norway, Peru and Portugal, we only found series for overall industrial production. Some of the activity series do not span the whole time period used in the analysis. To avoid excluding these variables from the sample, we have applied the EM algorithm, as described in Stock and Watson (2002), to construct the missing observations. However, experiments conducted on the data set excluding the series with missing observations, do not change our main conclusions.

To measure oil production and the real price of oil, we use world crude oil production, in millions of barrels per day, and US real refiner' acquisition cost for imported crude oil, respectively. The nominal oil price has been deflated using the US consumer price index. These are the same variables used in Kilian (2009) and many other papers.

3 Factors and correlations

Figure 1, Panels (a) and (b), display the two observable series: global oil production and the real price of oil. The figure shows significant growth in the real price of oil during the economic booms in 1999/2000 and 2006/2007 and a decrease in the real price of oil during the Asian crisis and the recent financial crisis, (see Panel (a)). The economic booms and busts are also evident in global oil production (see Panel (b)), where production slows down during the two recessions and increases during the two expansions. Furthermore, there is also evidence of a slowdown in global oil production during 2002/2003. The dates coincide with the Venezuelan unrest (strike) and US attack on Iraq (second Persian Gulf War).

Figure 1, Panels (c) and (d), display the two key activity variables used in the analysis: the emerging and developed economy factors. As the figure shows, the two factors capture features commonly associated with the business cycles in each region over the last 20 years.

Table 1: **Oil production and consumption by countries**

	Country	Production	Consumption	Net exporter
Developed	*Australia*	0.64	0.87	No
	Belgium	0.01	0.61	No
	Canada	2.83	2.03	Yes
	Denmark	0.29	0.20	Yes
	Finland	0.01	0.21	No
	France	0.09	1.97	No
	Germany	0.13	2.74	No
	Italy	0.14	1.83	No
	Japan	0.11	5.36	No
	Netherlands	0.06	0.89	No
	New Zealand	0.05	0.14	No
	Norway	2.93	0.22	Yes
	Portugal	0.00	0.31	No
	Spain	0.03	1.39	No
	Sweden	0.00	0.37	No
	Switzerland	0.00	0.27	No
	United Kingdom	2.32	1.78	Yes
	United States	9.04	19.19	No
Emerging	*Argentina*	0.80	0.50	Yes
	Brazil	1.59	2.06	No
	Chile	0.02	0.23	No
	China	3.45	5.17	No
	Hong Kong	0.00	0.25	No
	India	0.77	2.12	No
	Indonesia	1.38	1.04	Yes
	Korea, South	0.01	2.06	No
	Malaysia	0.75	0.46	Yes
	Mexico	3.42	2.00	Yes
	Peru	0.11	0.15	No
	Singapore	0.01	0.69	No
	South Africa	0.20	0.47	No
	Taiwan	0.00	0.83	No
	Thailand	0.21	0.77	No

Note: Column three to five reports oil production and oil consumption in millions of barrels per day, measured as averages for the period 1992-2009 (Source: EIA).

Both the booms and busts pre-dating and following the Asian crisis near the end of the 1990s, and the dot-com bubble around 2001 are evident in the emerging-country and the developed-country factors, respectively.

There is, however, a notable difference in how the recent financial crisis has affected the two factors. The decline in the activity factor representing the developed economies is much larger than any other previous decline in that factor. For the emerging-country activity factor, the recent financial crisis also caused large negative movements. However, compared to earlier downturns, the recent crisis does not seem particularly different. Ad-

Figure 1: **Observable variables and estimated factors**

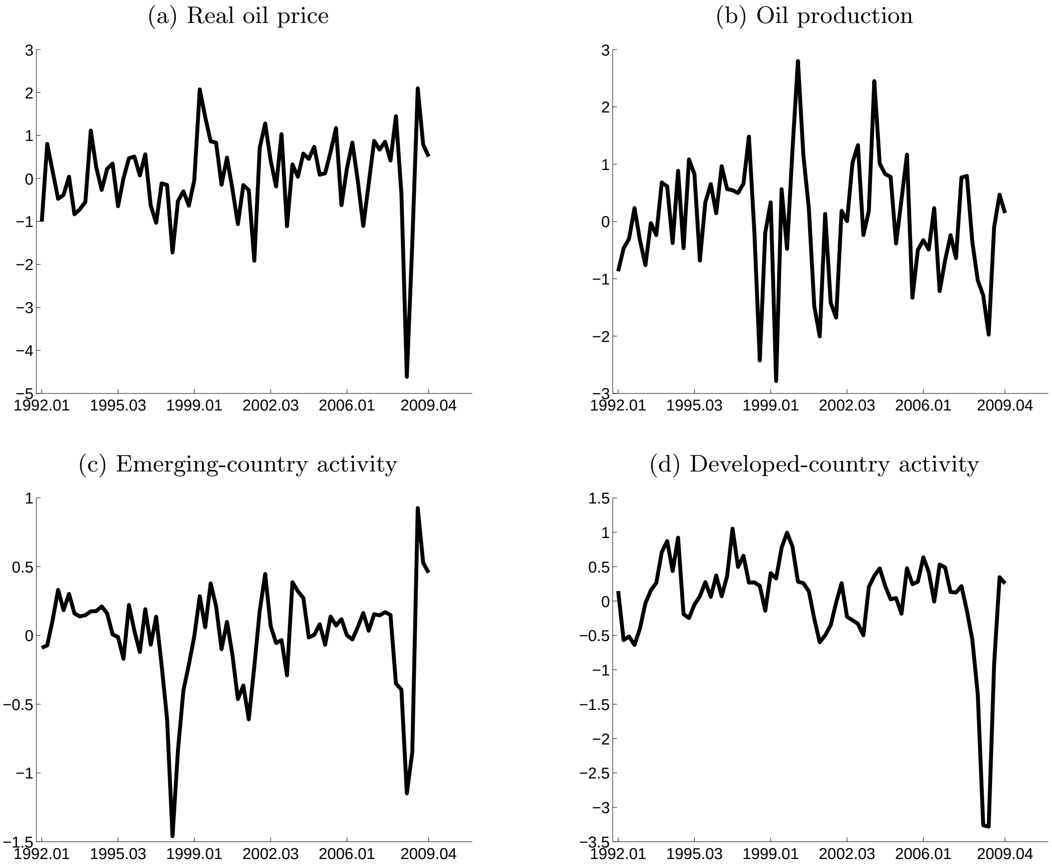

Note: The figure shows the standardized values of the first differences of the logs of each observable variable, i.e. the real price of oil and global oil production, and the estimated activity factors (the median). The sample used in the VAR is 1992:Q1 to 2009:Q4, while we use information from 1991:Q1 to 2009:Q4 to estimate the unobserved factors.

ditionally, the recovery in the emerging-country activity factor has been stronger than in the developed-country economy factor.

Although the factors should capture common movements among the countries in each group, the various countries may still contribute differently to the factor estimate. In particular, some countries may be more correlated with their respective factor than others. To illustrate this (and to further interpret the factors), Table 2 displays the correlation between the activity variables in each country and the developed-country and emerging-country factors. First, regarding the developed-country factor, the table indicates that with the exception of Australia, Japan, New Zealand and Norway, all developed countries are highly correlated with the factor (as expected). For Japan and New Zealand, however, the correlation with the emerging-country factor (that contains many Asian countries) is slightly higher than with the developed-country factor. Clearly, location is important. For Norway, and to some extent Australia, the correlation between GDP and either the developed-country or the emerging-country factors is low, suggesting a more idiosyncratic pattern in these countries.

Regarding the emerging-country factor, the results for the Asian and the South American countries are more diverse. While the Asian countries are highly correlated with the emerging-country factor, three of the South American countries (Argentina, Chile, Mexico) and South Africa are slightly more correlated with the developed-country factor than with the emerging-country factor. This indicates that the Asian countries account for the majority of variation in the emerging-country factor.

