CENTRE FOR APPLIED MACROECONOMICS AND COMMODITY PRICES (CAMP)

CAMP Working Paper Series
No 11/2018

On the China factor in international oil markets: A regime switching approach

Jamie L. Cross, Chenghan Hou and Bao H. Nguyen

NORWEGIAN BUSINESS SCHOOL

On the China factor in international oil markets: A regime switching approach[1]

Jamie L. Cross[a,b], Chenghan Hou[c,*], Bao H. Nguyen[d,e,b]

[a] *Centre of Applied Macroeconomics and Commodity Prices (CAMP), BI Norwegian Business School*
[b] *Centre for Applied Macroeconomic Analysis (CAMA)*
[c] *Center for Economics, Finance and Management Studies, Hunan University*
[d] *Development Policy Centre, Crawford School of Public Policy, Australian National University*
[e] *School of Economics, University of Economics Ho Chi Minh City (UEH)*

Abstract

We investigate the relationship between world oil markets and China's macroeconomic performance over the past two decades. Our analysis starts by proposing a simple method for disentangling real economic activity stemming from China and the rest of the world. We then consider a sufficiently large set of dynamic VAR models to distinguish between abrupt and gradual changes in the macroeconomic relationships and volatility clustering in the shocks. A model exercise shows that a Markov-switching model is preferred to previously used models in the literature. When investigating the role of oil market shocks on China's output, we find that oil supply shocks tend to elicit a positive response, while the response of oil demand shocks is negative. Next, when analyzing world oil price dynamics, we find that demand shocks have had significant positive impacts over the past two decades. The average proportion of oil price variation explained by demand from China and rest of world demand are around 30 percent over the sample period. Importantly, while China specific effects are relatively constant, rest of world aggregate demand shocks are found to have larger impact during times of global macroeconomic downturn. This highlights the importance of our model comparison exercise. Finally, we find that the recent 2014/15 oil price drop was due to a combination of increased oil supply and decreased demand from China.

JEL-codes: C32, E31, E32
Keywords: Oil prices, China, Vector autoregression (VAR), Markov-switching, Sign restrictions.

---

[1] We thank Francesco Ravazzolo, Hilde Bjørnland, Ippei Fujiwara, Knut Aastveit, Leif Anders Thorsrud, Tatsuyoshi Okimoto, Renee Fry-McKibbin, Warwick McKibbin and members of the workshop on Energy Economics hosted by the Free University of Bozen-Bolzano for their comments in the development of this research.

*Corresponding author
Email addresses: `jamie.cross@bi.no` (Jamie L. Cross), `chenghan.hou@hotmail.com` (Chenghan Hou), `bao.nguyen@anu.edu.au` (Bao H. Nguyen)

1. Introduction

Over the past three decades China has risen from being the 9th largest economy by world share of GDP, to the second largest economy in 2017. This rapid expansion of macroeconomic growth lead the country to be the world's largest energy consumer in 2011 and 2013 it became the biggest net importer of petroleum and other liquid fuels.[2] Despite these facts, there is surprisingly little research examining the relationship between China's macroeconomic growth and world oil markets. Moreover, the small existing literature that does exist has provided contrasting results.

On one side of the literature, researchers have speculated that recent oil price fluctuations have been predominately driven by growth in emerging markets, including China (Hamilton, 2009; **?**; Baumeister and Peersman, 2013). Recent empirical evidence in support of this claim was provided in Aastveit et al. (2015). Using a factor analysis with 33 countries, they find that approximately 40% of the 1-2 year variation in the oil price is explained by demand shocks from emerging markets (including China), while demand shocks from developed countries explain approximately 15%. Despite this result, direct evidence on the role of China specifically, has not been clear. For instance, while Liu et al. (2016) found that China specific demand accounted for 51 percent of the variation between 2000 and 2014, others have suggested that China's oil demand has little or zero impact on the global oil price (Mu and Ye, 2011; Wu and Zhang, 2014; Lin and Li, 2015; Cross and Nguyen, 2017). On the other side of the literature, researchers who have examined the effects of oil market shocks on China are yet to reach a consensus on the general nature of the impacts. For instance, Tang et al. (2010) and Zhao et al. (2016) find that positive oil demand shocks negatively impact China's output. In contrast, Du et al. (2010) and Herwartz and Plödt (2016) find that such shocks elicit positive growth. More recently, Cross and Nguyen (2017) suggest that such shocks may elicit either positive or negative movements, depending on the period in which they occurred. In particular, they find that such shocks had negative impacts in the early 1990's to 2000's and positive effects after the 2007/08 financial crisis.

A potential shortcoming of the studies that have focused on China is that they do not explicitly control for aggregate demand effects stemming both China and the rest of the world. This is critical because **?** shows that different oil market shocks have very different effects on the real price of oil. For example, increases in precautionary demand for crude oil elicit

[2]The data source is obtained through the US Energy Information Administration (EIA) website: https://www.eia.gov/beta/international/analysis.cfm?iso=CHN, and the IMF: https://www.imf.org/external/pubs/ft/weo/2017/02/weodata/download.aspx

an immediate, persistent, and large increase in the real price of crude oil, while an increase in aggregate demand causes a somewhat delayed, but sustained, increase. Moreover, he finds that the recent 2003-08 oil price boom was predominantly driven by the cumulative effects of positive global demand shocks.

This paper breaks ground in this area by proposing a simple two step method for modeling real economic activity (REA) stemming from China and the rest of the world.[3] In the first step, we define the *rest of world real economic activity* as the residuals from a simple (time varying coefficients) regression model in which world REA is regressed on China's REA. In the second step, we estimate a structural vector autoregression (VAR) model with these residuals as a proxy for rest of the world REA, along with China's REA, and the international price and production of crude oil, to answer our research question.

In conducting step two we build on the recent analysis of Cross and Nguyen (2017), by considering a class of non-linear VAR models. In that paper, Cross and Nguyen (2017) investigate the role of China in world oil markets by employing a class of time varying parameter VAR models (Primiceri, 2005; Chan and Eisenstat, 2018). The justification for this modeling approach stems from noting that changes in China's macroeconomic conditions and demand for oil in particular have been gradually occurring over time. Despite this fact, it is well known that energy prices tend to exhibit erratic jumps in volatility (as opposed to gradual changes). Indeed, multiple scholars have shown that world energy market dynamics are well captured by Markov-switching models (Raymond and Rich, 1997; Clements and Krolzig, 2002; Bjørnland et al., 2018a; Holm-Hadulla and Hubrich, 2017; Hou and Nguyen, 2018). With this in mind, we extend the set of models proposed in Cross and Nguyen (2017) to include a class of Markov-switching models.

While our methodological contributions are relatively simple, they yield significant insights into the relationship between China's recent macroeconomic performance and world oil markets. These results are summarized as follows. First, from a purely statistical perspective we find that a constant coefficients VAR model with Markov-switching in the covariance matrix provides better in-sample fit when compared to the time varying VAR in Cross and Nguyen (2017) and the traditional constant parameter VAR used in other studies (Du et al., 2010; Tang et al., 2010; Mu and Ye, 2011; Wu and Zhang, 2014; Liu et al., 2016; Herwartz and Plödt, 2016). Using the same set of sign restrictions in Cross and Nguyen (2017), we find that oil supply shocks tend to elicit a positive response in China's output, while the response of oil demand shocks is

[3]Following (?) real economic activity is not itself a measure of aggregate demand, however the structural VAR model introduced in Section 3.1 allow us to identify the demand driven component in real activity.

negative. Interestingly, the signs of these responses then switch during times of global economic downturn, highlighting the importance of our modeling choice. Next, when analyzing world oil price dynamics, we find that demand shocks have had significant positive impacts over the past two decades. To quantify these impacts, we compute a regime-dependent forecast error variance decomposition. The results show that the proportion of oil price variation explained by demand from China is around 30 percent over the sample period. Our results therefore provide empirical support to the conjectured claims in recent studies (Hamilton, 2009; **?**; Baumeister and Peersman, 2013). Importantly, while China specific effects are relatively constant, rest of world aggregate demand shocks are found to have larger impact during times of global macroeconomic downturn. This result highlights the importance of our model selection exercise. In addition to these insights, we also find evidence that the recent 2014/15 oil price drop was due to a combination of increased oil supply and decreased demand from China.

The rest of this paper is structured as follows. In Section 2 we provide the data sources and propose a method for disentangling world and China specific demand. In Section 3 we outline the set of VAR models and conduct a formal Bayesian model comparison exercise to select the best model for the main analysis. In Section 4 we discuss the identification procedure used to recover the structural shocks. In Section 5, we address the research questions. Finally, we present some robustness checks in Section 6 and conclude in Section 7.

2. Data

2.1. Data Sources and Transformations

We use quarterly data between 1992Q1 and 2015Q2 on four variables of interest in the model: The oil price and quantity, real economic activity from China and real activity from the rest of the world. In line with existing literature we consider two alternative measures of the oil price: the US refiners' acquisition cost (RAC) for imported crude oil and the West Texas Intermediate (WTI) price of crude oil. Since it is generally considered to be the best proxy for the free global oil price market (Baumeister et al., 2010), we use the former price for the benchmark model and the latter as a robustness check. The real oil price is obtained by deflating the nominal price by the US Consumer Price Index. Next, the quantity of oil is measured by the amount of world crude oil production (thousand barrels per day) as provided by the US Energy Information Administration (EIA).

Next, we measure world real economic activity using the global industrial production (IP) index provided by the Netherlands Bureau of Economic Policy Analysis (CPB).[4] There are

[4]The choice of an appropriate measure of global economic activity has recently received a lot of attention;

three reasons that we prefer the use of the CPB index over the commonly used OECD + 6 series. First, the CPB series has a more general coverage; containing IP data from 85 countries worldwide which account for approximately 97% of global IP. Second, unlike the OECD + 6 index, the CPB measure includes data on a range of emerging markets. This is important given the aforementioned results in Aastveit et al. (2015). Third, the use of the CPB series allows us to remain consistent with Cross and Nguyen (2017), which provides the benchmark results for our analysis.

Finally, also consistent with Cross and Nguyen (2017), China's real economic activity is measured by its real quarterly GDP published by Center of Quantitative Economic Research of Federal Reserve Bank of Atlanta. The reason for using China's real quarterly GDP, instead of the IP index, is that this data is unavailable. Instead, China's government officials only report *real value added in industry*, which is defined as gross output in industry minus the costs of factor inputs; see Kilian and Zhou (2018) for a discussion.

Prior to the analysis we conducted the following data transformations. First, since both the oil price and production are of monthly frequency, we transform them to non-annualized quarter-on-quarter growth rates by taking the first difference of the natural logarithm. Next, the IP index is converted to non-annualized month-on-month rates of growth by taking the first difference of the natural logarithm, seasonally adjusted, and then converted to quarterly values by taking the appropriate averages. Finally, the real GDP index is converted to quarter-on-quarter rates of growth by taking the first difference of the natural logarithm. The resulting series are plotted in Figure 1. A few features are worth highlighting: the high production and price volatility around 2000 and 2007, the gradual increase in production since 2012, the 2014/15 volatility in the oil price volatility and the slowdown of China's economy since 2012. We will return to each of these points when discussing the results in Section 6.

2.2. Disentangling China's demand and World demand for oil

To analyse the impacts of demand shocks stemming from China and the rest of the world it is first necessary to disentangle the real economic activity components. To this end, we regress the measure of world real economic activity on China specific real economic activity

see Kilian and Zhou (2018) for a survey of this literature. By far, the two most commonly used proxies for this activity are the global *real economic activity* (REA) index constructed by **?** (Kilian's index) and global *industrial production* (IP). For our analysis, we prefer latter measure for two reasons: First, recent evidence reported in Kalouptsidi (2017) shows that China has intervened and reduced shipyard costs by 13-20%. Since it relies on international shipping costs, the REA index may not fully reflect the true costs in China. Second, Hamilton (2018) has recently highlighted that the REA index has not captured well real economic activity in recent years, while world Industrial Production (IP) does.

and use the residual as a proxy for demand of oil by the rest of the world. Let $y_{w,t}$ and $y_{c,t}$ respectively denote the real growth of world and China specific demand for oil at date t. Then, the regression is given by

$$y_{w,t} = \beta_t y_{c,t} + e_t, \quad e_t \sim \mathcal{N}\left(0, \sigma^2_{y_w}\right), \tag{1}$$

$$\beta_t = \beta_{t-1} + u_t, \quad u_t \sim \mathcal{N}\left(0, \sigma^2_{\beta}\right), \tag{2}$$

where we account for the changing proportion of world demand explained by China through the use of time varying regression coefficients. To facilitate the readability of our paper, we defer specific details of the sampler to Appendix A.1.

The resulting posterior mean of the time varying coefficients and associated 16th and 84th percentiles of the posterior distribution are provided in Figure 2. To provide a benchmark, we also plot the associated least squares estimates for a constant coefficient regression. The primary takeaway is that the proportion of total world activity explained by China has substantially changed over the past two decades, with the average contribution ranging from 10 to 40 percent. The fact that this is a relatively large proportion highlights the importance of disentangling the two series before any econometric analysis. Also note that the least squares estimate is 33%. Thus, the least squares estimate would over-predict the contribution at the start of the sample, and under-predict it towards the end.

[Insert Figure 2 here]

3. Model Selection

3.1. Models

We consider two classes of time varying VARs. The first class models time variation in the parameters through a Markov-switching process, while the second class assumes that the parameters evolve according to an autoregressive process. Despite both classes of VARs being time varying parameter models, in what follows we adopt the convention of referring to the first class as Markov-switching VARs and the second class as time varying VARs. From an economic perspective, Markov-switching models are useful because they can capture abrupt changes in the endogenous relationship between the variables of interest or volatility clustering in the shocks. In contrast, time varying VARs are useful because they can capture gradual changes in these phenomena. Since there are many permutations of each model, we facilitate the presentation by first discussing the most general model and then noting how other models are specified as restricted versions of this general model. In both cases we let $\mathbf{y}_t$ denote an $n \times 1$ vector of time series of interest and p denote the number of lags. Since we wish to compare our results directly with the existing literature, we follow Cross and Nguyen (2017) and set $p = 4$.

