

Oil Price Elasticities and Oil Price Fluctuations

Caldara, Dario, Michele Cavallo, and Matteo Iacoviello

Please cite paper as:

Caldara, Dario, Michele Cavallo, and Matteo Iacoviello (2016).
Oil Price Elasticities and Oil Price Fluctuations.
International Finance Discussion Papers 1173.

<http://dx.doi.org/10.17016/IFDP.2016.1173>



International Finance Discussion Papers

Board of Governors of the Federal Reserve System

Number 1173

July 2016

Board of Governors of the Federal Reserve System
International Finance Discussion Papers

Number 1173

July 2016

Oil Price Elasticities and Oil Price Fluctuations

Dario Caldara, Michele Cavallo, and Matteo Iacoviello

NOTE: International Finance Discussion Papers are preliminary materials circulated to stimulate discussion and critical comment. References to International Finance Discussion Papers (other than an acknowledgment that the writer has had access to unpublished material) should be cleared with the author or authors. Recent IFDPs are available on the Web at www.federalreserve.gov/pubs/ifdp/. This paper can be downloaded without charge from the Social Science Research Network electronic library at www.ssrn.com.

Oil Price Elasticities and Oil Price Fluctuations*

Dario Caldara[†] Michele Cavallo[‡] Matteo Iacoviello[§]

July 11, 2016

Abstract

We study the identification of oil shocks in a structural vector autoregressive (SVAR) model of the oil market. First, we show that the cross-equation restrictions of a SVAR impose a nonlinear relation between the short-run price elasticities of oil supply and oil demand. This relation implies that seemingly plausible restrictions on oil supply elasticity may map into implausible values of the oil demand elasticity, and vice versa. Second, we propose an identification scheme that restricts these elasticities by minimizing the distance between the elasticities allowed by the SVAR and target values that we construct from a survey of relevant studies. Third, we use the identified SVAR to analyze sources and consequences of movements in oil prices. We find that (1) oil supply shocks and global demand shocks explain 50 and 35 percent of oil price fluctuations, respectively; (2) a drop in oil prices driven by supply shocks boosts economic activity in advanced economies, whereas it depresses economic activity in emerging economies; and (3) the selection of oil market elasticities is essential for understanding the source of oil price movements and to measuring the multipliers of oil prices on economic activity.

KEYWORDS: Oil Prices. Vector Autoregressions. Commodity Prices.

JEL CLASSIFICATION: Q43. C32. E32.

*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. E-mail: dario.caldara@frb.gov

[‡]FEDERAL RESERVE BOARD. E-mail: michele.cavallo@frb.gov

[§]FEDERAL RESERVE BOARD. E-mail: matteo.iacoviello@frb.gov

1 Introduction

Major swings in the price of oil draw a lot of attention among policymakers, academics, and practitioners. A common theme in the academic literature and in the financial press is that these swings originate from a variety of sources, such as changes in global demand, supply disruptions and precautionary motives. Empirically distinguishing among competing sources of fluctuations, however, has proved to be an elusive task. To achieve this goal, we implement a structural vector autoregressive (SVAR) model of the oil market using a new identification methodology and new data. We use the model to provide a full characterization of the two-way interaction between the oil market and the global economy, with a particular focus on the determinants of oil prices and on the differential effect of oil shocks on advanced and emerging economies.

Our methodological contribution aims at identifying the role of supply and demand shocks in driving oil prices and oil production. Doing so requires estimating supply and demand curves in the oil market. An illustration of this identification problem is depicted in figure 1, which shows the scatter plot between monthly surprises in oil prices and oil production implied by simple univariate AR(1) regressions. The dots clearly show that oil prices and oil production are uncorrelated.¹ The lack of correlation between prices and production could be the outcome of very different oil market configurations. As shown in figure 1, the supply curve could be inelastic, while the demand curve could be very elastic. As a result, fluctuations in oil prices and oil production would be decoupled, with prices driven uniquely by demand shocks and production driven uniquely by supply shocks. A market characterized by a very elastic oil supply curve and a very inelastic demand curve would also lead to a decoupling of movements in oil prices and oil production. In between these two extremes lies an oil market with a downward-sloping demand curve and an upward-sloping supply curve, which would imply that demand and supply shocks jointly affect oil prices and production. These market configurations, which we picked among many for illustrative purposes, are equally consistent with the data but have very different implications for the causes and the consequences of oil prices fluctuations.

The starting point of our analysis is to show that the cross-equation restrictions embedded in SVAR models impose a nonlinear relation between the short-run price elasticities of oil supply and oil demand—that is, on the slopes of the curves depicted in figure 1. This relation implies that seemingly plausible restrictions on the oil supply elasticity may map into implausible values of the oil demand elasticity, and vice versa. We build on this finding and propose an identification scheme that jointly restricts the slope of supply and demand schedules to plausible values. The identification scheme, in particular, minimizes the distance between the elasticities allowed by the SVAR on the one hand, and the target values indicated by a comprehensive analysis of the relevant literature, on the other. These targets are

¹This observation has also been noted by [Baumeister and Hamilton \(2015b\)](#).

consistent with the conventional view that the short-run price elasticities of both oil supply and demand are small.

Even with this identification strategy on hand, an additional challenge is to disentangle demand shocks that are specific to the oil market from demand shocks that originate from changes in global economic activity. The main difficulty is that global activity is hard to measure. Our novel approach is to use three indicators of real economic activity that provide a rich and comprehensive characterization of oil demand. We construct two separate indicators based on industrial production, one for emerging economies, and another for advanced economies, dating back to the mid 1980s. Our third indicator is an index of industrial metal prices, which are often viewed by policymakers and practitioners as early indicators of swings in economic activity and global risk sentiment.² Using a simple forecasting exercise, we show that industrial production in emerging economies and metal prices help predict fluctuations in future oil prices.³

With the identification strategy and the data on hand, we turn to the analysis of the causes and consequences of oil price movements. Our analysis delivers the following results.

1. Using our identification scheme, the short-run oil supply elasticity is about 0.1 and the oil demand elasticity is about -0.1 . Under these elasticities, oil supply shocks are the main driving force of oil market movements, accounting for 50 and 40 percent of the volatility of oil prices and oil production, respectively. Shocks to global economic conditions also play an important role, explaining about 35 percent of the volatility of oil prices, and 25 percent of the volatility of oil production.
2. A drop in oil prices driven by oil market shocks, regardless of whether the shock originates from supply or demand, depresses economic activity in emerging economies. A drop in oil prices boosts economic activity in advanced economies only if it is supply-driven. Given the historical realization of these oil shocks, these findings help rationalize the muted effects of changes in oil prices on global economic activity.
3. The selection of elasticities is essential for understanding sources and consequences of oil market movements. For instance, when we impose a smaller oil supply elasticity, say 0.02, the resulting demand elasticity is -0.4 , a value which is in the tail of the estimates in the literature. The smaller supply elasticity, in turn, implies that oil supply shocks explain about 80 and 15 percent of oil production and oil prices, respectively, and that oil-market specific demand shocks explain

²See, for instance, [Pindyck and Rotemberg \(1990\)](#) and [Cuddington and Jerrett \(2008\)](#).

³Industrial production in advanced economies does not provide any marginal improvement in forecasting power. We include advanced economies' activity in the SVAR to study how oil market shocks influence activity in advanced economies and we show that the lack of predictive power mostly reflects endogeneity bias.

5 and 55 percent of oil production and oil prices, respectively. Such decoupling between shocks that move oil prices and shocks that move oil production is avoided by our identification scheme.

Our paper contributes to the expanding literature that uses SVARs to study the interaction between the oil market and macroeconomic variables. Several studies have emphasized the importance of imposing restrictions on the oil supply elasticity to identify oil shocks.⁴ Unlike these papers, we jointly select oil supply and demand elasticities in a way that is coherent with both external information and the observed equilibrium relationships in the data. Our approach is also related to the recent study of [Baumeister and Hamilton \(2015b\)](#), who impose plausible prior distributions on both the oil supply and demand elasticities to identify oil shocks. Compared with this study, we show that small differences in the elasticities have a substantial effect on inference, in particular for decomposing movements in oil prices. We also show that, in the language of [Baumeister and Hamilton \(2015b\)](#), one could discipline the identification of the SVAR by imposing joint, rather than independent, prior distributions on oil market elasticities.

Regarding the data, our analysis builds on the seminal paper by [Kilian \(2009\)](#), who constructs a monthly measure of global economic activity based on data for dry cargo freight rates. As explained earlier, we replace his measure with three indicators to provide a richer characterization of global demand. These indicators boost the estimated contribution of global business cycle fluctuations to oil price movements. For instance, our identified global demand shocks explain about 35 percent of the historical fluctuations in oil prices, compared with the 8 percent estimate obtained by [Kilian \(2009\)](#) and the 4 percent estimate obtained by [Baumeister and Hamilton \(2015b\)](#). The use of IP follows [Aastveit et al. \(2015\)](#), who analyze the contribution of demand from advanced and emerging economies to movements in the price of oil using a factor-augmented vector autoregressive (FAVAR) model. The use of metal prices as an indicator of real activity has also been proposed by [Arezki and Blanchard \(2014\)](#), whose analysis exploits the idea that metal prices typically react to global activity even more than oil prices. Similarly, [Bernanke \(2016\)](#) argues that prices of commodities such as copper can act as 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.

Section 2 describes the indicators of global activity and their ability to capture the demand for oil. Section 3 discusses the econometric framework and the identification strategy. Section 4 presents the main results. Section 5 investigates the importance of selecting plausible oil elasticities. Section 6 extends our methodology to a setup that models explicitly oil inventories. Section 7 concludes.

⁴See, for instance, [Kilian and Murphy \(2012\)](#) as well as [Lippi and Nobili \(2012\)](#).

