Cross, J. L., Hou, C., & Nguyen, B. H. (2021). On the China factor in the world oil market: A regime switching approach. Energy Economics, 95, 105119.

On the China factor in the world oil market: A regime switching approach*

Jamie L. Cross† Chenghan Hou‡ Bao H. Nguyen§

October 21, 2020

Abstract

We investigate the relationship between China's macroeconomic performance and the world oil market over the past two decades. Unlike existing studies, we allow for possible regime changes by utilizing a class of Markov-switching vector autoregression (MS-VAR) models. The model identifies key regime changes in the structural shocks when the oil market experiences low and high volatility. We find that demand shocks from China and the rest of the world have a larger impact on the real price of crude oil during periods of high volatility. Supply shocks, in contrast, have a large effect on the price in the low volatility regime. A similar state-dependent phenomenon is observed for the impact of oil price shocks on China economic activity, however the size of these responses is relatively small. Thus, despite China being a major player in international oil markets, we conclude that oil market shocks tend to have little impact on China's real GDP growth.

JEL-codes: C32, E31, E32

Keywords: Oil prices, China, Markov-switching VARs, Sign restrictions.

*We thank Hilde Bjørnland, Tatsuyoshi Okimoto, Ippei Fujiwara, Knut Aastveit, Leif Anders Thorsrud, Francesco Ravazzolo, 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.

†Centre of Applied Macroeconomics and Commodity Prices (CAMP), BI Norwegian Business School and Centre for Applied Macroeconomic Analysis (CAMA). Email: jamie.cross@bi.no bozhangyc@gmail.com

‡*Corresponding author*: Center for Economics, Finance and Management Studies, Hunan University. Email: chenghan.hou@hotmail.com

§Tasmanian School of Business and Economics, University of Tasmania and CAMA. Email: b.nguyen@utas.edu.au

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 macroeconomic expansion has resulted in the country becoming the world's largest energy consumer since 2011, the biggest net importer of petroleum and other liquid fuels since 2013 and the top crude oil importer since 2017. Despite these facts, the existing literature on the effect of China's macroeconomic growth on the world price of crude oil has provided conflicting results. For instance, Liu et al. (2016) finds that China specific demand accounted for half of the real oil price variation between 2000 and 2014, while many others report that China's demand had little or zero impact in the same period (Mu and Ye, 2011; Wu and Zhang, 2014; Lin and Li, 2015; Cross and Nguyen, 2017).

A possible reason for these conflicting results is that most studies utilize linear models. This is surprising for at least two reasons. First, the growth from China has been unexpectedly large over the past few decades (Kilian and Hicks, 2013), suggesting a possible regime change in the demand for crude oil over this period. Second, multiple scholars have shown that world energy market dynamics are best modeled by Markov-switching models (Raymond and Rich, 1997; Clements and Krolzig, 2002; Fong and See, 2002; Vo, 2009; Bjørnland et al., 2018; Basher et al., 2016; Holm-Hadulla and Hubrich, 2017; Hou and Nguyen, 2018), suggesting that linear models may be misspecified.

This paper breaks ground in this area by estimating a class of Markov-switching vector autoregression (MS-VAR) models to study the relationship between China's economic growth and the world market for crude oil. The importance of specifying this class of models is demonstrated both statistically and economically. Statistically, we find that the MS-VAR provides superior in-sample fit compared to a linear VAR as well as a class of time-varying latent-parameter VARs. We find clear evidence of regime changes in the size and frequency of exogenous shocks, and that the price of oil responds differently to the shocks in low and high volatility states. Demand shocks from both China and the rest of world have a larger impact on the real price of oil during times of high oil market volatility compared to low volatility. In contrast, supply shocks tend to have a larger impact during the low volatility state. A similar state-dependent phenomenon was also observed for the impact of oil price shocks on China economic activity. In that case we observed that China's real GDP growth responds more to economic activity shocks from the rest of the world in the high volatility state, but is less responsive to supply shocks. That being said, the size of the responses are relatively small. Thus, despite China being a major player in international oil markets, we conclude that oil market shocks tend to have little impact on China's real GDP growth.

