# Oil Price Elasticities and Oil Price Fluctuations *

Dario Caldara † Matteo Iacoviello ‡ Michele Cavallo §

March 7, 2016

**Abstract**

We study the identification of oil shocks in a structural VAR (SVAR) model of the oil market. First, we show that the cross-equation restrictions of a SVARs impose a nonlinear relation between the short-run price elasticities of oil demand and oil supply. This relation implies that seemingly plausible restrictions on oil supply elasticity may map into implausible values of the oil demand elasticity, and viceversa. The selection of these elasticities is consequential for inference, in particular for the multipliers of oil prices on economic activity. Second, we propose an identification scheme that minimizes the distance between the SVAR-implied elasticities and targets constructed by a meta-analysis of relevant studies. Third, we use the identified SVAR to analyze sources and consequences of movements in oil prices, using a newly-constructed dataset on aggregate industrial production for both Advanced Economies and Emerging Economies.

We find that (i) oil supply shocks and oil-specific demand shocks are, on average, equal drivers of oil price fluctuations, while shocks to global demand play a minor role; (ii) a drop in oil prices driven by oil-market shocks boosts economic activity in advanced economies, while it depresses economic activity in emerging economies, thus accounting for the muted global economic effects of changes in oil prices; (iii) oil supply shocks caused about 70% of the decline in oil prices between 2014 and 2015.

*[EXTREMELY PRELIMINARY AND INCOMPLETE.] J Seymour and Lucas Husted provided outstanding research assistance. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of anyone else associated with the Federal Reserve System.

†Federal Reserve Board of Governors. Email: dario.caldara@frb.gov

‡Federal Reserve Board of Governors. Email: matteo.iacoviello@frb.gov

§Federal Reserve Board of Governors. Email: michele.cavallo@frb.gov

# 1 Introduction

This paper studies the identification of oil shocks in a structural VAR (SVAR) model of the oil market, using a new methodology and new data. Our goal is to provide a simple characterization of the two-way interaction between the oil market and the macroeconomy. Our main finding is that oil supply shocks are an important driver of oil prices, and that supply–driven increases in oil prices lead to decline in activity among advanced economies, and to an expansion in activity among emerging economies.

Our new methodology is as follows. First, we show that the cross-equation restrictions of a SVARs impose a nonlinear relation between the short-run price elasticities of oil demand and oil supply. This relation implies that seemingly plausible restrictions on oil supply elasticity may map into implausible values of the oil demand elasticity, and viceversa. The selection of these elasticities is consequential for inference, in particular for the multipliers of oil prices on economic activity. Second, we propose an identification scheme that minimizes the distance between the SVAR-implied elasticities and targets constructed by a meta-analysis of relevant studies. Third, we use the identified SVAR to analyze sources and consequences of movements in oil prices.

As for the data, our analysis builds on the classic paper by Kilian (2009) in setting up a VAR of the oil market that contains oil market variables – *crude oil production* and *real oil prices* – on the one hand, and three indicators of global economic activity on the other. Kilian constructed a new monthly measure of global real economic activity based on data for dry cargo bulk freight rates. We replace his measure with novel indicators: first, we construct two separate indicators of aggregate activity dating back to the 1980s, one based on *industrial production in emerging economies*, one based on *industrial production in advanced economies.*[1] Moreover, we include *industrial metals prices* as an additional endogenous variable in our VAR: including metals prices in the VAR improves our inference about the driving forces of oil prices. Metals prices are often viewed by practitioners and policy makers as early indicators of swings in economic activity and global risk sentiment (see e.g. Pindyck and Rotemberg (1990) and Cuddington and Jerrett (2008)). Along the same vein, Hamilton (2014) has more recently put forward the notion that concurrent swings in the prices for oil and metals could be indicative of changes in global demand. The analysis in Arezki and Blanchard (2014) has also exploited the idea that metals prices typically react to global activity even more than oil prices, while Bernanke (2016) has made use of the argument that prices of commodities such as copper can act indicators of underlying global growth as they are likely to respond to investors' perceptions of global demand, and not so much to changes in oil supply.

[1] The use of IP follows recent work by Aastveit et al. (2015) who extend Kilian's three-variable VAR to a FAVAR model that explicitly analyzes the contribution of demand from developed and emerging countries to movements in the real price of oil (from 1992 to 2009).

We find that (i) oil price fluctuations are primarily driven by oil supply and oil-specific demand shocks, whose quantitative contributions are, on average, equally important, whereas shocks to global demand play a minor role; (ii) a drop in oil prices driven by oil-market shocks, regardless of whether the shock originates from supply or demand, boosts economic activity in advanced economies, while it depresses activity in emerging economies. This finding thus helps understand the muted global economic effects of changes in oil prices; (iii) oil supply shocks explained about 70 percent of the decline in oil prices between 2014 and 2015.

Our approach has two attractive features relative to the existing literature. First, and different from Kilian (2009), the inclusion of a broader set of indicators of global demand enhances the contribution of business cycle fluctuations in explaining oil price movements. For instance, our identified global demand shocks explain about [33] % of the historical fluctuations in oil market variables, compared to [20] % using Kilian's 2009 preferred model. Additionally, our approach also yields responses to oil market shocks that largely match those of much more structured and detailed models.

The paper is structured as follows. Section 2 describes the identification and the estimation of the econometric model. Sections 3 and 4 show the main empirical findings. Section 5 explores robustness to alternative specifications of the VAR model. Section 6 concludes.

# 2 Measuring Global Demand for Oil

## 2.1 Industrial Production as a Coincident Indicator

An empirically compelling model of oil prices must take into account the drivers of oil demand in order to properly isolate exogenous oil market shocks from the endogenous response of oil markets to global economic developments. Yet global demand for oil is difficult to quantify and, as a result, the empirical literature uses several indicators. Rather than taking a stand on any particular indicator, we consider four different proxies that can be constructed at monthly frequency.

Our first two indicators are based on international industrial production data. We use industrial production for several reasons: (i) it is widely available across countries; (ii) it is historically a very reliable business cycle indicator;[2] (iii) it can capture demand for oil better than other indicators covering government and service sectors, since oil is an important factor of production in the industrial sector. Accordingly, our first two indicators are manufacturing industrial production in advanced economies ($ya$) and manufacturing industrial production in emerging market economies ($ye$). The split between advanced and emerging economies is dictated by the evidence that: (i) emerging economies produce relatively more energy–intensive goods than advanced economies, hence the impact on oil demand from

[2]See e.g. `http://www.nber.org/cycles/dec2008.html`.

Table 1: Correlations Among Global Activity Indicators
(1975M1–2015M12)

| | Pairwise Correlations | | | |
|---|---|---|---|---|
| | $ya_t$ | $ye_t$ | $m_t$ | $rea_t$ |
| $ya_t$ | 1.00 | $-0.25^{***}$ | $-0.35^{***}$ | $0.40^{***}$ |
| $ye_t$ | | 1.00 | $0.89^{***}$ | $0.40^{***}$ |
| $m_t$ | | | 1.00 | $0.29^{***}$ |
| $rea_t$ | | | | 1.00 |

NOTE: $^{*}$ $p < .10$, $^{**}$ $p < .05$, and $^{***}$ $p < .01$. $ya_t$ = Industrial production in advanced economies; $ye_t$ = Industrial production in emerging economies; $m_t$ = IMF Metal Price Index; $rea_t$ = Index of global activity based on.

a given change in industrial production from either group could be different; (ii) emerging economies are as a whole oil indipendent, while advanced economies are a whole net oil importers, hence the impact on aggregate activity of each group of a given oil shock could be different too.

While intuitive and readily available, industrial production may be a lagging indicator of global demand, as it may take time to adjust production due, for instance, to various types of adjustment costs. For this reason, we also consider two leading indicators of the world business cycle. The first indicator is the IMF Metal Price Index ($m$). Base metals – such as iron, copper, aluminum – are the lifeblood of global industrial activity: as a consequence, they are an ideal bellweather of shifts in cuurent and expected economic activity at the world level, especially in developing economies. The second indicator is the index of global real economic activity ($rea$) proposed in Kilian (2009) based on dry cargo, single voyage, ocean freight rates. This index is designed to capture shifts in the demand for industrial commodities and, similarly to the metal price index, should capture current and expected changes in world activity. All series are made stationary by removing a linear trend. Additional details about the data construction are provided in the Appendix.

Table 1 shows the contemporaneous correlations between these variables over the period 1985M1–2015M12, which corresponds to the sample period in our VAR analysis that follows. IP in advanced economies, $ya_t$, is negatively correlated with both $ye_t$ and $m_t$. The highest degree of comovement is between $ye_t$ and $m_t$. Taken together, these correlations suggest that metal prices might be a better indicator of economic activity in emerging economies than in advanced economies. Finally, $rea_t$ comoves with all other indicators, in line with Kilian (2009) interpretation that it is a broad indicator of global demand.

Table 2 displays the pairwise cross-correlations between oil prices ($p$) and the different global activity indicators at various leads and lags. The cross-correlation between $p_t$ and $ya_t$ is negative and becomes

Table 2: Cross-Correlations Between Oil Prices and Indicators of Global Activity

| Lag/Lead ($h$) | $ya_t$ | $ye_t$ | $m_t$ | $rea_t$ |
|---|---|---|---|---|
| $-12$ | $-0.30^{***}$ | $0.51^{***}$ | $0.61^{***}$ | $0.21^{***}$ |
| $-6$ | $-0.33^{***}$ | $0.63^{***}$ | $0.66^{***}$ | $0.25^{***}$ |
| $-3$ | $-0.34^{***}$ | $0.68^{***}$ | $0.70^{***}$ | $0.25^{***}$ |
| 0 | $-0.35^{***}$ | $0.71^{***}$ | $0.71^{***}$ | $0.26^{***}$ |
| 3 | $-0.40^{***}$ | $0.68^{***}$ | $0.68^{***}$ | $0.16^{***}$ |
| 6 | $-0.49^{***}$ | $0.63^{***}$ | $0.64^{***}$ | 0.03 |
| 12 | $-0.63^{***}$ | $0.53^{***}$ | $0.57^{***}$ | $-0.09^{*}$ |

NOTE: Cross-correlations between oil prices in month $t$ and the specified indicator of global activity in month $t+h$. $ya_t$ = Industrial production in advanced economies; $ye_t$ = Industrial production in emerging economies; $m_t$ = IMF Metal Price Index; $rea_t$ = Index of global activity based on Kilian (2009). $^{*}$ $p < .10$, $^{**}$ $p < .05$, and $^{***}$ $p < .01$.

larger at larger horizons, with oil prices leading negatively $ya$. Instead, the cross-correlation between $p_t$ and $ye_t$ and between $p_t$ and $m_t$ is positive and large, while between $p_t$ and $rea_t$ is positive but low. All told, the descriptive statistics suggest that the four indicators contain different information about global activity and have different relationships with oil prices.

## 2.2 Leading Indicators of Global Activity

To understand whether $m_t$ and $rea_t$ are leading indicator of global activity, we estimate (using OLS) the following forecasting regression:

$$y_{t+h} = \alpha + \beta_1 m_t + \beta_2 rea_t + \sum_{i=1}^{12} \delta_{1,i} y_{t-i} + \sum_{i=1}^{12} \delta_{2,i} m_{t-i} + \sum_{i=1}^{12} \delta_{3,i} rea_{t-i} + \epsilon_{t+h}, \tag{1}$$

where $h \geq 0$ is the forecast horizon, $y_t = \{ya_t, ye_t\}$, and $\epsilon_{t+h}$ is the forecast error.

To facilitate the comparison of the predictive power of metal prices and the real economic activity index, we report the standardized estimates of the coefficients $\beta_1$ and $\beta_2$. The results are tabulated in Table 3. Both $m_t$ and $rea_t$ have predictive power over advanced economies industrial output, as shown by panel (a); while $m_t$ predicts $ya_t$ for up to a year, $rea_t$ is informative only up to six months and has a lower predictive power than $m_t$. Both indicators are also very reliable predictors of near-term economic developments in emerging economies (panel (b)). At horizons up to one year, the estimated coefficients imply an economically and statistically significant positive relationship between $m$ and $rea$ and industrial production. For example, using a coefficient estimate of 1, an increase of 10 percent in

Table 3: Leading Indicators of Global Activity

| Forecast Horizon | $h = 1$ | $h = 6$ | $h = 12$ | $h = 24$ |
|---|---|---|---|---|
| *(a) Advanced Economies Industrial Production* | | | | |
| $\mathrm{m}_t$ | $0.18^{***}$ | $0.48^{**}$ | $0.41^{**}$ | 0.00 |
| | [3.23] | [2.40] | [2.14] | [0.02] |
| $\mathrm{rea}_t$ | $0.08^{***}$ | $0.20^{**}$ | 0.07 | $-0.22^{*}$ |
| | [2.87] | [1.98] | [0.96] | [1.67] |
| Adj. $R^2$ | 0.99 | 0.89 | 0.74 | 0.73 |
| *(b) Emerging Economies Industrial Production* | | | | |
| $\mathrm{m}_t$ | $0.56^{***}$ | $0.93^{***}$ | $1.02^{***}$ | $1.15^{***}$ |
| | [3.89] | [3.97] | [4.09] | [3.84] |
| $\mathrm{rea}_t$ | $0.25^{***}$ | $0.38^{***}$ | $0.27^{*}$ | $-0.03$ |
| | [3.30] | [4.12] | [1.91] | [0.23] |
| Adj. $R^2$ | 0.94 | 0.74 | 0.51 | 0.32 |

NOTE: The dependent variable in each specification is $ya_{t+h}$ in panel (a) and $ye_{t+h}$ in panel (b). Each row associated with $\mathrm{m}_t$ and $\mathrm{rea}_t$ reports the standardized estimates of the OLS coefficients associated with the indicator in month $t$: $m_t$ = IMF Metal Price Index; $rea_t$ = Index of global real economic activity based on Kilian (2009). $^{*}$ $p < .10$, $^{**}$ $p < .05$, and $^{***}$ $p < .01$. Each specification also includes a constant, 12 lags of the endogenous variables and 12 lags of $\mathrm{m}_t$ and $\mathrm{rea}_t$ (not reported). Absolute $t$-statistics reported in brackets are based on the heteroskedasticity- and autocorrelation-consistent asymptotic covariance matrix computed according to Newey and West (1987) with the automatic lag selection method of Newey and West (1994): * $p < .10$; ** $p < .05$; and *** $p < .01$.

metal prices in month $t$ implies a nearly 1.25 percent increase in emerging economies industrial output over the following six months.[3] At the two year horizon, the only significant relation is between $m_t$ and $ye_t$.

## 2.3 Global Activity as Predictor of Oil Prices

To better gauge whether our indices of global activity are good proxies for global demand for oil, we explore their relative roles in predicting oil prices. Specifically, we estimate the following forecasting

[3] Similarly, using a coefficient of 0.4, a 10 percent increase in $\mathrm{rea}_t$ predicts a 0.6 percent increase in emerging economies industrial output six-month ahead.