2 Measuring Global Demand for Oil

2.1 Overview

An empirically compelling structural model of the oil market must be able to separate the responses of oil market variables to oil-specific shocks from those to global economic developments. With this goal in mind, we select global economic indicators that fulfill two requirements: First, they must capture key salient features of the global business cycle, and, second, they must be good predictors of the global demand for oil. We consider four different monthly indicators. In this section we describe the indicators, their properties, and their roles as predictors of global oil demand.

2.2 Indicators of Global Activity

Our first two indicators of global activity are based on cross-country data on IP. We use data on IP for three reasons: (1) IP is widely available across a large number of economies, advanced and emerging alike; (2) IP is historically a very reliable business cycle indicator;⁵ and (3) IP can capture the demand for oil better than other indicators of real activity covering government and service sectors, as oil is an important factor of production in the industrial sector. We construct two indexes that aggregate IP data for advanced economies (*ya*) and for emerging economies (*ye*).⁶

The split between advanced and emerging economies is dictated by two pieces of evidence, which are summarized in detail in Appendix A.⁷ First, emerging economies use relatively more oil than advanced economies. In particular, even if emerging economies produce a smaller share of global GDP than advanced economies, their share of global oil consumption is similar. As a consequence, the effect on oil demand of a given change in IP from either group could be different. Second, emerging economies are—as a whole—oil independent, while advanced economies are—as a whole—net oil importers. In particular, emerging economies in our sample produce about 35 percent and absorb about 40 percent of global oil availability. By contrast, the advanced economies produce 20 percent and consume 40 percent of global oil availability. As a consequence, the effect on IP in advanced and emerging economies of a given oil shock could be different too.

Although intuitive and readily available, IP may be a lagging indicator of global oil demand, as

⁵The NBER treats industrial production as one of the main coincident indicators as part of its business cycle dating methodology. See, for instance, <http://www.nber.org/cycles/dec2008.html>.

⁶The Netherlands Bureau for Economic Policy Analysis also constructs IP indices for advanced and emerging economies from country-specific data, but their indices only start in 1991.

⁷Appendix A describes in the detail the construction of the IP indexes. Because of the availability of monthly IP data, our sample includes economies that account for 90 percent of global GDP (excluding both advanced and emerging economies somewhat proportionally), 55 percent of global oil production (because many oil producers do not have reliable IP data), and 80 percent of global oil consumption.

it may take time to alter production schedules in reaction to changing market conditions because of the presence of various types of adjustment costs. For this reason, we also consider two potential leading indicators of the global business cycle. The first indicator is the International Monetary Fund (IMF) metal price index (m). Base metals are the lifeblood of industrial activity, as they are critical material inputs in sectors whose cycles are normally synchronized such as capital equipment and heavy manufacturing, construction, infrastructure investment, and consumer durables. Moreover, production plans in these sectors are typically set up in advance, and, thus, purchase decisions for metal inputs are necessarily based on expectations of future market conditions. As a consequence, metal prices represent an ideal bellwether of shifts in current and expected economic activity at the global level.⁸

There is an extensive body of literature that supports our choice to include metal prices as indicators of the health of the global economy. The notion that movements in a set of commodity prices may be informative about expectations of global demand and macroeconomic conditions in general was originally investigated by [Pindyck and Rotemberg \(1990\)](#), who found that that co-movements in commodity prices can be indicative of shifts in market participants' expectations about the future paths of macroeconomic variables. Similarly, using factor analysis, [Labys et al. \(1999\)](#) show that movements in the common metal price factor appear to be led by fluctuations in the overall OECD industrial production index.⁹ In more recent related work, [Lombardi et al. \(2012\)](#) document that global economic activity responds positively to higher metal prices as captured by shocks to the common metal factor, thus lending empirical credibility to the idea that co-movements in metal prices reflect anticipations about the pace of expansion in global economic activity.¹⁰

The second indicator is the index of global real economic activity (rea) proposed by [Kilian \(2009\)](#) based on dry cargo 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 expected changes in global activity. All series are linearly detrended.

Table 1 shows the contemporaneous correlations among these variables over the 1985:M1–2015:M12 period corresponding to the sample used in the vector autoregressive (VAR) analysis that follows. We set 1985 as the starting date because of incomplete data availability for emerging economies IP prior to this date.¹¹ The IP indexes of advanced and emerging economies are negatively correlated. Metal prices

⁸[IMF \(2015a\)](#) and [Arezki and Matsumoto \(2014\)](#) provide an in-depth analysis of metal markets in the global economy.

⁹[Barsky and Kilian \(2001\)](#) note that, since the early 1970s, episodes of sustained increases in non-oil industrial commodity prices were consistent with a picture of rising global demand for all commodities, as those increases were too broad based to reflect commodity-specific supply shocks in individual markets.

¹⁰Other studies that find a meaningful link between cyclical fluctuations in global economic activity and movements in metal prices include [Cuddington and Jerrett \(2011\)](#), [Issler et al. \(2014\)](#), [Delle Chiaie et al. \(2015\)](#), and [Stuermer \(2016\)](#).

¹¹Although IP data are available for a handful of emerging economies starting in the 1960s, 1985:M1 is the earliest observation when our sample includes enough emerging economies to account, using today's GDP weights, for at least 25 percent of emerging economies' aggregate GDP.

display a highly positive correlation with emerging economies’ IP but are negatively correlated with advanced economies’ IP. Finally, *rea* co-moves with all other indicators, in line with the interpretation in Kilian (2009) that it represents a broad indicator of global activity. All told, the different correlations suggest that these indicators may capture distinct aspects of the global business cycle.

Do metal prices and the global real activity index lead the business cycle? To test this hypothesis, we look at the ability of *m* and *rea* to predict activity in advanced and emerging economies by estimating the following forecasting regression:

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

where $h \geq 0$ is the forecast horizon, $y_t = \{ya_t, ye_t\}$, ϵ_{t+h} is the forecast error, and ρ and δ control for past observations. Our interest is in understanding the different predictive power of *m* and *rea* for either *ya* or *ye*, as measured by β_1 and β_2 , once the past state of the economy is conditioned for.

To facilitate the comparison between metal prices and the real economic activity index, we report in table 2 the standardized estimates of the coefficients β_1 and β_2 . Both *m* and *rea* have predictive power for advanced economies’ IP, as shown by panel A. In particular, *m* predicts *ya* for up to a year, while *rea* is informative for up to six months and has a lower predictive power than *m*. Both indicators are also reliable predictors of near-term economic developments in emerging economies (panel B). At horizons up to one year, the coefficients imply an economically and statistically significant relationship between *m*, *rea*, and industrial production, particularly for *m*. For example, using the standardized estimate of $\beta_1 \simeq 1$ from the regression, an increase of about 30 percent in metal prices in month t —corresponding to a one standard-deviation increase—implies a nearly 4 percent increase in *ye*—corresponding to a one-standard deviation increase—over a period ranging from six months to about two years. Similarly, using the standardized estimate of $\beta_2 \simeq 0.2$ (an average across the 6, 12 and 24-month horizons), a one standard-deviation increase in *rea* predicts a 0.8 percent increase in *ye* six months ahead. In sum, both *m* and *rea* appear to be powerful predictors of economic activity, with the predictive power of *m* looking more stable across horizons and quantitatively larger. The finding that the predictive power of *rea* for *ya* and *ye* is less stable and quantitatively smaller than that of *m* is in line with the evidence presented in Beidas-Strom and Pescatori (2014) that swings in *rea* and in global IP do not always coincide, particularly at times of structural change in the shipping sector.¹²

¹²The economic and financial press has also emphasized how the surge in ship delivery and the resulting glut of large vessels has led to a plunge in dry bulk cargo shipping rates that has hampered the ability of indexes based on those rates to represent a valid gauge of global trade and global economic activity more in general. See, for example, the *Financial Times* articles by Robert Wright, “Dry bulk rates tumble on ship delivery surge” (January 11, 2012), and “Shipping costs fall to 25-year low” (February 1, 2012).

2.3 Global Activity as a Predictor of Oil Prices

We now assess the role of the various activity indicators—both the contemporaneous ones, that is, ya and ye , and the “leading” ones, that is, m and rea —in predicting oil prices. Table 3 displays pairwise cross-correlations between oil prices (p) and the activity indicators at various leads and lags. The cross-correlation between p and ya is negative and becomes larger at longer horizons, with oil prices negatively leading ya . The correlations between p and ye and between p and m are positive and large, while the correlations between p and rea are positive but small.

To better gauge whether our indexes of global activity capture global demand for oil, we explore their relative roles in predicting oil prices. Specifically, we estimate the following forecasting regression:

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

where $h \geq 0$ is the forecast horizon, ϵ_{t+h} is the forecast error, and the ρ and γ are coefficients that control for past observations. As in the previous set of forecasting regressions, our interest is in understanding the different predictive power of the various economic indicators for oil prices, as measured by the standardized estimates of the coefficients β_1 , β_2 , β_3 , and β_4 (once the past state of the economy is conditioned for).

The results are reported in table 4. Reading across rows shows how the indicators differ sharply in their predictive abilities. The coefficients on ya are not statistically significant, and have the “wrong” sign in that higher IP in advanced economies predicts lower oil prices over a 2-year horizon. Instead, the coefficients on ye and m are positive, and, in general, statistically significant. In particular, the coefficients on m indicate that metal prices are the only indicator with predictive power that extends beyond one year. The coefficients on rea are economically small and not statistically significant.

Taken together, results from the forecasting exercises show that the marginal contribution of rea to IP, on the one hand, and to oil prices dynamics on the other, is small once we condition on the other activity indicators in the regressions. For this reason, we leave rea out of the baseline VAR model that follows, and we report the results from a VAR model that includes rea in Appendix C.

The evidence on the predictive ability of various measures of global activity for oil prices is mixed, thus suggesting the opportunity to include a broader set of indicators in the VAR model rather than a single one. In particular, the index of metal prices is included in the set of endogenous variables in the VAR also for the purpose of controlling for information that economic agents may have about future global demand and that is not adequately captured by the other activity variables in the system.