Our results contribute to the growing literature on the relationship between China and the world market for crude oil. Interest in this relationship started when multiple researchers suggested that oil price dynamics since the turn of the century have been largely driven by growth in emerging markets, including China (Hamilton, 2009; Kilian, 2009; Baumeister and Peersman, 2013). Empirical evidence in support of these claims was then provided in Kilian and Hicks (2013) and Aastveit et al. (2015). For instance, Aastveit et al. (2015) 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. Liu et al. (2016) found that China specific demand was an important driver of the oil price between 2000 and 2014, however 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). Our results contribute towards this disagreement by noting that the effect is state-dependent, with demand shocks from both China having a larger impact on the real price of oil during times of high oil market volatility compared to low volatility.

A similar disagreement also exists among researchers who have examined the effects of oil market shocks on China's macroeconomic growth. Tang et al. (2010) and Zhao et al. (2016) find that unanticipated oil price increases 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) estimate a time-varying latent-parameter VAR model, and suggest that such shocks may elicit either positive or negative movements, depending on the type of shock and the period in which they occurred. Our results contribute towards this literature by first showing that the class of Markov-switching models is better suited to modeling the effects of oil market shocks on China's macroeconomic growth relative to linear or time-varying latent-parameter VAR models. Moreover, we find that while China's real GDP growth response is state dependent, the effects of supply shocks and economic activity related demand shocks tend to be positive in either state, while the residual demand component elicits a negative response. That being said, the size of the responses are relatively small, which leads us to conclude that oil market shocks tend to have little impact on China's real GDP growth.

The rest of the paper is organized as follows. In Section 2 we present the data and Section 3 discuss the methodology, including the MS-VAR, model selection and identification but defer technical details on estimation to the Appendix. In Section 4 we present the empirical results, followed by various robustness checks in Section 5. We conclude in Section 6.

2 Data

We use quarterly data between 1992Q1 and 2019Q2 on four variables of interest: oil production, the real price of oil, real global economic activity and real economic activity from China. The particular sample period is chosen mainly due to the availability of the time-series data of China's GDP.

In line with existing literature, we consider three alternative measures of the oil price: the US refiners' acquisition cost (RAC) for imported crude oil, the West Texas Intermediate (WTI) and the Brent price of crude oil. Since it is generally considered to be the best proxy for the free global oil price market (Kilian and Zhou, 2020), we use RAC for the benchmark model and the the WIT and the Brent price as a robustness check. The real oil price is obtained by deflating the nominal price by the US Consumer Price Index. Next, oil production is measured by the amount of world crude oil production (thousand barrels per day) as provided by the US Energy Information Administration (EIA). We measure world real economic activity by using the OECD+6 industrial production index provided by Baumeister and Hamilton (2019).[1] Finally, as in Cross and Nguyen (2017), real economic activity in China is measured by quarterly real GDP published by the 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.[2] All series are converted to quarter-on-quarter rates of growth by taking the first difference of the natural logarithm.

The series used in this paper are plotted in Figure 1. As can be seen from the figure, both of the global oil market and Chinese economy behave differently between good and bad times. We see that, during bad times, e.g. the Asian Financial Crisis in 1997 and the Global Financial Crisis in 2008-09, the volatility of all series are significantly larger than good times. However, the responses of the real price were different. The oil price is seen to increase during the former event while it is found to fall over the latter one. These features are suggestive of possible regime changes in the structural shocks. We will refer

[1]The choice of an appropriate measure of global economic activity has recently received much attention; see Kilian and Zhou (2018) for a survey of this literature. The two most commonly used proxies for this activity are the global *real economic activity* (REA) index constructed by Kilian (2009) and global *industrial production* (IP). For our analysis, we prefer the 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, Baumeister et al. (2020) and Hamilton (2019) has recently highlighted that the REA index has not captured well real economic activity in recent years, while world IP (IP) does.