3.1.1. Time varying VARs

The structural version of the time varying parameter VAR with stochastic volatility (TVP-VAR-SV) model, with n variables and p lags, has a state space representation with observation equation defined by:

$$\mathbf{B}_{0,t}\mathbf{y}_t = \mathbf{b}_t + \mathbf{B}_{1,t}\mathbf{y}_{t-1} + \cdots + \mathbf{B}_{p,t}\mathbf{y}_{t-p} + \boldsymbol{\epsilon}_t, \quad \boldsymbol{\epsilon}_t \sim \mathcal{N}\left(\mathbf{0}, \boldsymbol{\Sigma}_t\right), \tag{3}$$

where $\mathbf{b}_t$ is an $n \times 1$ vector of time varying intercepts, $\mathbf{B}_{j,t}$, $j = 1, \ldots, p$, are $n \times n$ matrices of time varying VAR coefficients, $\mathbf{B}_{0,t}$ is a lower triangular $n \times n$ matrix with ones along the main diagonal and $\boldsymbol{\Sigma}_t = \mathrm{diag}\left(e^{h_{1,t}}, \ldots, e^{h_{n,t}}\right)$ in which $e^{h_{i,t}}$, $i = 1, \ldots, n$ are latent time varying volatilities. We highlight the fact that (3) presents the structural form of the TVP-VAR-SV model and is therefore different from the reduced form representation in Primiceri (2005), which was utilized in Cross and Nguyen (2017). The reason that we prefer this representation is that it allows for more efficient posterior simulation (Eisenstat et al., 2016). Of course, prior to model comparison, the reduced-form coefficients can be easily recovered from the structural-form coefficients.

For estimation purposes (3) can be written as:

$$\mathbf{y}_t = \widetilde{\mathbf{X}}_t\boldsymbol{\beta}_t + \mathbf{W}_t\boldsymbol{\gamma}_t + \boldsymbol{\epsilon}_t, \quad \boldsymbol{\epsilon}_t \sim \mathcal{N}\left(\mathbf{0}, \boldsymbol{\Sigma}_t\right) \tag{4}$$

where $\widetilde{\mathbf{X}}_t = \mathbf{I}_n \otimes [\mathbf{1} \quad \mathbf{y}'_{t-1} \ldots \quad \mathbf{y}'_{t-1}]$ is a $n \times k_\beta$ matrix, $\boldsymbol{\beta}_t = \mathrm{vec}([\mathbf{b}_t \quad \mathbf{B}_{1,t} \ldots \mathbf{B}_{p,t}]')$ is a $k_\beta \times 1$ vector with $k_\beta = n\,(1 + np)$, $\boldsymbol{\gamma}_t$ is a $k_\gamma \times 1$ vector of the contemporaneous time varying coefficients in $\mathbf{B}_{0,t}$ (collected by rows) with $k_\gamma = \frac{n(n-1)}{2}$, and $\mathbf{W}_t$ is an $n \times k_\gamma$ matrix that contains appropriate elements of $-\mathbf{y}_t$. For example, when $n = 4$ it follows that $k_\gamma = 6$, i.e. $\boldsymbol{\gamma}_t = (b_{0,21,t}, b_{0,31,t}, b_{0,32,t}, b_{0,41,t}, b_{0,42,t}, b_{0,43,t})'$ and $\mathbf{W}_t$ has the form:

$$\mathbf{W}_t = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ -y_{1,t} & 0 & 0 & 0 & 0 & 0 \\ 0 & -y_{1,t} & -y_{2,t} & 0 & 0 & 0 \\ 0 & 0 & 0 & -y_{1,t} & -y_{2,t} & -y_{3,t} \end{pmatrix},$$

where $y_{i,t}$ is the i-th element of $\mathbf{y}_t$, and $b_{0,ji,t}$ is the ji-th element in the matrix $\mathbf{B}_{0,t}$, with

$i = 1, 2, 3$ and $j = 1, 2, 3, 4$.

Finally, (4) can be written more compactly as:

$$\mathbf{y}_t = \mathbf{X}_t \boldsymbol{\theta}_t + \boldsymbol{\epsilon}_t, \quad \boldsymbol{\epsilon}_t \sim \mathcal{N}\left(\mathbf{0}, \boldsymbol{\Sigma}_t\right), \tag{5}$$

where $\mathbf{X}_t = \left(\widetilde{\mathbf{X}}_t, \mathbf{W}_t\right)$ is an $n \times k_\theta$ matrix and $\boldsymbol{\theta}_t = (\boldsymbol{\beta}_t', \boldsymbol{\gamma}_t')'$ is a $k_\theta \times 1$ vector where $k_\theta = k_\beta + k_\gamma$.

Following Cross and Nguyen (2017), the state equations for the latent parameters and log volatilities are assumed to follow the subsequent independent random walk processes:

$$\boldsymbol{\theta}_t = \boldsymbol{\theta}_{t-1} + \boldsymbol{\nu}_t, \quad \boldsymbol{\nu}_t \sim \mathcal{N}\left(\mathbf{0}, \boldsymbol{\Omega}_{\boldsymbol{\theta}}\right), \tag{6}$$

$$\mathbf{h}_t = \mathbf{h}_{t-1} + \boldsymbol{\eta}_t, \quad \boldsymbol{\eta}_t \sim \mathcal{N}\left(\mathbf{0}, \boldsymbol{\Omega}_{\mathbf{h}}\right), \tag{7}$$

where $\boldsymbol{\Omega}_{\boldsymbol{\theta}} = \mathrm{diag}\left(\omega_{\theta 1}^2, \ldots, \omega_{\theta k}^2\right)$, $\boldsymbol{\Omega}_{\mathbf{h}} = \mathrm{diag}\left(\omega_{h1}^2, \ldots, \omega_{hn}^2\right)$ and we treat the initial conditions $\boldsymbol{\theta}_0$ and $\mathbf{h}_0$ as values to be estimated.

The model specification is completed by selecting priors for the state variances and the initial states. The priors for $\boldsymbol{\theta}_0$ and $\mathbf{h}_0$ are both multivariate Gaussian distributions: $\boldsymbol{\theta}_0 \sim \mathcal{N}\left(\mathbf{a}_{\boldsymbol{\theta}}, \mathbf{V}_{\boldsymbol{\theta}}\right), \mathbf{h}_0 \sim \mathcal{N}\left(\mathbf{a}_{\mathbf{h}}, \mathbf{V}_{\mathbf{h}}\right)$, while the diagonal elements of the covariance matrices $\boldsymbol{\Omega}_{\mathbf{h}}$ and $\boldsymbol{\Omega}_{\boldsymbol{\theta}}$ are independently distributed as: $\omega_{\theta i}^2 \sim \mathcal{IG}\left(\nu_{\theta i}, S_{\theta i}\right)$ and $\omega_{hj}^2 \sim \mathcal{IG}\left(\nu_{hj}, S_{hj}\right)$, with $i = 1, \ldots, k_\theta$ and $j = 1, \ldots, n$. The hyperparamters for these distributions are set as follows: $\mathbf{a}_{\boldsymbol{\theta}} = \mathbf{0}$, $\mathbf{V}_{\boldsymbol{\theta}} = 10 \cdot \mathbf{I}_{k_\theta}$, $\mathbf{a}_{\mathbf{h}} = \mathbf{0}$, $\mathbf{V}_{\mathbf{h}} = 10 \cdot \mathbf{I}_n$ and $\nu_{\theta i} = \nu_{hj} = 5$. The scale parameters for the volatilities, i.e. S_{hj}, are then chosen so that the prior means are 0.1^2 respectively. Finally, we distinguish between VAR coefficients and intercept terms, by choosing the scale parameters so that the prior means are set to 0.01^2 for the former and 0.1^2 for the latter.

3.1.2. Markov-switching VARs

Let $s_t \in \{1, \ldots, M\}$ denote a regime indicator variable at date t, where M denote the number of regimes. Then, the structural M-state Markov-Switching VAR (MS-VAR) can be expressed as:

$$\mathbf{B}_{0,s_t}\mathbf{y}_t = \mathbf{b}_{s_t} + \mathbf{B}_{1,s_t}\mathbf{y}_{t-1} + \cdots + \mathbf{B}_{p,s_t}\mathbf{y}_{t-p} + \boldsymbol{e}_t, \quad \boldsymbol{e}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Omega}_{s_t}), \tag{8}$$

where the collection $(\mathbf{B}_{0,s_t}, \mathbf{b}_{s_t}, \mathbf{B}_{1,s_t}, \ldots, \mathbf{B}_{p,s_t})$ and $\boldsymbol{\Omega}_{s_t}$ are regime-specific VAR parameters, which are of the same dimension and structure as their time varying counterparts in (3). The regime indicator variable s_t is assumed to follow a M-state Markov process with transition probabilities $Pr(s_t = j | s_{t-1} = i) = p_{ij}$, $i, j = 1, \ldots, M$. In practice we set $M = 2$ to allow for the natural interpretation of high and low volatility regimes.[5]

Since there is no computational efficiency to be gained through direct estimation of this model, we instead directly estimate the reduced form variant:

$$\mathbf{y}_t = \mathbf{c}_{s_t} + \mathbf{A}_{1,s_t}\mathbf{y}_{t-1} + \cdots + \mathbf{A}_{p,s_t}\mathbf{y}_{t-p} + \boldsymbol{\epsilon}_t, \quad \boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{s_t}), \tag{9}$$

where $\mathbf{c}_{s_t} = \mathbf{B}_{0,s_t}^{-1}\mathbf{b}_{s_t}$, $\mathbf{A}_{i,s_t} = \mathbf{B}_{0,s_t}^{-1}\mathbf{B}_{i,s_t}$, $i = 1, \ldots, p$, $\boldsymbol{\epsilon}_t = \mathbf{B}_{0,s_t}^{-1}\boldsymbol{e}_t$ and $\boldsymbol{\Sigma}_{s_t} = \left(\mathbf{B}'_{0,s_t}\boldsymbol{\Omega}_{s_t}^{-1}\mathbf{B}_{0,s_t}\right)^{-1}$.

For estimation purposes, it is useful to rewrite (9) as:

$$\mathbf{y}_t = \mathbf{X}_t\boldsymbol{\beta}_{s_t} + \boldsymbol{\epsilon}_t, \tag{10}$$

where $\boldsymbol{\beta}_{s_t} = \text{vec}\left((\mathbf{c}_{s_t}, \mathbf{A}_{1,s_t}, \ldots, \mathbf{A}_{p,s_t})'\right)$ is a $k_\beta \times 1$ vector of VAR coefficients and $\mathbf{X}_t = \mathbf{I}_n \otimes (1, \mathbf{y}'_{t-1}, \ldots, \mathbf{y}'_{t-p})$.

To complete the model specification, we assume the following independent priors:

$$\boldsymbol{\beta}_i \sim \mathcal{N}(\boldsymbol{\beta}_0, \mathbf{V}_0), \quad \boldsymbol{\Sigma}_i \sim \mathcal{IW}(\mathbf{S}_0, \nu_0), \quad (p_{i1}, \ldots, p_{iM}) \sim \mathcal{D}(\alpha_{i1}, \ldots, \alpha_{iM}), \quad \text{for } i = 1, \ldots, M,$$

where $\mathcal{IW}(\mathbf{S}, \nu)$ denotes the Inverse Wishart distribution with scale matrix $\mathbf{S}$ and the degree of freedom ν, and $\mathcal{D}(a_1, \ldots, a_M)$ denotes the Dirichlet distribution with concentration parameters $(a_1, \ldots, a_M)$. Since the time series of interest exhibit high persistence, frequent switching among regimes over time is empirically implausible. To incorporate this fact, we implement an informative prior on the regime transition probability in which the concentration matrix is constrained to be

$$\begin{pmatrix} \alpha_{11} & \alpha_{12} & \ldots & \alpha_{1M} \\ \alpha_{21} & \alpha_{22} & \ldots & \alpha_{2M} \\ \vdots & \vdots & \ddots & \vdots \\ \alpha_{M1} & \alpha_{M2} & \ldots & \alpha_{MM} \end{pmatrix} = \mathbf{1}_M + \rho\mathbf{I}_M,$$

where $\mathbf{1}_M$ is a $M \times M$ matrix with its entries all equal to one and $\rho > 0$ governs the degree of regime persistence. For instance, it is easily verified that the expected value of probability for two subsequent periods belonging in the same regime is $\mathbb{E}(p_{ii}) = \frac{1+\rho}{\rho+M}$, which implies that a higher value of ρ indicates a high regime persistence.

The hyperparameters for these distributions are set as follows. First, we utilize a Minnesota

[5]We provide a statistical justification for this assumption in Section 3.2.

prior for $\boldsymbol{\beta}$ in which $\boldsymbol{\beta}_0 = \mathbf{0}$ and $\mathbf{V}_0 = \text{diag}(v_1, \ldots, v_{k_\beta})$, where the entries in $(v_1, \ldots, v_k)$ correspond to those in $\text{vec}((\mathbf{c}_0, \mathbf{A}_{10}, , \ldots, \mathbf{A}_{p0})')$. We distinguish between the intercepts and VAR coefficients by setting the v_i associated with the former to be 10 and those with the latter as

$$v_i = \begin{cases} \frac{\lambda_1^2}{r^2} & \text{for coefficients on own lags,} \\ \frac{\lambda_1^2 \lambda_2}{r^2} \frac{\sigma_i}{\sigma_j} & \text{for coefficients on cross lags,} \end{cases}$$

where σ_j is set equal to the standard deviation of the residual from AR(p) model for the variable $i = 1, \ldots, n$ and $r = 1, \ldots, p$. The remaining hyperparameters are set as $\lambda_1 = 0.2$ and $\lambda_2 = 0.5$. This specification allows us to capture a number of economic ideas. For instance, the fact that recent lags are more important than older ones is captured by the geometric rate of decay on the term r^2. Also, by setting $\lambda_2 < 1$ we incorporate the belief that own lags are likely to be more important than cross lags. Next, for the covariance matrix, we set $\nu_0 = n + 4$ and $\mathbf{S}_0 = (\nu_0 - n - 1) \times \mathbf{I}_n$. Finally, for the transition probabilities, we set $\rho = 10$ implying that the prior probability of transitioning between two states is approximately 0.08.