Table 4: Global Activity and Oil Prices

| Forecast Horizon | $h = 1$ | $h = 6$ | $h = 12$ | $h = 24$ |
|---|---|---|---|---|
| $\text{ya}_t$ | $-0.09$ | $-0.39$ | $-0.87$ | $-0.76$ |
| | [0.38] | [0.80] | [1.47] | [1.52] |
| $\text{ye}_t$ | 0.22** | 0.77*** | 0.36** | 0.16 |
| | [2.11] | [3.68] | [2.48] | [0.77] |
| $\text{m}_t$ | 0.68*** | 0.56** | 0.46 | 0.75*** |
| | [4.05] | [1.98] | [1.35] | [2.92] |
| $\text{rea}_t$ | 0.10 | 0.07 | $-0.03$ | $-0.04$ |
| | [0.98] | [0.77] | [0.23] | [0.20] |
| Adj. $R^2$ | 0.89 | 0.66 | 0.63 | 0.42 |

NOTE: The dependent variable in each specification is $p_{t+h}$. Each row associated with a global activity indicator reports the standardized estimates of the OLS coefficients associated with the indicator in month $t$:ya = Manufacturing industrial production in advanced economies; ye = Manufacturing industrial production in emerging economies; m = IMF Metal Price Index; bea = Index of global real economic activity based on Kilian (2009). $^{*}$ $p < .10$, $^{**}$ $p < .05$, and $^{***}$ $p < .01$. Each specification also includes a constant, 12 lags of $p_t$ and 12 lags of $\text{GA}_t$ (not reported). Absolute $t$-statistics reported in brackets are based on the heteroskedasticity- and autocorrelation-consistent asymptotic covariance matrix computed according to Newey and West (1987) with the automatic lag selection method of Newey and West (1994): * $p < .10$; ** $p < .05$; and *** $p < .01$.

regression:

$$\begin{aligned} p_{t+h} = \alpha + \beta_1 ya_t + \beta_2 ye_t + \beta_3 m_t + \beta_4 rea_t + \\ + \sum_{i=1}^{12} \rho_i p_{t-i} + \sum_{i=1}^{12} \left(\gamma_{1,i} ya_{t-i} + \gamma_{2,i} ye_{t-i} + \gamma_{3,i} m_{t-i} + \gamma_{4,i} rea_{t-i}\right) + \epsilon_{t+h}, \end{aligned} \tag{2}$$

where $h \geq 0$ is the forecast horizon, $\epsilon_{t+h}$ is the forecast error, and the $\rho$ and $\gamma$ are coefficients that control for past observations. As before, our interest is in understanding the different predictive power of the various economic indicators for oil prices, once the past state of the economy is conditioned for, as measured by the standardized estimates of the coefficients $\beta_1$, $\beta_2$, $\beta_3$, and $\beta_4$.

The results are tabulated in Table 4. According to Table 4, the predictive content is uneven across economic indicators. The coefficients on $ya_t$ are all negative but not statistically significant. The two indicators that significantly predict oil prices are $ye_t$ and $m_t$. The latter is the only indicator that has predictive power beyond one year. The coefficients on $rea_t$ are economically small and not statistically significant. This result motivates our choice of including $m_t$ in our baseline VAR model.

The results from the above forecasting exercises are instructive for two reasons. First, they underscore the fact that information content of various measures of global activity for oil prices is mixed, and hence it is sensible to include a set of indicators in the VAR model rather than relying on a single one. Second, since the forecasting regression is completely silent on the causal relationship between global activity and oil prices, it is not clear what is the relative importance of shocks to global activity and to the oil market in shaping the correlations between these two blocks of variables. To answer this question empirically, however, one has to take a stand on the joint identification of these two types of shocks, the subject of the remainder of the paper.

# 3 Identifying Oil Market and Global Activity Shocks

## 3.1 Overview

The two building blocks of our estimation strategy are (i) to use a broad set of indicators to capture global demand for oil and (ii) to exploit outside information contained in several studies that have measured the price elasticity of the demand and the supply of oil to discipline the modelling of the oil market in a structural VAR model.

Suppose that the true economic structure describing the oil market and its relationship with the world economy is given by

$$\mathbf{A}\mathbf{X}_t = \sum_{j=1}^{p} \boldsymbol{\alpha}_j \mathbf{X}_{t-j} + \mathbf{u}_t \tag{3}$$

where $E\left[\mathbf{u}_t \mathbf{u}_t'\right] = \boldsymbol{\Sigma}_u$, a diagonal matrix, $\mathbf{X}$ is a vector of macroeconomic variables including oil production and oil prices, $\mathbf{u}_t$ is a vector describing the structural shocks, and $p$ is the number of lags in the VAR model. Without loss of generality, we normalize the elements on the main diagonal of $\mathbf{A}$ to 1. The reduced–form representation for $\mathbf{X}_t$ is:

$$\mathbf{X}_t = \sum_{j=1}^{p} \boldsymbol{\gamma}_j \mathbf{X}_{t-j} + \boldsymbol{\varepsilon}_t \tag{4}$$

where the reduced-form residuals $\boldsymbol{\varepsilon}_t$ are related to the structural shocks $\mathbf{u}_t$ by the following:

$$\boldsymbol{\varepsilon}_t = \mathbf{B}\mathbf{u}_t, \tag{5}$$

$$\boldsymbol{\Sigma}_\varepsilon = \mathbf{B}\boldsymbol{\Sigma}_u\mathbf{B}', \tag{6}$$

where $\mathbf{B} = \mathbf{A}^{-1}$, so that alternatively $\mathbf{u}_t = \mathbf{A}\boldsymbol{\varepsilon}_t$. Estimation of the reduced–form VAR allows recovering a consistent estimate of the $n(n+1)/2$ elements of $E[\boldsymbol{\varepsilon}_t\boldsymbol{\varepsilon}_t'] = \boldsymbol{\Sigma}_\varepsilon$. However, in order the recover the $n^2$ unknown elements of $\mathbf{B}$ and $\boldsymbol{\Sigma}_u$, additional identifying assumptions must be made.

To discuss our identification strategy, however, it is useful to list the set of endogenous variables. The benchmark monthly VAR specification consists of five endogenous variables ordered as follows: (1) log world crude oil production, $q_t$; (2) the log of advanced economies industrial production, $ya_t$; (3) the log of emerging economies industrial production, $ye_t$; (4) the log oil prices, $(p_t)$; and (5) the log of the IMF metal price index $(m_t)$.

## 3.2 Identification

We divide the discussion of the identification problem in two steps. First, we discuss the structure we impose on the oil market. Second, we discuss the interaction between the oil market and global activity.

**The Oil Market** Our structural VAR model specifies an oil supply and an oil demand equation, corresponding to the equations associated with the oil supply and oil demand shocks. After controlling for their own lags and for other determinants of oil demand and supply, these equations define a short-run oil supply and a short-run oil demand curve. If we consider a simple bivariate system in $q$ and $p$, these curves are defined as follows:

$$q_t = \eta_S p_t + u_{s,t}, \tag{7}$$

$$q_t = \eta_D p_t + u_{d,t}, \tag{8}$$

where $\eta_S$ and $\eta_D$ are the short-run price elasticity of oil supply and oil demand, respectively. For ease of interpretation, and in line with standard demand theory, we impose the restriction that the price elasticity of supply (which is ordinarily positive) is greater than the price elasticity of demand (which is ordinarily negative), so that $\eta_S > \eta_D$: this restriction rules out solutions that simultaneoulsy assume that (i) increases in demand lead to a reduction of prices; and (ii) increases in supply lead to an increase in prices.[4]

Given values for the elements of $\boldsymbol{\Sigma}_\varepsilon$, the covariance structure of the VAR imposes a unique relationship between $\eta_S$ and $\eta_D$. The red line in Figure 1 plots this relationship computed by setting $\boldsymbol{\Sigma}_\varepsilon$ at its

[4]The assumption that $\eta_S > \eta_D$ is a normalization that picks the economically most interesting solution (out of two) when solving the matrix quadratic equation that allows to recover $\eta_D$ and $\eta_S$ given the covariance structure of VAR residuals. In principle, if $\eta_S = \eta_S^*$, $\eta_D = \eta_D^*$, $\sigma_S = \sigma_S^*$, $\sigma_D = \sigma_D^*$ are a solution to the matrix quadratic equation, there is also another solution – corresponding to switching the labels on the demand and supply curves – where $\eta_S = \eta_D^*$, $\eta_D = \eta_S^*$, $\sigma_S = \sigma_D^*$ and $\sigma_D = \sigma_S^*$. Imposing that $\eta_S > \eta_D$ allows picking one solution only.

OLS estimate:[5] the larger the supply elasticity, the smaller the demand elasticity in absolute value.

The relationship between demand and supply elasticities implied by the VAR works as follows. By construction, the identification exercise in the VAR attempts to to match the volatility of oil prices and quantities $p$ and $q$ as well as their covariance, which in the data is close to zero. These three targets can be hit in principle with four parameters: the variance of the oil demand shocks, the variance of oil supply shocks, the oil demand elasticity, and the oil supply elasticity. Given shock variances, the oil demand elasticity restricts the oil supply elasticity, and vice versa.

Panels A and B in figure 2 plots two combinations of slopes of demand and supply schedules in the oil market that achieve the target, given the target values of the entire covariance matrix $\mathbf{\Sigma}_\varepsilon$. Panel A summarizes our benchmark identification scheme. For given slopes of oil demand and oil supply curve, the right mix of demand and supply shocks generates the "right" volatilities of prices and quantities and zero correlation between them. Panel B illustrates another observationally equivalent solution: more elastic demand relative to the benchmark, and more inelastic supply. In this case, large demand shocks move oil prices but have negligible effects on oil production, and large supply shocks move oil production but have negligible effects on oil prices. While the specification in B is observationally equivalent, it creates a disconnect between movements in quantitites and movements in prices in the oil market, not only unconditionally, but also conditional on each separate shock.

The middle panels (C and D) explain why the VAR restrictions rule out the possibility of both inelastic supply and inelastic demand. When both demand and supply are inelastic, either both demand and supply shocks are small, and the VAR matches the volatility of prices but underestimates the volatility of oil production (panel C). Or both shocks are large, and the VAR can match the volatility of production but overpredicts the volatility of oil prices (panel D). Finally, the panels at the bottom (E and F) illustrate why the VAR restrictions also rule out both elastic supply and elastic demand. Either both shocks are large, and the VAR can match the volatility of prices but overpredicts the volatility of oil production (panel E). Or both shocks are small, and the VAR can match the volatility of production but underestimates the volatility of oil prices (panel F).[6]

The two observationally equivalent possibilities of panels A and B and all their convex combinations suggest that various combinations of oil market structures can deliver volatilities and comovements between $p_t$ and $q_t$ consistent with the data. That is, these two pairs of elasticities are part of the loci

[5] We extract the elements of $\mathbf{\Sigma}_\varepsilon$ related to oil prices and quantities from the 5 equation VAR. The estimation of the bivariate system returns a similar covariance matrix of residuals.

[6] One can always find a solution where both demand and supply elasticity have the "right" sign, and the estimated variances of demand and supply shocks match any desired correlation between prices and production. There are however cases in which the solution can call for: (i) an upward sloping demand curve if the correlation between prices and production is large and positive and the supply elasticity is assumed to be "large"; (ii) a downward sloping supply curve if the correlation between prices and production is large and negative and the demand elasticity is assumed to be "large" in absolute value.

graphically represented by the curve plotted in Figure 1. Along this curve, the likelihood of the VAR is constant, hence there is no statistical basis for preferring one point over another. To choose the "right" elasticities, we need to rely on extra model information.

Our identification proceeds in two steps. First, we search the empirical literature on oil price elasticities and compile a list of studies that estimate short-run demand and supply elasticities. We list the papers and the associated elasticities in Table A.4. The consensus in the literature, as summarized by the median elasticity, is that the demand elasticity is $-0.16$, while the supply elasticity is 0.09. This pair of elasticities is summarized by the blue circle in Figure 1. Importantly, the blue circle does not lie on the red line, which means they are not admitted by our SVAR.

Second, to select a pair of admissible elasticities, our identification strategy selects the VAR restriction that minimizes the Euclidean distance between the VAR-admissible elasticities and the elasticities obtained by the meta-analysis. Without loss of generality, let us assume that we identify the VAR by imposing the restriction on $\eta_S$. If we denote $\eta_D$ as a function of $\eta_S$ and the covariance matrix of residuals, $\eta_D\left(\eta_S; \mathbf{\Sigma}_\varepsilon\right)$, our identification strategy solves the following problem:[7]

$$\min_{\eta_S} \begin{bmatrix} \eta_S - \eta_S^* \\ \eta_D\left(\eta_S; \mathbf{\Sigma}_\varepsilon\right) - \eta_D^* \end{bmatrix} V^{-1} \begin{bmatrix} \eta_S - \eta_S^* \\ \eta_D\left(\eta_S; \mathbf{\Sigma}_\varepsilon\right) - \eta_D^* \end{bmatrix}, \tag{9}$$

where $\eta_S^*$ and $\eta_D^*$ are the target supply and demand elasticities, respectively, and $V$ is a matrix of weights.[8] Our identification scheme selects the green circle in Figure 1 corresponding to the point on the curve that is as close as possible to the blue circle. Such identification yields, for $\mathbf{\Sigma}_\varepsilon$ evaluated at its OLS estimate, an oil supply elasticity of 0.07, and an oil demand elasticity of $-0.14$. The bottom right panel in Figure 2 illustrates the significance of these elasticities for gauging price movements in the oil market, by plotting the oil demand and oil supply curves implied by our estimated VAR. The elasticities imply that an exogenous decline in world crude oil supply of 1 percent (approximately 0.8 million barrels per day) should lead, over a 1-month horizon, to an increase in oil prices of 5 percent and to an offsetting second-round increase in oil supply of about 0.25 percent (equivalent to 0.2 million barrels per day).

**The Oil Market and Global Activity** With a modelling of the oil market at hand, the next step is to model its interaction with the macroeconomy. Current and expected developments in the global economy can lead to shifts in the demand for oil, and hence affect the identification of exogenous oil

[7]Alternatively, we could have parametrized the problem either by imposing the restriction on $\eta_D$, in which case $\eta_S\left(\eta_D; \mathbf{\Sigma}_\varepsilon\right)$, or, following Uhlig (2005) and Rubio-Ramírez et al. (2010) among many, in terms of an orthonormal matrix $Q$, so that $\eta_S\left(\Omega; \mathbf{\Sigma}_\varepsilon\right)$ and $\eta_D\left(\Omega; \mathbf{\Sigma}_\varepsilon\right)$.

[8]V is a diagonal matrix whose elements are the variance of the estimated elasticities reported in Table A.4

market specific shocks. In turn, such shocks can lead to changes in economic conditions, which is typically what researchers are interested in. The following five equations decribe the joint modelling of the oil market and the global economy:[9]

$$
\begin{aligned}
q_t &= \eta_S p_t + u_{s,t}, && (10)\\
ya_t &= \nu_Q q_t + u_{ya,t} && (11)\\
ye_t &= \mu_Q q_t + \mu_A ya_t + u_{ye,t} && (12)\\
q_t &= \eta_A ya_t + \eta_E ye_t + \eta_D p_t + u_{d,t} && (13)\\
m_t &= \psi_Q q_t + \psi_A ya_t + \psi_E ye_t + \psi_P p_t + u_{m,t}. && (14)
\end{aligned}
$$

Equations (10)-(14) summarize the parametric restrictions we impose on elements of the matrix $\mathbf{A}$. These parametric restrictions are guided, where possible, by economic theory. Equation (10) describes the oil supply schedule and we assume that oil production $q_t$ responds contemporaneously only to changes in oil prices. Importantly, the exclusion restrictions imposed on the oil supply equation do not rule out the possibility that movements in global activity have a contemporaneous effect on oil production. Global activity can affect production by inducing changes in oil prices.