[2]See Kilian and Zhou (2018) for a discussion.

to these periods when discussing the main results.

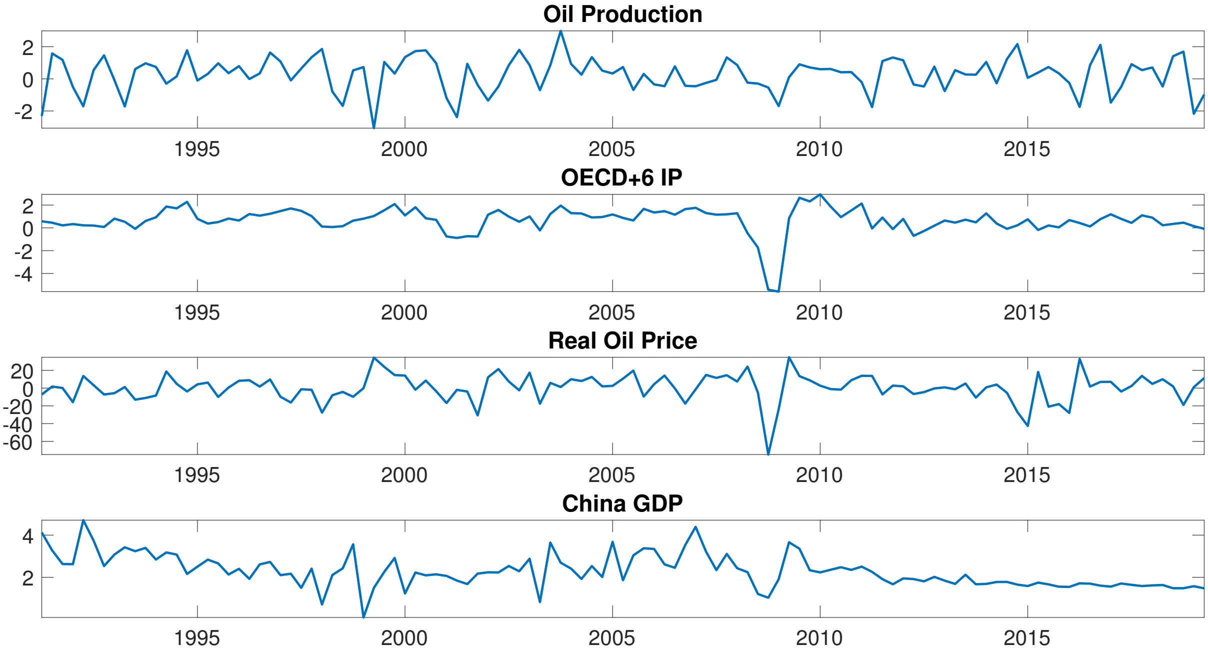

Figure 1: Historical evolution of the series (1992Q2-2019Q4).

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.

3 Empirical Methodology

To identify the relevant modeling features of the observed data, we first compare the in-sample fit of the well known constant parameter VAR (CVAR) with the more flexible MS-VAR models. A complete list of the considered models is provided in Table 1. We estimate each of these models using MCMC methods, details of which are provided in Appendix B.

Table 1: A list of competing models.

Model	Description
CVAR	A VAR with constant coefficients & constant 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

In each case, we set $p = 4$ and $M = 2$. The former is in line with the BIC, while setting the number of regimes to be two provides the natural interpretation of high and low volatility regimes.[3]

3.1 Markov Switching VAR

We consider a broad class of MS-VAR models that allow for regime changes in either, or both of, the VAR coefficients and the structural shocks. To conserve space, we discuss the most general model and note how other models arise as restricted versions.

The reduced form M-state Markov-Switching VAR (MS-VAR) is given by:

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

where $s_t \in \{1, \ldots, M\}$ denotes a regime indicator variable at date t, which is assumed to follow a M-state Markov process with transition probabilities $Pr(s_t = j | s_{t-1} = i) = p_{ij}$, $i, j = 1, \ldots, M$. The well known constant parameter VAR (CVAR) can be viewed as a MS-VAR without any regime dependence. 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, and the C-VAR-MS makes the restriction that the coefficients are constant across regimes.