Table 2: **Correlation with factors**

Developed				**Emerging**			
Country	**Var.**	*emeAct*	*devAct*	**Country**	**Var.**	*emeAct*	*devAct*
Australia	*GDP*	0.09 (0.12)	0.35 (0.12)	*Argentina*	*GDP*	0.20 (0.15)	0.25 (0.18)
	IP	0.34 (0.13)	0.48 (0.12)		*IP*	0.36 (0.13)	0.38 (0.13)
Belgium	*GDP*	0.41 (0.14)	0.78 (0.14)	*Brazil*	*GDP*	0.44 (0.12)	0.41 (0.12)
	IP	0.32 (0.12)	0.63 (0.12)		*IP*	0.51 (0.13)	0.41 (0.12)
Canada	*GDP*	0.16 (0.14)	0.75 (0.17)	*Chile*	*GDP*	0.24 (0.13)	0.32 (0.15)
	IP	0.32 (0.14)	0.70 (0.16)		*IP*	0.37 (0.13)	0.41 (0.13)
Denmark	*GDP*	0.19 (0.11)	0.56 (0.11)	*China*	*GDP*	0.36 (0.11)	0.14 (0.13)
	IP	0.04 (0.11)	0.43 (0.09)		*IP*	0.27 (0.11)	0.10 (0.10)
Finland	*GDP*	0.26 (0.14)	0.80 (0.16)	*Hong Kong*	*GDP*	0.78 (0.14)	0.47 (0.15)
	IP	0.15 (0.14)	0.62 (0.14)		*IP*	0.43 (0.13)	0.33 (0.14)
France	*GDP*	0.25 (0.15)	0.83 (0.16)	*India*	*GDP*	*N/A*	*N/A*
	IP	0.35 (0.15)	0.75 (0.15)		*IP*	0.19 (0.11)	0.24 (0.12)
Germany	*GDP*	0.29 (0.13)	0.74 (0.14)	*Indonesia*	*GDP*	0.54 (0.13)	-0.02 (0.14)
	IP	0.20 (0.14)	0.68 (0.15)		*IP*	0.58 (0.14)	0.01 (0.15)
Italy	*GDP*	0.42 (0.15)	0.80 (0.16)	*Korea*	*GDP*	0.59 (0.13)	0.49 (0.14)
	IP	0.43 (0.15)	0.83 (0.16)		*IP*	0.70 (0.13)	0.45 (0.13)
Japan	*GDP*	0.63 (0.12)	0.52 (0.13)	*Malaysia*	*GDP*	0.49 (0.11)	0.27 (0.11)
	IP	0.66 (0.14)	0.46 (0.15)		*IP*	0.69 (0.11)	0.48 (0.15)
Netherlands	*GDP*	0.17 (0.14)	0.78 (0.16)	*Mexico*	*GDP*	0.26 (0.14)	0.67 (0.14)
]	*IP*	0.29 (0.12)	0.56 (0.12)		*IP*	0.16 (0.15)	0.61 (0.15)
New Zealand	*GDP*	0.47 (0.14)	0.43 (0.16)	*Peru*	*GDP*	0.31 (0.12)	0.08 (0.11)
	IP	*N/A*	*N/A*		*IP*	0.45 (0.13)	0.33 (0.13)
Norway	*GDP*	0.08 (0.10)	0.33 (0.11)	*Singapore*	*GDP*	0.75 (0.13)	0.43 (0.14)
	IP	0.17 (0.12)	0.51 (0.13)		*IP*	0.54 (0.12)	0.32 (0.11)
Portugal	*GDP*	0.10 (0.12)	0.66 (0.15)	*South Africa*	*GDP*	0.24 (0.14)	0.56 (0.18)
	IP	-0.07 (0.10)	0.24 (0.10)		*IP*	0.40 (0.13)	0.61 (0.14)
Spain	*GDP*	-0.02 (0.15)	0.75 (0.19)	*Taiwan*	*GDP*	0.56 (0.12)	0.52 (0.13)
	IP	0.31 (0.13)	0.76 (0.13)		*IP*	0.61 (0.13)	0.27 (0.13)
Sweden	*GDP*	0.32 (0.14)	0.83 (0.16)	*Thailand*	*GDP*	0.48 (0.13)	0.22 (0.14)
	IP	0.32 (0.14)	0.78 (0.15)		*IP*	0.63 (0.13)	0.42 (0.14)
Switzerland	*GDP*	0.17 (0.15)	0.69 (0.16)				
	IP	0.33 (0.12)	0.62 (0.12)				
United Kingdom	*GDP*	0.23 (0.15)	0.84 (0.18)				
	IP	0.37 (0.14)	0.80 (0.15)				
United States	*GDP*	0.27 (0.14)	0.71 (0.16)				
	IP	0.36 (0.16)	0.81 (0.18)				
	Mean	***0.27***	***0.65***		***Mean***	***0.45***	***0.35***

Note: Column three to four, and seven to eight report the correlation between observable activity variables and the identified emerging-country and developed-country activity factors. IP is an abbreviation for industrial production. N/A are missing values. Standard errors in parenthesis.

4 Implementation of sign restrictions

We implement the following algorithm for each draw of the reduced form covariance matrix Ω:

1. Let $\Omega = PP'$ be the Cholesky decomposition of the VAR covariance matrix Ω, and $\tilde{A}_0 = P$.

2. Draw an independent standard normal 2×2 matrix J. Let $J = QR$ be the "economy size" QR decomposition of J with the diagonal of R normalized to be positive.

3. Compute a candidate structural impact matrix $A_0 = \tilde{A}_0\tilde{Q}$, where $\tilde{Q}$ is constructed from a 4×4 identity matrix, with the 2×2 matrix Q placed in the 2nd and 3rd row and column of $\tilde{Q}$.

If the candidate matrix satisfies the sign restrictions, we keep it. Otherwise the procedure above is repeated. The imposed signs can also be restricted to hold for many periods, in which case the candidate matrix must be past into the impulse response function before validation.

5 Country details

In the paper we report the median response at the two year horizon across countries within the same geographical region: Asia, Europe, North America (NA) and South America (SA), respectively. Here, Figures 2 and 3 report the country specific details. That is, the individual countries' responses to a oil supply and oil-specific demand shock.