Of course, the forecasting regressions offer little guidance about the causal relationship between global activity and oil prices. For instance, they are silent on the relative importance of shocks to global activity and shocks to the oil market in shaping the dynamic correlations between these oil and non-oil variables. To address this question empirically, we move to the SVAR analysis, the subject of the remainder of the paper.

3 Identifying Oil Market and Global Activity Shocks

3.1 Overview

Suppose that the economic structure describing the oil market and its relationship with the global economy is given by the following SVAR model:

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

where \mathbf{X} is a vector of oil market and macroeconomic variables, \mathbf{u}_t is the vector of structural shocks, p is the lag length, and \mathbf{A} and $\boldsymbol{\alpha}_j$ for $j = 1, \dots, p$ are matrices of structural parameters. The vector \mathbf{u}_t is Gaussian with mean zero and variance-covariance matrix $E[\mathbf{u}_t \mathbf{u}_t'] = \boldsymbol{\Sigma}_u$. Without loss of generality, we normalize the entries on the main diagonal of \mathbf{A} to 1 and we assume that $\boldsymbol{\Sigma}_u$ is a diagonal matrix. The reduced-form representation for \mathbf{X}_t is the following:

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

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

$$\boldsymbol{\varepsilon}_t = \mathbf{B}\mathbf{u}_t, \quad (5)$$

$$\boldsymbol{\Sigma}_\varepsilon = \mathbf{B}\boldsymbol{\Sigma}_u\mathbf{B}', \quad (6)$$

where $\mathbf{B} = \mathbf{A}^{-1}$, so that, alternatively, $\mathbf{u}_t = \mathbf{A}\boldsymbol{\varepsilon}_t$. Estimation of the reduced-form VAR allows to obtain a consistent estimate of the $n(n+1)/2$ distinct entries of $E[\boldsymbol{\varepsilon}_t \boldsymbol{\varepsilon}_t'] = \boldsymbol{\Sigma}_\varepsilon$. Hence, to recover the n^2 unknown entries of \mathbf{B} and $\boldsymbol{\Sigma}_u$, we need to make $(n-1)n/2$ identifying assumptions.

To discuss our identification strategy, it is useful to list the set of endogenous variables that we include in the model. The benchmark monthly VAR specification consists of five endogenous variables ordered as follows: (1) log global 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 of oil prices, p_t ; and (5) the log of the IMF metal price index, m_t .

We next discuss our identification strategy 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.

3.2 The Oil Market

Our structural VAR model specifies an oil supply and an oil demand equation, and we interpret innovations to these equations as oil supply and oil demand shocks, respectively. After controlling for their own lags and for other determinants of oil supply and demand, these equations define a short-run oil supply curve and a short-run oil demand curve. If we consider a simple bivariate system of order zero 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 η_S and η_D are the short-run price elasticity of oil supply and oil demand, and $u_{s,t}$ and $u_{d,t}$ are the oil supply and oil demand shocks, assumed to be mutually orthogonal, with standard deviations given by σ_S and σ_D .

The identification problem can be summarized as follows. By construction, the identification of the SVAR has to select parameters that match the estimated variance-covariance matrix of the reduced-form VAR residuals Σ_ε . In the data, the volatility of oil prices is relatively large, the volatility of oil production is relatively small, and their covariance is close to zero. These three target values can be hit with four structural parameters: the standard deviations of the oil supply and demand shocks as well as the oil supply and demand elasticities. In what follows, we explore how assumptions on the oil demand elasticity restrict the oil supply elasticity, and vice versa. As routine in the literature, we do not directly impose restrictions on the variance of the structural shocks, which are instead pinned down by the restrictions imposed on the elasticities.

Figure 2 presents six combinations of slopes of supply and demand schedules in the oil market that could be used to match the data. Panels A and B plot two combinations that achieve the targets. Panel A shows that, given oil supply and oil demand curves, the right mix of supply and demand shocks can generate the “right” volatility of prices and quantities and zero correlation between them. Panel B illustrates another observationally equivalent solution: a highly inelastic supply curve and a highly elastic demand curve. 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.

Although the oil market configurations plotted in panels A and B are observationally equivalent, their economic implications are different. When both curves are elastic, oil supply shocks generate a negative co-movement between prices and quantities, while oil demand shocks generate a positive co-movement between prices and quantities. That is, the unconditional zero correlation between prices and quantities observed in the data is accounted for by two shocks that generate non-zero conditional correlations. By contrast, when the supply curve is inelastic and the oil demand curve is very elastic, there is a disconnect between movements in quantities and movements in prices in the oil market, not only unconditionally, but also conditional on each separate shock.

The remaining panels show why other combinations of supply and demand elasticities are inadmissible. The middle panels (C and D) show that the estimated variance-covariance matrix Σ_ϵ rules out the possibility of both highly inelastic supply and inelastic demand. When both the supply and demand curves are inelastic, two configurations may arise. With supply and demand shocks that are both small, the SVAR matches the volatility of prices but underestimates the volatility of oil production (panel C). Alternatively, when both shocks are large, the SVAR can match the volatility of production but overpredicts the volatility of oil prices (panel D).

The lower panels (E and F) illustrate why the estimated variance-covariance matrix Σ_ϵ also rules out both highly elastic supply and highly elastic demand. Two configurations may also arise in this case. With supply and demand shocks that are large, the SVAR can match the volatility of prices but overpredicts the volatility of oil production (panel E). Alternatively, with supply and demand shocks that are small, the SVAR can match the volatility of production but underestimates the volatility of oil prices (panel F).¹³

More formally, given values for the entries of Σ_ϵ , the variance-covariance structure of the VAR imposes a unique relationship between η_S and η_D . The red solid line in figure 3 plots this relationship in our model computed by setting Σ_ϵ at its OLS estimate.¹⁴ The two observationally equivalent configurations depicted in panels A and B of figure 2 denote just two points along a continuum of models that are graphically represented by the curve plotted in figure 3. 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 information outside the model.

Our identification proceeds as follows. First, we search the related empirical literature on oil market

¹³One can always find a solution where both supply and demand elasticity have the “right” sign, and the estimated variances of supply and demand shocks match any desired correlation between prices and production. There are, however, cases in which the solution can call for (1) 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” or (2) 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.

¹⁴We extract the entries of Σ_ϵ related to oil prices and quantities from the five-equation VAR. The estimation of the bivariate system returns a similar variance-covariance matrix of residuals.

price elasticities and compile a list of studies that estimate short-run supply and demand elasticities. We list these studies and the associated elasticities in table A.4. Six studies estimate the short-run price elasticity of oil supply: Half of them estimate a supply elasticity of about 0.25, two of them found elasticities near zero, and one study estimates a negative supply elasticity. By contrast, thirty studies estimate the short-run price elasticity of demand. Estimates of the demand elasticity range from -0.9 to -0.03 , with the bulk of estimates between -0.3 and -0.1 .

The “consensus” in the literature, as summarized by the median elasticity, is that the supply elasticity is 0.13 and the demand elasticity is -0.13 . The blue circle in figure 3 graphically represents this pair of elasticities. Importantly, the blue circle does not lie on the red solid line, which means that, on the basis of our SVAR, there is a slight tension between what the literature finds and what our SVAR admits. This tension between supply and demand elasticities arises for any selection of the value for the supply elasticity in table A.4: Through the lenses of the SVAR, values for the supply elasticity on the low and the high ends of the estimates presented in table A.4 are incompatible with a demand elasticity of around -0.2 . To emphasize this point, in Appendix C we report results for a SVAR model identified by targeting a supply elasticity of 0.02, as in Kilian and Murphy (2012).

Second, our identification selects a pair of admissible elasticities by minimizing 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 η_S . If we denote η_D as a function of η_S and the variance-covariance matrix of residuals, $\eta_D(\eta_S; \Sigma_\varepsilon)$, our identification strategy solves the following problem:¹⁵

$$\min_{\eta_S} \begin{bmatrix} \eta_S - \eta_S^* \\ \eta_D(\eta_S; \Sigma_\varepsilon) - \eta_D^* \end{bmatrix} V^{-1} \begin{bmatrix} \eta_S - \eta_S^* \\ \eta_D(\eta_S; \Sigma_\varepsilon) - \eta_D^* \end{bmatrix}, \quad (9)$$

where η_S^* and η_D^* are the target values for the supply and demand elasticities, respectively, and V is a diagonal matrix of weights given by the variances of the elasticities reported in table A.4. Our identification scheme selects the green square in figure 3 corresponding to the point on the curve that is as close as possible to the blue circle. Such identification yields, for Σ_ε evaluated at its OLS estimate, an oil supply elasticity of 0.11 and an oil demand elasticity of -0.11 . The green square in figure 3 graphically represents this pair of elasticities.

Panel A in figure 2 illustrates the importance of the values of these elasticities for characterizing movements in the oil market by plotting the oil supply and oil demand curves implied by our VAR estimates. The elasticities imply that an exogenous decline in global crude oil supply of 1 percent

¹⁵Alternatively, we could have parameterized the problem either by imposing the restriction on η_D , in which case $\eta_S(\eta_D; \Sigma_\varepsilon)$, or, following Uhlig (2005) and Rubio-Ramírez et al. (2010) among many, in terms of an orthonormal matrix Q , so that $\eta_S(\Omega; \Sigma_\varepsilon)$ and $\eta_D(\Omega; \Sigma_\varepsilon)$.

(approximately 0.8 million barrels per day) should lead, over a one-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).

Comparing our Identification Scheme with Existing Studies. Our characterization of the identification problem builds on several studies in the literature. [Lippi and Nobili \(2012\)](#) identify oil shocks by imposing sign restrictions on the impulse responses of models variables to oil and non-oil shocks. As already noted by [Kilian and Murphy \(2012\)](#), SVARs identified using the sign restriction approach are likely to include models with counterfactual implications about the oil supply elasticity. This observation motivated the subsequent strategy of combining sign restrictions on impulse responses with bounds on the oil supply elasticity ranging from 0 to about 0.02 ([Kilian and Murphy, 2012, 2014](#)). The relationship between η_S and η_D depicted in our figure 3 shows that the set of models identified by imposing sign restrictions in combination with these bounds might include SVARs with (1) extremely different implications for the oil market and (2) with potentially counterfactual values for the elasticity of oil demand. In fact, we show that an upper bound on the supply elasticity imposes a lower bound on the demand elasticity. Moreover, because the mapping is very steep in this region, a tight bound on the supply elasticity results in a very large set for the values of demand elasticities that, despite its size, might exclude plausible low values of the demand elasticity.