Equation (11) is the behavioural equation for advanced economies activity. We assume that $ya_t$ responds only to oil production within the period. Similarly, Equation (12) is the behavioural equation for emerging economies activity. We assume that $ye_t$ responds to $ya_t$ and to oil production within the period. We allow both $ya_t$ and $ye_t$ to react contemporaneously to changes in $q_t$ because oil is an input in the production of manifacturing goods. $ye_t$ reacts contemporaneously to $ye_t$ because exports to advanced economies are an important component of aggregate demand in emerging economies.

Equation (13) describes an oil demand schedule in which oil demand is allowed to respond contemporaneously to $ya_t$ and $ye_t$ and to changes in oil prices. Importantly, the demand price elasticity $\eta_D$ is defined as the change in desired demand $q$ for a given change in oil prices $p$, holding global activity constant within the period.

Finally, equation (14) describes the behavioral equation for metals prices, which are allowed to respond contemporaneously to all variables in the system. The idea is that $u_{m,t}$ is a shock that mainly captures news about global activity, but also current developments in global activity that are not adequately accounted for by IP indices.

To summarize, in matrix notation, our set of restrictions assumes that the relationship between oil

[9] For expositional convenience we omit the lagged terms, which instead appear in equation (15).

and the macroeconomy is described by the following set of equations:

$$\underbrace{\begin{bmatrix} 1 & 0 & 0 & -\eta_S & \mathbf{0} \\ -\nu_Q & 1 & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ -\mu_Q & -\mu_A & 1 & \mathbf{0} & \mathbf{0} \\ 1 & -\eta_A & -\eta_E & -\eta_D & \mathbf{0} \\ -\psi_Q & -\psi_A & -\psi_E & \psi_P & 1 \end{bmatrix}}_{\mathbf{A}} \begin{bmatrix} q_t \\ ya_t \\ ye_t \\ p_t \\ m_t \end{bmatrix} = \sum_{j=1}^{p} \boldsymbol{\alpha}_j \mathbf{X}_{t-j} + \begin{bmatrix} u_{s,t} \\ u_{ya,t} \\ u_{ye,t} \\ u_{d,t} \\ u_{m,t} \end{bmatrix}, \tag{15}$$

where $\nu_Q$ and $\mu_Q$ can be interpreted as the short–run elasticities of economic activity to oil production, and $\eta_A, \eta_E$ denotes the elasticity of oil demand to economic activity.

The restrictions above – shown in bold – in the matrix **A** above allow to uniquely identify the structural parameters $(\nu_Q, \mu_Q, \mu_A, \eta_A, \eta_E, \eta_D, \psi_Q, \psi_A, \psi_E, \psi_P, \mathbf{\Sigma}_u)$ given information from $\mathbf{\Sigma}_\varepsilon$, the variance–covariance matrix of the reduced–form VAR residuals.

**Comparison with Existing Studies.** We compare our identification strategy with two prominent papers in the literature.

Baumeister and Hamilton (2015) (BH) offers an alternative approach to "filling in" the elements of the matrix **A** in a VAR model with oil prices, world IP, oil production, and oil inventories.

BH specify a number of exclusion restrictions for some of the elements of **A** which are less than the number that is sufficient to achieve exact identification. Next, they specify priors for all the remaining elements of the **A** matrix in order to help distinguishing between alternative representations of the **A** matrix which would otherwise all yield the same likelihood. While this approach is tempting, it is not immune from risks. First, by specifying priors for a large number of elements of **A**, the researcher easily ends up ceding to the temptations of assuming his own conclusions. To give an example, BH impose that economic activity does not depend on oil production ($\nu_Q = \mu_Q = 0$, in our notation), but depend negatively on oil prices, with a coefficient which is bounded between 0 and $-1$. If one thinks of the equations behind a VAR as coming from a structural model, there should be no reason to assume that economic activity $y$ depends directly on prices but not production. Obviously, to the extent that all shocks move simultaneously oil prices and economic activity, the reduced-form relationship between economic activity and oil prices will involve a non-zero response of activity to prices, but this response cannot form the basis for a prior on the very parameter that one is interested in estimating. Second, coming up with a full set of credible short-run identifying restrictions is difficult. Third, the more priors one specifies, the higher is the likelihood that the VAR is overidentified, and thus incapable of explaining the full variance–covariance structure of the data.

The paper by Kilian (2009) is, under some specific assumptions, a special case of our VAR. Kilian works with a three equation VAR where his measure of economic activity $y$ is based on data for dry cargo bulk freight rates. Compared to his approach, we use both IP and metals prices to measure economic activity $y$. Additionally, we do not restrict $\widehat{\mathbf{e}}_S$ to be equal to zero, which is Kilian's working assumption, which implicitly implies a very large oil demand elasticity.

# 4 Are All Shocks Alike?

## 4.1 Overview

This section first presents the impulse responses implied by our Structural VAR. We then examine the historical contribution of the identified shocks to oil market variables and economic activity. To estimate the model, we employ Bayesian estimation techniques. In particular, we impose a Minnesota prior on the reduced-form VAR parameters by using dummy observations (??, del) with hyper-parameters $\lambda = [1, 2, 1, 3, 3]$. The resulting specification, which includes a constant, is estimated over the 1985:M1–2015:M12 period using twelve lags of the endogenous variables.[10]

## 4.2 Impulse Responses

The solid lines in the left column of Figure 3 show the median impulse responses of the five endogenous variables to a one–standard deviation oil supply shock, while the shaded bands represent the corresponding 90-percent (light blue) and 68-percent (dark blue) pointwise credible bands. An unanticipated disruption in oil supply reduces production by about 0.75 percent and elicits a persistent increase in oil price, which rise by 5 percent on impact and remains elevated thereafter.

On the activity side, the response of activity between advanced and emerging economies is markedly different. Industrial production in advanced economies declines gradually, bottoming out at 0.3 percent two and a half years after the shock. In contrast, industrial production in emerging economies raises after the shock, peaking after six months 0.15 percent above baseline. The heterogeneous responses of advanced and emerging economies to oil supply shocks are consistent with the fact that, on average, our group of emerging economies are, on net, oil independent, while our group of advanced economies are oil dependent. As shown by Iacoviello (2016), following a surprise increase in oil prices, real activity and private consumption increases in exporting countries, while it declines in oil importing countries.

[10]We use the first year of the sample as a training sample for the Minnesota prior. All the results reported in the paper are based on 50,000 draws from the posterior distribution of the structural parameters, where the first 10,000 draws were used as a burn-in period.

The right column of Figure 3 shows the responses to an oil demand shock. The shock leads to an increase in oil prices of 6 percent and induces a rise in oil production of about 0.5 percent. The near-term response of activity in advanced and emerging economies is similar. Industrial production increases mildly in both groups of economies for six months. Thereafter, real activity contracts in advanced economies while remains elevated in emerging economies, even though the responses are economically small and only marginally significant.[11]

Figure 4 traces out the effects of the three global activity shocks. The left column plots the responses to a shock to activity in the advanced economies; the middle column plots the responses to a shock to activity in emerging economies activity shock; and the right column plots the responses to a metal price shock. All shocks have the characteristics of a business–cycle, demand–driven increase in oil price, and lead to a rise in oil prices and oil production.

In line with results from the forecasting regression discussed in Section 2.3, positive shocks to activity in emerging economies and positive shocks to metal prices induce a persistent increase in oil prices. An increase in advanced economies activity induces a mild and short-lived increase in oil prices. Nonetheless, in constrast to the forecasting regressions—which predict a decline in oil price following an increase in advanced economies activity—the Structral VAR attributes the negative correlation between oil price and advanced economies industrial production mostly to oil supply shocks, with all three global activity shocks generating a positive co-movement between these two variables.

Finally, the responses of industrial production in both advanced and emerging economies to a shock to the metal price index support the view that metal prices are a good leading indicator of global activity.

## 4.3 Historical Decomposition

This section presents the historical decomposition of the actual paths of the VAR variables that is attributable to the oil market and global activity shocks. We first give an overview over the 1986–2015 period presenting the data at an annual frequency. We then zoom in three important episodes in the sample involving large changes in the price of oil. The first episode is centered around the Asian financial crisis; the second episodes focuses on the period of the global financial crisis; the third episode corresponds to the dramatic fall in the price of oil that started in July 2014 and lasts through the end of the sample. [12]

[11] The response of the metal price index to both shocks is positive and persistent. The immediate jump in metal prices followed by a a response with a profile similar to industrial production—in particular in emerging economies—is in line with our mantained assumption that metal prices are a leading indicator of global activity.

[12] We compute the historical variance decomposition at the OLS estimates of the reduced-form parameters. The sample used for the estimation includes actual data and the dummy observations used to implement the Minnesota prior (see Appendix **??** for details on the prior specification).

Figure 5 shows the decompositions across the whole sample for the price of oil, oil production, and industrial production in both advanced economies and emerging economies. The solid black lines depict the actual paths of the VAR variables, whereas the areas coloured in green, blue, cyan, red, and orange represent the contributions to the actual paths made by metals prices, advanced economies activity, emerging economies activity, oil supply, and oil demand shocks, respectively.

Overall, our results imply that, across the whole sample period, fluctuations in the price of oil and in oil production were mostly determined by shocks to the oil market, and in particular by supply shocks. Nonetheless, global activity contributed significantly to shape the up and downs of oil market variables. Shocks to emerging economies and to metal prices played a prominent role, while shocks to advanced economies activity played a modest role. Shocks emanating from the oil market contributed to fluctuations in advanced and emerging economies real activity. However, the significance of oil shocks, and in particular of oil supply shocks, has varied considerably over the past three decades. Such shocks played an important role in the late 1990s and since 2012, while they had a distinctly secondary role in the remaining of the sample.

While the aggregation to annual data offers a useful broad overveiw about the relative importance of the drivers in economic fluctuations, it misses some important aspects of the interaction between the oil market and the macroeconomy that can be better appreciated by a monthly narrative, which we discuss next for three episodes of large swings in oil prices.

**The Asian Financial Crisis.** In Figure 6, we plot the estimated historical decomposition of the model variables for the period of the Asian Financial Crisis from July 1997 to December 1998. All panels display the change in the log of the variable since June 1997. The defining feature of this event was a sharp contraction in real activity in emerging economies, as shown in the bottom right panel. The decline in the demand of oil from emerging economies induced some downward pressure on oil prices, the upper left panel which, if we include the contribution of shocks to metal prices, accounted for about one third of the decline in oil prices.

Importantly, throughout this period, despite a lower demand for oil from emerging countries, oil exporters did not cut on production with some producers, most notably Iraq, increasing production until early 1999. Accordingly, our model attributes a major role in the decline of the price of oil to supply shocks. Oil-specific demand shocks constituted an offsetting factor for oil prices and account for about 15 percent of the increase in oil production in late 1997.

The impact of disturbances originated in the oil market on economic activity is postive but economically small in advanced economies, while is negative and more significant in emerging economies. The drop in oil prices caused by “excess production” in some oil exporting countries resulted in a decline of almost 1 percent in emerging economies real activity.

**The Global Financial Crisis.** In Figure 7, we plot the estimated historical decomposition of the model variables for the period of the Global Financial Crisis ranging from July 2008 to December 2009. All panels display the change in the log of the variable since June 2008. As indicated by the solid black line in the upper left panel, the real price of oil experienced a dramatic plunge throughout this period. At the onset of the crisis, the model attributes much of the decline in the price of oil to negative oil-specific demand shocks, due to the simultaneous decline oil production and the relatively small movements in global activity factors. The role of global activity shocks becomes more prominent drivers of the price of oil towards the end of 2008, as the sharp decline in real activity materializes both in advanced and emerging economies.

The model attributes some of the decline in the price of oil to positive supply shocks even though oil production is consistently below trend. Through the lense of the model, the sharp contraction in global activity at the end of 2008 should have led to an even larger reduction in oil production than observed in the data. The higher-than expected reading in oil production is razionalized by the model through the only shock that generates a negative co-movement between oil production and the price of oil, an oil supply shock.

The decomposition of real activity variables plotted in the bottom row of Figure 7 provide a good reality check for the model. In accordance to conventional wisdom, shocks originating in the oil market had nearly no role in the collapse of economic activity associated with the Great Finacial Crisis.

**The 2014-2015 slump.** Finally, Figure 8 displays the estimated historical decompositions for the July 2014–December 2015 period, characterized by a major slump in the real price of oil. All panels display the change in the log of the variable since June 2014. Throughout the episode, our identification attributes most of the decline in the price of oil to supply shocks. Oil-specific demand shocks contributed to the acceleration in the decline of the price of oil in early 2015. On one hand, positive shocks to global supply, as detected by the decomposition of global production plotted in the upper right panel, likely resulted from the enduring expansion in unconventional shale oil production. On the other hand, the negative shocks to oil-specific demand were likely due to waning concerns about future availability of oil supplies and thus heightened expectations of future excess supply in global oil markets. These expectations, in turn, presumably reflected a few main factors, that is, the return to production of oil fields in Iraq and Libya following the end of military threats from extremists, greater market confidence that the expansion in shale oil production would not suddenly lose momentum following the price slump, and OPEC's unwillingness to cut production and thus follow through with a policy of stabilizing prices, as typically done in earlier decades.[13] Since early 2015 the decline in the price of oil also began to reflect negative shocks in emerging economies and to expectations of global activity as captured by the shock to metal prices.

As shown in the bottom row of Figure 8, oil shocks had no role in shaping economic activity in advanced economies until mid-2015. Since June 2015, oil supply shocks added about 1 percent to growth in industrial production. On the contrary, oil shocks were part of the headwinds faced by emerging economies since late 2014. The contribution to the decline in industrial output was a little over a 0.1 percent in October 2014, and it grew to about 1 percent at the end of 2015. The boosts in economic activity in advanced economies and the drag on economic activity in emerging economies of about equal size might explain the muted response of global activity to developments in oil markets.

# 5 Discussion

## 5.1 Oil Price Elasticity and the Decoupling Puzzle

Figure 9 and 10 explore drivers of oil prices during the 2014-2015 oil slump and the Asian Financial crisis, comparing our baseline identification scheme with two alternative identification schemes and datasets.

In Figure 9, Panel (a) shows again how the 2014–2015 oil slump is mostly driven by oil supply shocks, with a non–negligible role coming from global activity shocks. By contrast, panel (b) shows how a model with zero oil supply elasticity requires positive oil supply shocks to explain the rise in oil production, and negative oil demand shocks to explain the fall in oil prices. The zero supply elasticity model, in other words, decouples oil supply and oil demand shocks, which seems at odds, at least from a real world perspective, with the recent events. Finally, panel (c) show the results of the model that retains the zero oil supply elasticity, but replaces the three global activity indicators with Kilian (2009)'s global activity index. Such indicator explains a smaller portion of the decline in oil prices than our three combined indicators.