3.1.3. Competing Models

We identify the relevant modeling features of the observed data by comparing the in-sample fit of the TVP-VAR-SV as MS-VAR models with various restricted versions. A complete list of the considered models is provided in Table 1. Note that the C-VAR-SV is a restricted version of the TVP-VAR-SV model with $\boldsymbol{\theta}_t = \boldsymbol{\theta}_0$ for all dates $t = 1, \ldots, T$, while the TVP-VAR-C makes the restriction that $\mathbf{h}_t = \mathbf{h}_0$. Similarly, the MS-VAR-C is a restricted version of the MS-VAR model in which the error covariance matrix is restricted to be constant across regimes, while the C-VAR-MS makes the restriction that the coefficients are constant across regimes. Finally, the CVAR can be viewed as either a TVP-VAR-SV or MS-VAR with the appropriate restrictions on both the VAR coefficients and the error covariance matrix. We estimate each of these models using MCMC methods, details of which are provided in Appendix A.2.

Table 1: A list of competing models.

Model	Description
CVAR	A VAR with constant coefficients & constant error covariance
C-VAR-SV	A VAR with constant coefficients & time varying error covariance
TVP-VAR-C	A VAR with time varying coefficients & constant error covariance
TVP-VAR-SV	A VAR with time varying coefficients & time varying error covariance
MS-VAR-C	A VAR with regime-dependent coefficients & constant error covariance
C-VAR-MS	A VAR with constant coefficients & regime-dependent error covariance
MS-VAR	A VAR with joint regime-dependence in both coefficients & error covariance

3.2. Model Comparison

Suppose we are interested in comparing the in-sample fit of two distinct models M_i and M_j. In a Bayesian framework, each model is formally defined by a likelihood function, denoted by $p(\mathbf{y}|\boldsymbol{\theta}_k, M_k)$, $k = i, j$, and a prior probability distribution on the model-specific parameter vector $\boldsymbol{\theta}_k$, denoted by $p(\boldsymbol{\theta}_k|M_k)$. Given this information, a formal method of model comparison is the Bayes factor of M_i against M_j, which is defined as:

$$\mathrm{BF}_{ij} = \frac{p(\mathbf{y}|M_i)}{p(\mathbf{y}|M_j)}, \tag{11}$$

where

$$p(\mathbf{y}|M_k) = \int p(\mathbf{y}|M_k, \boldsymbol{\theta}_k)\, p(\boldsymbol{\theta}_k|M_k)\, d\boldsymbol{\theta}_k, \tag{12}$$

is the marginal likelihood of M_k , $k = i, j$.

As discussed in Geweke and Amisano (2011), the marginal likelihood can be interpreted as a density forecast of the observed data $\mathbf{y}$ under the M_k. This implies that the marginal likelihood value will be relatively large for models in which the observed data are more likely. Thus, if the observed data are more likely under M_i as compared to M_j, then $\mathrm{BF}_{ij} > 1$. In this case, posterior inference would then be conducted with M_i. More generally, given a set of m models, $\mathcal{M} = \{M_1, \ldots, M_m\}$, the model with the largest marginal likelihood value will be used to generate posterior inference. The model comparison results for the set of models presented in Table 1 is reported in Table 2.[6]

Table 2: Log marginal likelihoods for the set of VAR models

CVAR	C-VAR-SV	TVP-VAR-C	TVP-VAR-SV	C-VAR-MS	MS-VAR-C	MS-VAR
-706.446	-834.10	-2518.80	-987.56	**-696.43**	-699.17	-702.71

Note: Log marginal likelihood for the best model is in bold.

Three points are worth noting. First, in contrast to Cross and Nguyen (2017) we find that the CVAR model provides superior in-sample fit compared to the class of time varying parameter

[6]These (log-) marginal likelihood values are computed with a recently developed importance sampling based algorithm proposed in Chan and Eisenstat (2017). We use this algorithm because recent work has shown traditional methods such as the (modified) harmonic mean of Gelfand and Dey (1994) or the method in Chib (1995) to be extremely inaccurate. Specifically, see Chan and Grant (2015) for a discussion of the former and Frühwirth-Schnatter and Wagner (2010) for a similar discussion regarding the latter. Since it would lead us too far astray, we do not list the algorithm here, but instead refer the interested reader to Chan and Eisenstat (2017) for details.

models. The reason for this discrepancy is that in this paper we have explicitly disentangled the real GDP components between China and the rest of the world. This highlights the importance of this step in our estimation procedure. Second, the class of MS-VAR models outperforms both the CVAR and the class of time varying parameter models. Third, the C-VAR-MS is the best model. Statistically speaking, this suggests that allowing for regime dependence in the error covariance matrix is a significant modeling feature of the data. Economically, this means that volatility clustering at high and low levels is a key feature regarding the relationship between China's economy and world oil markets. Moreover, the use of alternative models may lead to biased estimates of the underlying shocks. For instance, the use of a traditional (homoscedastic) VAR models would erroneously average the shocks in both regimes, resulting in a positive bias during the low volatility regime and a negative bias during the high volatility regime. While concrete evidence of this claim requires the identification of the underlying structural shocks, suggestive evidence can be found in Figure 3, which presents the posterior means and 68 percent credible sets for the respective elements in the reduced form covariance matrix.[7]

[Insert Figure 3 here]

Since the credible sets of the variance terms (those elements on the main diagonal) do not contain zero, they are each statistically significant over the sample period. Hence the above claim seems to hold. Interestingly, the same can not be said for the covariance terms. In those cases, the credible sets contain zero in the low volatility regime (i.e., the entire sample excluding the periods 2000-04 and 2007-09). While this suggests that contemporaneous shocks may not be important during normal economic conditions, it does not imply that they are never important over the entire sample. This is because the credible sets do not contain zero in the high volatility regime, and are therefore statistically significant. Taken together, these result suggests that while *own shocks* are likely to be of primary importance in periods of low volatility, *contemporaneous interactions* seem to become important in high volatility regimes. Moreover, the periods of high volatility coincide with the US recessions. We return to this point when discussing the main results in Section 5.

Finally, before moving on to the structural analysis, it is worth investigating whether the assumption of two regimes is a reasonable one. To this end, we follow Koop and Potter (1999) and compute the marginal likelihood for various C-VAR-MS models which differ in the number

[7]Under the assumption of normality, a 68 percent credible interval shows the set of values within one standard deviation of the mean.

of possible regimes. The results contained in Table 3 show that the two regime C-VAR-MS is the best model.

Table 3: Log marginal likelihoods for C-VAR-MS with various regimes

No. of regimes	2	3	4	5
Log-ML	**-696.43**	-701.17	-708.01	-706.50

Note: Log marginal likelihood for the best model is in bold.

4. Identification

It is well known that the identification of a structural VAR model is a subject of ongoing research.[8] Since a key objective of our paper is to provide a set of results that are comparable with the existing literature, we adopt the same set of agnostic sign restrictions used in Cross and Nguyen (2017).[9] These restrictions are based on the comparative statics of a simple supply and demand model for the global oil market, in which the quantity is measured by world oil production and the price is given by the real international price of crude oil.

The first type of shock are *supply shocks*, denoted by $\epsilon_{Q,t}$. Such shocks represent an exogenous disruption of global oil production that may be caused by, for example, geopolitical turmoil. Under this interpretation, positive oil supply shocks simultaneously cause positive responses of global oil production and world economic activities but reduces the real oil price. The second type of shock arises from the fact that increases in aggregate global economic activity, after excluding the China factor, tend to generate higher commodity prices, and are therefore called *global oil demand shocks*, $\epsilon_{Y_W,t}$. These shocks are associated with increases in both global oil production and the real price of oil. The third type of shocks originates from specific factor generated demand and are therefore called *oil specific demand shocks*, $\epsilon_{P,t}$. This idea comes from **?** who finds that increased in precautionary demand for crude oil which are associated with changes in market expectations about the availability of future oil supply relative

[8]See Kilian (2013) for a general overview of various identification strategies in VAR models; including the sign restriction method used in this paper. For a critical review of sign restrictions, we refer the reader to Fry and Pagan (2011).

[9]In a seminal paper, **?** proposed a recursive identification strategy for a trivariate VAR in which the observed data series are available at monthly frequency. The justification for these restrictions is that the short-run oil supply curve is perfectly inelastic, thus implying that global oil production does not respond to oil demand shocks instantaneously, but only with a delay of one month. Since this argument is based on the observation of a one month lag, such a restriction should not be imposed at the quarterly frequency. Instead, a large number of studies have sought to relax **?**'s identifying assumption and utilized sign restrictions to study the effect of oil supply and demand shocks on macroeconomic aggregates. See, for instance, Lippi and Nobili (2012); Kilian and Murphy (2012); Baumeister and Peersman (2013); Lütkepohl and Netšunajev (2014); Kilian and Murphy (2014); Cross and Nguyen (2017), among others.

to demand is an important factor causing oil price shocks. This implies that such shocks induce a positive correlation between the oil production and its real price but reduces global economic activity. Notice that we have not specified any directional responses of China's GDP growth to each of these shocks. This is because a key objective of this paper is to study the effects of global oil market shocks on China's macroeconomic growth. Instead, we remain agnostic and allow the reactions of these variables to be completely determined by the data.

The final type of shock that we consider is a *China specific demand shock*, $\epsilon_{Y_C,t}$. Since the primary purpose of this paper is to study the effects of such shocks on the world oil market, we do not impose any sign restrictions on the market response. By doing this, the direction and magnitude of the responses are purely determined by the data. The directional signs for these restrictions of the impact matrix are summarized in Table 4. A practical guide to the implementation of these restrictions is provided in Appendix B.

Table 4: Sign restrictions

	$\epsilon_{Q,t}$	$\epsilon_{Y_W,t}$	$\epsilon_{P,t}$	$\epsilon_{Y_C,t}$
Oil production	$+$	$+$	$+$	$\times$
World economic growth (excluding China)	$+$	$+$	$-$	$\times$
Real oil price	$-$	$+$	$+$	$\times$
China's GDP growth	$\times$	$\times$	$\times$	$+$

Note: $+$ and $-$ respectively indicate positive and negative responses, while $\times$ leaves the effect unrestricted.

The cost of providing such an agnostic identification procedure is that the structural shocks may be partially identified. For instance, under the restrictions in Table 4, if China specific demand shocks elicit a positive response in all of the variables in the system, and the aggregate demand shocks has a positive impact on China specific demand, then these two shocks are indistinguishable. To overcome this issue we add an additional elasticity restriction in which we assume that own shocks yield a greater responsiveness than contemporaneous interactions. This assumption has been adopted in a range of papers, e.g. Peersman (2005), Aastveit et al. (2015) and Cross and Nguyen (2017). For completeness, in Section 6 we provide estimation results with no elasticity restrictions.

5. Results

5.1. The Effect of Oil Market Shocks on China

In this section we investigative the effects of oil market shocks on China's output by considering the propagation of the identified intertemporal structural shocks. The supply shock is normalized as a negative shock implying a reduction in world oil production that would lead to increase the oil price, while reducing world economic growth (excluding China). In each case, the shock size is one standard deviation. The resulting impulse response functions are displayed in Figure 4. Following Bjørnland et al. (2018b), we use a box plot to illustrate the cumulative three year ahead impacts of the oil market shocks on China' GDP growth. In what follows, we discuss two broad conclusions.

First, in normal times, negative oil supply shocks have positive effects, while both positive oil demand and oil specific demand shocks have negative effects. Interestingly, the signs of these responses then switch during times of US economic crisis. This is a new result in the oil market literature which therefore builds on existing results. The finding that negative oil supply shocks can have positive effects complements those in Zhao et al. (2016), but is in contrast to those in Herwartz and Plödt (2016) and Cross and Nguyen (2017), who find that such shocks tend to have no impact. Interestingly, during the 2000 recession, we also find that such shocks elicit an almost zero response on China's real GDP growth. The result that positive oil demand shocks tend to have a negative effect on China's growth is also consistent with those in Tang et al. (2010) and Zhao et al. (2016) but in contrast with those in Du et al. (2010) and Herwartz and Plödt (2016). The responses are mostly consistent with Cross and Nguyen (2017) who also provide evidence that demand shocks had negative impacts in the 1990s and positive effects in the 2007/08 GFC, however they find no evidence of the earlier switch during the 2000 recession. Instead, their result is that China had no impact during this period. This suggests that the Markov-switching model used in this study is better capable of capturing this abrupt event as compared to the autorgeressive models used in that paper.

Second, the magnitude response of China's output to the oil price shocks are found to be small and economically insignificant. This result is consistent with evidence in Herwartz and Plödt (2016) and Cross and Nguyen (2017, 2018) who independently observe that the reaction of Chinese real GDP to different global oil price shocks is relatively flat. As discussed in Hamilton (2009); Aastveit et al. (2015) and Cross and Nguyen (2018), a likely reason for these small effects is the structure of China's energy expenditure. More specifically, Cross and Nguyen (2018) document that coal provides the dominant proportion of China's total energy expenditure share, with oil expenditure contributing between 24% and 35% of this total. The main takeaway from this point is that despite China being a major player in international oil markets, oil market shocks have historically had little impact on China's real GDP growth.

5.2. The China Factor in World Oil Markets

In this section we assess the impacts of China specific demand shocks on the oil price. To this end, we construct impulse response functions and counterfactual historical decompositions. The impulse response functions allow us to investigate how such shocks transmit through the oil market market, while the historical decompositions allow us to quantify their effects over the sample period.

Figure 5 displays the response of the oil price to one standard deviation shocks in the remaining three variables. The demand shocks exert increasing positive pressure on the oil price in times of US macroeconomic crisis, while the supply shock impact decreases. With respect to our research question, the responses show that the China factor has had a significant positive impact on the oil price, however this impact is likely substantially less than demand shocks from the rest of the world. Our findings therefore contrasts those in Mu and Ye (2011); Wu and Zhang (2014) and Cross and Nguyen (2017) who each find no evidence to support the hypothesis that China's demand for oil has impacted the world price.