In recent work, [Baumeister and Hamilton \(2015b\)](#) acknowledge the importance of specifying plausible priors on both the oil supply and demand elasticities.¹⁶ What emerges from our analysis is that one could use the joint distribution of the oil supply and demand elasticities to narrow down the set of admissible models. Moreover, under the assumption implicitly made in the literature that VAR models provide a good description of the dynamic interaction between oil prices and quantities, our results highlight the importance of understanding the cross-equation restrictions imposed by the data in the context of prior selection. For instance, priors centered around low values for both elasticities—say, $\eta_S = -0.02$ and $\eta_D = 0.02$ —would describe the oil market structures depicted in panels C and D of figure 2, which are inconsistent with the data. If the prior standard deviations are sufficiently large, the cost of imposing these prior distributions is to give weight to implausible models in the posterior, at the expense of sharp inference. By contrast, if the prior standard deviations are tight, the resulting SVAR would generate a covariance between prices and quantities that is different from what observed in the data, akin to imposing over-identifying restrictions.

¹⁶See also [Baumeister and Hamilton \(2015a\)](#).

3.3 The Interaction between the Oil Market and Global Activity

With a modeling of the oil market on hand, the next step is to model its interaction with macroeconomic variables. 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 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 describe the joint modeling of the oil market and the global economy:¹⁷

$$q_t = \eta_S p_t + u_{s,t}, \quad (10)$$

$$q_t = \eta_A y a_t + \eta_E y e_t + \eta_D p_t + u_{d,t}, \quad (11)$$

$$y a_t = \nu_Q q_t + u_{y a,t}, \quad (12)$$

$$y e_t = \mu_Q q_t + \mu_A y a_t + u_{y e,t}, \quad (13)$$

$$m_t = \psi_Q q_t + \psi_A y a_t + \psi_E y e_t + \psi_P p_t + u_{m,t}. \quad (14)$$

Equations (10)–(14) summarize the restrictions we impose on the entries of matrix \mathbf{A} . Equation (10) describes the oil supply schedule, and we assume that oil production q_t responds contemporaneously only to changes in oil prices. The short-run supply elasticity η_S is the solution to the minimization problem described in equation (9). The supply shock $u_{s,t}$ captures disturbances to oil supply due, for instance, to wars, sanctions, and technological progress in the oil industry. Equation (11) describes the oil demand schedule: oil demand is allowed to respond contemporaneously to $y a_t$ and $y e_t$ and to changes in oil prices. The demand price elasticity η_D is defined as the change in desired demand q for a given change in oil prices p , holding global activity constant. The oil-specific demand shock $u_{d,t}$ captures changes in oil prices due to speculation, changes in oil intensity, and precautionary demand for oil due to oil price volatility.¹⁸ Our restrictions imply that global activity can only shift contemporaneously the oil demand equation. That is, for a given oil price, demand rises as global activity increases, while producers' decisions are not affected within the period. Hence, 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 quantities because oil demand shifts along a positively sloped supply curve.

Equation (12) determines activity in advanced economies. We assume that $y a_t$ responds within the period only to oil production. Similarly, equation (13) determines activity in emerging economies, $y e_t$. We assume that $y e_t$ responds within the period to $y a_t$ and to oil production. Our assumption that $y e_t$

¹⁷For expositional convenience we omit the lagged terms, which instead appear in equation (15).

¹⁸The studies by Juvenal and Petrella (2015) and Beidas-Strom and Pescatori (2014) examine more explicitly the role of speculation in driving oil price fluctuations.

reacts contemporaneously to ya_t is meant to capture the idea that exports to advanced economies are an important component of aggregate demand in emerging economies. Finally, equation (14) determines metal 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 and also captures current developments in global activity that are not adequately accounted for by IP indexes.¹⁹

To summarize, in matrix notation, our restrictions assumes that the relationship between oil and macroeconomic variables is described by the following set of equations:

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

where ν_Q and μ_Q can be interpreted as the short-run elasticities of economic activity to oil production, and η_A, η_E denotes the elasticity of oil demand to economic activity.

Together with the knowledge of either η_D or η_S , the additional zero restrictions in the matrix \mathbf{A} satisfy the necessary and sufficient coefficients for identification of (Rubio-Ramírez et al., 2010), and allow us to identify uniquely the structural parameters ($\nu_Q, \mu_Q, \mu_A, \eta_A, \eta_E, \eta_D, \psi_Q, \psi_A, \psi_E, \psi_P, \boldsymbol{\Sigma}_u$) given information from $\boldsymbol{\Sigma}_\varepsilon$, the variance-covariance matrix of the reduced-form VAR residuals.

4 Main Results

4.1 Overview

In this section, we first present the impulse responses implied by our SVAR. We then examine the contribution of the identified shocks to fluctuations in oil market variables and economic activity by presenting results from the forecast error variance decomposition and the historical decomposition. 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 Negro and Schorfheide, 2011). The resulting specification, which includes a constant, is estimated over the 1985:M1–2015:M12

¹⁹Our model assumes that the oil market has a contemporaneous and direct effect on both ya_t and ye_t only through changes in q_t , as oil is an input in the production of manufacturing goods. Note that changes in oil prices have an indirect contemporaneous effect on real activity by inducing changes in oil production. The assumption that metal prices react to oil prices, and not vice versa, is less rooted in economic theory. For this reason, we explored a model that imposes the alternative ordering between p and m and found very similar results to those reported below.

period using 12 lags of the endogenous variables.²⁰

4.2 Impulse Responses

The solid lines in the left column of figure 4 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 prices, which rise by 6 percent on impact and remain elevated thereafter. On the activity side, the response of activity between advanced and emerging economies is markedly different. IP in advanced economies declines gradually, bottoming out at -0.35 percent two and a half years after the shock. The negative response is consistent with the fact that, on average, our group of advanced economies are oil dependent. In contrast, IP in emerging economies rises rapidly, peaking after six months at 0.2 percent above baseline. Hence, emerging economies, which are on average oil independent, experience positive effects from a supply-driven increase in oil prices.²¹

The right column of figure 4 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, with IP increases mildly in both groups of economies for six months. Thereafter, real activity contracts in advanced economies while remaining elevated in emerging economies, even though the responses are economically small and only marginally significant.

Figure 5 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, and the right column the responses to a metal price shock. The three shocks generate correlations that are typical of demand-driven business cycle fluctuations: The increase in real activity in advanced and emerging economies, which is accompanied by a persistent increase in metal prices, is associated with a rise in both 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. In contrast to the forecasting regressions—which predict a decline in oil prices following an increase in advanced economies’ activity—the SVAR attributes the negative correlation between oil prices and

²⁰The vector of hyperparameters of the Minnesota prior is $\lambda = [1, 2, 1, 3, 3]$. 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.

²¹Iacoviello (2016) finds that oil-exporting countries experience a rise in consumption and GDP following supply-driven increases in oil prices. Peersman and Van Robays (2012) also find that the effects of exogenous oil supply shocks on economic activity are negative for net oil importing countries and insignificant or even positive for net oil exporters.

advanced economies’ activity documented in table 3 to oil supply shocks, with all three global activity shocks generating a positive co-movement between these two variables. Nonetheless, an increase in advanced economies’ activity induces only a mild and short-lived increase in oil prices.

The positive responses of IP in both advanced and emerging economies to a shock to the metal price index support the view that metal prices are a valid leading indicator of global activity. Additionally, the finding that activity in emerging economies reacts positively to oil price and metal price increases corroborates our identification strategy. Additionally, our findings are also consistent with the literature that emphasizes commodity prices as drivers of business cycles in emerging economies.²²

4.3 Forecast Error Variance Decomposition

Up to now, we have characterized the dynamic effects of oil market and global activity shocks. This section explores the relative importance of oil and non-oil shocks for fluctuations in oil prices and global economic activity. We ask whether movements in oil prices are an important indicator of shifts in economic activity and whether oil shocks are an important driver, on average, of global business cycles.

Table 5 shows the variation in the two-year-ahead forecast error variance in oil prices, oil production, advanced economies’ IP and emerging economies’ IP that is attributable to the five structural shocks identified by our SVAR model. We choose to focus on the two-year horizon, as we consider it the typical horizon for business cycle analysis. As shown in the first row of table 5, about three-fourths of the fluctuations in oil prices are due to shocks that originate in the oil market. In particular, shocks to oil supply account for more than 50 percent of the forecast error variance, while oil-specific demand shocks explain about 15 percent. Movements in global demand are also a significant source of oil price fluctuations, explaining the remaining 25 percent of oil price fluctuations.²³ Shocks to activity in emerging economies and in metal prices are equally important drivers of oil prices, while the contribution of shocks to activity in advanced economies is almost zero.²⁴ The second row of table 5 shows that oil production is equally driven by oil supply and oil demand shocks, which together account for 70 percent of its volatility.

²²The evidence in IMF (2015b) suggests that fluctuations in commodity prices and commodity terms-of-trade can lead to sizable output fluctuations in commodity exporters. Fernandez et al. (2015) also find that shocks to the prices of exported commodities explain a considerable portion of business cycle fluctuations in emerging economies. Furthermore, our findings are also consistent with those from the literature on emerging economies’ business cycle that terms-of-trade and world price shocks explain a significant fraction of cyclical fluctuations and that, in turn, variations in terms-of-trade are largely explained by variations in the prices of a few exported commodities. See for instance Mendoza (1995).

²³Because the median contribution is calculated for each shock independently, the rows do not sum to 100 percent. The medians of the total contributions to oil price fluctuations of oil-specific shocks and of global demand shocks are 72.2 percent and 25.3 percent, respectively.