Figure 2: **Impulse responses: Oil supply shock**

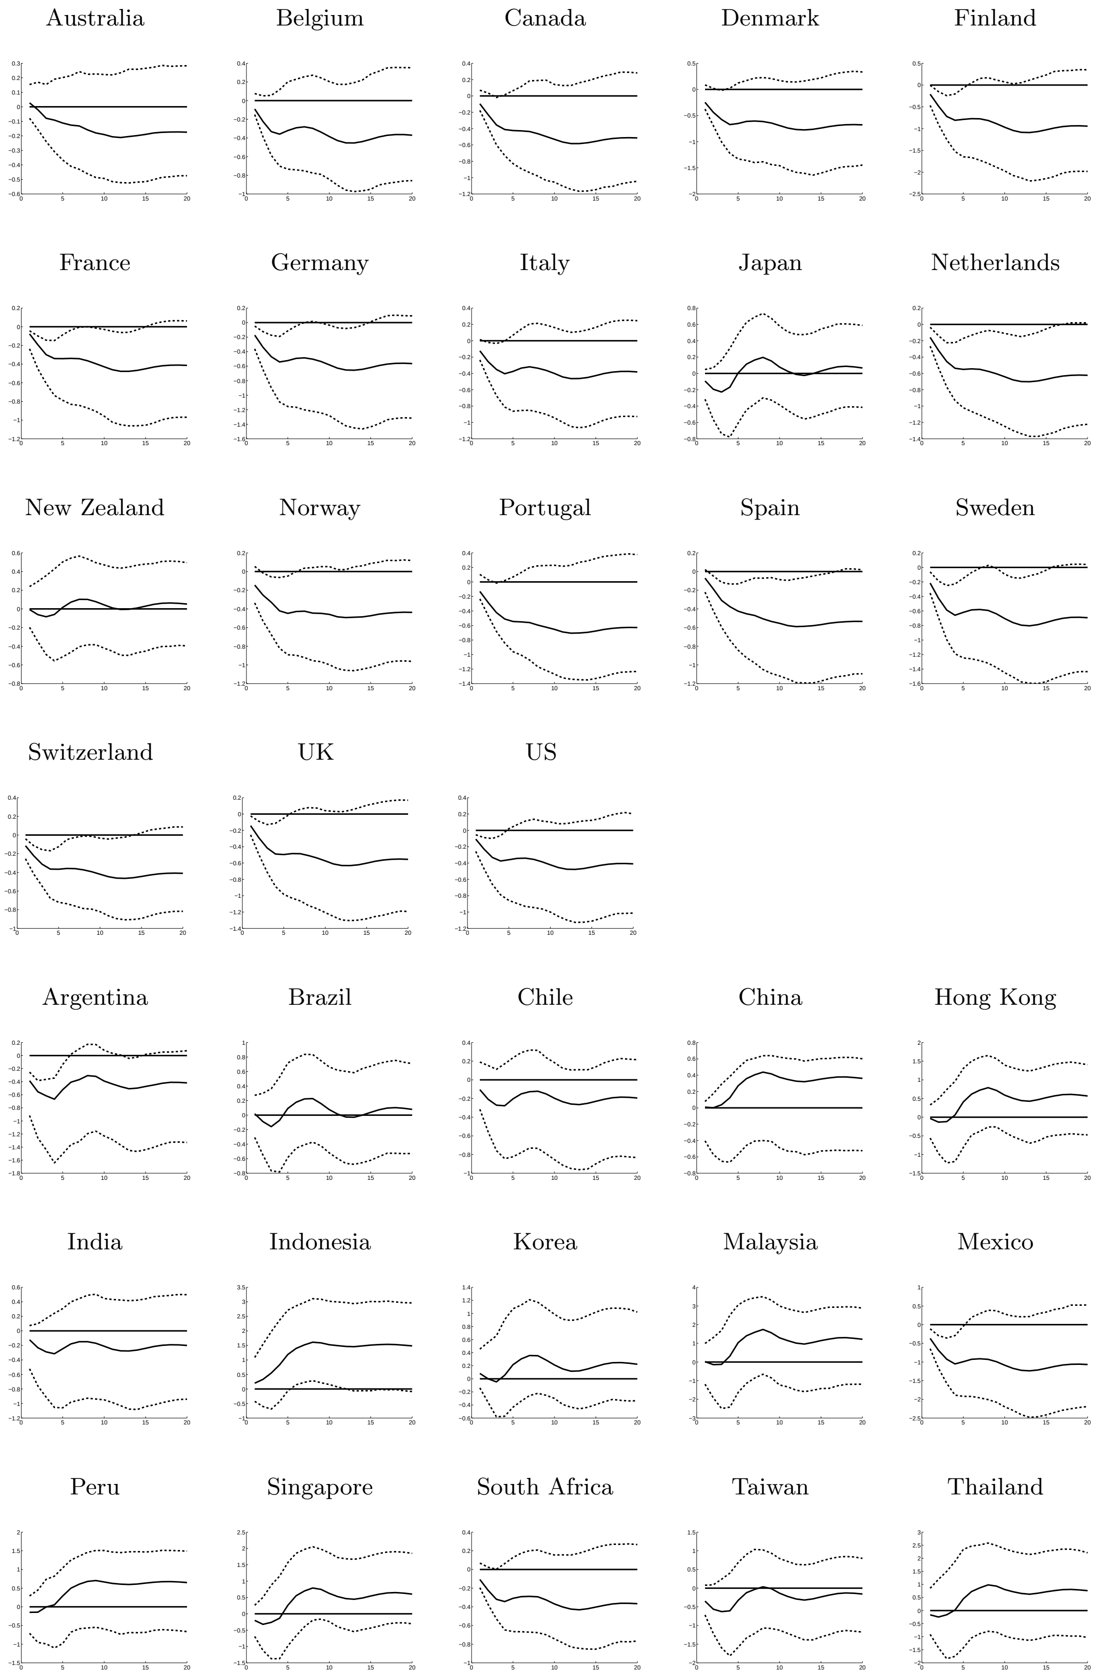

Note: The figures show the responses of GDP (in percent) in a given country after a oil supply shock that is normalized to decrease oil production by one percent. The responses are displayed in levels of the variables. The dotted lines represent 68 percent confidence bands (bootstrapped), while the black solid lines are the point estimates.