To further investigate the role of each variable in driving the price of oil, we now construct counterfactual historical decompositions. Such decompositions were first proposed in Kilian and Lee (2014) as a method of indicating how the variable of interest—in this case the price of oil—would have evolved, if we replaced all realizations of shock j by zero, while preserving the remaining structural shocks in the model. Thus, for interpretation purposes, a positive value indicates that the structural shock increased the real price of oil and vice versa. The results for each counterfactual are presented in Figure 6. The two panels in the first row focus on the oil market variables, while those in the second row focus on the specific demand variables.

The top left panel shows that in the absence of supply shocks. Four observations are worth noting: the negative effects in the lead up to the recession in 2000, the sharp positive effects in the early 2000s, the positive effects throughout most of the 2000s, and the noticeable negative pressure around the time of the recent 2013 and 2015 price drops. To the best of our knowledge, the last observation is new in the literature. It suggests the recent price drop in the price of oil was largely due to positive oil supply shocks. Interestingly, a quick glance at the bottom right panel reveals that decreased demand from China has also played a roughly equivalent role.

The effects of the demand shocks more generally are somewhat similar to the supply shocks, however key differences do exist around the 2000 recession and the 2007/08 GFC. More specifically, China specific demand seems to have had a large positive impact on oil price movements in the year 2000, while each of the demand shocks had a large negative impact on the price of oil in the 2007/08 GFC. The largest contributor in the latter crisis was the aggregate demand shocks, whose effects are found to be approximately twice as large as those from China. Thus, while both supply and demand shocks have positively contributed to the oil price since the turn of the century, our results support **?** in the sense that demand shocks are found to be important causes of oil price movements around the time of the 2007/08 crisis. Finally, to address our key research question in some more detail, we elaborate on the results in the bottom right panel. Generally speaking, with the exception of the negative shock around the 2007/08 crisis, the results clearly suggest that the China factor has had a positive impact on the oil price over the past two decades.

Finally, to better relate these results to the variance decompositions in Aastveit et al. (2015) and Liu et al. (2016), in Figure 7 we report the 12-step ahead variance decomposition for the oil price across.

The results suggests that the proportion of variation explained by aggregate demand hocks from the rest of the world is around 30 percent during normal times and close to 60 percent during times of US macroeconomic crisis. In contrast, the proportion of variation explained by oil specific demand and oil supply shocks decreases duirng times of US macroeconomic crisis. This latter result is consistent with the findings in **?**. Interestingly, the proportion of variance explained by China's demand shocks is relatively constant over the sample period, at around 30 percent. Thus, while the China factor certainly explains a non-negligible proportion of the variance, our result is substantially less than that in Liu et al. (2016) who found such shocks accounted for 51% of the variation in the oil price between 2000 and 2014. A likely reason for the discrepant quantitative results is that their study limits the real economic activity variables in the VAR model to be US industrial production and China value-added. The issue with this is that any variation from Asia, which have been noted as an essential contributor to the price of oil , will be captured by the China factor. Evidence for this hypothesis can be found by comparing the variance decompositions between the two studies (Aastveit et al., 2015). Specifically, Aastveit et al. (2015) find that economic activity from developing economies can explain around 40% of the oil prices two year ahead forecast error variance, which is remarkably close to the 46% of variation that Liu et al. (2016) attribute to China at the same impulse horizon. In contrast, the CPB index used in this paper controls for the effects of Asia.

6. Robustness Checks

To assess the necessity of imposing elasticity assumptions on the various model relations we simply re-estimate the model without them. The results from this exercise are provided in Figures 8—10. It suffices to note that they are almost identical to those in our main analysis. This suggests that partial identification is not an issue in this study.

Next, we test the robustness of the results to the use of the WTI oil price. The results are provided in Figures 11—13. Again, the results are almost identical. The only difference is that the oil supply shocks now have an even smaller impact on China's real GDP growth during both the recession of the 2000s and 2007/08. Otherwise, our findings are robust across the two data sets.

7. Conclusion

Our primary objective in this paper was to answer two questions: (1) how demand shocks from China impact the world oil market; and (2) how shocks from the world oil market affect China's macroeconomic output. To address these questions, we first disentangled oil specific demand stemming from China and the rest of the world. We then proposed two classes of flexible VAR models: Markov-switching and time varying parameter VARs, that are capable of capturing either gradual or abrupt changes in the relationships among variables of interest and volatility clustering in the shocks. To distinguish the relevant modeling features of the data, we conducted a formal Bayesian model comparison exercise, based on the Bayes factor. This exercise revealed that the constant coefficient Markov-switching switching VAR with regime dependent covariance matrix (C-VAR-MS) was preferred to all other models. This model allowed us to detect key regime changes, which a traditional constant parameter VAR model would be unable to detect. In particular, we found evidence to support the idea of regime changes in times of global economic downturns.

Allowing for regime changes was shown to be particularly important when investigating our research questions. First, we found that oil supply shocks tend to elicit a positive response in China's output, while the response of oil demand shocks is negative. Importantly, the signs of these responses then switch during times of global economic downturn, highlighting the importance of our modeling choice. Next, when analyzing world oil price dynamics, we found that demand shocks have had significant positive impacts over the past two decades. Using a regime-dependent forecast error variance decomposition, we showed that the proportion of oil price variation explained by demand from China is around 30 percent over the sample period. Importantly, while China specific effects are relatively constant, the impact of rest of world aggregate demand shocks were found to be larger during times of global macroeconomic downturn. Finally, we also found evidence that the recent 2014/15 oil price drop was due to a combination of increased oil supply and decreased demand from China. Taken together, our results provide empirical support to the conjectured claims that China has been influential in driving oil price dynamics over the past two decades (Hamilton, 2009; ?; Baumeister and Peersman, 2013).

Figures

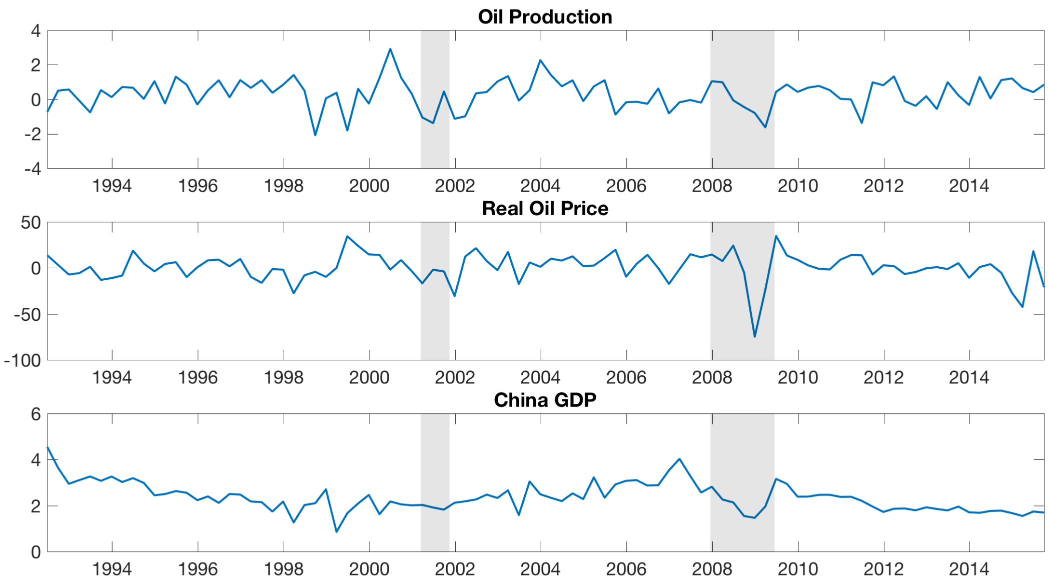

Figure 1: Historical evolution of the series (1992Q2-2015Q3).

Note: The raw data of crude oil production and prices are collected from EIA. China's GDP is sourced from Fed of Atlanta. All series are expressed in quarter-on-quarter percent changes. The shaded region shows recessions as defined by the NBER.

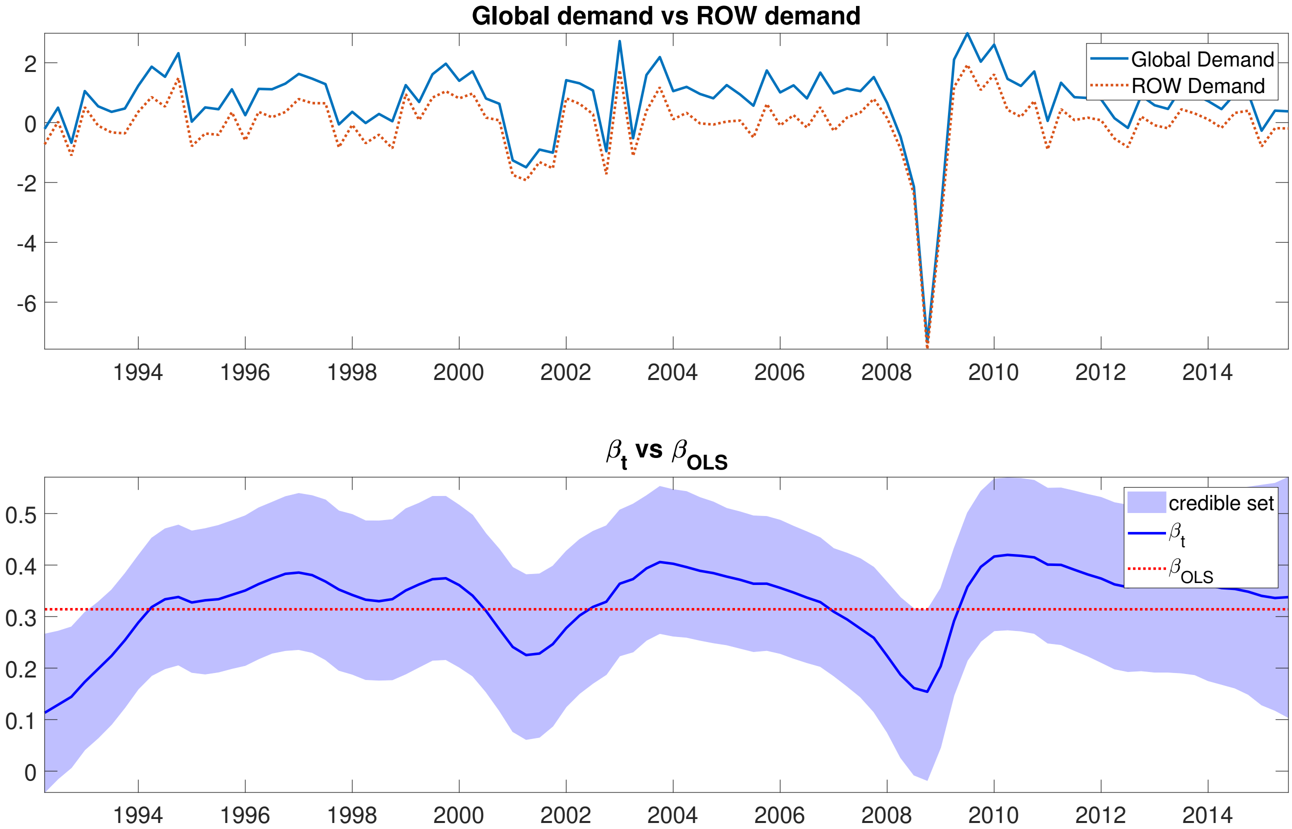

Figure 2: Global demand and ROW demand after excluding China's demand.

Note: β_t and ROW demand (e_t) derived from equation 1-2 and β_{OLS} obtained from the associated constant coefficient regression. The shaded region shows the 68% credible intervals.

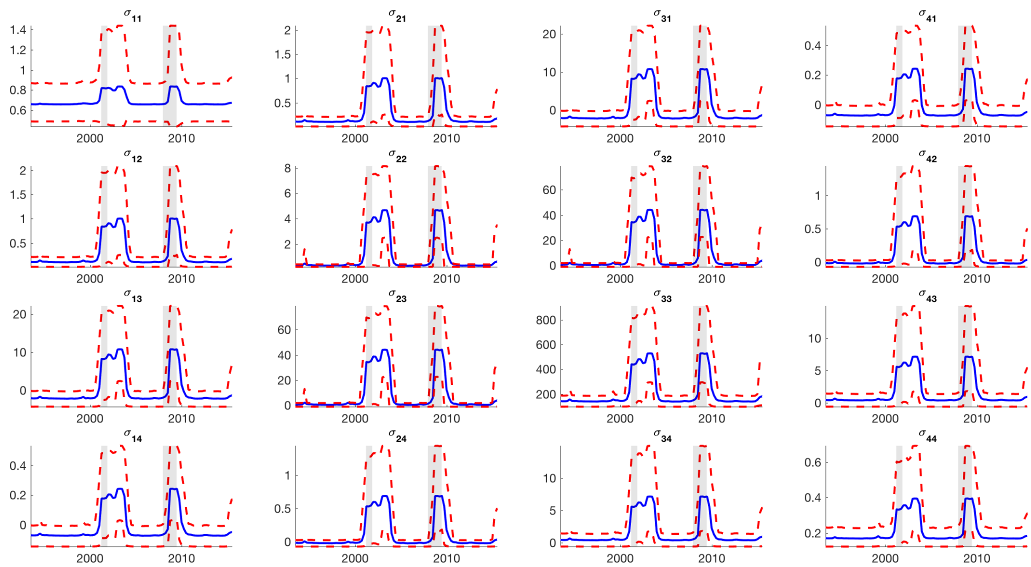

Figure 3: Reduced Form Covariance Matrix

Note: The figure plots time-varying reduced form covariance matrix of the C-VAR-MS model along with the 68% credible intervals. The shaded region shows recessions as defined by the NBER.