In Figure 10, the comparison between panel (b) and panel (a) again highlights how the decoupling between oil supply and oil demand shocks is apparent also during the Asian crisis when oil supply elasticity is assumed to be zero. Additionally, the comparison between (a) and (c) illustrates the difference between our world global activity measures and Kilian (2009)'s index. Such index fails to find any connection between the sharp slowdown in Emerging Economies activity and the concurrent decline in oil prices.

[13] Even though many commentaries remarked that the decision of OPEC to keep its production target unchanged was largely anticipated, between the end of November and the beginning of December in 2014, that is, in the days immediately following OPEC's decision, the Brent price of oil lost a remarkable amount of about $5 per barrel.

## 5.2 Global Effects of Oil Market Shocks

Figure 11 illustrates the importance of selecting between alternative elasticities in order to gauge the aggregate effects of oil price shocks. To make our point, we consider the hypothetical question of how much activity in either Advanced Economies or Emerging Economies would respond to a 1% oil price change, by plotting the two–year elasticity of IP to oil prices as a function of the oil supply elasticity. Under our benchmark identification scheme, a 1 percent rise in oil prices driven by an oil supply shock leads to a 0.05 percent drop in advanced economies IP, and to 0.03 percent rise in emerging economies IP. By contrast, the same oil price rise, when driven by shocks to oil demand, has more muted effects on activity in the AEs, even if its effects on the EEs are roughly unchanged. This result favor an interpretation of oil demand shocks that are driven, at least in part, by higher demand for energy–intensive goods (such as cars or heavy equipment) driven by better growth prospects in the EEs: under this interpretation, it is plausible to find that, all else equal, the effects on world economic activity of an oil price rise should be more benign.

[ TBD: Say more about effects when elasticity is different from benchmark ]

## 5.3 Oil Price Elasticities and Volatilities

Up to now we have focused on answering the question, what are the dynamic effects of oil market and global activity shocks? A related question is: how much do oil and non–oil shocks matter, on average, for oil prices and world economic activity? The answer to this question is of interest for two reasons. First, it sheds light of the issue of whether oil shocks are an important driver, on average, of global business cycles. Second, it sheds light on whether movements oil prices are an important indicator of shifts in economic activity.

Figure 12 plots the percentage of the variance of the 2–year ahead forecast errors for oil prices, IP in advanced economies, and IP in emerging economies that is attributable to oil shocks, with the remaining percentage due to global activity shocks. We report these percentages as a function of the supply elasticity. According to our benchmark identification scheme, oil shocks explain about 70 percent of the two–year forecast error variance of oil prices, with oil supply shocks accounting for the largest share (about 50 percentage points out of 70). By contrast, oil demand shocks explain all of the oil–market driven volatility in oil prices when $\eta_S$, the oil supply elasticity, is close to zero.

At two–year horizon, oil shocks explain between 5 and 10 percent of the volatility in AE and EE activity.

# 6 Conclusions

[ ]

# References


Aastveit, K. A., H. C. Bjrnland, and L. A. Thorsrud (2015). What drives oil prices? emerging versus developed economies. *Journal of Applied Econometrics 30*(7), 1013–1028.

Altinay, G. (2007). Short-run and long-run elasticities of import demand for crude oil in turkey. *Energy Policy 35*(11), 5829 – 5835.

Arezki, R. and O. J. Blanchard (2014). Seven questions about the recent oil price slump. https://blog-imfdirect.imf.org/2014/12/22/seven-questions-about-the-recent-oil-price-slump/.

Asali, M. (2011). Income and price elasticities and oil-saving technological changes in ardl models of demand for oil in g7 and bric. *OPEC Energy Review 35*(3), 189–219.

Baranzini, A. and S. Weber (2013). Elasticities of gasoline demand in switzerland. *Energy policy 63*, 674–680.

Baumeister, C. and J. D. Hamilton (2015). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. *Unpublished*.

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.

Bentzen, J. (1994). An empirical analysis of gasoline demand in denmark using cointegration techniques. *Energy Economics 16*(2), 139–143.

Bernanke, B. S. (2016). The relationship between stocks and oil prices. http://www.brookings.edu/blogs/ben-bernanke/posts/2016/02/19-stocks-and-oil-prices?cid=00900015020089101US0001-02201.

Brons, M., P. Nijkamp, E. Pels, and P. Rietveld (2008). A meta-analysis of the price elasticity of gasoline demand. a SUR approach. *Energy Economics 30*(5), 2105–2122.

Chang, D. and A. Serletis (2014). The demand for gasoline: Evidence from household survey data. *Journal of Applied Econometrics 29*(2), 291–313.

Cooper, J. C. (2003). Price elasticity of demand for crude oil: Estimates for 23 countries. *OPEC review 27*(1), 1–8.

Coyle, D., J. DeBacker, and R. Prisinzano (2012). Estimating the supply and demand of gasoline using tax data. *Energy Economics 34*(1), 195–200.

Crotte, A., R. B. Noland, and D. J. Graham (2010). An analysis of gasoline demand elasticities at the national and local levels in mexico. *Energy Policy 38*(8), 4445–4456.

Cuddington, J. T. and D. Jerrett (2008). Super cycles in real metals prices? *IMF staff Papers*, 541–565.

Dahl, C. (1993). A survey of oil demand elasticities for developing countries. *OPEC Review 17*(4), 399–420.

Dahl, C. and J. Ko (1998). The effect of deregulation on us fossil fuel substitution in the generation of electricity. *Energy policy 26*(13), 981–988.

Dahl, C. A. (1982). Do gasoline demand elasticities vary? *Land Economics 58*(3), 373–382.

Eltony, M. N. (2008). Estimating energy price elasticities for the non-oil manufacturing industries in kuwait. *OPEC Energy Review 32*(2), 184–195.

Espey, M. (1998). Gasoline demand revisited: an international meta-analysis of elasticities. *Energy Economics 20*(3), 273–295.

Faris Al-Faris, A.-r. (1997). Energy-pricing policies in arab countries: Impacts of structural adjustment programmes. *OPEC Review 21*(4), 245–260.

Goodwin, P., J. Dargay, and M. Hanly (2004). Elasticities of road traffic and fuel consumption with respect to price and income: a review. *Transport Reviews 24*(3), 275–292.

Goodwin, P. B. (1992). A review of new demand elasticities with special reference to short and long run effects of price changes. *Journal of Transport Economics and Policy*, 155–169.

Graham, D. J. and S. Glaister (2004). Road traffic demand elasticity estimates: a review. *Transport Reviews 24*(3), 261–274.

Güntner, J. H. (2014). How do oil producers respond to oil demand shocks? *Energy Economics 44*, 1–13.

Hamilton, J. D. (2009). Causes and Consequences of the Oil Shock of 2007-08. *Brookings Papers on Economic Activity 40*(1 (Spring), 215–283.

Hamilton, J. D. (2014). Oil prices as an indicator of global economic conditions. http://econbrowser.com/archives/2014/12/oil-prices-as-an-indicator-of-global-economic-conditions.

Hughes, J. E., C. R. Knittel, and D. Sperling (2008). Evidence of a shift in the short-run price elasticity of gasoline demand. *The Energy Journal 29*(1).

Iacoviello, M. (2016). Oil Prices and Consumption across Countries and U.S. States. Federal Reserve Board.

IMF (2011). World economic outlook.

Javan, A. and N. Zahran (2015). Dynamic panel data approaches for estimating oil demand elasticity. *OPEC Energy Review 39*(1), 53–76.

Kayser, H. A. (2000). Gasoline demand and car choice: Estimating gasoline demand using household information. *Energy economics 22*(3), 331–348.

Kilian, L. (2009, June). Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market. *American Economic Review 99*(3), 1053–69.

Kilian, L. and D. P. Murphy (2012, October). Why agnostic sign restrictions are not enough: Understanding the dynamics of oil market var models. *Journal of the European Economic Association 10*(5), 1166–1188.

Kilian, L. and D. P. Murphy (2014). The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics 29*(3), 454–478.

Krichene, N. (2002). World crude oil and natural gas: a demand and supply model. *Energy economics 24*(6), 557–576.

Lin, C.-Y. C. and L. Prince (2013). Gasoline price volatility and the elasticity of demand for gasoline. *Energy Economics 38*, 111–117.

Liu, W. (2014). Modelling gasoline demand in the united states: A flexible semiparametric approach. *Energy Economics 45*, 244–253.

Liu, W. (2015). Gasoline taxes or efficiency standards? a heterogeneous household demand analysis. *Energy Policy 80*, 54–64.

McRae, R. (1994). Gasoline demand in developing asian countries. *The Energy Journal*, 143–155.

Newey, W. K. and K. D. West (1987). A simple, positive semi-definite, heteroskedasticity- and autocorrelation-consistent covariance matrix. *55*(3), 703–708.

Newey, W. K. and K. D. West (1994). Automatic lag selection in covariance matrix estimation. *61*(4), 631–653.

Nicol, C. J. (2003). Elasticities of demand for gasoline in canada and the united states. *Energy economics 25*(2), 201–214.

Pindyck, R. S. and J. J. Rotemberg (1990, December). The Excess Co-movement of Commodity Prices. *Economic Journal 100*(403), 1173–89.

Polemis, M. L. (2007). Modelling industrial energy demand in greece using cointegration techniques. *Energy Policy 35*(8), 4039–4050.

Ramanathan, R. (1999). Short-and long-run elasticities of gasoline demand in india: An empirical analysis using cointegration techniques. *Energy economics 21*(4), 321–330.

Rubio-Ramírez, J. F., D. F. Waggoner, and T. Zha (2010, 04). Structural vector autoregressions: Theory of identification and algorithms for inference. *Review of Economic Studies 77*(2), 665–96.

Sene, S. O. (2012). Estimating the demand for gasoline in developing countries: Senegal. *Energy economics 34*(1), 189–194.

Sita, B. B., W. Marrouch, and S. Abosedra (2012). Short-run price and income elasticity of gasoline demand: Evidence from lebanon. *Energy Policy 46*, 109–115.

Taghizadeh Hesary, F. and N. Yoshino (2014). Monetary policies and oil price determination: An empirical analysis. *OPEC Energy Review 38*(1), 1–20.

Uhlig, H. (2005, March). What are the effects of monetary policy on output? results from an agnostic identification procedure. *Journal of Monetary Economics 52*(2), 381–419.

Wadud, Z., D. J. Graham, and R. B. Noland (2009). Modelling fuel demand for different socio-economic groups. *Applied Energy 86*(12), 2740–2749.

Figure 1: Oil Demand and Supply Elasticities Implied by the SVAR

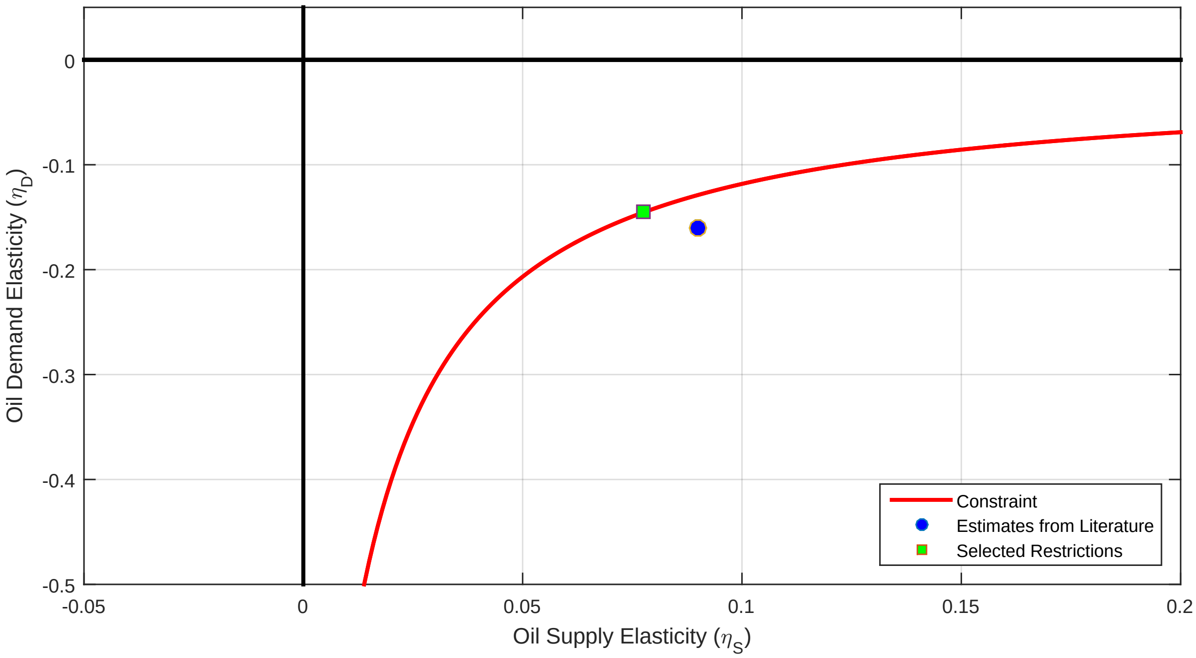


Note: The red solid line shows the relationship between the oil demand and supply elasticities imposed by the 4-equation VAR. The green square corresponds to the selected identification scheme, given the VAR (supply elasticity: 0.07, demand elasticity: -0.14). The blue circle corresponds to elasticities suggested by the meta-analysis of the literature (supply elasticity: 0.09, demand elasticity: -0.16).