Figure 3: **Impulse responses: Oil-specific demand shock**

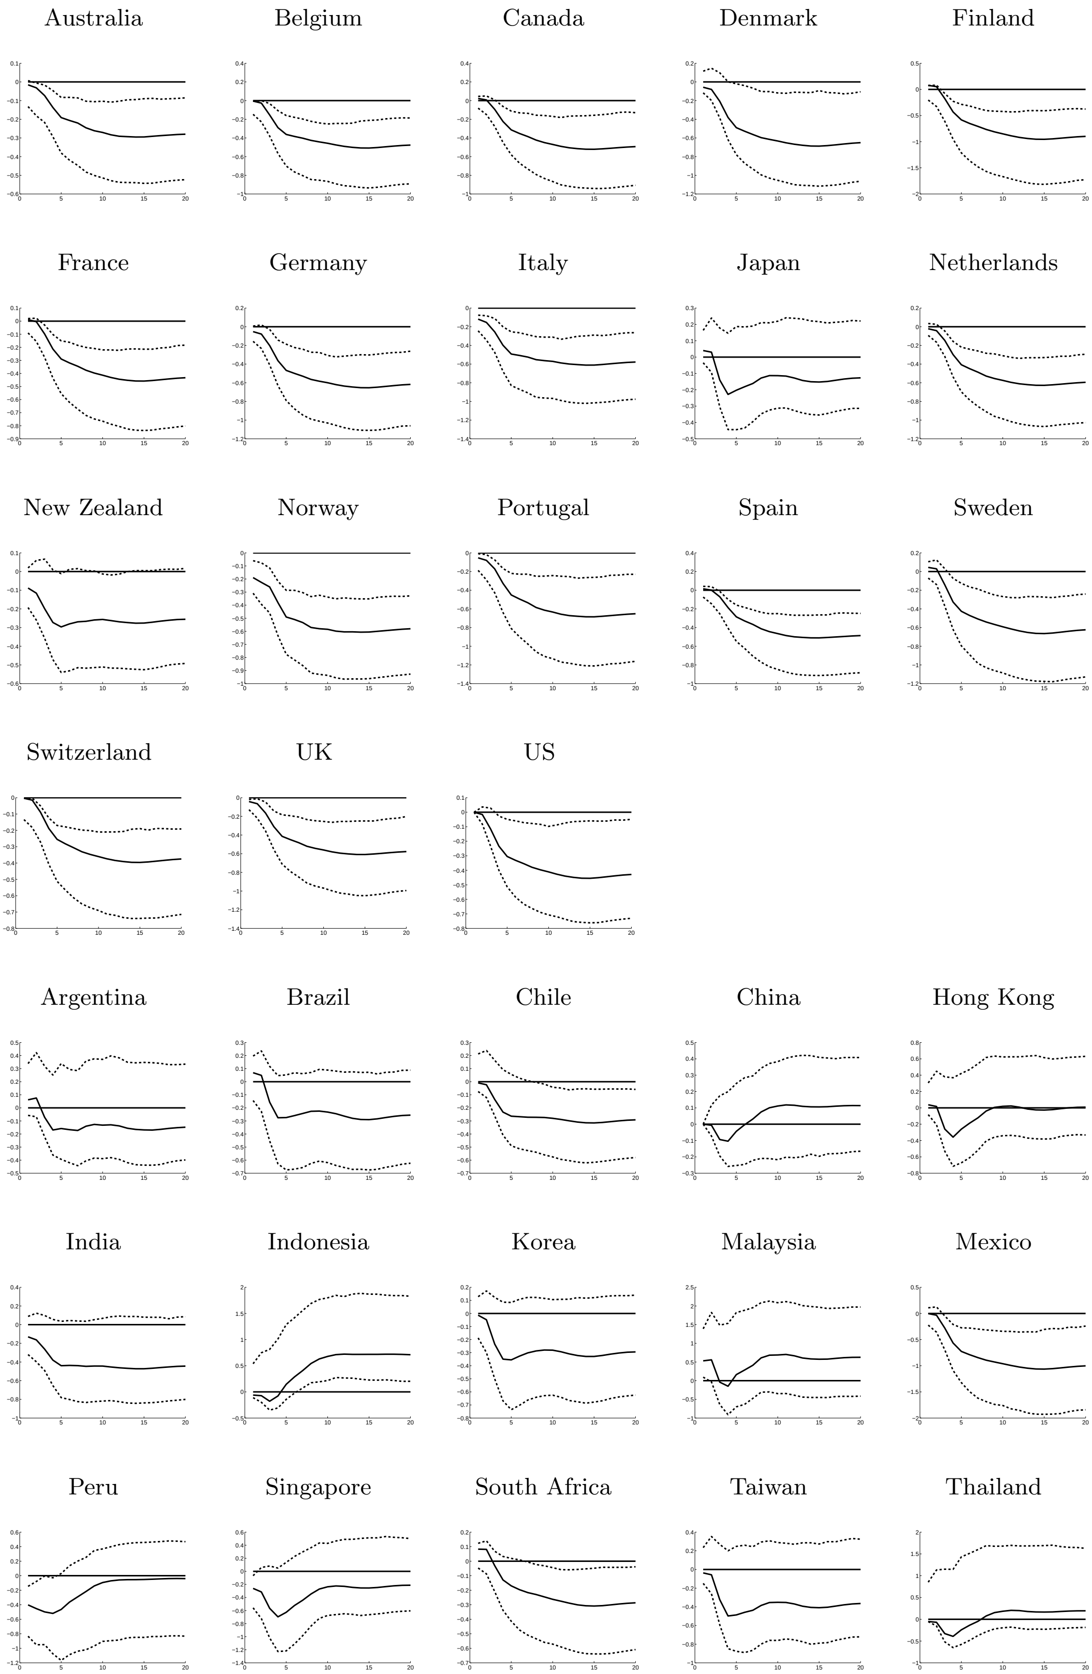

Note: The figures show the responses of GDP (in percent) in a given country after a oil-specific demand shock that increases oil prices with 10 percent. The responses are displayed in levels of the variables. The dotted lines represent 68 percent confidence bands (bootstrapped), while the black solid lines are the point estimates.

Figure 4: **Regression of oil shocks on observable GDP**

(a) Oil supply

(b) Oil-specific demand

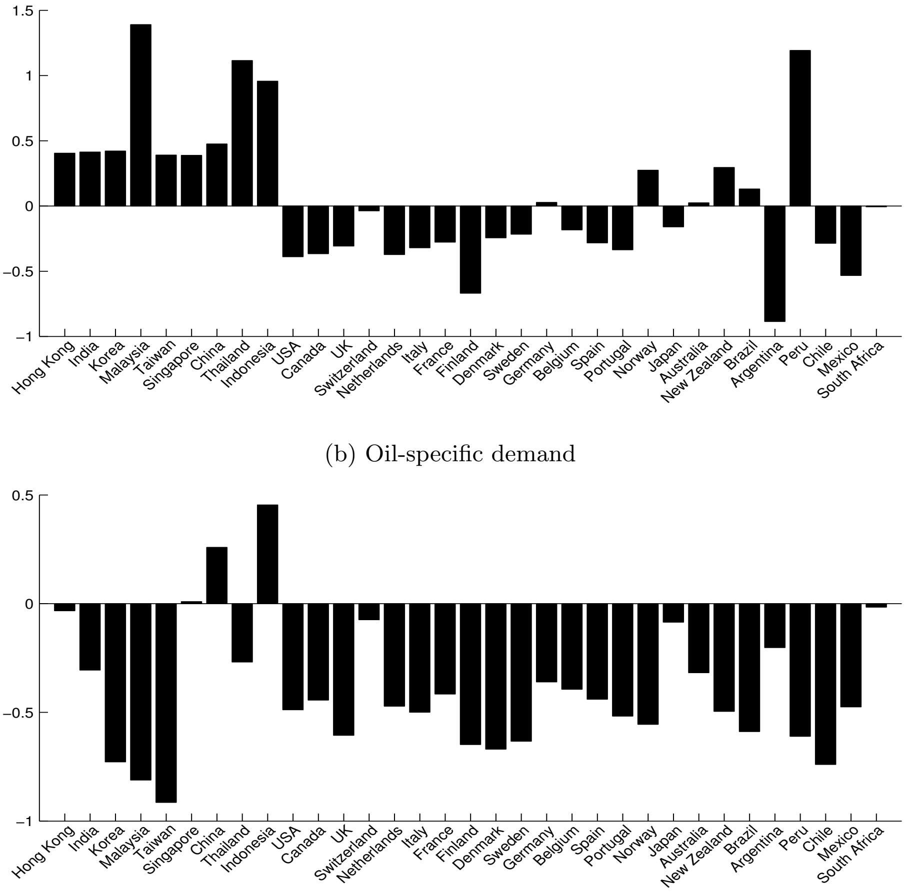

Note: The bars show for each country the accumulated regression coefficients from the following regressions:

$$\Delta X_{t,i} = \alpha_i + \sum_{p=1}^{4} \beta_{p,i} s_{t-p} + e_{t,i}$$

where $\Delta X_{t,i}$ is the observable GDP growth in country i at time t, α and β are coefficients, and s_{t-p} are lags of the structural shocks (oil supply or oil-specific demand) identified in our model.

One concern with our modeling strategy, and therefore the identified country specific responses, is related to the idiosyncratic part of our model. That is, the factors might explain very different proportions of the variance in each individual country's activity measure. For example, the correlation between Norwegian GDP and the developed-economy activity factor is only 0.3, while the correlation between US GDP and the developed-economy activity factor is as high as 0.7, see Table 2 above. Thus, to avoid a direct dependence on the factor loading structure imposed in the FAVAR, we regress the structural oil supply and oil-specific demand shocks estimated in the model on the individual countries' GDP growth rates using standard OLS. The results are plotted in Figure 4. This also serves as a robustness check for the individual countries' impulse responses, plotted in Figures 2 and 3.

The findings confirm the baseline results that oil supply shocks (that temporarily increase the oil price) stimulate GDP in all emerging countries in Asia as well as in Brazil and Peru, while for the remaining countries in South America and for most of the developed countries, GDP instead falls. There are, however, a few exceptions to this picture: In Australia, Germany, New Zealand and Norway, GDP also increases temporarily (as in Asia).

Regarding the oil-specific demand shock, most countries respond negatively as expected. However, as seen using the FAVAR model, some Asian countries (most notably Indonesia) respond positively, implying that the average response for the Asian countries is less severe than for the other countries.

6 Robustness

Our main results are not particularly sensitive to the number of lags used in the transition equation. In fact, when we estimate the model with two lags instead of four, the results are slightly stronger, implying that the emerging-country factor explains an even larger share of the variation in the real price of oil and oil production. Below we discuss in greater detail the results for a number of additional robustness checks.