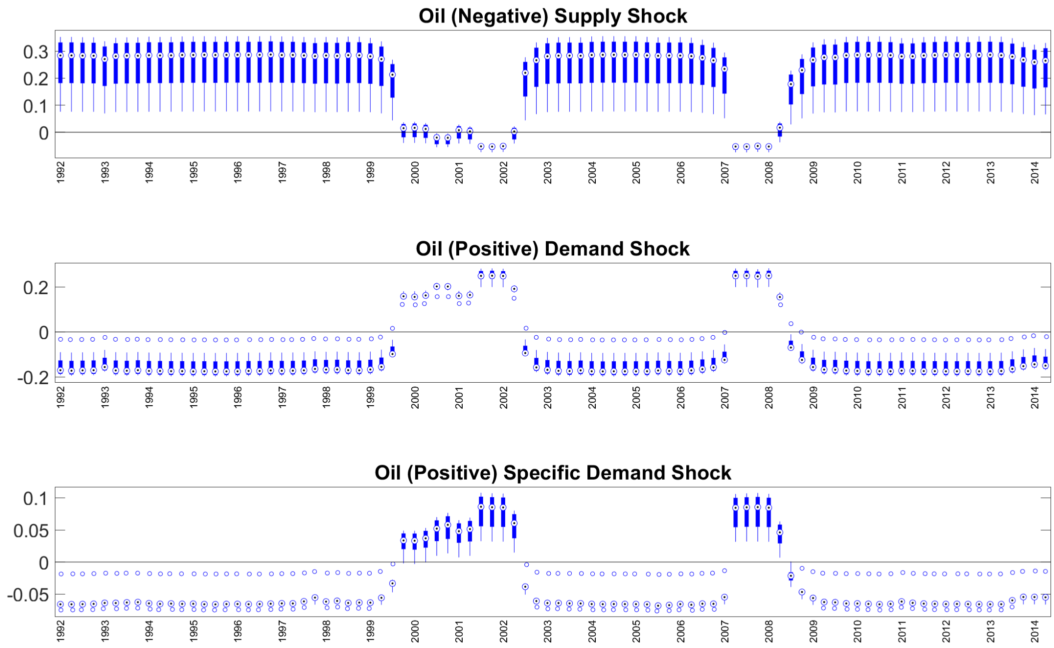

Figure 4: China's GDP response to the global oil market shocks.

Note: The boxes report the inter-quantile range of the cumulative three year ahead impacts. The circles inside each box is the median estimate.

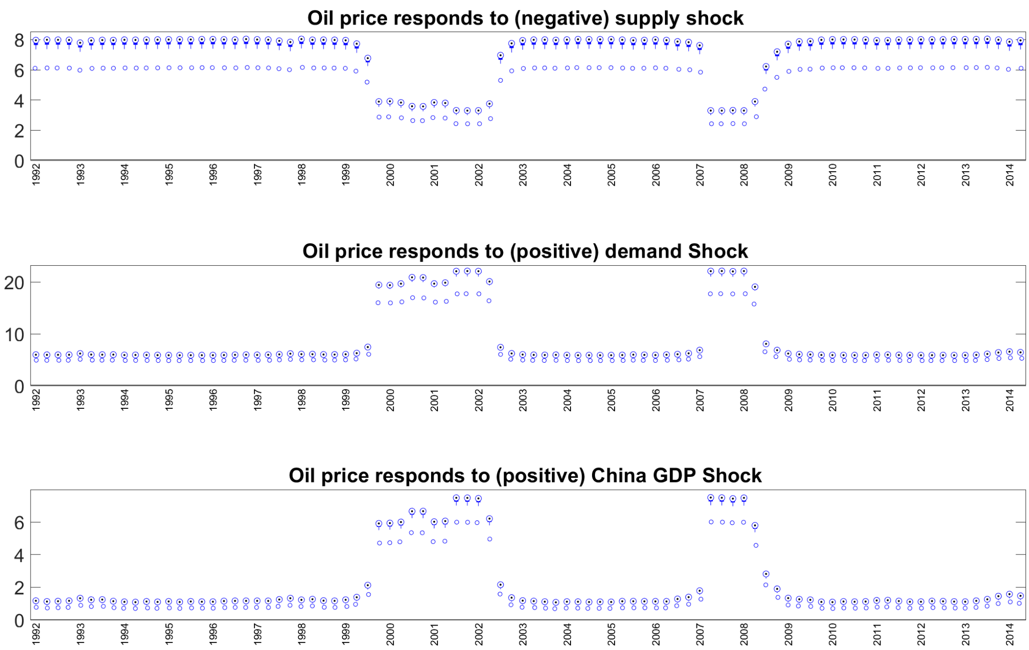

Figure 5: Oil price response to supply and demand shocks.

Note: See Figure 4.

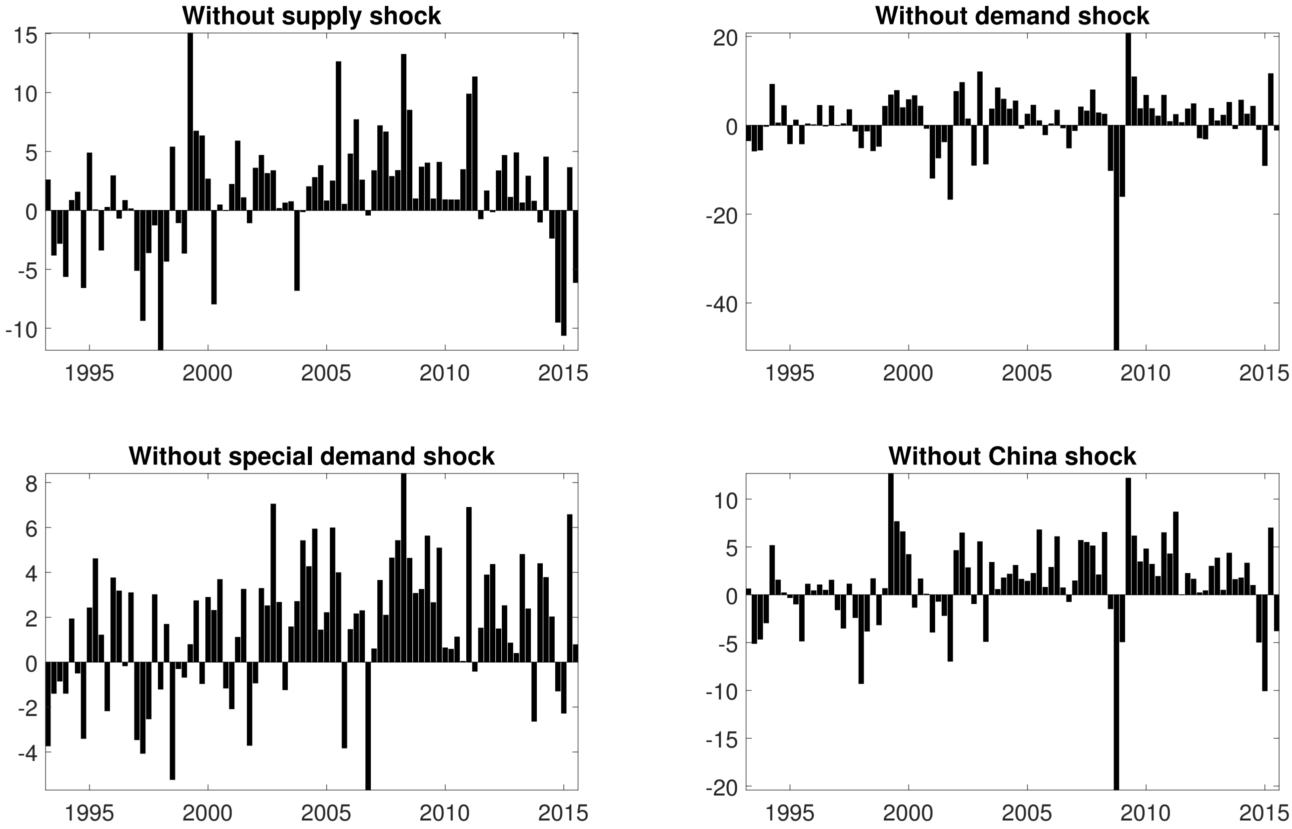

Figure 6: Historical decomposition of factors contributing to oil price fluctuations.

Note: The figure shows the difference between the actual oil real price and the total contribution of the structural shocks without the shock of interest. A positive value indicates that the structural shock contributed to increasing the price of oil, and vice versa.

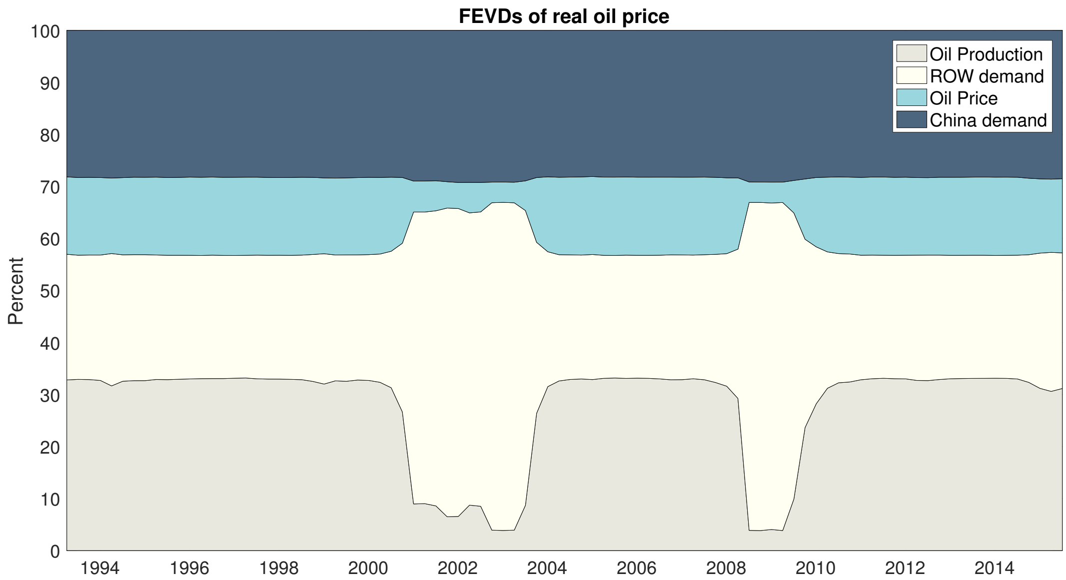

Figure 7: FEDVs of the real oil price

Note: The figure plots 12 step ahead regime conditional variance decomposition for the oil price over the sample period.

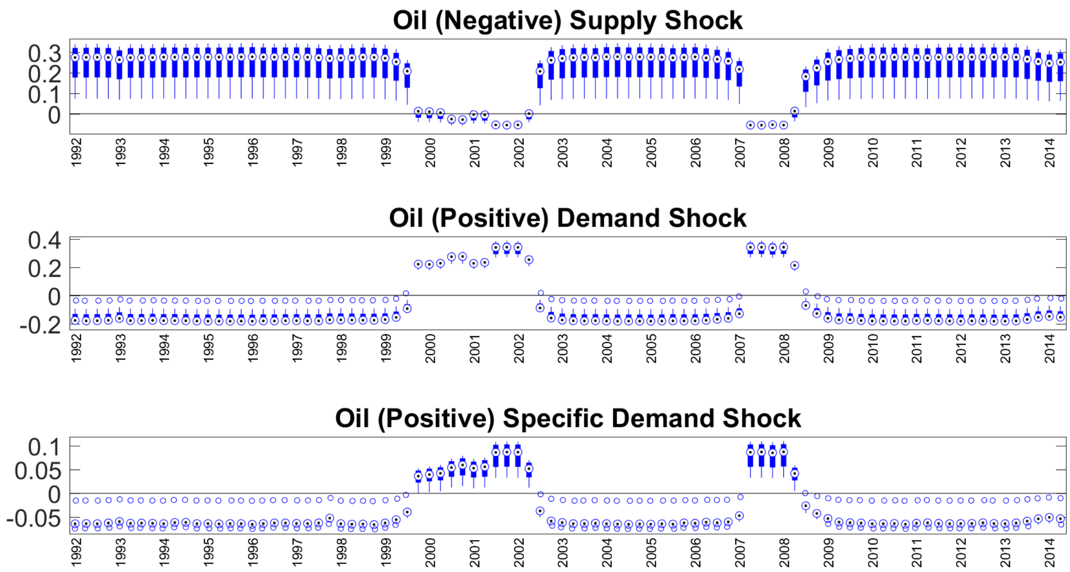

Figure 8: Robustness: China's GDP response to the global oil market shocks with no elasticity constraints.

Note: See Figure 4.

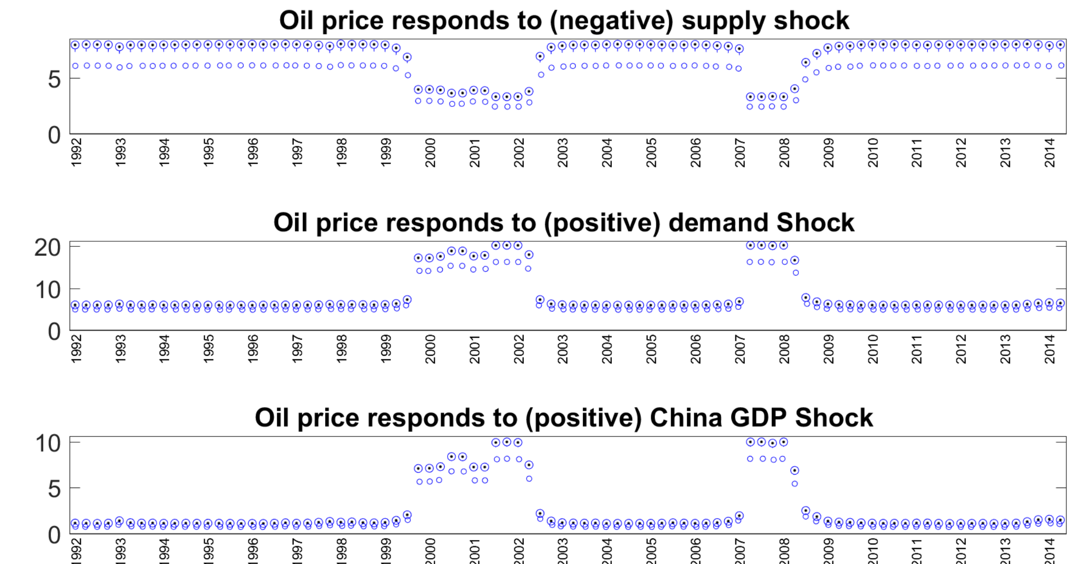

Figure 9: Robustness: Oil price response to supply and demand shocks with no elasticity constraints.