Figure 2: Alternative Oil Market Structures

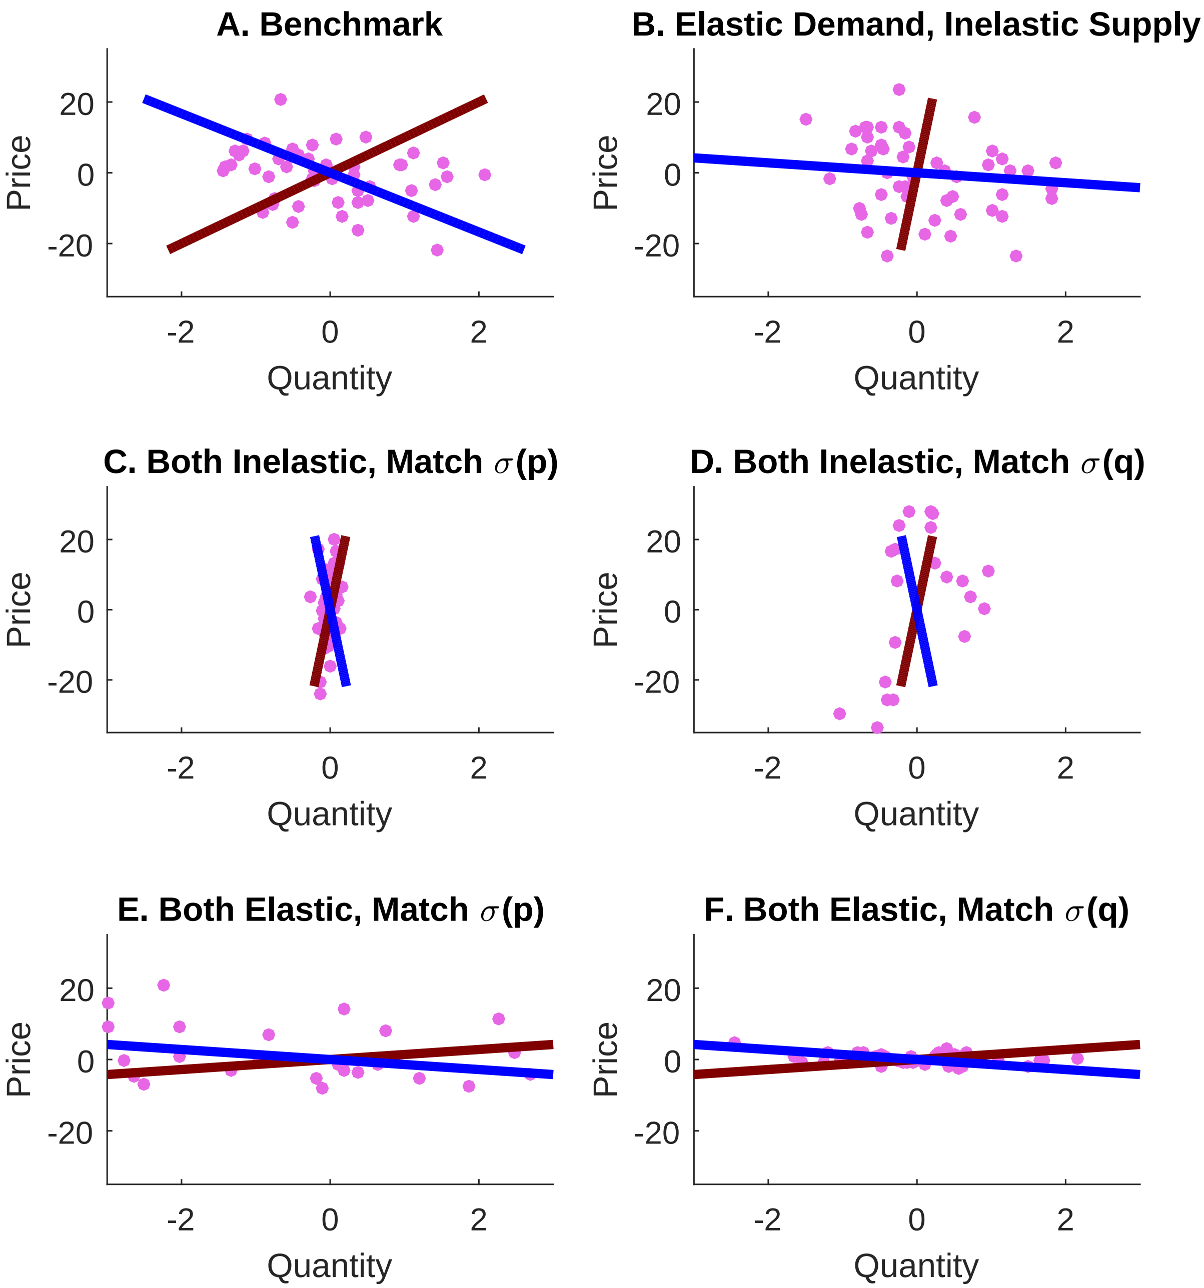


Note: The solid lines are oil demand curves and the oil supply curves. The blue dots represent simulated data implied by a given oil market structure.

Figure 3: Impulse Responses to Oil Market Shocks

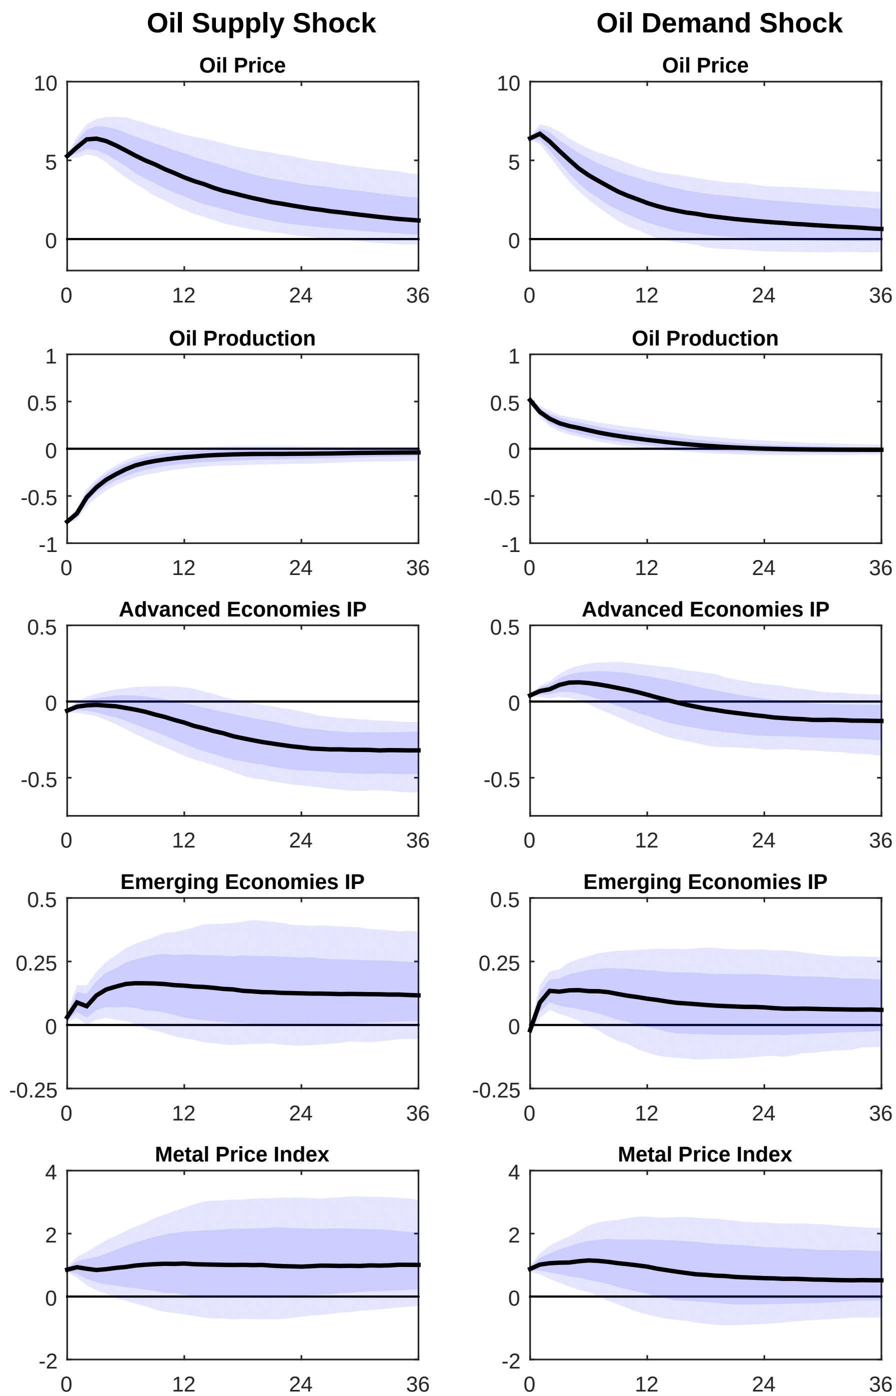


Note: The solid lines in the left column depict median responses of the specified variable to an oil supply shock of size 1 standard deviation, while those in the right column depict median responses to an oil demand shock of size 1 standard deviation; the light shaded bands represent the 90-percent pointwise credible sets and the dark shaded bands represent the 68-percent pointwise credible sets.

Figure 4: Impulse Responses to Global Activity Shocks

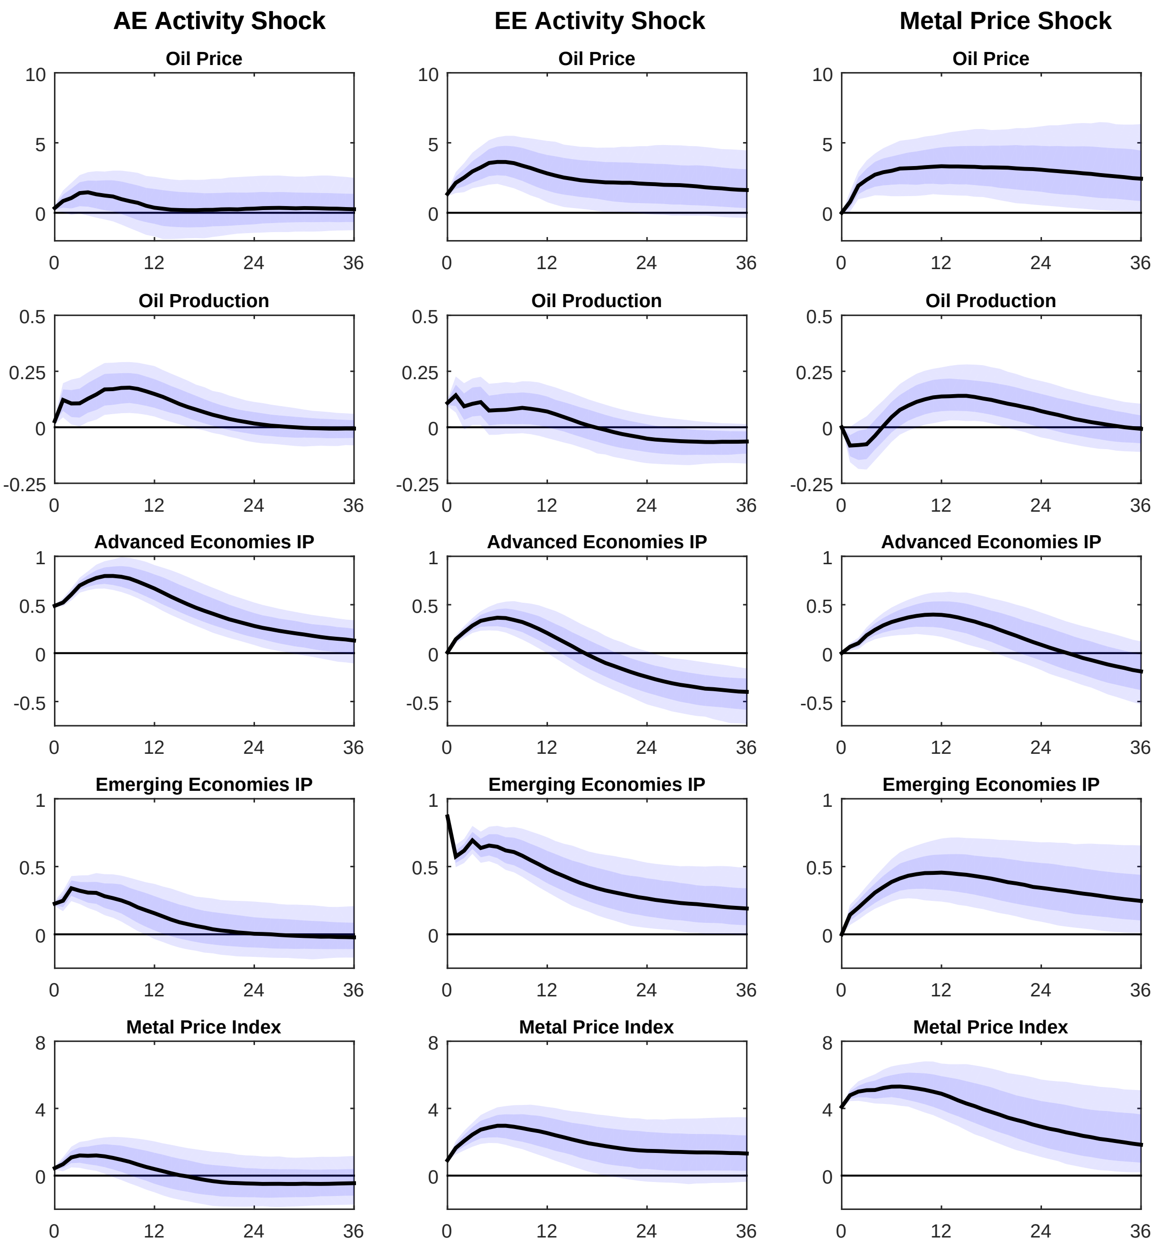


Note: The solid lines in the left column depict median responses of the specified variable to an advanced economies' activity shock of size 1 standard deviation, those in the middle column depict median responses to emerging economies' activity shock of size 1 standard deviation, and those in the right column depict median responses to a metal price shock of size 1 standard deviation; the light shaded bands represent the 90-percent pointwise credible sets and the dark shaded bands represent the 68-percent pointwise credible sets.

Figure 5: Historical Variance Decomposition of Selected Variables
(Full Sample)

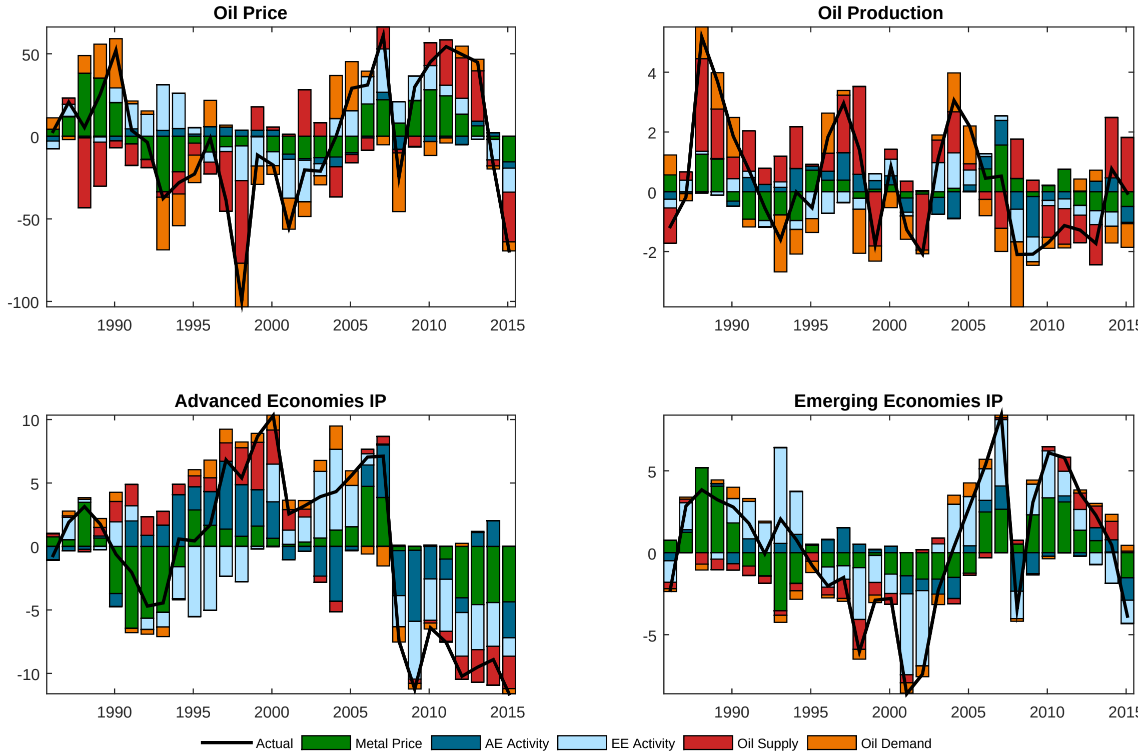


Note: Sample period: 1986 to 2015. The shaded regions in each panel depict the historical contributions of oil market and global activity shocks to the specified variable, while the solid lines depict the actual series.