6.1 Asia or South America?

We examine whether Asia or South America (or a combination of both) drives the relationship between the oil-market and the macro economy presented in our paper. To do so, we split the sample of emerging countries into two blocks and estimate two different factors: one consisting of emerging Asian countries and one consisting of emerging South American countries (including South Africa). When estimating the South American factor we use GDP growth in Brazil to normalize the factor loadings. Then, we sequentially use these new factor estimates in our main model, as a replacement for the original emerging-country factor.

Figure 5: **Variance decomposition: Asia and South America separately**

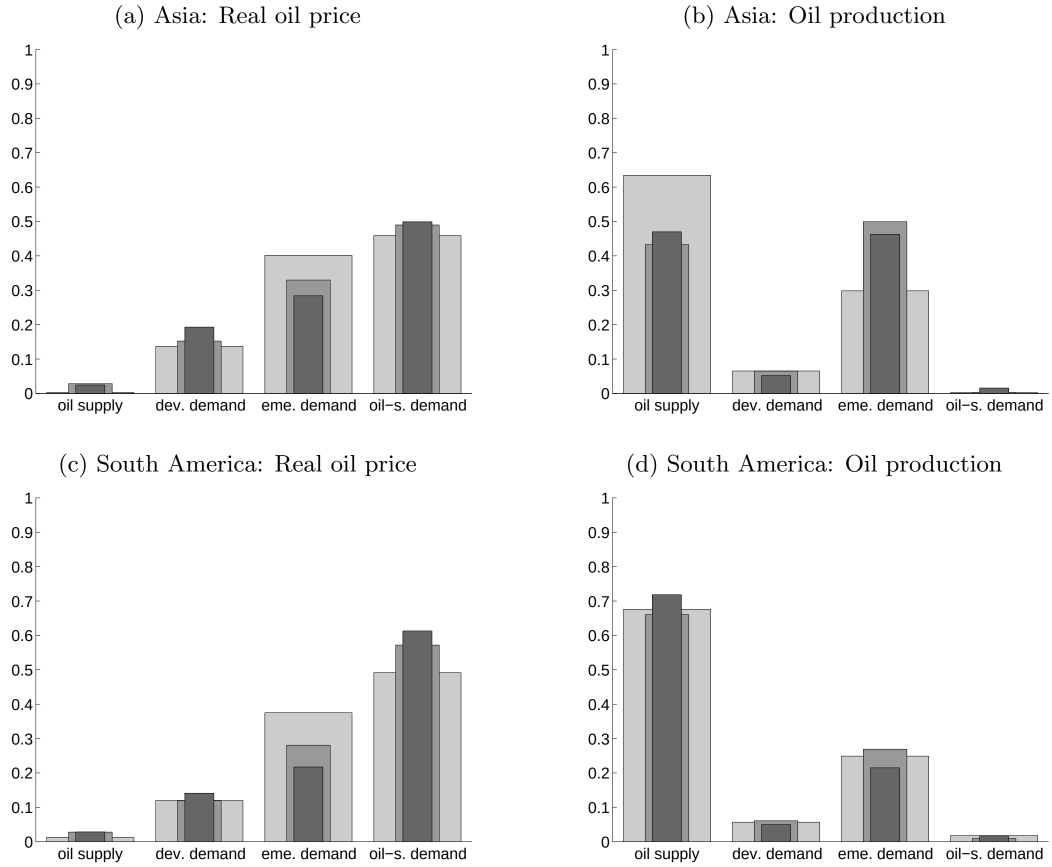

Note: In the model for Asia, only Asia is contained in the emerging factor, while in the model for South America, only South America is contained in the emerging factor. The bars display the variance decomposition with respect to the shocks for horizons 4, 8 and 12 quarters. The widest bars correspond to the shorter horizon.

The results using the emerging Asian factor are similar to our baseline results, while the results change when we use the emerging South American factor, see Figure 5. In particular, the emerging South American demand shock explains slightly less of the variance in oil prices and almost half of the variance in oil production compared to the results reported in the paper. This confirms that Asia is the main driver of the emerging factor, but the role of South America is far from negligible.

6.2 The impact of US demand and monetary policy

In this section we do two separate exercises. First, we exclude demand from the US altogether from the baseline model and second, we augment the model with the US interest rate. The exercises are motivated by the fact that many recent studies have argued that US

Figure 6: **Variance decomposition: Without the US**

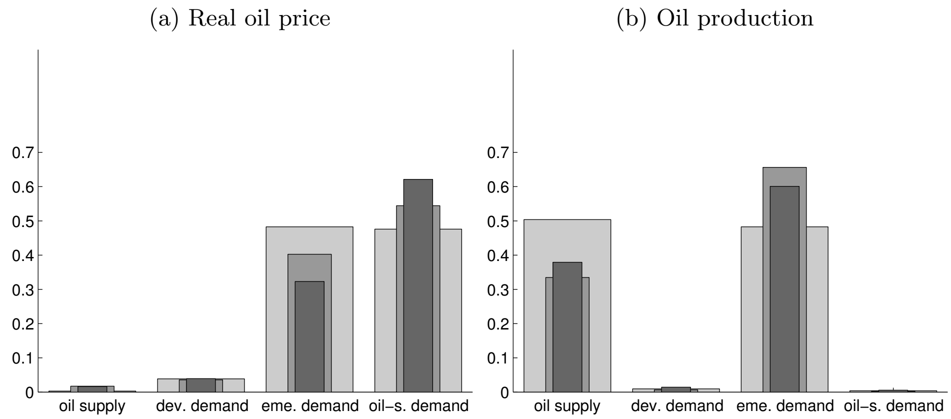

Note: The bars display the variance decomposition with respect to the shocks for horizons 4, 8 and 12 quarters. The widest bars correspond to the shorter horizon.

monetary policy could be an important driver of commodity prices, see, e.g. Anzuini et al. (2013), who analyze the effect of expansionary US monetary policy shocks on commodity prices.

Excluding the US from the analysis gives an even smaller role to developed countries in explaining the oil price, see the variance decomposition in Figure 6.[1] This is not very surprising, given the importance of the US for developed countries, and given the importance of the US as the major consumer of oil.

Second, we find that augmenting the model with the US Federal Funds rate (the interest rate is ordered second to last to allow monetary policy to affect the oil price on impact) does not alter the main results.[2] Emerging countries are still the main drivers of the oil price, see Figure 7. The small effects of US monetary policy shocks on oil prices is also consistent with the findings in Anzuini et al. (2013).

However, although we find US monetary policy to have only a trivial effect on the oil price, we still find US monetary policy to explain a considerable share of economic activity in emerging countries on longer horizons. The latter finding is consistent with Canova (2005) and Maćkowiak (2007), analyzing the impact of US monetary policy shocks on some emerging countries.

[1]We now choose Germany to have a loading with one on the developed factor.
[2]Results are robust to alternative measures of US monetary policy (i.e., the US Treasury Bill rate) and to alternative orderings of the US monetary policy.