²⁴Stock and Watson (2016) use factor models and find that the contribution of global demand shocks to oil price fluctuations ranges between 20 and 40 percent, with their estimates reasonably close to our findings.

The third and fourth rows of table 5 show that, on average, oil-specific shocks contribute little to the volatility in real activity. Oil supply shocks account for about 5 percent of the forecast error variance of advanced economies’ and emerging economies’ IP, while the contribution of oil-specific demand shocks is a little below 2 percent in both advanced and emerging economies. Not surprisingly, activity variables are mostly driven by their own shocks. Yet, in line with the results from the forecast regressions, we find that shocks to metal prices, which are ordered last in our SVAR, account for 15 percent of the forecast error variance in advanced economies and for 30 percent in emerging economies.

Summing up, our results show that fluctuations in oil prices are primarily driven, on average, by oil supply shocks. In addition, movements in oil prices driven by global demand mostly reflect changes in oil demand from emerging rather than advanced economies. We also find that oil shocks have a significant but secondary role in driving IP in both advanced and emerging economies, accounting for slightly less than 10 percent of the variability at business cycle frequency.

4.4 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.²⁵ We first give an overview of the 1986–2015 period, presenting the data at an annual frequency. We then zoom in on three important episodes in the sample involving large changes in the price of oil. The first episode is centered on the Asian financial crisis, the second episode focuses on the period of the global financial crisis, and 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. For these three episodes, we calculate the cumulative effects on each variable of shocks that materialized from the onset of the event onwards.²⁶

Figure 6 shows the decompositions across the whole sample for the oil prices, oil production, and IP in both advanced economies and emerging economies. Overall, our decomposition show that, over the whole sample period, fluctuations in the price of oil and in oil production were mostly determined by oil-market specific shocks, in particular by supply shocks. Nonetheless, global activity contributed significantly to shaping fluctuations in 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 economies’ and emerging economies’ real activity. However, the quantitative importance of oil shocks, and, in particular, of oil supply shocks has varied considerably over the past three decades. Oil supply shocks played an

²⁵We compute the historical decomposition using the OLS estimates of the reduced-form parameters. The sample used for the estimation includes both actual data and the dummy observations used to implement the Minnesota prior.

²⁶For these episodes, our decomposition sets previous shocks to zero, which explains why the vertical bars in figures 7 to 9, 12, and 13 do not sum to the actual data depicted by the solid lines.

important role in the late 1990s and since 2012, but they had a smaller role in the rest of the sample.

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

The Asian financial crisis. Figure 7 plots the historical decomposition of the model variables for the period of the Asian Financial Crisis from July 1997 to December 1998. All the 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 lower-right panel. The decline in the demand of oil from emerging economies induced downward pressure on oil prices and, together with shocks to metal prices, accounted for about one-third of the decline in oil prices.²⁷

Throughout this period, despite a lower demand for oil from emerging countries, a few oil exporters, most notably Iraq, did not cut production and, on the contrary, increased production throughout early 1999. Accordingly, our model attributes a major role in the decline of the price of oil to supply shocks. Oil-specific demand shocks represented an offsetting factor for oil prices and, at the same time, account for about 15 percent of the increase in oil production in late 1997.

The effect of disturbances originated in the oil market on economic activity is positive but economically small in advanced economies, while it 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 8, we plot the estimated historical decomposition of the model variables for the period of the global financial crisis 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 in oil production and the relatively small movements in global activity factors. The role of global activity shocks as drivers of the price of oil becomes more prominent toward the end of 2008, as the sharp decline in real activity materializes in both 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 lens of the model, the sharp contraction in

²⁷Baumeister and Kilian (2016a) attribute a large portion of the slide to reduced demand for oil, arguably caused by the Asian financial crisis, which was followed by economic crises in other countries including Russia, Brazil, and Argentina.

global activity at the end of 2008 should have led to an even larger reduction in oil production than that observed in the data. The higher-than-expected oil production is rationalized by the model through the only shock that generates a negative co-movement between oil production and the price of oil, the oil supply shock.

Finally, the decomposition of the real activity variables plotted in the lower row of figure 8 shows that shocks originating in the oil market had nearly no role in the collapse of economic activity associated with the Global Financial Crisis.

The 2014–15 slump. Figure 9 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 the 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, as also acknowledged by [Kilian \(2016\)](#). 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—for instance, 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. Since early 2015, the decline in the price of oil also began to reflect negative shocks both in emerging economies and to expectations of global activity as captured by the shock to metal prices.²⁸

As shown in the lower row of figure 9, oil shocks had no role in shaping economic activity in advanced economies until mid-2015. Since June 2015, oil supply shocks have added about 1 percent to growth in industrial production in advanced economies. 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 0.1 percent as of October 2014, rising subsequently to about 1 percent at the end of 2015. The boost in economic activity in advanced economies in conjunction with the drag on economic activity in emerging economies of about equal size thus rationalizes the muted response of global activity to oil market developments.

²⁸The gap between the vertical bars and the black line in the upper-left panel of figure 9 shows that about 30 percent of the decline in oil prices was actually predictable, as it reflected the cumulative effects of earlier shocks, slightly smaller than the 50 percent value found by [Baumeister and Kilian \(2016b\)](#).

4.5 Summary of Findings

All told, the analysis in this section shows that (1) under our identification, oil supply and oil-specific demand shocks induce significant movements in oil prices and quantities; (2) shocks to global demand account, on average, for about 25 percent of fluctuations in oil prices, and are important drivers of oil prices during some historical episodes; and (3) shocks that originate in the oil market have asymmetric effects on activity in advanced and emerging economies.

5 The Role of Oil Price Elasticities

5.1 Overview

In this section we analyze the sensitivity of our results to changes in the assumptions about the oil supply elasticity. We first show how small differences in the oil supply elasticity can change the sign of the predicted effects on economic activity of shocks originating in the oil market. We then show how different assumption about the elasticity change the share of the variance of oil and activity variables explained by the various shocks. Moreover, we show how changes in the oil supply elasticity affect the inference about the driving forces of oil prices during two episodes of large oil price movements. This exercise also highlights how the use of multiple indicators of global demand for oil is instrumental toward the goal of disentangling oil market-specific shocks and global demand shocks.

5.2 Global Effects of Oil Market Shocks

Figure 10 illustrates the importance of selecting among alternative oil supply elasticities in order to gauge the aggregate effects of shocks originating in the oil market. To make our point, we consider the hypothetical question of how economic activity would respond to a 1 percent change in the price of oil. The upper panel plots the two-year elasticity of IP in advanced economies to oil prices changes driven by oil supply shocks (the red solid line) and by oil-specific demand shocks (the blue dashed line) as a function of the oil supply elasticity. Similarly, the lower panel plots the corresponding two-year elasticity of IP in emerging economies.

Under our benchmark identification scheme (shown by the vertical line corresponding to $\eta_S = 0.11$), a 1 percent rise in oil prices driven by an oil supply shock leads to a 0.04 percent drop in advanced economies' IP and to a 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 advanced economies, even as its effects on emerging economies are roughly unchanged. This result suggests that oil-specific

demand shocks capture, at least in part, higher demand for energy-intensive goods (such as cars or heavy equipment) driven by better growth prospects in emerging economies.

The upper panel shows that as the oil supply elasticity rises (and, consequently, the oil demand elasticity falls), the VAR predicts more benign effects on activity in advanced economies from oil-market driven changes in oil prices. For instance, the IP elasticity evaluated at $\eta_S = 0$, the value used in Kilian (2009), is -0.1, twice as large as under our identification. By contrast, as shown in the lower panel, as the oil supply elasticity rises, supply-driven increases in oil prices have more positive effects on emerging economies' activity, whereas demand-driven increases in oil prices have more negative effects.

5.3 Oil Price Elasticities and Volatilities

Figure 11 plots the share of the variance of the two-year-ahead forecast errors for oil prices, IP in advanced economies, and IP in emerging economies that is attributable to oil shocks as a function of the supply elasticity. The red solid line plots the contribution of oil supply shocks, the blue solid line plots the contribution of oil-specific demand shocks, and the black dashed line plots their sum. The remaining share of variability is due to global activity shocks.

The left panel of figure 11 shows the forecast error variance decomposition for oil prices. The relative importance of oil-specific demand and supply shocks changes considerably as the oil supply elasticity varies. As reported in table 5, according to our benchmark identification scheme (shown by the vertical line), oil-specific shocks explain about 65 percent of the two-year forecast error variance of oil prices, with oil supply shocks accounting for the largest share, about 50 percent. By contrast, if the supply elasticity were close to zero, oil supply shocks would explain none of the oil-market driven volatility in prices, and, if the supply elasticity were to rise to 0.2, would explain nearly all of the oil-market driven volatility in prices. On the contrary, the assumption on the oil supply elasticity has only a minor effect on the overall share of volatility that is explained by oil shocks, which rises from 65 percent for a supply elasticity equal to zero to 70 percent for a supply elasticity equal to 0.2.

The middle and right panels of figure 11 show the share of two-year-ahead volatility in advanced economies' and emerging economies' activity that is explained by oil market shocks, again as a function of the oil supply elasticity. As the two panels show, while the relative importance of oil demand and oil supply shocks changes considerably as the oil supply elasticity varies, the overall share of volatility in advanced and emerging economies that is explained by oil shocks is small regardless of the value of elasticity, ranging between 5 and 10 percent. Nonetheless, and in line with the comparison of the IP elasticities to oil prices, we find that an increase in the oil supply elasticity decreases the share of fluctuations in advanced economies' IP explained by oil market-specific shocks, whereas it increases the share of fluctuations in emerging economies' IP generated by the oil shocks.

5.4 Oil Price Elasticities and the Decoupling Puzzle

Figures 12 and 13 compare the historical decomposition for oil prices and oil production during the Asian financial crisis and during the 2014–15 oil slump obtained with three different specifications of the SVAR model.