Figure 6: Historical Variance Decomposition of Selected Variables
(Asian Financial Crisis)

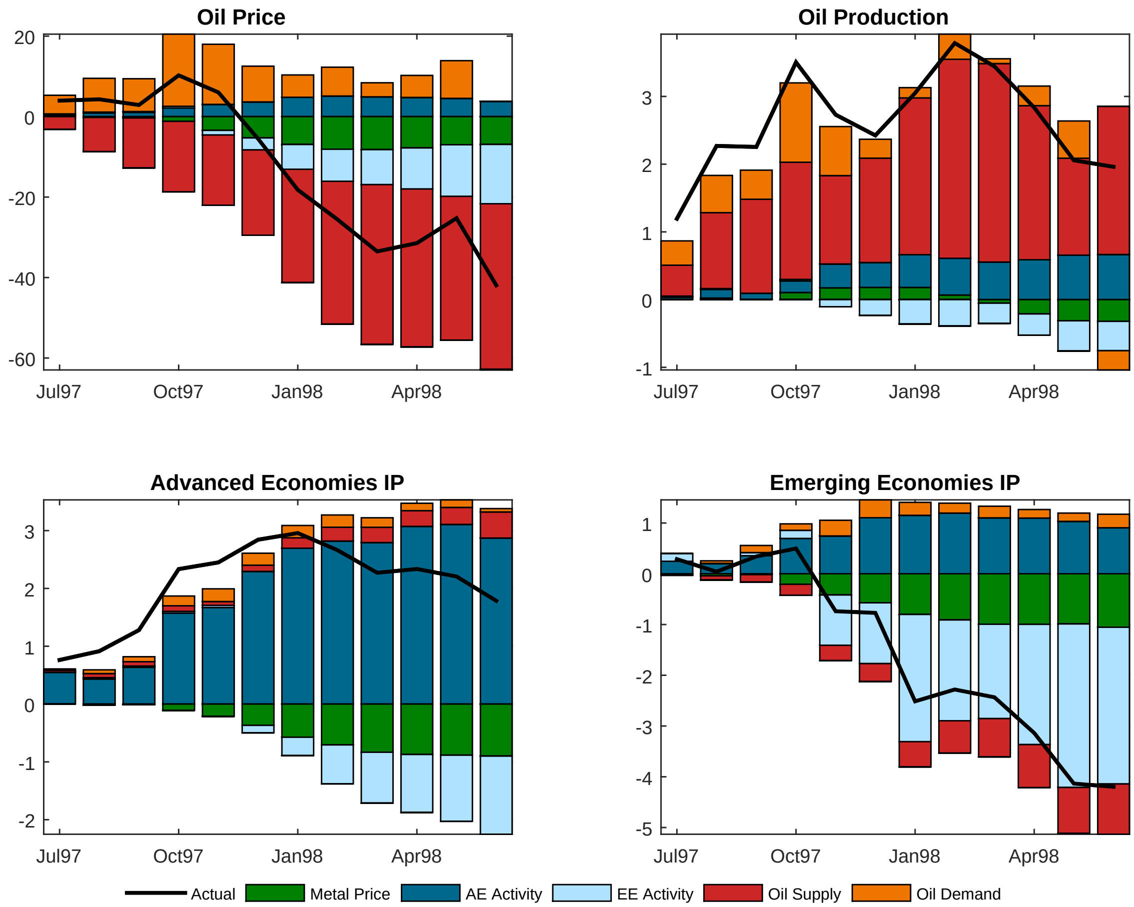


Note: Sample period: July 1997 to December 1998. The shaded regions in each panel depict the historical contributions of oil market and global activity shocks to the specified variable, while the solid lines depict the actual series. All variables are expressed in deviations from June 1997.

Figure 7: Historical Variance Decomposition of Selected Variables
(Global Financial Crisis)

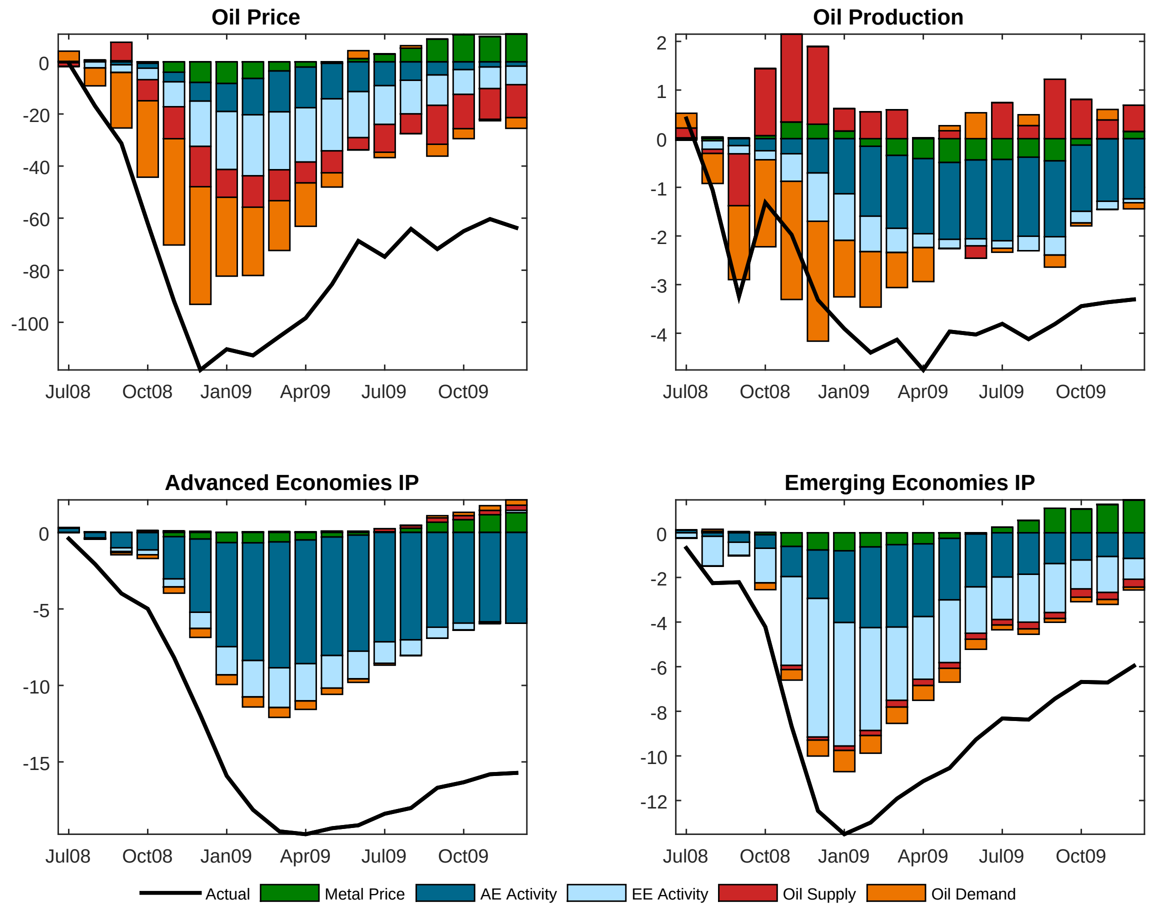


Note: Sample period: July 2008 to December 2009. The shaded regions in each panel depict the historical contributions of oil market and global activity shocks to the specified variable, while the solid lines depict the actual series. All variables are expressed in deviations from June 2008.

Figure 8: Historical Variance Decomposition of Selected Variables
(2014-2015 Oil Slump)

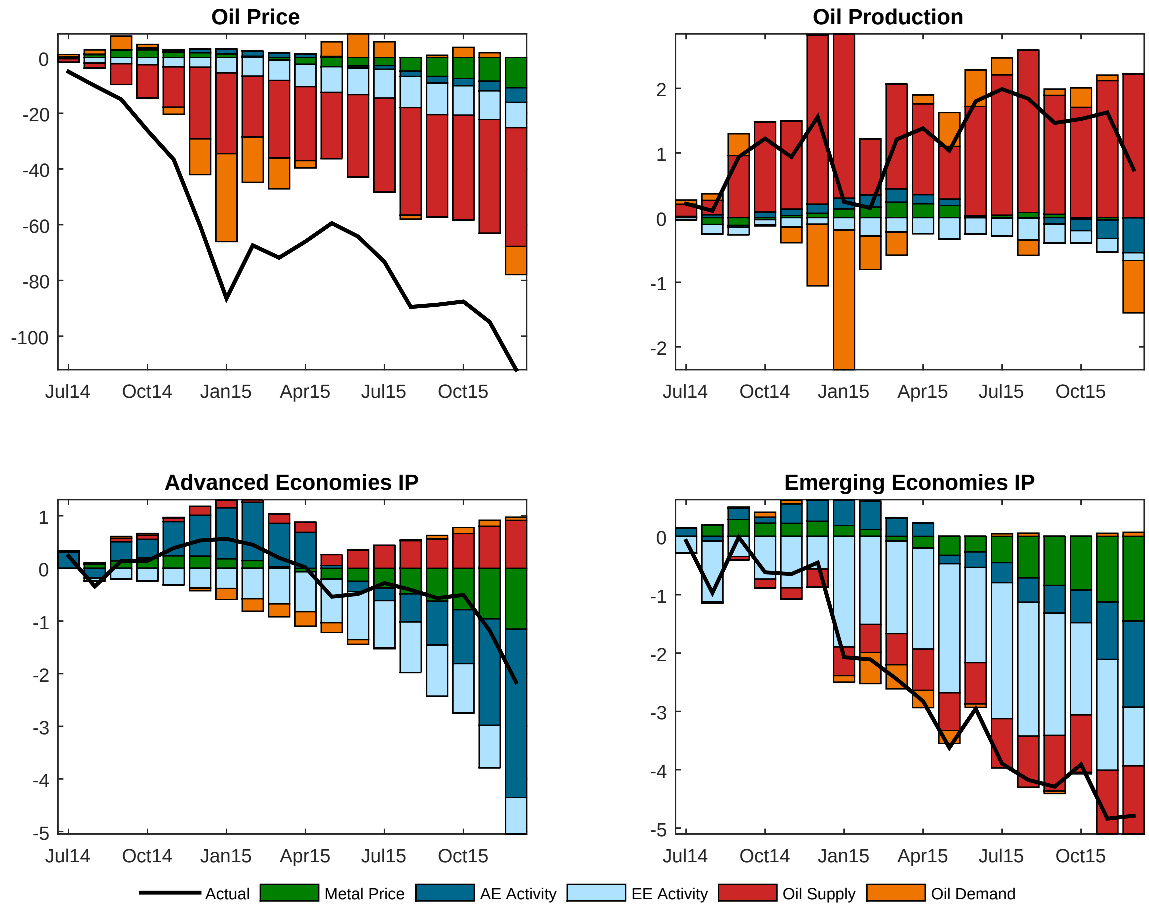


Note: Sample period: July 2014 to December 2015. The shaded regions in each panel depict the historical contributions of oil market and global activity shocks to the specified variable, while the solid lines depict the actual series. All variables are expressed in deviations from June 2014.

Figure 9: HISTORICAL DECOMPOSITION OF SELECTED VARIABLES
(2014–2015 Oil Slump: Model Comparison)

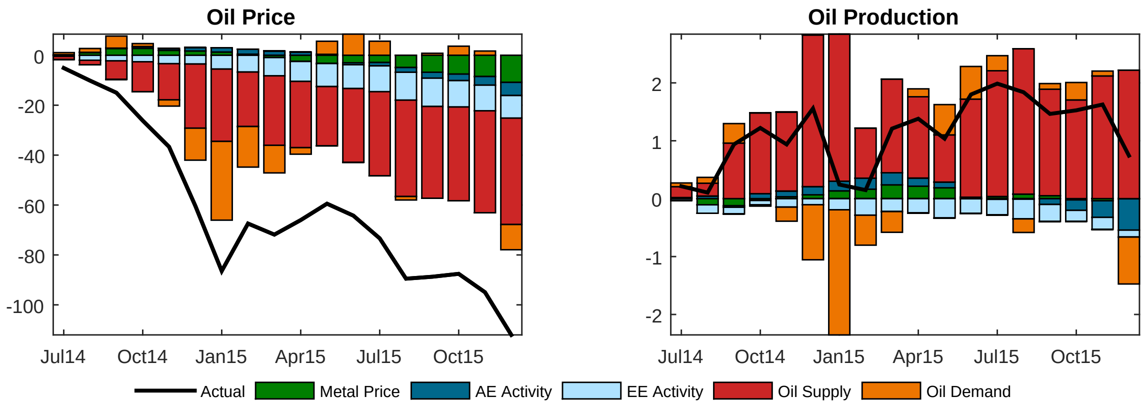


(a) 5-equation SVAR model: Baseline Identification

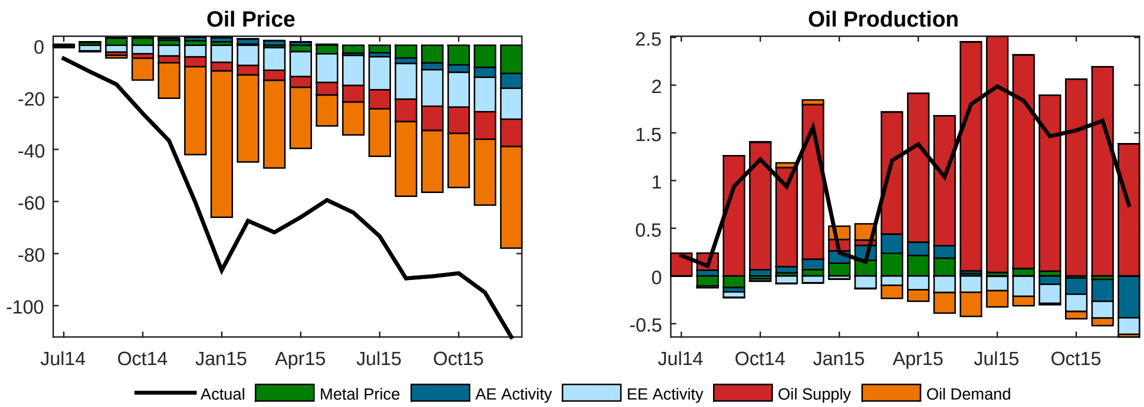


(b) 5-equation SVAR model: $\eta_S = 0$

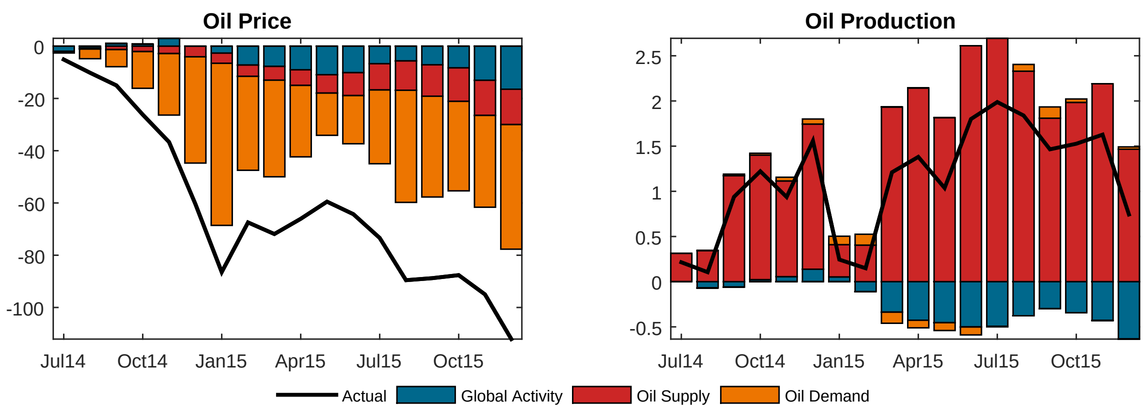


(c) 3-equation SVAR model: $\eta_S = 0$

NOTE: Sample period: July 2014 to December 2015. The shaded regions in panel (a) depict the historical contributions of oil market and global activity shocks to the specified variable based on our baseline identification; panel (b) depicts the historical decomposition based on the identification that sets the oil supply elasticity $\eta_S$ to zero; panel (c) reports the historical decomposition based on a 3-equation VAR that also sets $\eta_S = 0$. The solid lines depict the actual series. All variables are expressed in deviations from June 2014. See text for details