Figure 7: **Variance decomposition: US monetary policy included**

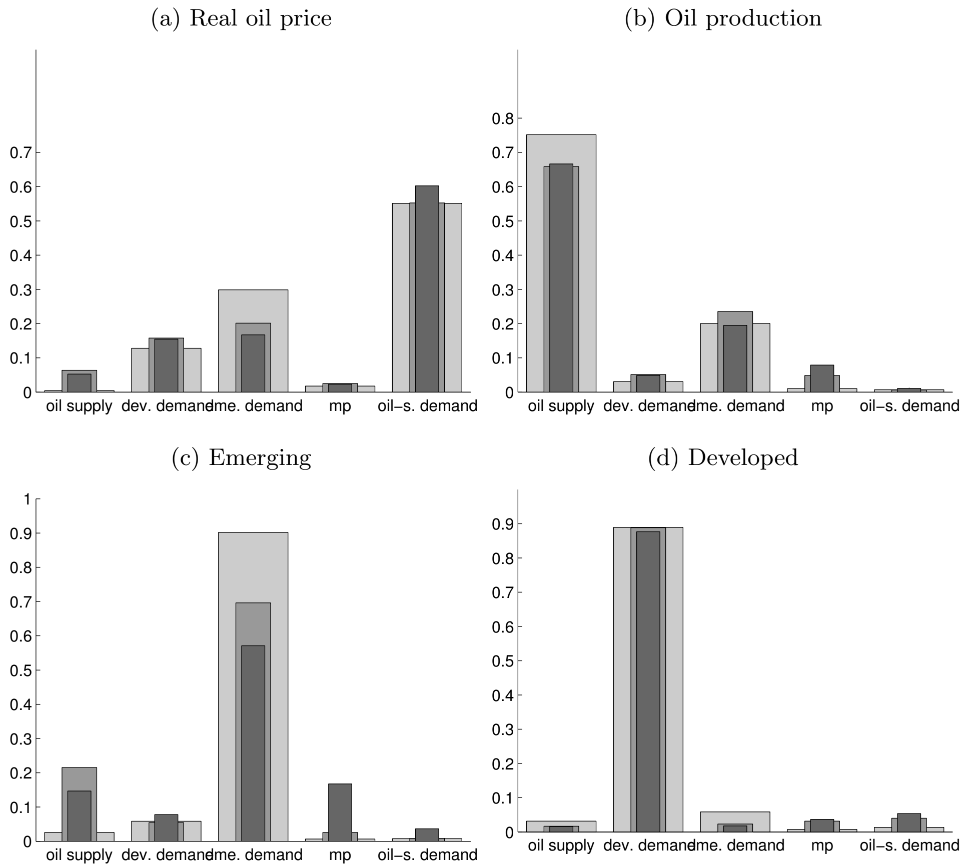

Note: The bars display the variance decomposition with respect to the shocks for horizons 4, 8 and 12 quarters. The widest bars correspond to the shorter horizon.

6.3 Recursive identification

An advantage with our identification strategy is that we can identify distinct demand shocks that affect both the developed and the emerging factors simultaneously. If simultaneity was unimportant, however, then the FAVAR model could be identified using a standard recursive identification strategy, ordering the developed-economy factor above the emerging-economy factor or vice versa. Identifying such a recursive model, however, yields as expected, very different results from our baseline model. In particular, now the activity factor that is ordered first will always explain more of the variation in the oil price than the activity factor that is ordered second. Thus, simultaneity matters, which a standard recursive identification strategy does not adequately capture. Despite this, recursive identification strategies nonetheless reveal that the emerging-country factor plays an important role, see Table 3. First, the emerging-country activity factor will always explain

Table 3: Variance decompositions: Alternative identification schemes

Horizon	**VAR A - Cholesky**		**VAR B - Cholesky**		**VAR C - Sign**	
& demand shock	dev.	eme.	dev.	eme.	dev.	eme.
Real price of oil						
4	0.02	0.58	0.34	0.22	0.13	0.43
8	0.03	0.50	0.31	0.17	0.13	0.36
12	0.04	0.43	0.31	0.14	0.15	0.30
Oil production						
4	0.00	0.41	0.12	0.25	0.03	0.36
8	0.00	0.58	0.13	0.41	0.03	0.55
12	0.02	0.52	0.11	0.39	0.02	0.50

Note: VAR A contains the variables: $\begin{bmatrix}\Delta prod & F^{eme} & F^{dev} & \Delta rpo\end{bmatrix}$*, and is identified using the Cholesky decomposition, i.e., a recursive ordering. VAR B is equal to VAR A, except from the ordering of the variables:* $\begin{bmatrix}\Delta prod & F^{dev} & F^{eme} & \Delta rpo\end{bmatrix}$*. VAR C is identical to our main model, except that we only enforce the sign restriction to hold for one period.*

relatively more of the variation in the oil price than the developed-country activity factor, irrespective of where it is ordered. That is, when the emerging-country factor is ordered first, it explains nearly twice as much of the variation in the real oil price compared to when the developed-country factor is ordered first. Similarly, when the emerging-country factor is ordered last, it explains more than twice as much of the variation in the real oil price compared to what the developed-country factor does when ordered last. Second, the emerging-country factor always explain more of the variance in oil production than the developed-country factor, independent of the ordering of the emerging-country and the developed-country factors.

6.4 Set identification and alternative estimation strategies

Neither the frequentist nor the Bayesian theory is conclusive on how confidence bands should be presented when structural disturbances are generated from sign restrictions. Moon et al. (2011) analyze the problem for classical VARs, in a frequentist setting. This do not apply directly to us, as we estimate a FAVAR where uncertainty in both parameter and factor estimates should be taken into account. Further, the sign restrictions we employ are not those that are studied in Moon et al. (2011).

Generally, our sign restrictions are very informative, meaning that the set of admissible impulse responses is rather narrow. This is illustrated in Figure 8. The figure reports the set of admissible impulse responses for 5000 draws based on the point estimates of β and Σ.

Importantly, variance decompositions of the set of impulse responses reported in Figure 8 are highly conclusive and supportive of the results already reported. Demand shocks originating in developed economies never explain more than 0.15 percent of the variation in global oil production. At horizons exceeding 5 quarters ahead, demand shocks originating in emerging economies never explain less than around 40 percent of the same variation. Further, for horizons up to 12 quarters ahead, demand shocks originating in emerging economies are almost unambiguously more important than demand shocks originating in developed economies in explaining the variation in the real price of oil. For example, at horizon 4, the set of variance decompositions attributed to emerging demand is in the range 20 to 60 percent, while the same range for developed demand is 0.02 to 30 percent.

We stress that for a particular (bootstrapped) parameter set, β and Σ, any draw of the sign-restricted structural disturbances will not cause the impulse responses in columns one and four in Figure 8 to change. Only impulse responses in columns two and three are affected. Thus, these are set identified, while the former are point identified. That is, the confidence bands reported in columns one and four of the impulse response figure in the main text are purely a function of parameter and factor uncertainty.