Panel A in figure 12 reproduces the historical decomposition of the Asian financial crisis under our identification scheme. As already discussed, our decomposition is characterized by (1) the joint movement in oil prices and oil production induced by positive oil supply shocks, and (2) the importance of the decline in emerging economies activity in accounting for the decline in oil prices. 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. That is, the zero supply elasticity model decouples fluctuations in oil supply and oil demand shocks. Finally, panel C shows the results of a SVAR that retains the zero oil supply elasticity but replaces the three global activity indicators with the real global activity index. The comparison between panels A and C illustrates the difference between our global activity indicators and Kilian (2009)’s index. It is remarkable how such an index fails to find any connection between the sharp slowdown in emerging economies activity and the concurrent decline in oil prices.

In figure 13 we apply a similar model comparison to the 2014–15 oil price slump. The comparison between panels A and B again highlights how the decoupling between oil supply and oil demand shocks introduced by setting the supply elasticity to zero is also apparent during this episode. This decoupling seems particularly at odds with the conventional narrative that attributes much of the decline in oil prices to positive surprises in shale production in the United States and to the resilient production in OPEC countries. Finally, panel C shows that the real global activity index explains a smaller portion of the decline in oil prices relative to our three indicators, but that the difference is not as significant as for the Asian financial crisis.

6 Adding Oil Inventories to the VAR

Our baseline specification abstracts from inventories, by assuming that oil production is fully absorbed by consumption in every period. Total oil inventories are, across countries, a small multiple of the flow of oil production within a month, thus making it hard to conjecture that changes in inventories can absorb large gaps between production and consumption at monthly frequency. For instance, U.S. crude oil stocks are less than twice monthly U.S. crude oil production.²⁹ However, inventories could in principle

²⁹In 2015, average stocks of crude oil excluding the Strategic Petroleum Reserve were 481 million barrels (Source: US Department of Energy). Average monthly crude oil production was about 285 million barrels (Source: Oil and Gas Journal).

quickly move to absorb differences between oil production and oil consumption, in turn affecting the dynamics of the oil market. To account for this possibility, we follow [Kilian and Murphy \(2014\)](#) and [Baumeister and Hamilton \(2015b\)](#) and extend our model of the oil market presented in Section 3 to include oil inventories. Define the change in inventories (ΔI) as the difference between oil production (Q) and oil consumption (C):

$$\Delta I_t = Q_t - C_t. \quad (16)$$

Recall that our benchmark VAR includes $q_t = 100 \log(Q_t/\bar{Q}_t)$, where \bar{Q}_t is the trend level of production constructed assuming a constant growth rate. Given this scaling, we can express detrended oil consumption as $c_t = q_t - \Delta i_t$, where $\Delta i_t = 100\Delta I_t/\bar{Q}_t$. This leads to the following structural model:

$$q_t = \eta_S p_t + u_{s,t}, \quad (17)$$

$$q_t - \Delta i_t = \eta_A y a_t + \eta_E y e_t + \eta_D p_t + u_{d,t}, \quad (18)$$

$$y a_t = \nu_Q q_t + u_{ya,t}, \quad (19)$$

$$y e_t = \mu_Q q_t + \mu_A y a_t + u_{ye,t}, \quad (20)$$

$$m_t = \psi_Q q_t + \psi_A y a_t + \psi_E y e_t + \psi_P p_t + u_{m,t}, \quad (21)$$

$$\Delta i_t = \phi_Q q_t + \phi_A y a_t + \phi_E y e_t + \phi_P p_t + \phi_M m_t + u_{i,t}. \quad (22)$$

There are three differences compared to our baseline model described in Equations (10)–(14). First, the inclusion of inventories in the VAR changes the dynamics of the model and consequently the estimation of the reduced-form residuals ε_t . Second, inventories enter the oil demand curve described by Equation (18). Third, there is an additional equation (Equation 22) describing how inventories react to movements in oil-market and macroeconomic variables. The disturbance $u_{i,t}$ denotes inventory-demand shocks that, in concert with oil-specific demand shocks $u_{d,t}$, drive oil demand holding macroeconomic conditions unchanged. The proposed model of inventories follows closely [Baumeister and Hamilton \(2015b\)](#), except that we do not allow for measurement error in inventories.³⁰

Figure 14 summarizes the key results. Panel A replicates Figure 3 for the model with inventories. The set of VAR-admissible demand and supply elasticities is very similar to the baseline model. Consequently, the oil supply and demand elasticities selected by our identification strategy are 0.11 and -0.1, nearly identical to those in the baseline model. Panel B plots the impulse responses to the identified inventory-demand shock (one standard deviation). The shock leads to an increase in inventories and a

³⁰There are no data on global crude oil inventories. Hence, as standard in the literature, we proxy for global oil inventories by using data on total U.S. oil inventories scaled by the ratio of OECD petroleum stocks over US petroleum stocks. All data are provided by the Energy Information Administration. [Kilian and Lee \(2014\)](#) construct a proxy for global crude oil inventories by using data compiled by the Energy Intelligence Group (EIG). Their measure of global stocks amounts to about three months of oil production.

modest increase in oil prices. As oil production increases less than the change in inventories, the shock also results in a temporary decline in oil consumption. The response of economic activity is statistically and economically not significant. Finally, panel C plots the historical decomposition of oil prices and oil production implied by the model with inventories. The contribution of inventories-demand shocks to historical fluctuations in oil market variables (the yellow bars) has been small throughout the sample.³¹ Moreover, the inclusion of inventories in the model does not alter the relative importance of the remaining shocks compared to our baseline specification. In particular, shocks to oil supply remain the key drivers of oil price fluctuations.

7 Conclusion

In this paper we identify a SVAR of the oil market by imposing plausible restrictions on the values of the short-run price elasticities of both oil supply and oil demand. In the estimating framework, global demand for oil is jointly measured by industrial production in advanced and emerging economies as well as by an index of metal prices. We find that oil supply shocks account for half of oil price fluctuations at business cycle frequencies, while shocks to global demand account for 30 percent. We also show that a drop in oil prices driven by oil supply shocks boosts economic activity in advanced economies, while it depresses economic activity in emerging economies, thus helping explain the muted effects of changes in oil prices on global economic activity. Our analysis shows how the selection of the oil elasticities is consequential to understand the nature of oil price movements and to measure the size of the multipliers of oil prices on economic activity.

³¹The decomposition plotted in Panel C is based on end-of-the-year values. Inventories typically decrease in December as oil companies delay imports, increase runs at refineries, or sell crude oil in order to reduce the local taxes that are levied on their stock of inventory. This seasonal pattern explains why the contribution to oil prices of inventory-demand shocks plotted in Panel C is mostly negative.

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

	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: The entries in the table denote the pairwise correlations between the specified global activity indicators. 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 *** $p < .01$.

Table 2: Leading Indicators of Global Activity

Forecast Horizon	$h = 1$	$h = 6$	$h = 12$	$h = 24$
<i>(a) Advanced Economies' IP</i>				
m_t	0.18*** [3.23]	0.48** [2.40]	0.41** [2.14]	0.00 [0.02]
rea_t	0.08*** [2.87]	0.20** [1.98]	0.07 [0.96]	-0.22* [1.67]
Adj. R^2	0.99	0.89	0.74	0.73
<i>(b) Emerging Economies' IP</i>				
m_t	0.56*** [3.89]	0.93*** [3.97]	1.02*** [4.09]	1.15*** [3.84]
rea_t	0.25*** [3.30]	0.38*** [4.12]	0.27* [1.91]	-0.03 [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 m_t and 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). Each specification also includes a constant, 12 lags of the endogenous variables and 12 lags of m_t and rea_t (not reported). Absolute t -statistics reported in brackets are based on the heteroskedasticity- and autocorrelation-consistent asymptotic variance-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 *** $p < .01$.

Table 3: 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 *** $p < .01$.

Table 4: Global Activity and Oil Prices

Forecast Horizon	$h = 1$	$h = 6$	$h = 12$	$h = 24$
ya_t	-0.09 [0.38]	-0.39 [0.80]	-0.87 [1.47]	-0.76 [1.52]
ye_t	0.22** [2.11]	0.77*** [3.68]	0.36** [2.48]	0.16 [0.77]
m_t	0.68*** [4.05]	0.56** [1.98]	0.46 [1.35]	0.75*** [2.92]
rea_t	0.10 [0.98]	0.07 [0.77]	-0.03 [0.23]	-0.04 [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_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). Each specification also includes a constant, 12 lags of p_t and 12 lags of each global activity indicator (not reported). Absolute t -statistics reported in brackets are based on the heteroskedasticity- and autocorrelation-consistent asymptotic variance-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 *** $p < .01$.

Table 5: Forecast Error Variance Decomposition of Selected Variables
(24-Month Ahead)

Shock	Oil Supply	Oil Demand	AE Activity	EE Activity	Metal Prices
Oil Prices	47.6 [34.7; 59.5]	14.3 [9.5; 22.3]	2.2 [0.8; 5.3]	13.4 [6.2; 23.0]	16.8 [8.3; 28.3]
Oil Production	38.2 [30.9; 46.8]	33.2 [26.7; 40.8]	10.2 [5.0; 16.9]	5.9 [3.4; 10.6]	8.3 [3.3; 15.7]
AE Activity	5.5 [1.8; 12.9]	1.8 [0.7; 4.5]	63.8 [51.7; 74.3]	10.7 [6.7; 15.5]	14.6 [6.6; 25.4]
EE Activity	5.8 [1.7; 13.5]	1.7 [0.5; 5.1]	8.4 [4.8; 13.7]	52.0 [40.2; 65.5]	29.1 [16.6; 41.1]


NOTE: The entries in the table denote the posterior median estimate of the portion of the forecast error variance of a specified variable at the 24-month horizon that is attributable to five structural shocks. The 16th and 84th percentiles of the posterior distributions are reported in bracket. See subsection 4.3 for details.

Figure 1: QUANTITIES AND PRICES IN THE OIL MARKET




NOTE: The figure depicts the scatter plot between the residuals from a regression of oil prices and oil production on their own lag and a constant. The solid red, black dashed, and blue dotted lines represent alternative configurations of the oil demand and the oil supply curves that are consistent with the data.