Figure 10: Historical Decomposition of Selected Variables
(Asian Financial Crisis: Model Comparison)

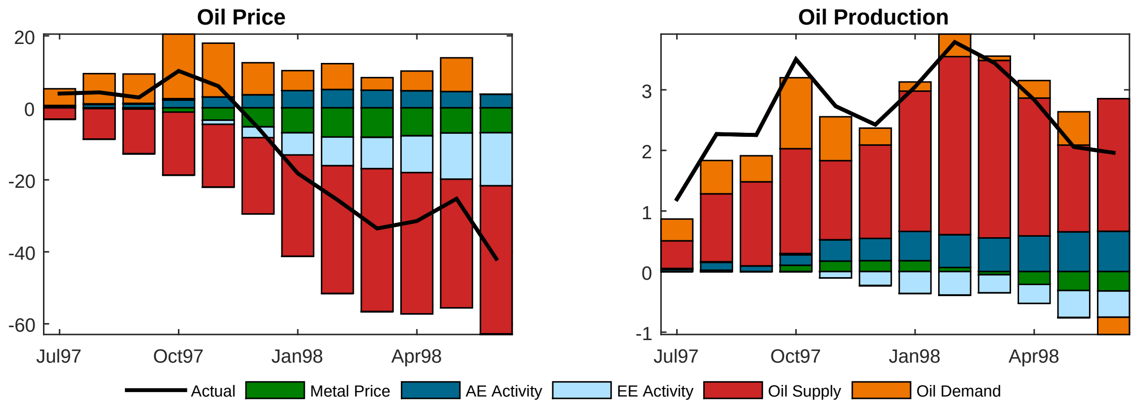


(a) 5-equation SVAR model: Baseline Identification

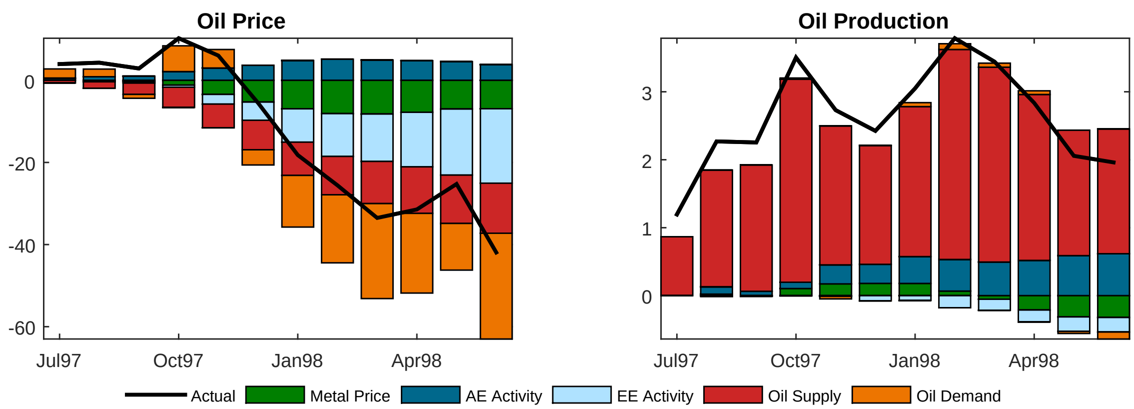


(b) 5-equation SVAR model: $\eta_S = 0$

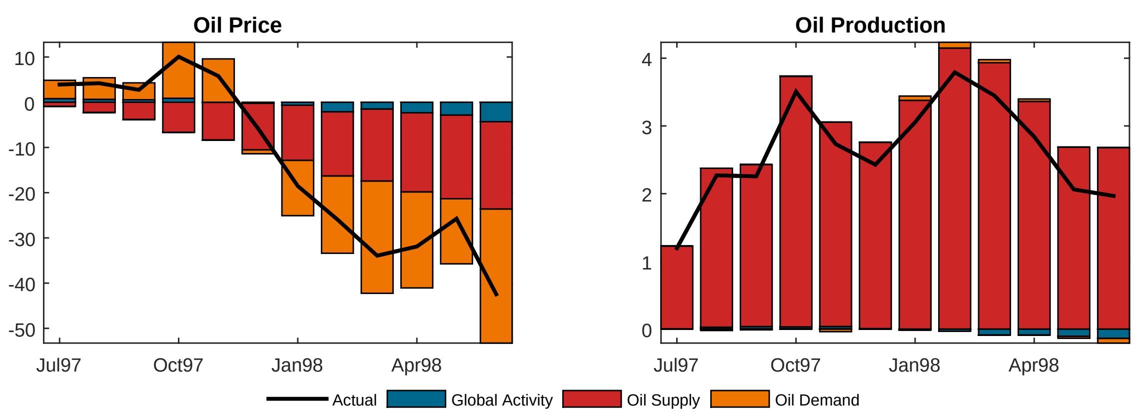


(c) 3-equation SVAR model: $\eta_S = 0$

Note: Sample period: July 1997 to December 1998. The shaded regions in panel (a) depict the historical contributions of oil market and global activity shocks to the specified variable based on our baseline identification; panel (b) depicts the historical decomposition based on the identification that sets the oil supply elasticity $\eta_S$ to zero; panel (c) reports the historical decomposition based on a 3-equation VAR that also sets $\eta_S = 0$. The solid lines depict the actual series. All variables are expressed in deviations from June 1997. See text for details.

Figure 11: How Activity Responds to Oil Price Changes

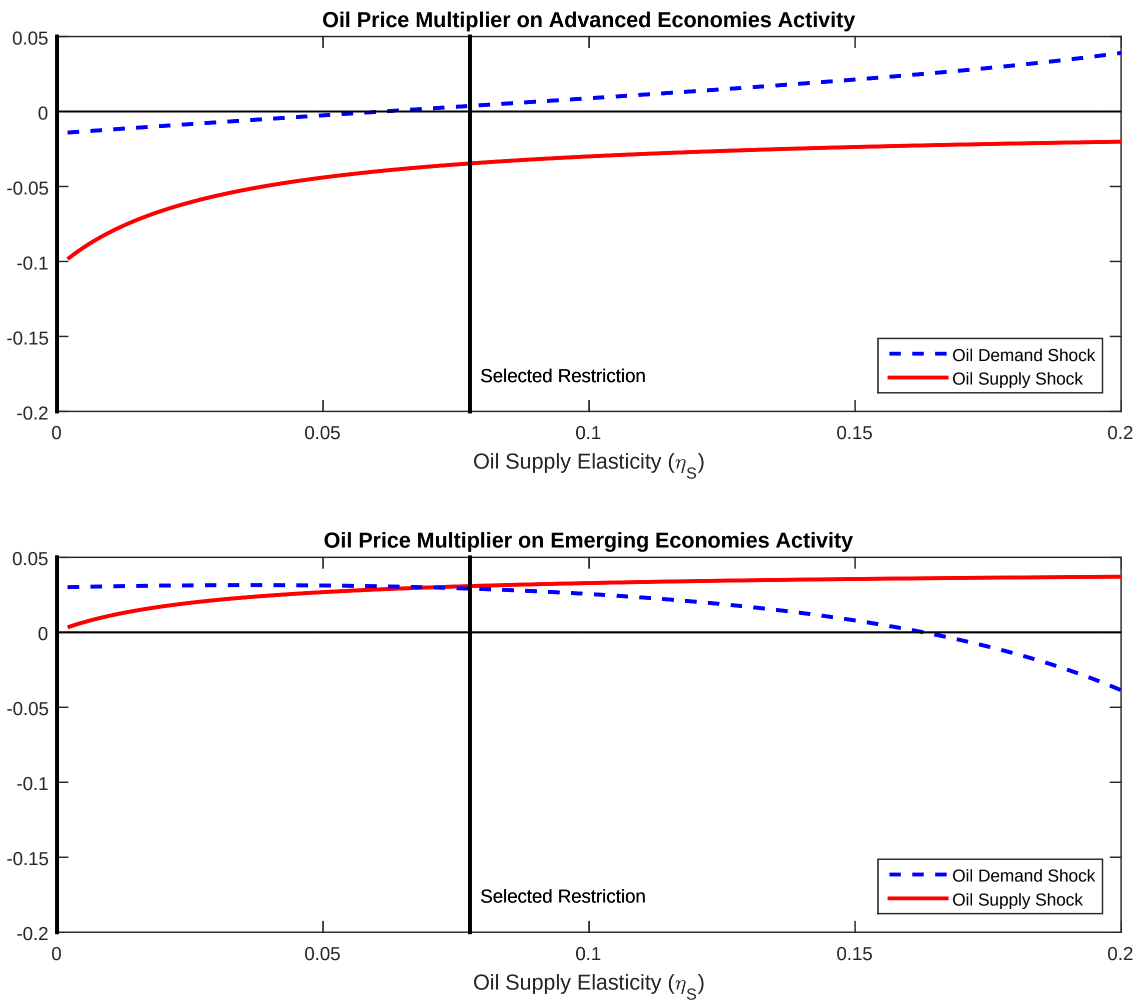


Note: The lines plot the two-year elasticity of AFE and EME IP to oil prices (average response of activity in the first two years divided by average response of prices in the first two years) implied by the estimated VAR model for the identified shocks.

Figure 12: Forecast Error Variance Decomposition of Selected Variables
(Share Explained by Oil Shocks)

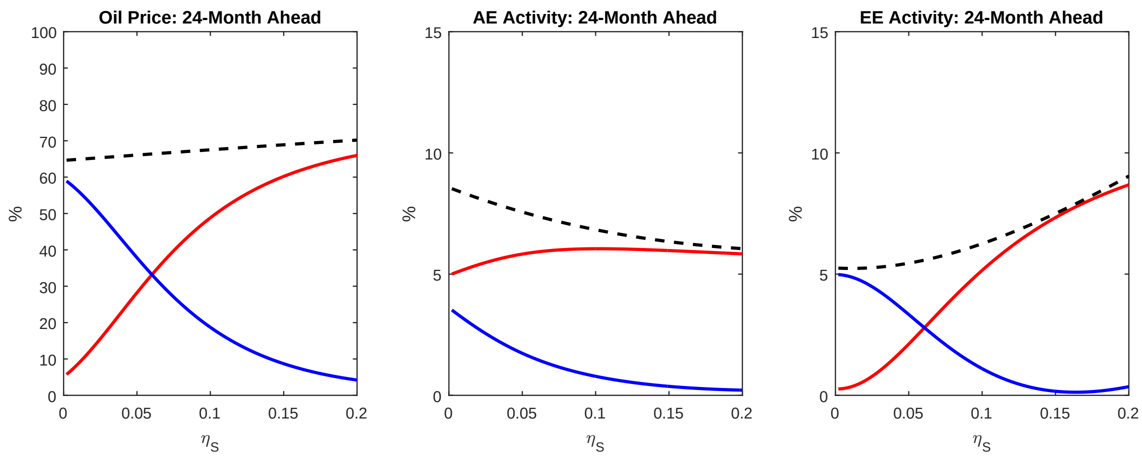


Note: Fraction of forecast error variance in oil prices (left column), advanced economies activity (middle column), and emerging economies activity (right column) explained by oil supply (red line) and oil demand (blue line) shocks.

# A The Data

**Industrial Production.** We construct a monthly index of industrial production for advanced economies and emerging economies aggregating country-level data. We take monthly, seasonally-adjusted, total industrial production excluding constructions. For countries where this series is not avaiable, we use monthly, seasonally-adjsuted, manufacturing industrial production.[14] The resulting unbalanced dataset run from 1960:M1 to 2015:M12 for advanced economies and from 1963:M1 to 2015:M12 for emerging economies. To construct the indexes, we first compute the growth rate of industrial production for each individual country. For both advanced and emerging economies, we then aggregate the country-specific growth rates by calculating annual weights based on gross domestic production (GDP) in current US dollars from the World Bank World Development Indicators. We then obtain the level of industrial production by cumulating the resulting monthly growth series. Both indexes are normalized to take value 100 in January 2007.

Table A.1 lists the countries included in the advanced economies index, while Table A.2 lists the countries included in the emerging economies index. For each country we report the weight in the total index as of 2013, as well as the sample availability. For advanced economies, since 1985—the first observation we use in the estimation—data are available for all countries but Finland, and Greece, and Portugal, all countries with a small weight in the overall index. Data availability is more scattered for emerging economies. From 1985 to the mid 1990s, the emerging economies index is driven mostly by India, Korea and Mexico. Data for Russia, the third largest country in the panel, become available in 1993 while data for China are available since 1997.

The countries in the sample account, in 2013, for 87 percent of world GDP in current dollars, with 53 and 34 percentage points of GDP accruing to advanced and emerging economies respectively. Because of the lack of monthly IP data, the largest economies missing from the sample are Australia, Saudi Arabia, Switzerland, Nigeria, Iran and United Arab Emirates, which together account for about 6 percent of world GDP.

The advanced economies in the sample account for 20 percent of world oil production and 41 percent of world oil consumption. The emerging economies in the sample account for 34 percent of world oil production and 39 percent of world oil consumption. (As a consequence, the missing countries account for 13 percent of world GDP, 46 percent of world oil production, 20 percent of world oil consumption)

**Metals Prices.** Metal prices are measured from the IMF Metals Price Index, linearly log detrended and expressed in real terms dividing by the U.S. CPI Index (Haver code: PCURS@USECON).