Figure 8: **Impulse responses: Point estimates and admissible sets**

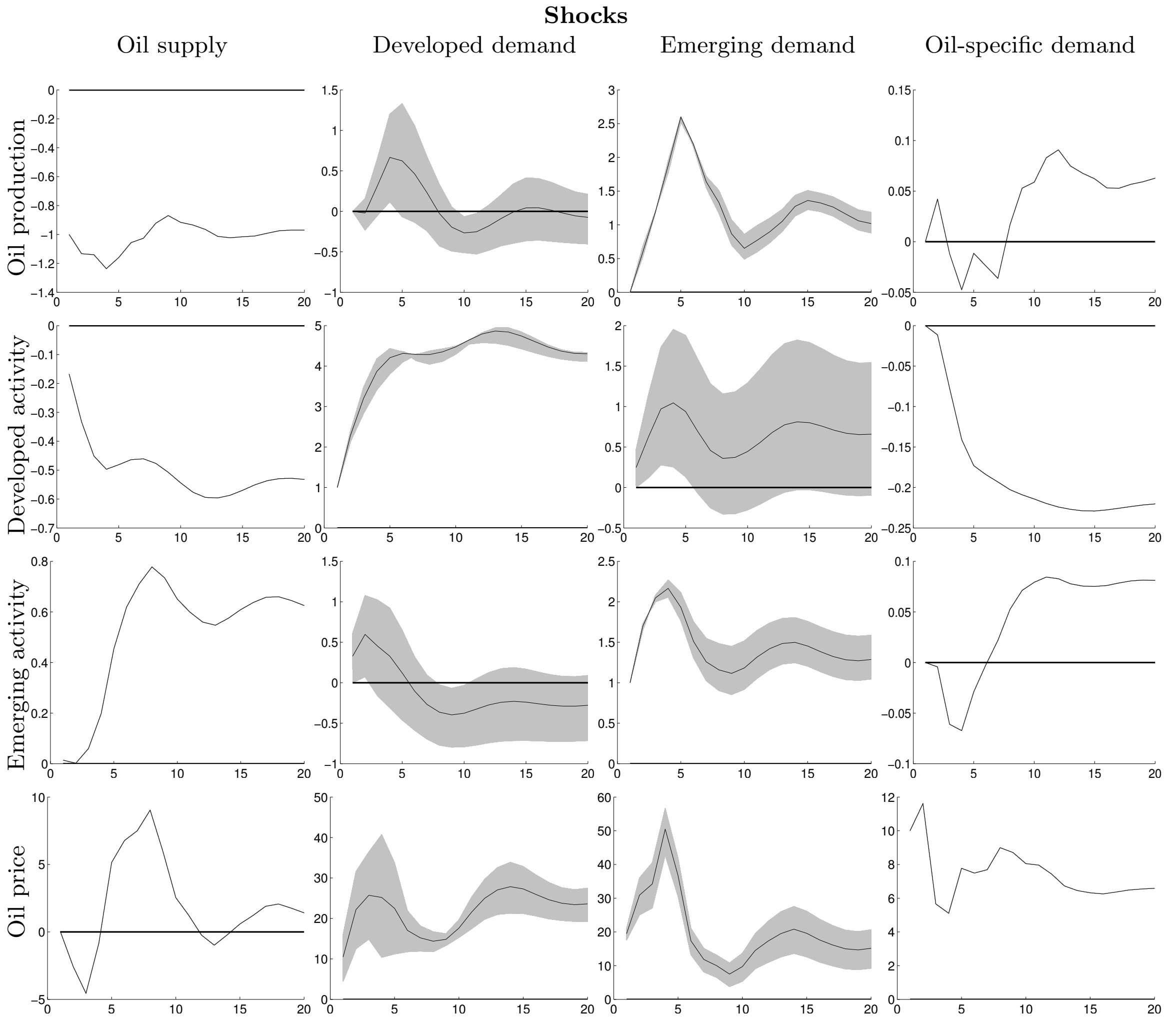

Note: The responses are displayed in levels of the variables. The developed-country and emerging-country demand shocks are normalized to increase activity in developed and emerging countries by one percent, respectively. To facilitate comparison with earlier studies, the oil supply shock is normalized to decrease oil production by one percent, while the oil-specific demand shock is normalized to increase the real oil price by 10 percent. The grey shaded areas represent the whole set of admissible impulse responses simulated, while the solid lines are the point estimates.

Figure 9: **Impulse responses: Bayesian estimates**

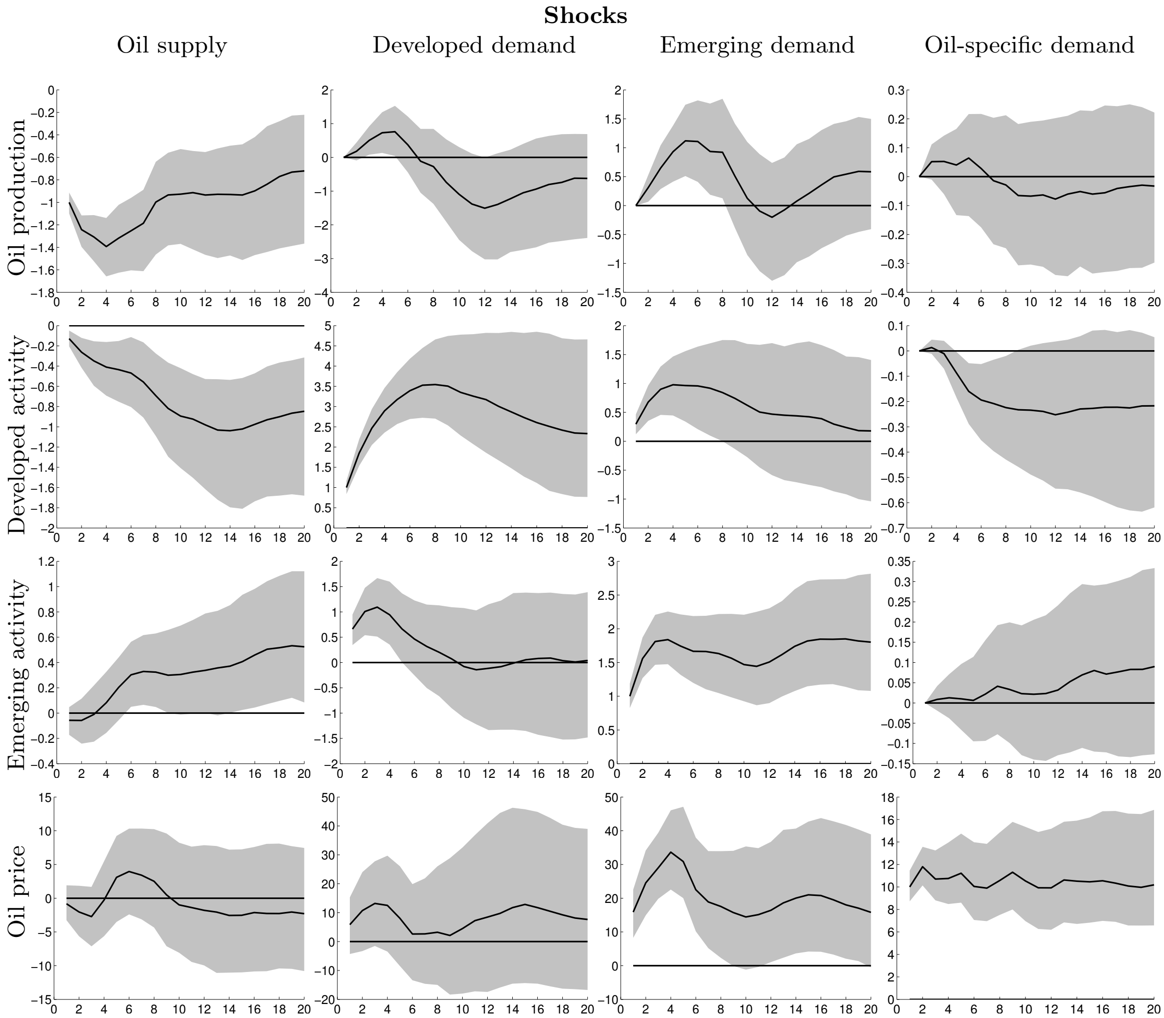

Note: The responses are displayed in levels of the variables. The developed-country and emerging-country demand shocks are normalized to increase activity in developed and emerging countries by one percent, respectively. To facilitate comparison with earlier studies, the oil supply shock is normalized to decrease oil production by one percent, while the oil-specific demand shock is normalized to increase the real oil price by 10 percent. The grey shaded areas represent 68 percent probability bands, while the black solid lines are the median estimates.

We have also estimated the FAVAR model using Bayesian techniques. Bernanke et al. (2005) show that a joint estimation of a related factor model, using likelihood-based Gibbs sampling techniques, yields very similar results to the ones applying a two-step procedure. This also holds in our application, when we specify and estimate the model as a Bayesian Dynamic Factor Model (BDFM). In particular, as shown in Figure 9, the impulse responses and the uncertainty estimates from the BDFM are very similar to the once presented in the main paper.

In this application we still prefer the FAVAR specification because of its simplicity, as estimating, identifying and simulating the FAVAR model requires only a few lines of Matlab code and takes less than 120 seconds. On the other hand, the BDFM is much more time consuming. Furthermore, the distributional assumptions (on, e.g., the error terms in the model) that are needed to estimate the BDFM using likelihood-based Gibbs sampling techniques are much more restrictive than in our FAVAR case, where the simulation is non-parametric.