Note: See Figure 4.

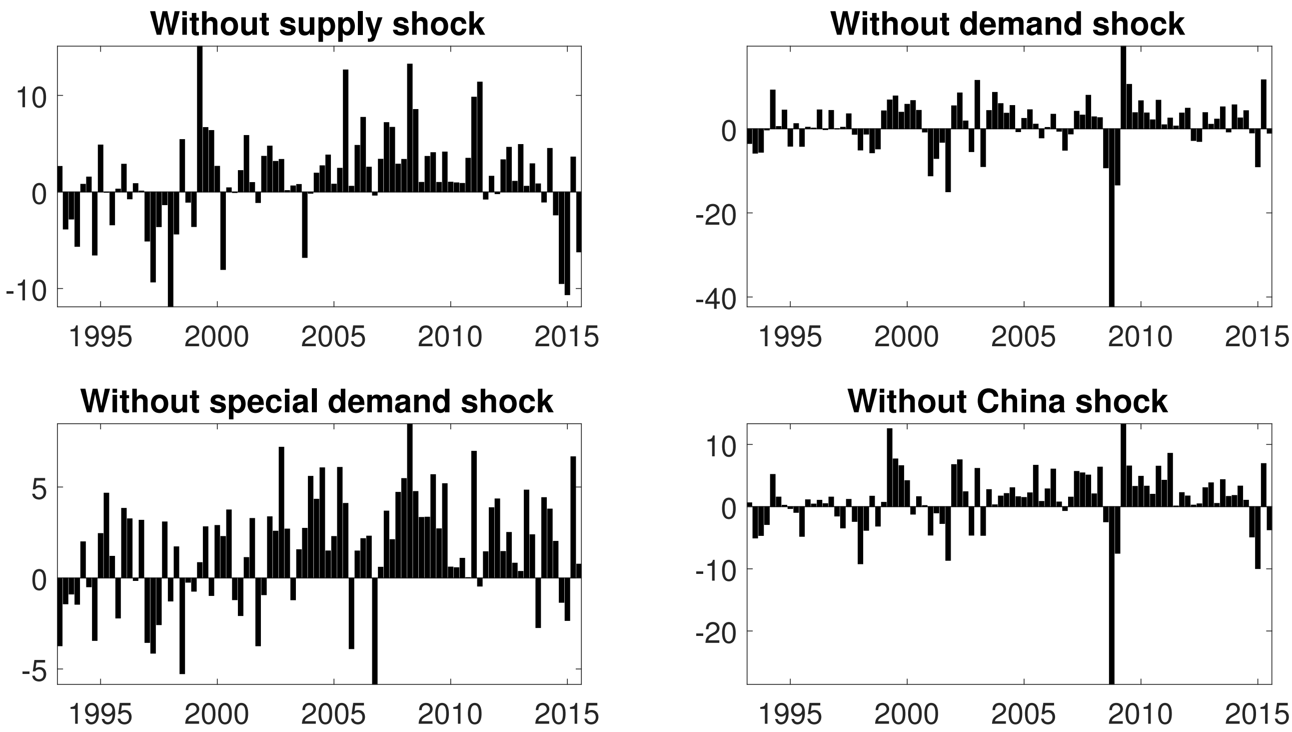

Figure 10: Robustness: Historical decomposion of factors contributing to oil price fluctuations with no elasticity constraints.

Note: See Figure 6

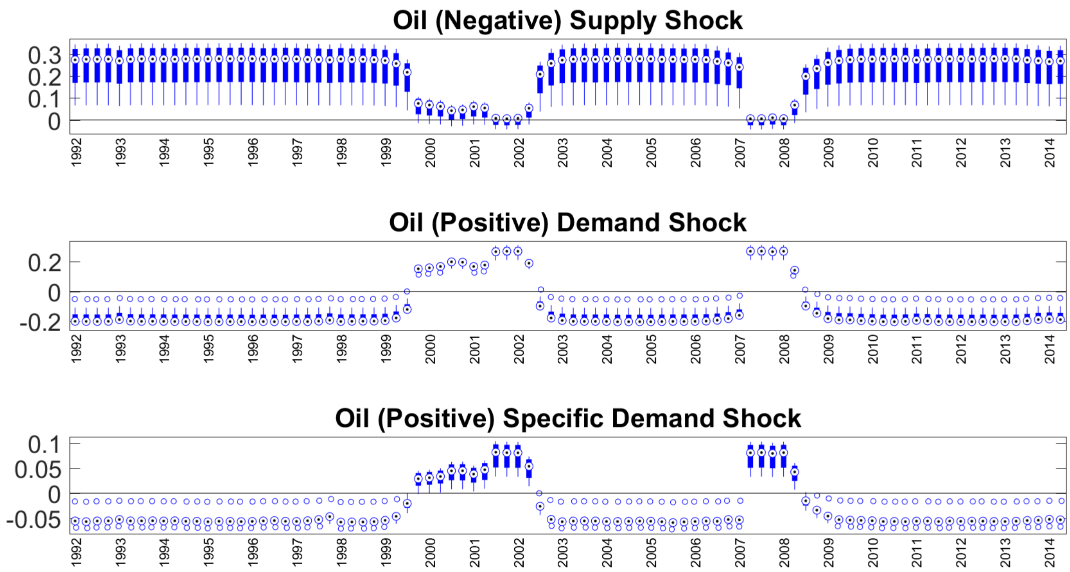

Figure 11: Robustness: China's GDP response to the global oil market shocks with WTI oil price data.

Note: See Figure 4.

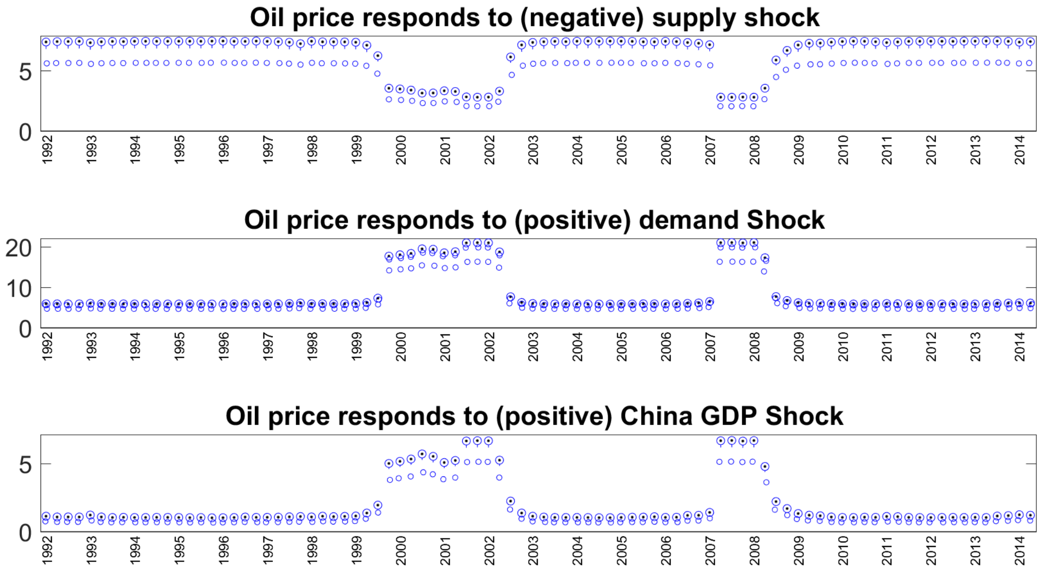

Figure 12: Robustness: Oil price response to supply and demand shocks with WTI oil price data.

Note: See Figure 4.

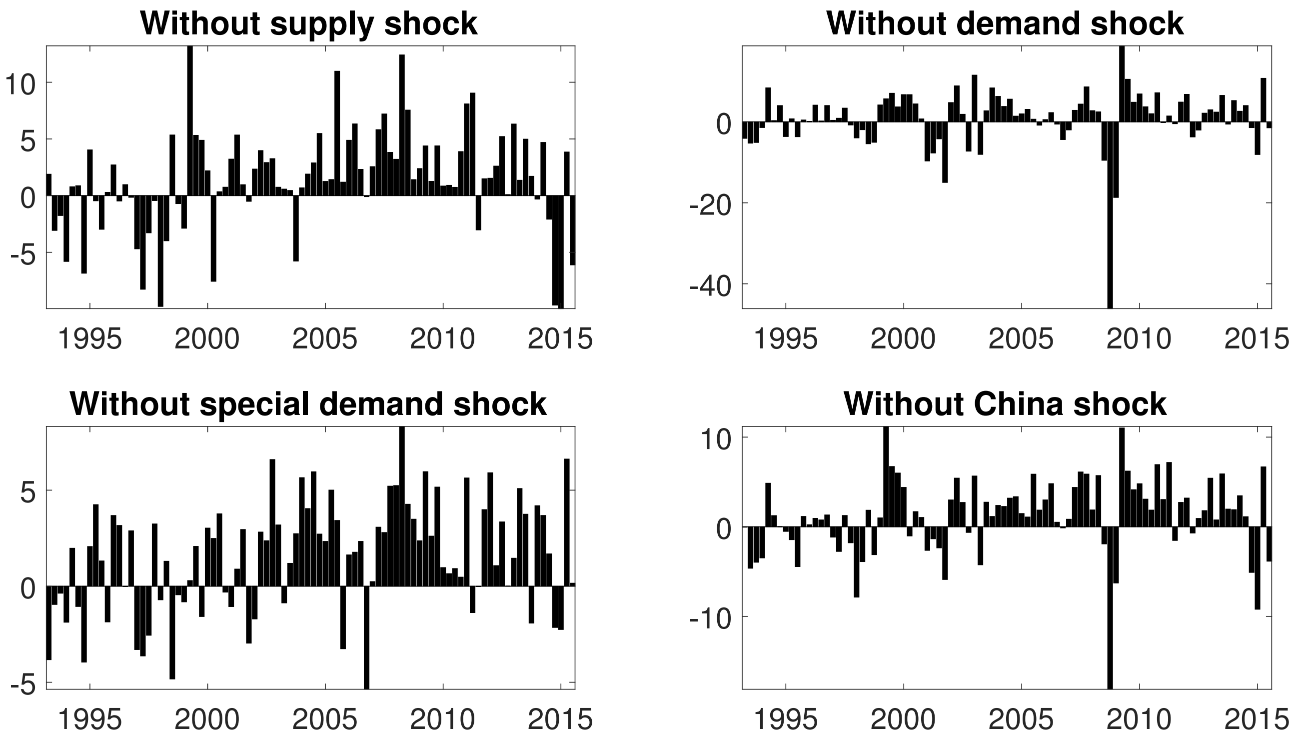

Figure 13: Robustness: Historical decomposion of factors contributing to oil price fluctuations with WTI oil price data.

Note: See Figure 6

Appendix A. Bayesian Estimation

In this appendix we outline the estimation details for fitting the TVP-VAR-SV and MS-VAR models. The remaining specifications are all nested versions of these models and can thus be estimated with straight forward modifications of the proceeding approach. In what follows we present these steps across two subsections. In each case we obtain $15,000$ posterior draws, discarding the first $5,000$ as a burn-in.

Appendix A.1. Stage 1: Regression Estimation

We start by outlining stage 1 of the Gibbs sampler in which we draw the latent parameters, variance terms and demand series in the time varying parameter regression model, discussed in Section 2. The demand series will then be used in the VAR models discussed in the subsequent sections.

We assume independent priors of the form:

$$\sigma^2_{y_w} \sim \mathcal{IG}\left(\nu_{0y_w}, S_{0y_w}\right), \, \sigma^2_{\beta} \sim \mathcal{IG}\left(\nu_{0\beta}, S_{0\beta}\right),$$

where $\mathcal{IG}(\nu, S)$ denotes the Inverse-Gamma distribution with shape parameter ν and scale parameter S. To ensure that the prior is non-informative we set a large prior variance. Specifically, we choose relatively small values for the shape parameters, i.e. $\nu_{0y_w} = \nu_{0\beta} = 5$, and set the scale terms so that the (unconditional) expectation of $\sigma^2_{y_w}$ is one and σ^2_{β} is 0.01.[10] Finally, for the initial condition, we let $\beta_1 \sim \mathcal{N}\left(\beta_0, \sigma^2_{0\beta}\right)$, where the mean is set to zero, i.e. $\beta_0 = 0$, and our uncertainty about this parameter the value is reflected by setting a large variance, i.e. $\sigma^2_{0\beta} = 10$.

Stage 1 of the posterior simulator can be viewed as a two block sampler that successively generates draws from the following full conditional distributions:

1. $p\left(\boldsymbol{\beta} \mid \mathbf{y}_w, \mathbf{y}_c, \sigma^2_{\beta}, \sigma^2_{y_w}\right)$,
2. $p\left(\sigma^2_{\beta}, \sigma^2_{y_w} \mid \mathbf{y}_w, \mathbf{y}_c, \boldsymbol{\beta}\right)$,

where $\mathbf{y}_w = (y_{w,1}, \ldots, y_{w,T})'$, $\mathbf{y}_c = (y_{c,1}, \ldots, y_{c,T})'$, $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_T)'$, $\mathbf{e} = (e_1, \ldots, e_T)'$, $\mathbf{u} = (u_1, \ldots, u_T)'$ and $\boldsymbol{\Theta}$ is the collection of VAR parameters that we estimate in subsequent blocks discussed in later sections. For instance, when estimating the TVP-VAR-SV model $\boldsymbol{\Theta} = (\boldsymbol{\beta}_t, \boldsymbol{\Sigma}_t)_{t=1}^{T}$. Similarly, when estimating the MS-VAR model $\boldsymbol{\Theta} = (\boldsymbol{\beta}_i, \boldsymbol{\Sigma}_i)_{i=1}^{M}$.