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 $(a_1, \ldots, a_M)$. The choice of these distributions is motivated by the support of the parameter space and conjugacy. For example, we use a Gaussian prior distribution for the BVAR coefficients since they can exist anywhere on the real line and result in a Gaussian (conditional) posterior distribution when combined with the likelihood function. Similarly, a Dirichlet prior distribution is used for the transition probabilities because they exist in the unit interval and result in a Dirichlet (conditional) posterior distribution when combined with the likelihood function.

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

[3] In the robustness section we also show that setting the number of regimes to be two is also in line with marginal likelihood computations.

is constrained such that $\alpha_{ij} = 1$ for $i \neq j$ and $\alpha_{ij} = \rho > 0$ for $i = j$, i.e.

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

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 prior for $\boldsymbol{\beta}$ in which $\boldsymbol{\beta}_0 = \mathbf{0}$ and $\mathbf{V}_0 = \text{diag}(v_1, \dots, v_{k_\beta})$, where the entries in $(v_1, \dots, v_k)$ correspond to those in $\text{vec}((\mathbf{c}_0, \mathbf{A}_{10}, , \dots, \mathbf{A}_{p0})')$. We distinguish between the intercepts and VAR coefficients by setting the v_i associated with the former to be 100 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 $i = 1, \dots, n$ and $r = 1, \dots, p$. A detailed discussion of the Minnesota prior can be found in Doan et al. (1984); Koop and Korobilis (2010); Karlsson (2013).

Following Carriero et al. (2019), we set the hyperparameters in the Minnesota prior as $\lambda_1 = 0.04$ and $\lambda_2 = 0.04^2$. 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 $\rho = 50$ implying that the prior probability of transitioning between two states is approximately 0.02.

3.2 Model Selection

To determine which of these reduced form models best represents the data, we conduct a formal Bayesian model comparison exercise. For concreteness, 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:

$$\mathrm{BF}_{ij} = \frac{p(\mathbf{y}|M_i)}{p(\mathbf{y}|M_j)}, \tag{2}$$

where

$$p(\mathbf{y}|M_k) = \int p(\mathbf{y}|M_k, \boldsymbol{\theta}_k)\, p(\boldsymbol{\theta}_k|M_k)\, d\boldsymbol{\theta}_k, \tag{3}$$

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

The marginal likelihood values for the general class of MS-VAR models listed in Table 1 are computed with the importance sampling based algorithm proposed in Chan and Eisenstat (2018) and the corresponding results are presented in Table 2. Similar to the estimation of the structure model we will discuss in the following section, the estimation results in this exercise are all based on 10000 posterior samples obtained after a burn-in period of 5000. Following standard practice, we set the lag length to p = 4.

For interpretation purposes, note that the marginal likelihood value will be relatively large for models in which the observed data are more likely, and vice versa. 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.

Table 2: Log marginal likelihoods for the class of MS-VAR models

CVAR	C-VAR-MS	MS-VAR-C	MS-VAR
-883.52	**-879.39**	-883.07	-885.83

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

Two lessons are learned. First, the preferred model is the C-VAR-MS. This means that allowing for a non-linear volatility process is crucial when examining the relationship between China's economy and world oil markets. It also provides some evidence to existing claims that unexpectedly substantial growth from China over the past two decades may have generated a possible regime change in China's effect on the price of crude oil (Kilian and Hicks, 2013; Aastveit et al., 2015). Second, the general class of MS-VAR models do not always outperform the CVAR. In particular, allowing for regime switching in the VAR coefficients does not improve upon the CVAR. This highlights the importance of conducting a formal model comparison procedure when choosing the best model to address our research question. In particular, the use of either a CVAR model or a MS-VAR-C model would lead to biased estimated of the underlying structural process. Specifically, there may be 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 is found in Figure 2, which presents the posterior means and 68 percent credible sets for the respective elements in the reduced form covariance matrix.