6.5 Using the nominal price of oil

One concerns is that the relative weak response in the real price of oil to the developed demand shock is due to the fact that CPI also increases (since we deflate the price of oil with US CPI). This is not the case. In fact, when we run the model using the nominal price of oil, the model predicts an even larger role for the emerging-country demand shocks. See, Figures 10 and 11.

Figure 10: **Impulse responses: Using the nominal price of oil**

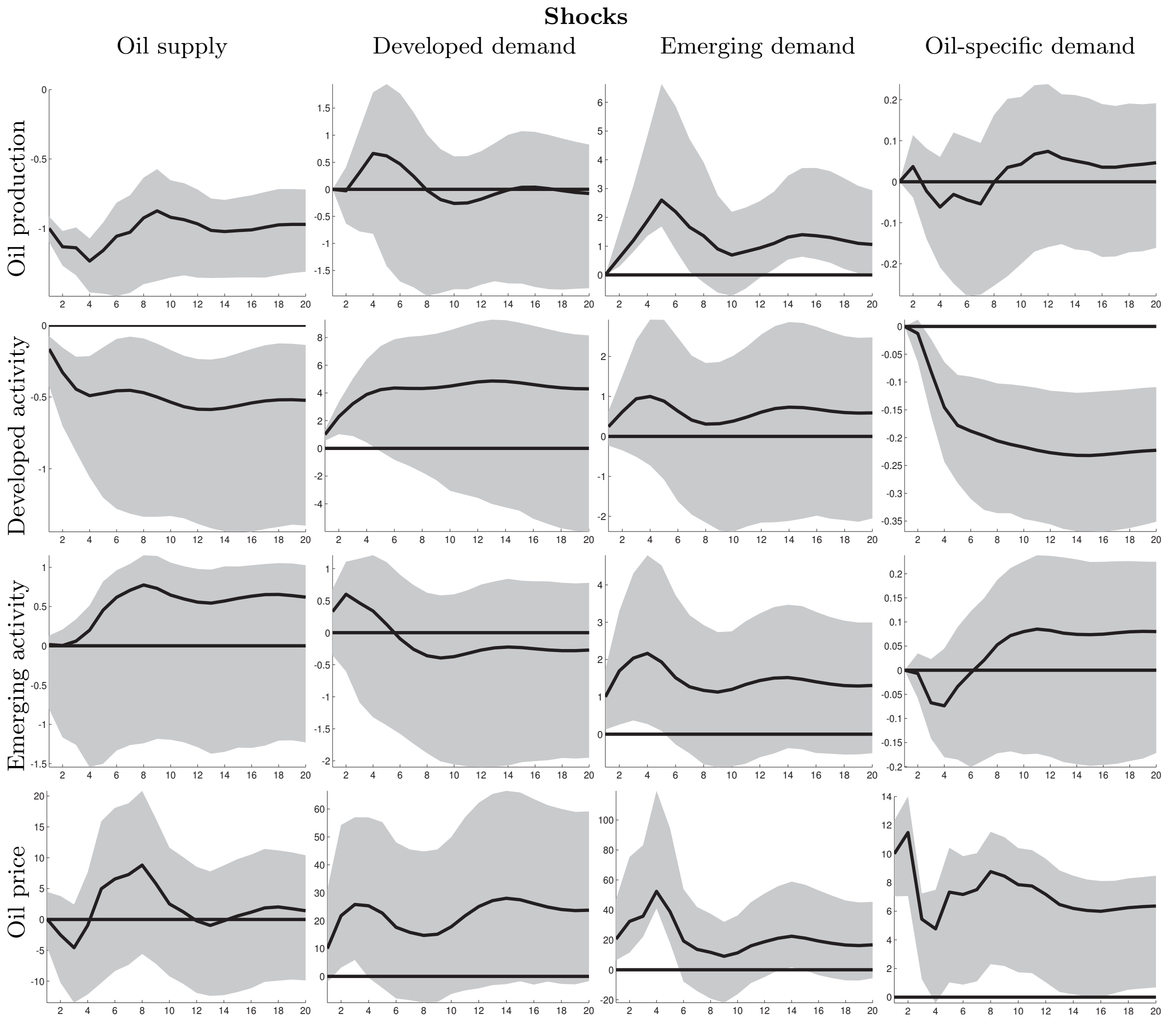

Note: The responses are displayed in levels of the variables. The developed-country and emerging-country demand shocks are normalized to increase activity in developed and emerging countries by one percent, respectively. To facilitate comparison with earlier studies, the oil supply shock is normalized to decrease oil production by one percent, while the oil-specific demand shock is normalized to increase the real oil price by 10 percent. The grey shaded areas represent 68 percent confidence bands, while the black solid lines are the point estimates.

Figure 11: **Variance decomposition: Using the nominal price of oil**

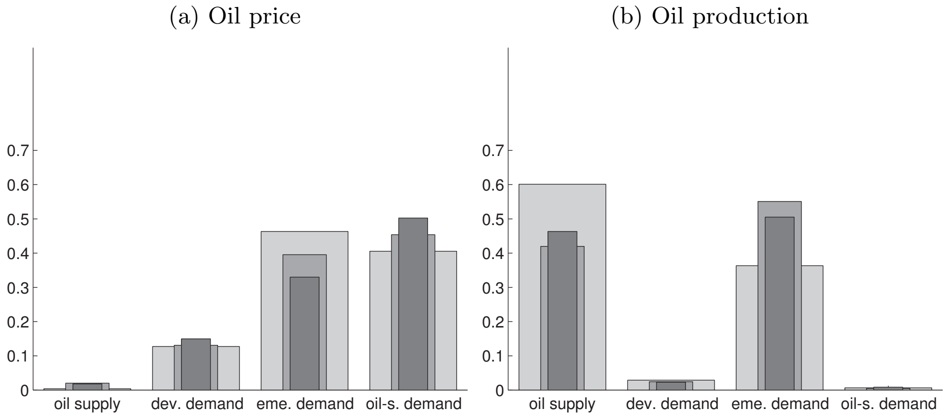

Note: The bars display the variance decomposition with respect to the shocks for horizons 4, 8 and 12 quarters. The widest bars correspond to the shorter horizon.

References

Anzuini, A., M. J. Lombardi, and P. Paganoa (2013). The impact of monetary policy shocks on commodity prices. *International Journal of Central Banking 9*, 119–144.

Bernanke, B., J. Boivin, and P. S. Eliasz (2005). Measuring the effects of monetary policy: A factor-augmented vector autoregressive (FAVAR) approach. *Quarterly Journal of Economics 120*(1), 387–422.

Canova, F. (2005). The transmission of US shocks to Latin America. *Journal of Applied Econometrics 20*(2), 229–251.

Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review 99*(3), 1053–1069.

Maćkowiak, B. (2007). External shocks, US monetary policy and macroeconomic fluctuations in emerging markets. *Journal of Monetary Economics 54*(8), 2512–2520.

Moon, H. R., F. Schorfheide, E. Granziera, and M. Lee (2011). Inference for VARs identified with sign restrictions. Working Papers 11-20, Federal Reserve Bank of Philadelphia.

Pesaran, M. H., T. Schuermann, and L. V. Smith (2009). Rejoinder to comments on forecasting economic and financial variables with global VARs. *International Journal of Forecasting 25*(4), 703–715.

Stock, J. H. and M. W. Watson (2002). Macroeconomic forecasting using diffusion indexes. *Journal of Business & Economic Statistics 20*(2), 147–162.

Centre for Applied Macro - and Petroleum economics (CAMP)
will bring together economists working on applied macroeconomic issues, with special emphasis on petroleum economics.

BI Norwegian Business School
Centre for Applied Macro - Petroleum economics (CAMP)
N-0442 Oslo

http://www.bi.edu/camp

CAMP Working Paper Series
ISSN: 1893-4811