The conditional distribution in Step 1 is obtained by combining the likelihood from (1) with the state equation in (2). To that end, first write (1) in matrix form to get

$$\mathbf{y}_w = \mathbf{X}_c\boldsymbol{\beta} + \boldsymbol{e}, \quad \boldsymbol{e} \sim \mathcal{N}\left(\mathbf{0}, \sigma^2_{y_w}\mathbf{I}_T\right), \tag{A.1}$$

[10] To be clear, this is attained by setting $S_{0y_w} = 4$ and $S_{0\beta} = 0.04$.

where $\mathbf{X}_c = \text{diag}(y_{w,1}, \ldots, y_{w,T})$. Thus, by a change of variable $\left(\mathbf{y}_w|\mathbf{y}_c, \boldsymbol{\beta}, \sigma^2_\beta, \sigma^2_{y_w}\right) \sim \mathcal{N}\left(\mathbf{X}_c\boldsymbol{\beta}, \sigma^2_{y_w}\mathbf{I}_T\right)$. Next, define $\mathbf{H}$ to be the $T \times T$ first difference matrix. Then, we can write (2) as

$$\mathbf{H}_\beta\boldsymbol{\beta} = \tilde{\boldsymbol{\alpha}}_\beta + \mathbf{u}, \quad \mathbf{u} \sim \mathcal{N}\left(\mathbf{0}, \sigma^2_\beta\mathbf{I}_T\right), \tag{A.2}$$

where $\tilde{\boldsymbol{\alpha}}_\beta = (\beta_0, 0 \ldots, 0)'$ is a $T \times 1$ matrix. Since the first difference matrix is invertible, (A.2) can be equivalently expressed as:

$$\boldsymbol{\beta} = \boldsymbol{\alpha}_{\boldsymbol{\beta}} + \mathbf{H}_\beta^{-1}\mathbf{u}, \tag{A.3}$$

where $\boldsymbol{\alpha}_{\boldsymbol{\beta}} = \mathbf{H}_\beta^{-1}\tilde{\boldsymbol{\alpha}}_\beta$. Thus, by a change of variable $\left(\boldsymbol{\beta}|\sigma^2_\beta\right) \sim \mathcal{N}\left(\boldsymbol{\alpha}_{\boldsymbol{\beta}}, \sigma^2_\beta\left(\mathbf{H}'_{\boldsymbol{\beta}}\mathbf{H}_\beta\right)^{-1}\right)$. By multiplying this density with the likelihood it is easy to see that

$$\left(\boldsymbol{\beta} \mid \mathbf{y}_w, \mathbf{y}_c, \sigma^2_\beta, \sigma^2_{y_w}\right) \sim \mathcal{N}\left(\widehat{\boldsymbol{\beta}}, \mathbf{K}_{\boldsymbol{\beta}}^{-1}\right),$$

where

$$\mathbf{K}_{\boldsymbol{\beta}} = \frac{1}{\sigma^2_\beta}\mathbf{H}'_{\boldsymbol{\beta}}\mathbf{H}_\beta + \frac{1}{\sigma^2_{y_w}}\mathbf{X}'_c\mathbf{X}_c, \quad \widehat{\boldsymbol{\beta}} = \mathbf{K}_{\boldsymbol{\beta}}^{-1}\left(\frac{1}{\sigma^2_\beta}\mathbf{H}'_{\boldsymbol{\beta}}\mathbf{H}_\beta\boldsymbol{\alpha}_{\boldsymbol{\beta}} + \frac{1}{\sigma^2_{y_w}}\mathbf{X}'_c\mathbf{y}_w\right).$$

To implement Step 2, first note that the variance terms are independent. Thus, we can sample them one at a time. Using the fact that the Inverse-Gamma prior distribution is conjugate to the Gaussian likelihood gives

$$\left(\sigma^2_\beta \mid \mathbf{y}_w, \mathbf{y}_c, \boldsymbol{\beta}, \sigma^2_{y_w}\right) \sim \mathcal{IG}\left(\nu_\beta + \frac{T}{2}, S_\beta + \frac{1}{2}\sum_{t=1}^{T}\left(\beta_t - \beta_{t-1}\right)^2\right),$$
$$\left(\sigma^2_{y_w} \mid \mathbf{y}_w, \mathbf{y}_c, \boldsymbol{\beta}, \sigma^2_\beta\right) \sim \mathcal{IG}\left(\nu_{y_w} + \frac{T}{2}, S_{y_w} + \frac{1}{2}\sum_{t=1}^{T}\left(y_{w,t} - y_{c,t}\right)^2\right).$$

Appendix A.2. Stage 2a: TVP-VAR-SV Estimation

Let $\mathbf{y}_t$ denote a 4×1 vector containing observations on oil production, real world economic activity excluding China, real oil price and real economic activity in China. We then estimate the TVP-VAR-SV model by successively sampling from the following full conditional distributions:

4. $p\left(\boldsymbol{\theta}|\mathbf{y}, \mathbf{h}, \boldsymbol{\theta}_0, \mathbf{h}_0, \boldsymbol{\Omega}_{\boldsymbol{\theta}}, \boldsymbol{\Omega}_{\mathbf{h}}\right)$
5. $p\left(\mathbf{h}|\mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\theta}_0, \mathbf{h}_0, \boldsymbol{\Omega}_{\boldsymbol{\theta}}, \boldsymbol{\Omega}_{\mathbf{h}}\right)$
6. $p\left(\boldsymbol{\Omega}_{\boldsymbol{\theta}}, \boldsymbol{\Omega}_{\mathbf{h}}|\mathbf{y}, \boldsymbol{\theta}, \mathbf{h}, \boldsymbol{\theta}_0, \mathbf{h}_0\right)$

7. $p\left(\boldsymbol{\theta}_0, \mathbf{h}_0 | \mathbf{y}, \boldsymbol{\theta}, \mathbf{h}, \boldsymbol{\Omega}_{\boldsymbol{\theta}}, \boldsymbol{\Omega}_{\mathbf{h}}\right)$

where $\boldsymbol{\theta} = (\boldsymbol{\theta}_1', \ldots, \boldsymbol{\theta}_T')'$, $\mathbf{y} = (\mathbf{y}_1', \ldots, \mathbf{y}_T')'$ and $\mathbf{h} = (\mathbf{h}_1', \ldots, \mathbf{h}_T')'$. To implement Step 1, we first show that the conditional distribution of $\boldsymbol{\theta}$ is multivariate Gaussian. To this end, first write (5) as

$$\mathbf{y} = \mathbf{X}\boldsymbol{\theta} + \boldsymbol{\epsilon} \quad \boldsymbol{\epsilon} \sim \mathcal{N}\left(\mathbf{0}, \boldsymbol{\Sigma}\right), \tag{A.4}$$

where $\mathbf{X} = \text{diag}\left(\mathbf{X}_1, \ldots, \mathbf{X}_T\right)$, $\boldsymbol{\epsilon} = (\boldsymbol{\epsilon}_1', \ldots, \boldsymbol{\epsilon}_T')'$ and $\boldsymbol{\Sigma} = \text{diag}\left(\boldsymbol{\Sigma}_1, \ldots, \boldsymbol{\Sigma}_T\right)$. Thus, by a change of variable $\left(\mathbf{y} | \boldsymbol{\theta}, \mathbf{h}, \boldsymbol{\theta}_0, \mathbf{h}_0, \boldsymbol{\Omega}_{\boldsymbol{\theta}}, \boldsymbol{\Omega}_{\mathbf{h}}\right) \sim \mathcal{N}\left(\mathbf{X}\boldsymbol{\theta}, \boldsymbol{\Sigma}\right)$. Next, define $\mathbf{H}_{\boldsymbol{\theta}}$ to be the $Tk_\theta \times Tk_\theta$ first difference matrix. Then, we can write (6) as

$$\mathbf{H}_{\boldsymbol{\theta}}\boldsymbol{\theta} = \tilde{\boldsymbol{\alpha}}_{\boldsymbol{\theta}} + \boldsymbol{\nu} \quad \boldsymbol{\nu} \sim \mathcal{N}\left(\mathbf{0}, \mathbf{S}_{\boldsymbol{\theta}}\right), \tag{A.5}$$

where $\tilde{\boldsymbol{\alpha}}_{\boldsymbol{\theta}} = (\boldsymbol{\theta}_0', \mathbf{0}', \ldots, \mathbf{0}')'$ is a $Tk_\theta \times 1$ vector, $\boldsymbol{\nu} = (\boldsymbol{\nu}_1', \ldots, \boldsymbol{\nu}_T')'$ and $\mathbf{S}_{\boldsymbol{\theta}} = \mathbf{I}_T \otimes \boldsymbol{\Omega}_{\boldsymbol{\theta}}$. Since the first difference matrix is invertible, (A.5) can be equivalently expressed as

$$\boldsymbol{\theta} = \boldsymbol{\alpha}_{\boldsymbol{\theta}} + \mathbf{H}_{\boldsymbol{\theta}}^{-1}\boldsymbol{\nu}, \tag{A.6}$$

where $\boldsymbol{\alpha}_{\boldsymbol{\theta}} = \mathbf{H}_{\boldsymbol{\theta}}^{-1}\tilde{\boldsymbol{\alpha}}_{\boldsymbol{\theta}}$. Thus, by a change of variable $\left(\boldsymbol{\theta} | \boldsymbol{\theta}_0, \boldsymbol{\Omega}_{\boldsymbol{\theta}}\right) \sim \mathcal{N}\left(\boldsymbol{\alpha}_{\boldsymbol{\theta}}, \left(\mathbf{H}_{\boldsymbol{\theta}}'\mathbf{S}_{\boldsymbol{\theta}}^{-1}\mathbf{H}_{\boldsymbol{\theta}}\right)^{-1}\right)$. Using the same type of steps as in the linear regression case, it's possible to show that

$$\left(\boldsymbol{\theta} | \mathbf{y}, \mathbf{h}, \boldsymbol{\theta}_0, \mathbf{h}_0, \boldsymbol{\Omega}_{\boldsymbol{\theta}}, \boldsymbol{\Omega}_{\mathbf{h}}\right) \sim \mathcal{N}\left(\widehat{\boldsymbol{\theta}}, \mathbf{K}_{\boldsymbol{\theta}}^{-1}\right),$$

where

$$\mathbf{K}_{\boldsymbol{\theta}} = \mathbf{H}_{\boldsymbol{\theta}}'\mathbf{S}_{\boldsymbol{\theta}}^{-1}\mathbf{H}_{\boldsymbol{\theta}} + \mathbf{X}'\boldsymbol{\Sigma}^{-1}\mathbf{X}, \quad \widehat{\boldsymbol{\theta}} = \mathbf{K}_{\boldsymbol{\theta}}^{-1}\left(\mathbf{H}_{\boldsymbol{\theta}}'\mathbf{S}_{\boldsymbol{\theta}}^{-1}\mathbf{H}_{\boldsymbol{\theta}}\boldsymbol{\alpha}_{\boldsymbol{\theta}} + \mathbf{X}'\boldsymbol{\Sigma}^{-1}\mathbf{y}\right).$$

Note that the precision matrix $\mathbf{K}_{\boldsymbol{\theta}}$ is a band matrix-i.e., all of the nonzero elements exist along a narrow band clustered around the main diagonal. Consequently, we use the precision sampler of Chan and Jeliazkov (2009) to sample from $\left(\widehat{\boldsymbol{\theta}}, \mathbf{K}_{\boldsymbol{\theta}}^{-1}\right)$. The algorithm is based on fast band and sparse matrix routines and is consequently more efficient than conventional Kalman filter-based algorithms.

To implement Step 2, we apply the auxiliary mixture sampler of Kim et al. (1998) and precision sampler to sequentially draw elements of $\mathbf{h}_{i\bullet} = (\mathbf{h}_{i1}, \ldots, \mathbf{h}_{iT})$, $i = 1, \cdots, n$.

Steps 3 can be implemented by respectively sampling from

$$(\omega_{\theta i}^2|\mathbf{y},\boldsymbol{\theta},\mathbf{h},\boldsymbol{\theta}_0,\mathbf{h}_0) \sim \mathcal{IG}\left(\nu_{\theta i}+\frac{T}{2}, S_{\theta i}+\frac{1}{2}\sum_{t=1}^{T}(\theta_{i,t}-\theta_{i,t-1})^2\right), i=1,\cdots,k_\theta$$
$$(\omega_{hj}^2|\mathbf{y},\boldsymbol{\theta},\mathbf{h},\boldsymbol{\theta}_0,\mathbf{h}_0) \sim \mathcal{IG}\left(\nu_{hj}+\frac{T}{2}, S_{hj}+\frac{1}{2}\sum_{t=1}^{T}(h_{j,t}-h_{j,t-1})^2\right), j=1,\cdots,n.$$

Finally, to implement Step 4 apply similar steps as in the linear regression model to get

$$(\boldsymbol{\theta}_0|\mathbf{y},\boldsymbol{\theta},\mathbf{h},\boldsymbol{\Omega}_{\boldsymbol{\theta}},\boldsymbol{\Omega}_{\mathbf{h}}) \sim \mathcal{N}\left(\widehat{\boldsymbol{\theta}}_0, \mathbf{K}_{\boldsymbol{\theta}_0}^{-1}\right),$$
$$(\mathbf{h}_0|\mathbf{y},\boldsymbol{\theta},\mathbf{h},\boldsymbol{\Omega}_{\boldsymbol{\theta}},\boldsymbol{\Omega}_{\mathbf{h}}) \sim \mathcal{N}\left(\widehat{\mathbf{h}_0}, \mathbf{K}_{\mathbf{h}_0}^{-1}\right),$$

where

$$\mathbf{K}_{\boldsymbol{\theta}_0} = \mathbf{V}_{\boldsymbol{\theta}}^{-1}+\boldsymbol{\Omega}_{\boldsymbol{\theta}}^{-1}, \quad \widehat{\boldsymbol{\theta}}_0 = \mathbf{K}_{\boldsymbol{\theta}_0}^{-1}\left(\mathbf{V}_{\boldsymbol{\theta}}^{-1}\mathbf{a}_{\boldsymbol{\theta}}+\boldsymbol{\Omega}_{\boldsymbol{\theta}}^{-1}\boldsymbol{\theta}_1\right),$$
$$\mathbf{K}_{\mathbf{h}_0} = \mathbf{V}_{\mathbf{h}}^{-1}+\boldsymbol{\Omega}_{\mathbf{h}}^{-1}, \quad \widehat{\mathbf{h}}_0 = \mathbf{K}_{\mathbf{h}_0}^{-1}\left(\mathbf{V}_{\mathbf{h}}^{-1}\mathbf{a}_{\mathbf{h}}+\boldsymbol{\Omega}_{\mathbf{h}}^{-1}\mathbf{h}_1\right).$$

Appendix A.3. Stage 2b: MS-VAR Estimation

We estimate the MS-VAR model by successively sampling from the following full conditional distributions:

4. $p(\mathbf{s}|\boldsymbol{\Theta},\mathbf{y})$,
5. $p(\boldsymbol{\Theta}|\mathbf{s},\mathbf{y})$,
6. $p(\mathbf{P}|\mathbf{s})$,

where $\mathbf{s}=(s_1,\ldots,s_T)'$ is a vector of regime indicators, $\boldsymbol{\Theta}=\{(\boldsymbol{\beta}_i,\boldsymbol{\Sigma}_i)\}_{i=1}^M$ denotes the collection of model parameters across the M regimes, and $\mathbf{P}$ be the $M\times M$ Markov transition matrix, i.e., $\mathbf{P}_{ij}=p_{ij}$. To simplify the notation, in what follows we define $x_{t_1:t_2}=(x_{t_1},\ldots,x_{t_2})$ for a general variable x.