Figure 2: ALTERNATIVE OIL MARKET STRUCTURES




NOTE: The figure presents six combinations of slopes of supply and demand schedules in the oil market that could be used to match data on oil prices and oil production. The blue dots represent simulated data implied by a given oil market structure. Panels A and B plot two combinations that match the data; Panels C plots a combination that underestimates the volatility of oil production; Panels D plots a combination that overestimates the volatility of oil prices; Panel E plots a combination that overestimate the volatility of oil production; Panel F plots a combination that underestimates the volatility of oil prices.

Figure 3: 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 five-equation VAR. The blue circle corresponds to elasticities suggested by the meta-analysis of the literature (supply elasticity: 0.13, demand elasticity: -0.13). The green square corresponds to the elasticities selected by identification scheme (supply elasticity: 0.11, demand elasticity: -0.11).

Figure 4: IMPULSE RESPONSES TO OIL MARKET SHOCKS




NOTE: The solid lines in the left column depict median responses of the specified variable to a one standard-deviation oil supply shock, while those in the right column depict median responses to a one standard-deviation oil demand shock; The light shaded bands represent the 90 percent pointwise credible sets and the dark shaded bands represent the 68-percent pointwise credible sets. All variables are expressed in log changes (multiplied by 100).

Figure 5: IMPULSE RESPONSES TO GLOBAL ACTIVITY SHOCKS




NOTE: The solid lines in the left column depict median responses of the specified variable to a one standard-deviation shock to advanced economies' activity, those in the middle column depict median responses to a one standard-deviation shock to emerging economies' activity, and those in the right column depict median responses to a one standard-deviation metal price shock; The light shaded bands represent the 90 percent pointwise credible sets and the dark shaded bands represent the 68 percent pointwise credible sets. All variables are expressed in log changes (multiplied by 100).

Figure 6: HISTORICAL 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. All variables are expressed in log changes (multiplied by 100).

Figure 7: HISTORICAL 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 log changes (multiplied by 100) from June 1997.

Figure 8: HISTORICAL 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 log changes (multiplied by 100) from June 2008.

Figure 9: HISTORICAL DECOMPOSITION OF SELECTED VARIABLES
(2014-15 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 log changes (multiplied by 100) from June 2014.

Figure 10: HOW ACTIVITY RESPONDS TO OIL PRICE CHANGES




NOTE: The lines plot the two-year elasticity of advanced economies' IP and emerging economies' 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 oil supply shocks (solid red line) and oil-specific demand shock (red dashed line). The vertical line denotes the restriction selected by our identification strategy.

Figure 11: 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. The vertical line denotes the restriction selected by our identification strategy.


Figure 12: HISTORICAL DECOMPOSITION OF SELECTED VARIABLES
(Asian financial crisis: model comparison)



(a) Five-equation SVAR model: baseline identification




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




(c) Three-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 η_S to zero; panel C reports the historical decomposition based on a three-equation VAR that also sets $\eta_S = 0$. The solid lines depict the actual series. All variables are expressed in log changes (multiplied by 100) from June 1997. See page 21 for details.


Figure 13: HISTORICAL DECOMPOSITION OF SELECTED VARIABLES
(2014–15 oil slump: model comparison)



(a) Five-equation SVAR model: Baseline Identification




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




(c) Three-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 η_S to zero; panel C reports the historical decomposition based on a three-equation VAR that also sets $\eta_S = 0$. The solid lines depict the actual series. All variables are expressed in log changes (multiplied by 100) from June 2014. See page 22 for details.


Figure 14: SELECTED RESULTS FROM MODEL WITH OIL INVENTORIES



(a) Oil Demand and Supply Elasticities Implied by the SVAR



(b) Impulse Responses to Inventories Shock



(c) Historical Decomposition — Full Sample

NOTE: Selected results from the VAR model with oil inventories. See Section 6 for details.

References

- Aastveit, K. A., H. C. Bjørnland, and L. A. Thorsrud (2015). What Drives Oil Prices? Emerging versus Developed Economies. *Journal of Applied Econometrics* 30(7), 1013–1028.
- 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/>.
- Arezki, R. and A. Matsumoto (2014). Metals and Oil: A Tale of Two Commodities. <https://blog-imfdirect.imf.org/2015/09/14/metals-and-oil-a-tale-of-two-commodities/>.
- Barsky, R. B. and L. Kilian (2001). Do We Really Know That Oil Caused the Great Stagflation? A Monetary Alternative. In B. S. Bernanke and K. Rogoff (Eds.), *NBER Macroeconomics Annual*, Volume 16, pp. 137–183. University of Chicago Press.
- Baumeister, C. and J. D. Hamilton (2015a). Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information. *Econometrica* 83(5), 1963–1999.
- Baumeister, C. and J. D. Hamilton (2015b). Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks. Manuscript, University of Notre Dame and UCSD.
- Baumeister, C. and L. Kilian (2016a). 40 Years of Oil Price Fluctuations: Why the Price of Oil May Still Surprise Us. *Journal of Economic Perspectives* 30(1), 139–160.
- Baumeister, C. and L. Kilian (2016b). Understanding the Decline in the Price of Oil since June 2014. *Journal of the Association of Environmental and Resource Economists* 3(1), 131–158.
- Beidas-Strom, S. and A. Pescatori (2014). Oil Price Volatility and the Role of Speculation. Working Paper 14/218, IMF.
- 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>.
- Cuddington, J. T. and D. Jerrett (2008). Super Cycles in Real Metals Prices? *IMF Staff Papers* 55, 541–565.
- Cuddington, J. T. and D. Jerrett (2011). Business Cycle Effects on Metal and Oil Prices: Understanding the Price Retreat of 2008-9. Manuscript, Colorado School of Mines.

- Del Negro, M. and F. Schorfheide (2011). Bayesian Macroeconometrics. In J. Geweke, G. Koop, and H. van Dijk (Eds.), *Handbook of Bayesian Econometrics*, pp. Chapter 7. Oxford University Press.
- Delle Chiaie, S., L. Ferrara, and D. Giannone (2015). Common Factors of Commodity Prices. Manuscript, Banque de France and Federal Reserve Bank of New York.
- Fernandez, A., A. Gonzalez, and D. Rodriguez (2015). Sharing a Ride on the Commodities Roller Coaster: Common Factors in Business Cycles of Emerging Economies. Working Paper 15/280, IMF.
- Iacoviello, M. (2016). Oil Prices and Consumption across Countries and U.S. States. Federal Reserve Board.
- IMF (2015a). Commodity Market Developments and Forecasts, with a Focus on Metals in the World Economy. In *World Economic Outlook*, Chapter 1, Special Feature, pp. 38–51. IMF.
- IMF (2015b). Where Are Commodity Exporters Headed? Output Growth in the Aftermath of the Commodity Boom. In *World Economic Outlook*, Chapter 2, pp. 65–104. IMF.
- Issler, J., C. Rodrigues, and R. Burkack (2014). Using Common Features to Understand the Behavior of Metal-Commodity Prices and Forecast Them at Different Horizons. *Journal of International Money and Finance* 42, 310–335.
- Juvenal, L. and I. Petrella (2015). Speculation in the Oil Market. *Journal of Applied Econometrics* 30(4), 621–649.
- 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), 1053–69.
- Kilian, L. (2016). The Impact of the Shale Oil Revolution on U.S. Oil and Gas Prices. *Review of Environmental Economics and Policy*, forthcoming.
- Kilian, L. and T. K. Lee (2014). Quantifying the Speculative Component in the Real Price of Oil: The Role of Global Oil Inventories. *Journal of International Money and Finance* 42(C), 71–87.
- Kilian, L. and D. P. Murphy (2012). Why Agnostic Sign Restrictions Are Not Enough: Understanding the Dynamics of Oil Market VAR Models. *Journal of the European Economic Association* 10(5), 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.

- Labys, W., A. Achouch, and M. Terraza (1999). Metal Prices and the Business Cycle. *Resources Policy* 25(4), 229–238.
- 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., C. Osbat, and B. Schnatz (2012). Global Commodity Cycles and Linkages: A FAVAR Approach. *Empirical Economics* 55, 541–565.
- Mendoza, E. G. (1995). The Terms of Trade, the Real Exchange Rate and Economic Fluctuations. *International Economic Review* 36(1), 101–137.
- Newey, W. K. and K. D. West (1987). A Simple, Positive Semi-Definite, Heteroskedasticity- and Autocorrelation-Consistent Covariance Matrix. *Econometrica* 55(3), 703–708.
- Newey, W. K. and K. D. West (1994). Automatic Lag Selection in Covariance Matrix Estimation. *Review of Economic Studies* 61(4), 631–653.
- Peersman, G. and I. Van Robays (2012). Cross-Country Differences in the Effects of Oil Shocks. *Energy Economics* 34(5), 1532–1547.
- Pindyck, R. S. and J. J. Rotemberg (1990). The Excess Co-movement of Commodity Prices. *Economic Journal* 100(403), 1173–89.
- Rubio-Ramírez, J. F., D. F. Waggoner, and T. Zha (2010). Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference. *Review of Economic Studies* 77(2), 665–96.
- Stock, J. H. and M. W. Watson (2016). Factor Models and Structural Vector Autoregressions in Macroeconomics. In J. B. Taylor and H. Uhlig (Eds.), *Handbook of Macroeconomics*, Volume 2, pp. forthcoming. Elsevier.
- Stuermer, L. (2016). 150 Years of Boom and Bust: What Drives Mineral Commodity Prices? *Macroeconomic Dynamics*, forthcoming.
- Uhlig, H. (2005). What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure. *Journal of Monetary Economics* 52(2), 381–419.