[14]For Japan total industrial production is available only starting 1998, while data on manufacturing industrial production go back to the 1950s. Since for the common sample the correlation between the two series is close to 1 and Japan is the second-largest country in our sample, we use the longest available series.

**Oil Market Variables.** The Real price of oil is the Brent Price of Oil (PZBRT@USECON), linearly log detrended and expressed in real terms dividing by the U.S. CPI Index (Haver code: PCURS@USECON). Oil Production is total production of crude oil (excluding natural gas liquids and nonconventional oils), from the International Energy Agency.

Table A.1: Advanced Economies Industrial Production

| Country | Share of World GDP | Share of World Oil Production | Share of World Oil Consumption | Sample |
|---|---|---|---|---|
| United States | 22.03% | 11.63% | 20.78% | 1985–2015 |
| Japan | 6.46% | 0.00% | 4.95% | 1985–2015 |
| Germany | 4.90% | 0.00% | 2.64% | 1985–2015 |
| France | 3.69% | 0.00% | 1.82% | 1985–2015 |
| United Kingdom | 3.52% | 1.00% | 1.64% | 1985–2015 |
| Italy | 2.81% | 0.13% | 1.41% | 1985–2015 |
| Canada | 2.42% | 4.59% | 2.61% | 1985–2015 |
| Spain | 1.83% | 0.00% | 1.31% | 1985–2015 |
| Netherlands | 1.12% | 0.00% | 0.98% | 2000–2015 |
| Sweden | 0.76% | 0.00% | 0.34% | 2000–2015 |
| Belgium | 0.69% | 0.00% | 0.69% | 1985–2015 |
| Norway | 0.69% | 2.12% | 0.27% | 1985–2015 |
| Austria | 0.56% | 0.00% | 0.29% | 1985–2015 |
| Denmark | 0.44% | 0.21% | 0.17% | 2000–2015 |
| Finland | 0.35% | 0.00% | 0.21% | 1995–2015 |
| Greece | 0.32% | 0.00% | 0.32% | 1995–2015 |
| Ireland | 0.30% | 0.00% | 0.15% | 1985–2015 |
| Portugal | 0.30% | 0.00% | 0.27% | 2010–2015 |
| Luxembourg | 0.08% | - | - | 1985–2015 |
| AFE total | 53.26% | 19.69% | 40.85% | |

Note: The entries in the table list the countries used in the construction of the index of manufacturing industrial production in advanced economies. The column titled 'Weight' reports the weight of each country in the total index for 2013. The column titled 'Sample' reports data availability for each individual country. Data for Japan is on manufacturing industrial production.

Table A.2: Emerging Economies Industrial Production

| Country | Share of World GDP | Share of World Oil Production | Share of World Oil Consumption | Sample |
|---|---|---|---|---|
| China | 12.47% | 4.87% | 11.69% | 1997–2015 |
| Brazil | 3.14% | 2.44% | 3.34% | 2002–2015 |
| Russia | 2.73% | 12.45% | 3.48% | 1993–2015 |
| India | 2.45% | 1.05% | 4.08% | 1985–2015 |
| Korea | 1.72% | 0.00% | 2.69% | 1985–2015 |
| Mexico | 1.66% | 3.32% | 2.21% | 1985–2015 |
| Indonesia | 1.20% | 1.02% | 1.77% | 1993–2015 |
| Turkey | 1.08% | 0.00% | 0.79% | 1985–2015 |
| Argentina | 0.82% | 0.73% | 0.73% | 1994–2015 |
| Poland | 0.69% | 0.00% | 0.57% | 1996–2015 |
| Thailand | 0.51% | 0.53% | 1.38% | 2000–2015 |
| Colombia | 0.50% | 1.16% | 0.33% | 1990–2015 |
| Venezuela | 0.49% | 3.10% | 0.90% | 1997–2013 |
| South Africa | 0.48% | 0.00% | 0.64% | 1985–2015 |
| Malaysia | 0.41% | 0.75% | 0.88% | 1985–2015 |
| Singapore | 0.40% | 0.00% | 1.35% | 1985–2015 |
| Israel | 0.38% | 0.00% | 0.25% | 1990–2015 |
| Chile | 0.36% | 0.00% | 0.39% | 1991–2015 |
| Philippines | 0.36% | 0.00% | 0.33% | 1998–2015 |
| Kazakhstan | 0.30% | 1.99% | 0.30% | 1999–2015 |
| Czech Republic | 0.27% | 0.00% | 0.20% | 2000–2015 |
| Peru | 0.27% | 0.12% | 0.25% | 1990–2015 |
| Romania | 0.25% | 0.10% | 0.19% | 2000–2015 |
| Ukraine | 0.24% | 0.00% | 0.28% | 2002–2015 |
| Hungary | 0.18% | 0.00% | 0.14% | 1993–2015 |
| Slovak Republic | 0.13% | 0.00% | 0.08% | 1998–2015 |
| Croatia | 0.08% | - | - | 2000–2015 |
| Bulgaria | 0.07% | 0.00% | 0.08% | 2000–2015 |
| Slovenia | 0.06% | - | - | 1998–2015 |
| Lithuania | 0.06% | 0.00% | 0.06% | 1995–2015 |
| Jordan | 0.04% | - | - | 2002–2015 |
| Latvia | 0.04% | - | - | 2000–2015 |
| Estonia | 0.03% | - | - | 2000–2015 |
| EME total | 33.86% | 33.62% | 39.39% | |

Note: The entries in the table list the countries used in the construction of the index of manufacturing industrial production in advanced economies. The column titled 'Weight' reports the weight of each country in the total index for 2013. The column titled 'Sample' reports data availability for each individual country. Data on manufacturing industrial production for the following countries: Indonesia, Thailand, Colombia, Venezuela, South Africa, Singapore, Israel, Chile, Philippines, Peru. Data for Mexico on total industrial production including construction.

Figure A.1: The Data

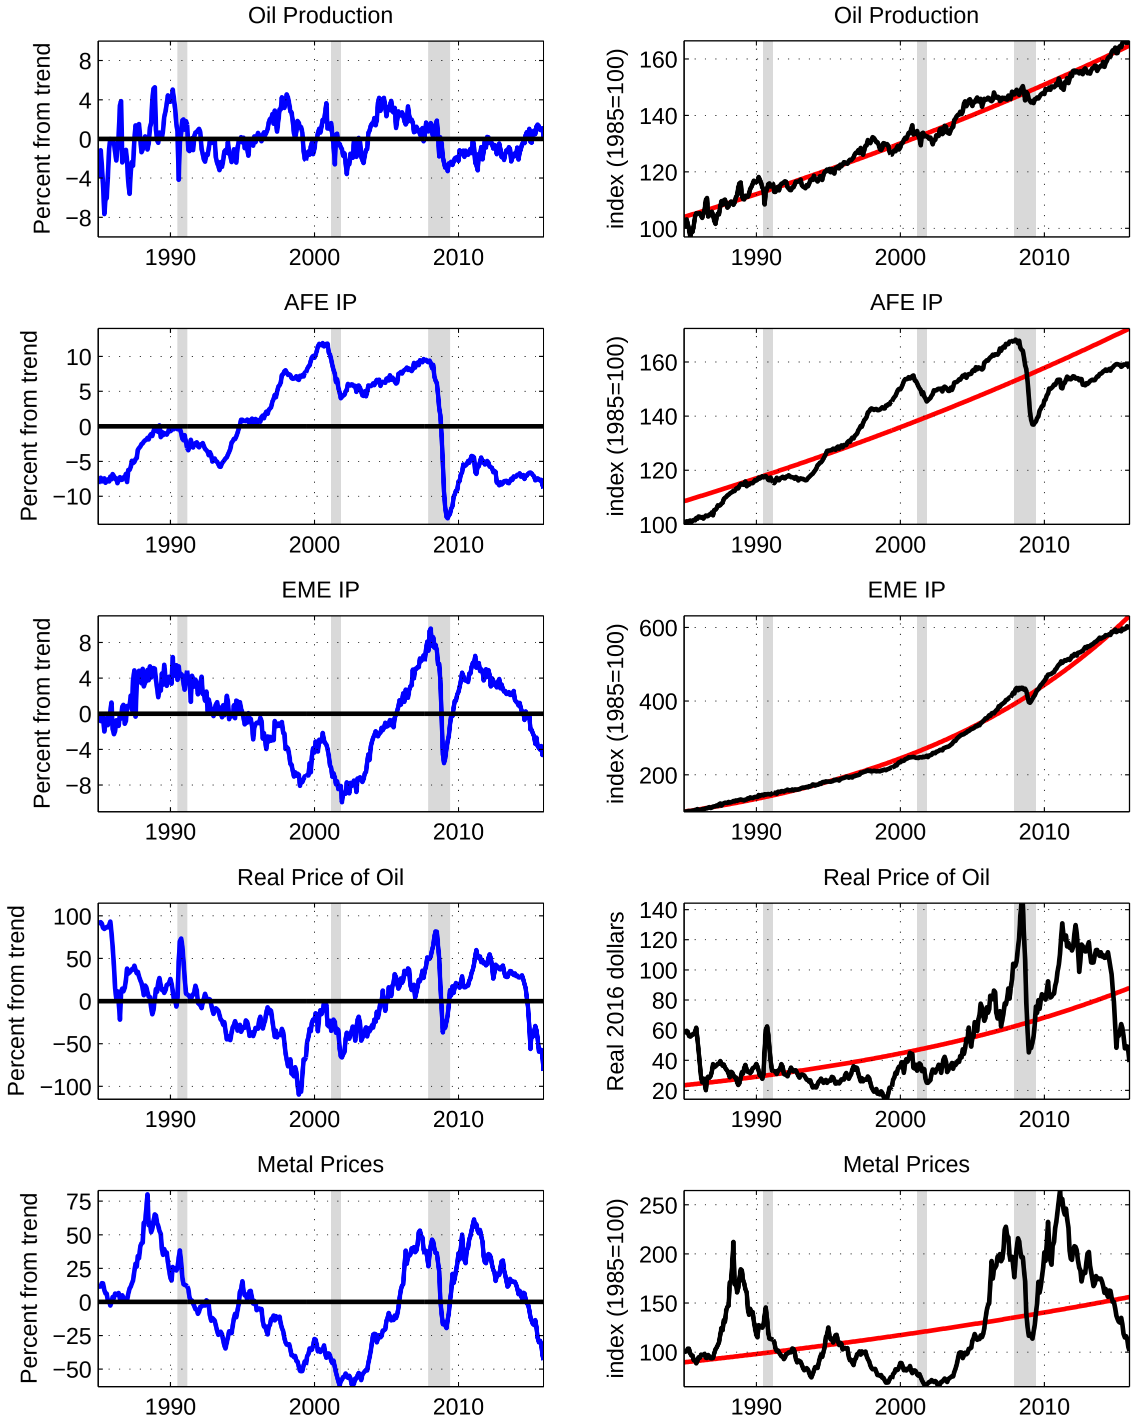


Note: Sample period: monthly data from 1985:M1 to 2015:M12. Data linearly detrended (left column) and undetrended (right column). The shaded vertical bars denote the NBER-dated U.S. recessions.

# B Search of the literature

We employed the following search strategy. We examined a narrative survey (Hamilton, 2009) and set up a search query able to capture most of the relevant studies. In particular, we searched for the words "oil demand elasticity", "oil supply elasticity", and "gasoline demand elasticity". We confined our searches to published papers and searched the following databases: EconLit, ScienceDirect, Wiley, and JSTOR. The search returned 125 unique entries. Other than obviously irrelevant studies (e.g. false positives and studies that did not contain estimates of oil elasticities), we excluded studies conducting meta-analysis, studies with estimates based only on pre-1985 data, and studies estimating only long-run elasticities. These inclusion criteria leave 32 studies in our database. We decided not to apply formal meta-analysis tools because (i) standard errors of the estimates were missing and (ii) since we span both the empirical macro and micro literature, there was a rich heterogeneity in the models and countries used for estimation but not enough studies to compute estimates conditional of specific characteristics. The search was terminated on March 18 2015.

Table A.3 tabulates the studies constituting our database. For comparison, Table A.4 tabulates oil elasticities documented in meta-analysis and survey papers excluded from our search. Both tables highlight how the literature concentrated on the estimation of the short-run price elasticity of demand, while there are only a handful of studies that estimate the short-run supply elasticity. In addition, the median short-run demand elasticity of −0.21 that emerges from our survey is in line with the majority of existing surveys and meta-analysis.

Table A.3: Oil Elasticities across Studies
(Literature Search)

| Study | **Demand Elasticity** | **Supply Elasticity** |
|---|---|---|
| Faris Al-Faris (1997) | -0.14 | |
| Asali (2011) | -0.05 | |
| Altinay (2007) | -0.10 | |
| Baranzini and Weber (2013) | -0.09 | |
| Baumeister and Peersman (2013a) | -0.26 | 0.27 |
| Baumeister and Peersman (2013b) | -0.35 | |
| Bentzen (1994) | -0.32 | |
| Chang and Serletis (2014) | -0.66 | |
| Cooper (2003) | -0.05 | |
| Coyle et al. (2012) | -0.08 | 0.26 |
| Crotte et al. (2010) | -0.06 | |
| Dahl (1982) | -0.20 | |
| Dahl and Ko (1998) | -0.72 | |
| Eltony (2008) | -0.18 | |
| Güntner (2014) | | -0.00 |
| Taghizadeh Hesary and Yoshino (2014) | -0.08 | 0.24 |
| Hughes et al. (2008) | -0.04 | |
| Javan and Zahran (2015) | -0.04 | |
| Kayser (2000) | -0.23 | |
| Kilian and Murphy (2012) | | 0.01 |
| Kilian and Murphy (2014) | -0.26 | |
| Krichene (2002) | -0.07 | -0.16 |
| McRae (1994) | -0.12 | |
| Nicol (2003) | -0.21 | |
| Lin and Prince (2013) | -0.03 | |
| Liu (2014) | -0.06 | |
| Liu (2015) | -0.30 | |
| Polemis (2007) | -0.13 | |
| Ramanathan (1999) | -0.21 | |
| Sene (2012) | -0.12 | |
| Sita et al. (2012) | -0.92 | |
| Wadud et al. (2009) | -0.09 | |
| *Mean* | **-0.13** | **0.13** |
| *Median* | **-0.21** | **0.10** |
| *Standard Deviation* | **0.21** | **0.18** |

Table A.4: Oil Elasticities across Studies
(Meta-Analysis & Surveys)

| **Authors** | **Oil Demand** | **Oil Supply** | **Google Cites** |
|---|---|---|---|
| Brons et al. (2008) | -0.34 | | 187 |
| Dahl (1993) | -0.07 | | 68 |
| Espey (1998) | -0.23 | | 349 |
| Goodwin (1992) | -0.27 | | 688 |
| Goodwin et al. (2004) | -0.25 | | 510 |
| Graham and Glaister (2004) | -0.25 | | 253 |
| Hamilton (2009) | -0.08 | | 833 |
| IMF (2011) | -0.02 | | |
| **MEDIAN** | **-0.24** | | |
| **AVERAGE** | **-0.19** | | |