To implement Step 1 we apply the forward-backward algorithm of Chib (1996). To be specific, given $p(s_{t-1}|\mathbf{y}_{1:t-1},\theta)$ we compute $p(s_t|\mathbf{y}_{1:t})$ by

$$\begin{aligned}
p(s_t|\mathbf{y}_{1:t},\theta) &= \frac{p(y_t|s_t,\boldsymbol{\Theta})p(s_t|\mathbf{y}_{1:t-1},\boldsymbol{\Theta})}{\sum_{s_t}p(y_t|s_t,\boldsymbol{\Theta})p(s_t|\mathbf{y}_{1:t-1},\boldsymbol{\Theta})}, \\
&= \frac{p(y_t|s_t,\boldsymbol{\Theta})\sum_{s_{t-1}}p(s_t,s_{t-1}|\mathbf{y}_{1:t-1},\boldsymbol{\Theta})}{\sum_{s_t}p(y_t|s_t,\boldsymbol{\Theta})p(s_t|\mathbf{y}_{1:t-1},\boldsymbol{\Theta})}, \\
&= \frac{p(y_t|s_t,\boldsymbol{\Theta})\sum_{s_{t-1}}p(s_t|s_{t-1})p(s_{t-1}|\mathbf{y}_{1:t-1},\boldsymbol{\Theta})}{\sum_{s_t}p(y_t|s_t,\boldsymbol{\Theta})p(s_t|\mathbf{y}_{1:t-1},\boldsymbol{\Theta})}.
\end{aligned}$$

until we get $p(s_T|\mathbf{y}_{1:T}, \mathbf{\Theta})$. Then we implement the backward sampling by first sample s_T from $p(s_T|\mathbf{y}_{1:T}, \mathbf{\Theta})$, then we sample s_t given s_{t+1} from

$$p(s_t|s_{t+1:T}, \mathbf{y}_{1:T}, \mathbf{\Theta}) = \frac{p(s_t|\mathbf{y}_{1:t}, \mathbf{\Theta})p(s_{t+1}|s_t)}{\sum_{s_t} p(s_t|\mathbf{y}_{1:t}, \mathbf{\Theta})p(s_{t+1}|s_t)}.$$

To implement Step 2, first note that given $s_{1:T}$ we can regroup data into M distinct regimes. That is, for $i = 1, \ldots, M$, the model in a regime i can be written as

$$\mathbf{y}_i = \mathbf{X}_i\boldsymbol{\beta}_i + \epsilon_i \quad \epsilon_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{T_i} \otimes \mathbf{\Sigma}_i),$$

where $\mathbf{y}_i$ and $\mathbf{X}_i$ collect the observations belonging to regime i, and T_i is the number of observations in regime i. Following the standard results for the linear regression model, we have

$$\boldsymbol{\beta}_i \sim \mathcal{N}(\widehat{\boldsymbol{\beta}}_i, \widehat{\mathbf{K}}_i^{-1}), \quad \mathbf{\Sigma}_i \sim \mathcal{IW}(\widehat{\mathbf{S}}_i, \widehat{\nu}_i),$$

where $\widehat{\mathbf{K}}_i = \mathbf{X}_i{}' \left(\mathbf{I}_{T_i} \otimes \mathbf{\Sigma}_i\right)^{-1} \mathbf{X}_i + \mathbf{V}_0^{-1}$, $\widehat{\boldsymbol{\beta}}_i = \widehat{\mathbf{K}}_i^{-1} \left(\mathbf{X}_i{}' \left(\mathbf{I}_{T_i} \otimes \mathbf{\Sigma}_i\right)^{-1} \mathbf{y}_i + \mathbf{V}_0^{-1}\boldsymbol{\beta}_0\right)$, $\widehat{\nu}_i = T_i + \nu_0$ and $\widehat{\mathbf{S}}_i = \left(\mathbf{y}_i - \mathbf{X}_i\right)\left(\mathbf{y}_i - \mathbf{X}_i\right)' + \mathbf{S}_0$.

To implement Step 3, we draw the jth row of $\mathbf{P}$ for $j = 1, \ldots, M$, given $s_{1:T}$, according to

$$(p_{j1}, \ldots, p_{jM}) \sim \mathcal{D}(\alpha_{j1} + n_{j1}, \ldots, \alpha_{jM} + n_{jM}),$$

where $n_{kl} = \sum_{j=1}^{T-1} \mathbb{1}\left(\{s_j = l, s_{j+1} = k\}\right)$ and $\mathbb{1}(A)$ is the indicator function that is equal to one if statement A is true and zero otherwise.

Appendix B. Identification by sign restrictions

The identification procedure discussed in Section 4 is implemented with the algorithm in Rubio-Ramirez et al. (2010), which is outlined as follows:

1. Take the eigenvalue-eigenvector decomposition of the reduced form covariance matrix: $\mathbf{\Sigma}_{s_t}$, so that $\mathbf{\Sigma}_{s_t} = \mathbf{P}_{s_t}\mathbf{D}_{s_t}\mathbf{P}'_{s_t}$ where $\mathbf{D}_{s_t}$ is a diagonal matrix of eigenvalues and $\mathbf{P}_{s_t}$ is a matrix of corresponding (right) eigenvectors.
2. Draw a random $n \times n$ matrix $\mathbf{K}$ with its entries following standard normal distribution.
3. Take the $\mathbf{QR}$ decomposition of $\mathbf{K}$ so that $\mathbf{K} = \mathbf{QR}$ where $\mathbf{Q}$ is an orthogonal matrix and $\mathbf{R}$ is an upper triangular matrix.
4. Compute the time varying impact matrix $\mathbf{A}_{s_t} := \mathbf{P}_{s_t}\mathbf{D}_{s_t}^{\frac{1}{2}}\mathbf{Q}'$.

5. Check that the proposed matrix satisfies the restrictions outlined in in Section 4. If yes, keep it. Otherwise, discard it and redraw **K**.

References

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

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

C. Baumeister, G. Peersman, I. Van Robays, et al. The economic consequences of oil shocks: differences across countries and time. *Inflation in an era of relative price shocks, Reserve Bank of Australia*, pages 91–128, 2010.

H. Bjørnland, V. Larsen, and J. Maih. Oil and macroeconomic (in) stability. *American Economic Journal: Macroeconomics*, 10(4):128–51, 2018a.

H. Bjørnland, L. Thorsrud, and R. Torvik. Dutch disease dynamics reconsidered. *CAMP Working Paper*, 2018b.

J. C. Chan and E. Eisenstat. Efficient estimation of Bayesian VARMAs with time-varying coefficients. *Journal of Applied Econometrics*, 32(7):1277–1297, 2017.

J. C. Chan and E. Eisenstat. Bayesian model comparison for time-varying parameter VARs with stochastic volatility. *Journal of Applied Econometrics*, 2018.

J. C. Chan and A. L. Grant. Pitfalls of estimating the marginal likelihood using the modified harmonic mean. *Economics Letters*, 131:29–33, 2015.

J. C. Chan and I. Jeliazkov. Efficient simulation and integrated likelihood estimation in state space models. *International Journal of Mathematical Modelling and Numerical Optimisation*, 1(1-2):101–120, 2009.

S. Chib. Marginal likelihood from the Gibbs output. *Journal of the American Statistical Association*, 90(432): 1313–1321, 1995.

S. Chib. Calculating posterior distributions and modal estimates in Markov mixture models. *Journal of Econometrics*, 75(1):79–97, 1996.

M. P. Clements and H.-M. Krolzig. Can oil shocks explain asymmetries in the US business cycle? In *Advances in Markov-Switching Models*, pages 41–60. Springer, 2002.

J. Cross and B. Nguyen. The relationship between global oil price shocks and China's output: A time-varying analysis. *Energy Economics*, 62:79–91, 2017.

J. Cross and B. H. Nguyen. Time varying macroeconomic effects of energy price shocks: A new measure for China. *Energy Economics*, 2018.

L. Du, H. Yanan, and C. Wei. The relationship between oil price shocks and China's macro-economy: an empirical analysis. *Energy Policy*, 38(8):4142–4151, 2010.

E. Eisenstat, J. C. Chan, and R. W. Strachan. Stochastic model specification search for time-varying parameter VARs. *Econometric Reviews*, 35(8-10):1638–1665, 2016.

S. Frühwirth-Schnatter and H. Wagner. Stochastic model specification search for Gaussian and partial non-Gaussian state space models. *Journal of Econometrics*, 154(1):85–100, 2010.

R. Fry and A. Pagan. Sign restrictions in structural vector autoregressions: A critical review. *Journal of Economic Literature*, 49(4):938–960, 2011.

A. E. Gelfand and D. K. Dey. Bayesian model choice: asymptotics and exact calculations. *Journal of the Royal Statistical Society. Series B*, pages 501–514, 1994.

J. Geweke and G. Amisano. Hierarchical Markov normal mixture models with applications to financial asset returns. *Journal of Applied Econometrics*, 26(1):1–29, 2011.

J. Hamilton. Measuring global economic activity. *Mimeo*, 2018.

J. D. Hamilton. Causes and Consequences of the Oil Shock of 2007-08. Technical report, National Bureau of Economic Research, 2009.

H. Herwartz and M. Plödt. The macroeconomic effects of oil price shocks: Evidence from a statistical identification approach. *Journal of International Money and Finance*, 61:30–44, 2016.

F. Holm-Hadulla and K. Hubrich. Macroeconomic implications of oil price fluctuations: a regime-switching framework for the euro area. Working Paper Series 2119, European Central Bank, Dec. 2017.

C. Hou and B. Nguyen. Understanding the US natural gas market: A Markov switching VAR approach. *Energy Economics*, forthcoming, 2018.

M. Kalouptsidi. Detection and impact of industrial subsidies: The case of Chinese shipbuilding. *The Review of Economic Studies*, 85(2):1111–1158, 2017.

L. Kilian. Structural vector autoregressions. In N. Hashimzade and M. A. Thornton, editors, *Handbook of research methods and applications in empirical macroeconomics*, chapter 22, pages 515–554. Oxford University Press, Oxford, 2013.

L. Kilian and T. K. Lee. Quantifying the speculative component in the real price of oil: The role of global oil inventories. *Journal of International Money and Finance*, 42:71–87, 2014.

L. Kilian and D. P. Murphy. 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, 2012.

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

L. Kilian and X. Zhou. Modeling fluctuations in the global demand for commodities. *Journal of International Money and Finance*, 88:54–78, 2018.

S. Kim, N. Shephard, and S. Chib. Stochastic volatility: likelihood inference and comparison with ARCH models. *The Review of Economic Studies*, 65(3):361–393, 1998.

G. Koop and S. M. Potter. Bayes factors and nonlinearity: Evidence from economic time series. *Journal of Econometrics*, 88(2):251–281, 1999.

B. Lin and J. Li. The Determinants of Endogenous Oil Price: Considering the Influence from China. *Emerging Markets Finance and Trade*, 51(5):1034–1050, 2015.

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

L. Liu, Y. Wang, C. Wu, and W. Wu. Disentangling the determinants of real oil prices. *Energy Economics*, 56: 363–373, 2016.

H. Lütkepohl and A. Netšunajev. Disentangling demand and supply shocks in the crude oil market: How to check sign restrictions in structural VARs. *Journal of Applied Econometrics*, 29(3):479–496, 2014.

X. Mu and H. Ye. Understanding the crude oil price: how important is the China factor? *Energy Journal*, 32 (4):10–050, 2011.

G. Peersman. What caused the early millennium slowdown? Evidence based on vector autoregressions. *Journal of Applied Econometrics*, 20(2):185–207, 2005.

G. E. Primiceri. Time varying structural vector autoregressions and monetary policy. *The Review of Economic Studies*, 72(3):821–852, 2005.

J. E. Raymond and R. W. Rich. Oil and the macroeconomy: A Markov state-switching approach. *Journal of Money, Credit, and Banking*, pages 193–213, 1997.

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

W. Tang, L. Wu, and Z. Zhang. Oil price shocks and their short-and long-term effects on the Chinese economy. *Energy Economics*, 32:S3–S14, 2010.

G. Wu and Y.-J. Zhang. Does China factor matter? An econometric analysis of international crude oil prices. *Energy Policy*, 72:78–86, 2014.

L. Zhao, X. Zhang, S. Wang, and S. Xu. The effects of oil price shocks on output and inflation in China. *Energy Economics*, 53:101–110, 2016.

Centre for Applied Macroeconomics and Commodity Prices (CAMP)
will bring together economists working on applied macroeconomic issues, with special emphasis on petroleum economics.

BI Norwegian Business School
Centre for Applied Macroeconomics and Commodity Prices (CAMP)
N-0442 Oslo

http://www.bi.no/camp

CAMP Working Paper Series
ISSN: 1892-2198