Appendix

A The Data

Industrial Production. We construct a monthly index of industrial production (IP) for advanced economies and emerging economies by aggregating country-level data. We take monthly, seasonally adjusted, total IP excluding the construction industries. For countries where this series is not available, we use monthly, seasonally adjusted, manufacturing industrial production. The initial unbalanced dataset runs 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 IP 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 U.S. dollars from the World Bank's World Development Indicators. Next, we obtain the level of industrial production by cumulating the resulting monthly growth series. Both indexes are normalized to take the value of 100 in January 2007. Although the IP data potentially start in the 1960s, we set 1985 as the starting date because 1985:M1 is the earliest observation when our sample includes enough emerging economies so that they account, using today's GDP weights, for at least 25 percent of emerging economies' GDP.

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 except Finland, Greece, and Portugal, which are 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 global GDP in current dollars, with 53 percentage points 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 global GDP.

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

Metal Prices. Metal prices are measured from the IMF Metal Price Index, linearly log de-trended and expressed in real terms dividing by the U.S. consumer price index (CPI). The metal price index (code: PMETA) is

available at <https://www.imf.org/external/np/res/commod/index.aspx>.

Oil Market Variables. The real price of oil is the Brent price of oil, linearly log de-trended and expressed in real terms dividing by the U.S. CPI index. Oil production is total oil supply from the International Energy Agency (Table 3, Oil Market Report).

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 phrases “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 (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 (1) standard errors of the estimates were missing and (2) because we span both the empirical macro and micro literature, there was a rich heterogeneity in the models and economies used for estimation but not enough studies to compute estimates conditional on specific characteristics. The search was terminated on March 18, 2015.

Table A.3 lists the studies that make up our database. For comparison, table A.4 displays the oil elasticities documented in meta-analysis and survey papers excluded from our search. Both tables highlight how the literature placed a strong focus on the estimation of the short-run price elasticity of demand, as there is 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.

C Additional Robustness Exercises

In this section we explore the robustness of our key findings to two modifications of our baseline SVAR model. First, we identify the SVAR model by imposing a lower target for the short-run price elasticity of oil supply. In particular, we set $\eta_S^* = 0.02$, the value selected by Kilian and Murphy (2012) using the Iraq’s invasion of Kuwait as an event study. Exploring the robustness of our findings to the choice of η_S^* is important as there are only few studies that estimate this elasticity. Second, we add to the model *rea* as a fourth indicator of global economic activity. In particular, we assume that *rea* is allowed to react to all remaining variables in the model, while a shock to *rea* can affect the remaining variables only with a one-period lag.

The red line in Figure A.2 plots the oil demand and supply elasticities implied by the baseline model (top panel) and by the baseline model augmented with *rea* (bottom panel). The inclusion of *rea* does not affect the

set of admissible oil elasticities. As shown in the top panel, when the target supply and demand elasticities are 0.02 and -0.13, respectively, our identification strategy selects a supply elasticity of 0.07—about 40 percent smaller than the elasticity selected by our baseline identification—and a demand elasticity of -0.17—about 50 percent larger than the elasticity selected by our baseline identification. By contrast, the addition of *rea* does not affect the selection of the oil elasticities, which are identical to those in the baseline model.

Table A.5 shows the amount of variation in the two-year-ahead forecast error variance in oil prices that is attributable to the structural shocks identified by our baseline SVAR model (first row) and by the two variants of the baseline model discussed above. As shown in the second row of table A.5, a configuration of the oil market characterized by a lower supply elasticity and a high demand elasticity implies that oil supply shocks and oil-specific demand shocks are equally important drivers of oil prices, explaining 32 percent and 27 percent of the fluctuations in oil prices, respectively. Hence, even if oil supply shocks contribute less to oil price fluctuations, they remain the most important driver. The addition of *rea* to the model has only a minor impact on the decomposition of oil price fluctuations. A shock to *rea* explains only about 3 percent of movements in oil prices, while the relative importance of other shocks is substantially unchanged.

Table A.1: ADVANCED ECONOMIES' INDUSTRIAL PRODUCTION

Country	Share of Global GDP	Share of Global Oil Production	Share of Global 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 calculation of the industrial production index in advanced economies. The underlying country indexes refer to total industrial production excluding construction. Data for Japan are on manufacturing industrial production.

Table A.2: EMERGING ECONOMIES' INDUSTRIAL PRODUCTION

Country	Share of Global GDP	Share of Global Oil Production	Share of Global 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 calculation of the industrial production index in emerging economies. The underlying country indexes refer to total industrial production excluding construction. Data for Chile, Colombia, Indonesia, Israel, Peru, Philippines, Singapore, South Africa, Thailand, and Venezuela refer to manufacturing industrial production. Data for Mexico are for total industrial production including construction.

Table A.3: OIL ELASTICITIES ACROSS STUDIES: LITERATURE SEARCH

Study	Demand Elasticity	Supply Elasticity
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	
Faris Al-Faris (1997)	-0.14	
Güntner (2014)		-0.00
Hughes et al. (2008)	-0.04	
Javan and Zahran (2015)	-0.04	
Kayser (2000)	-0.23	
Kilian and Murphy (2012)		0.02
Kilian and Murphy (2014)	-0.26	
Krichene (2002)	-0.07	-0.16
Lin and Prince (2013)	-0.03	
Liu (2014)	-0.06	
Liu (2015)	-0.30	
McRae (1994)	-0.12	
Nicol (2003)	-0.21	
Polemis (2007)	-0.13	
Ramanathan (1999)	-0.21	
Sene (2012)	-0.12	
Sita et al. (2012)	-0.92	
Taghizadeh Hesary and Yoshino (2014)	-0.08	0.24
Wadud et al. (2009)	-0.09	
<i>Mean</i>	-0.22	0.10
<i>Median</i>	-0.13	0.13
<i>Standard Deviation</i>	0.22	0.18

NOTE: The table reports a list of papers that estimate short-run price elasticities of demand and supply. The papers were found searching for the phrases "oil demand elasticity," "oil supply elasticity," and "gasoline demand elasticity." in the following databases: EconLit, ScienceDirect, Wiley, and JSTOR. We excluded studies conducting meta-analysis, studies with estimates based only on pre-1985 data, and studies estimating only long-run elasticities.

Table A.4: OIL ELASTICITIES ACROSS SURVEY STUDIES

Authors	Demand Elasticity	Supply Elasticity
Brons et al. (2008)	-0.34	
Dahl (1993)	-0.07	
Espey (1998)	-0.23	
Greene and Leiby (2006)	-0.07	0.07
Goodwin (1992)	-0.27	
Goodwin et al. (2004)	-0.25	
Graham and Glaister (2004)	-0.25	
Hamilton (2009)	-0.08	
IMF (2011)	-0.02	
MEDIAN	-0.24	
AVERAGE	-0.19	


NOTE: The table reports a list of papers that survey the empirical literature on short-run price elasticities of demand and supply.

Table A.5: Forecast Error Variance Decomposition of Selected Variables
(Oil Prices — 24-Month Ahead)

Shock	Oil Supply	Oil Demand	AE Activity	EE Activity	Metal Prices	Rea
Baseline	47.6 [34.7; 59.5]	14.3 [9.5; 22.3]	2.2 [0.8; 5.3]	13.4 [6.2; 23.0]	16.8 [8.3; 28.3]	
$\eta_S^* = 0.02$	32.3 [22.0; 44.6]	27.2 [18.7; 37.1]	2.3 [0.8; 5.3]	14.1 [6.8; 23.3]	17.1 [8.0; 28.6]	
With <i>rea</i>	49.1 [37.3; 59.7]	13.2 [8.5; 20.6]	2.1 [0.8; 4.9]	9.0 [4.2; 17.5]	16.1 [7.3; 28.8]	2.8 [0.6; 8.0]


NOTE: The entries in the table denote the posterior median estimate of the portion of the forecast error variance of oil prices at the 24-month horizon from three models that is attributable to the structural shocks. The 16th and 84th percentiles of the posterior distributions are reported in bracket. See subsection 4.3 for details.

Figure A.1: DATA




NOTE: Sample period: monthly data from 1985:M1 to 2015:M12. The solid lines in the left column show the linearly detrended data. The black and red lines in the right column show the undetrended data and the trends, respectively. The shaded vertical bars denote U.S. recessions as dated by the National Bureau of Economic Research (NBER).

Figure A.2: OIL DEMAND AND SUPPLY ELASTICITIES IMPLIED BY THE SVAR



(a) Baseline Model with $\eta_S^* = 0.02$



(b) Extended Baseline Model with rea

NOTE: The red solid line in both panels shows the relationship between the oil demand and supply elasticities imposed by the VAR. The blue circle corresponds to the target elasticities, while the green square corresponds to the elasticities selected by the identification scheme. See text for more details. Panel A: baseline identification with $\eta_S^* = 0.02$; Panel B: Baseline model augmented with rea_t .

References for Appendix

- Altinay, G. (2007). Short-Run and Long-Run Elasticities of Import Demand for Crude Oil in Turkey. *Energy Policy* 35(11), 5829 – 5835.
- 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 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 U.S. 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.
- 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.
- 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 U.S. 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* 26(2), 155–169.
- Graham, D. J. and S. Glaister (2004). Road Traffic Demand Elasticity Estimates: A Review. *Transport Reviews* 24(3), 261–274.
- Greene, D. L. and P. N. Leiby (2006). The Oil Security Metrics Model. ORNL/TM-2006/505, Oak Ridge National Laboratory, Oak Ridge, Tennessee.
- 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.
- 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), 93–114.
- IMF (2011). World Economic Outlook. Technical report, International Monetary Fund.
- 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. and D. P. Murphy (2012). Why Agnostic Sign Restrictions Are Not Enough: Understanding the Dynamics of Oil Market VAR Models. *Journal of the European Economic Association* 10(5), 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* 15(1), 143–155.
- Nicol, C. J. (2003). Elasticities of Demand for Gasoline in Canada and the United States. *Energy Economics* 25(2), 201–214.
- 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.
- 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.
- Wadud, Z., D. J. Graham, and R. B. Noland (2009). Modelling Fuel Demand for Different Socio-Economic Groups. *Applied Energy* 86(12), 2740–2749.