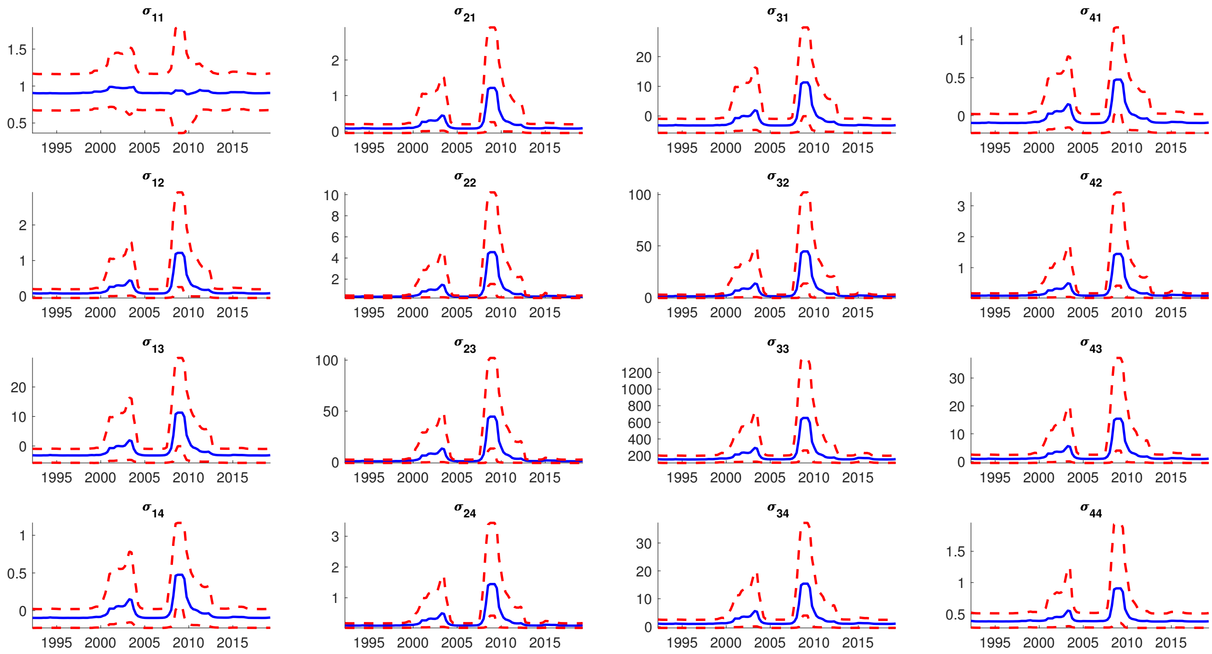

Figure 2: Reduced Form Covariance Matrix

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.

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. However, 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). This suggests that while *own shocks* are likely to be of primary importance during normal economic conditions, *contemporaneous interactions* become extremely important during global economic downturns, when the oil market is likely highly volatile. We return to this point when discussing the main results in Section 4.

To provide a more nuanced picture of regime clustering over the sample period, we follow Song (2014) and Hou (2017) plot the estimates of $P(s_i = s_j | y_{1:T})$ as a heat map in Figure 3. The heatmap can be viewed as a table in which the colors represent different probabilities over time. More precisely, the clustering of regimes is presented through a $T \times T$ matrix in which $i = 1, \ldots, T$ and $j = 1, \ldots, T$; therefore the figure is symmetric against the 45^0 line. For interpretation purposes, the light color on the main diagonal of the figure indicates a regime that occurs in the period $i = j$. Set in this manner, lighter colors off the main diagonal indicate regime recurrence, while darker colors represent regime change. Presented in this manner, we clearly observe periods of regime change, especially around the early 2000s and the Great Recession, which further supports the existence of structural changes in the oil market.

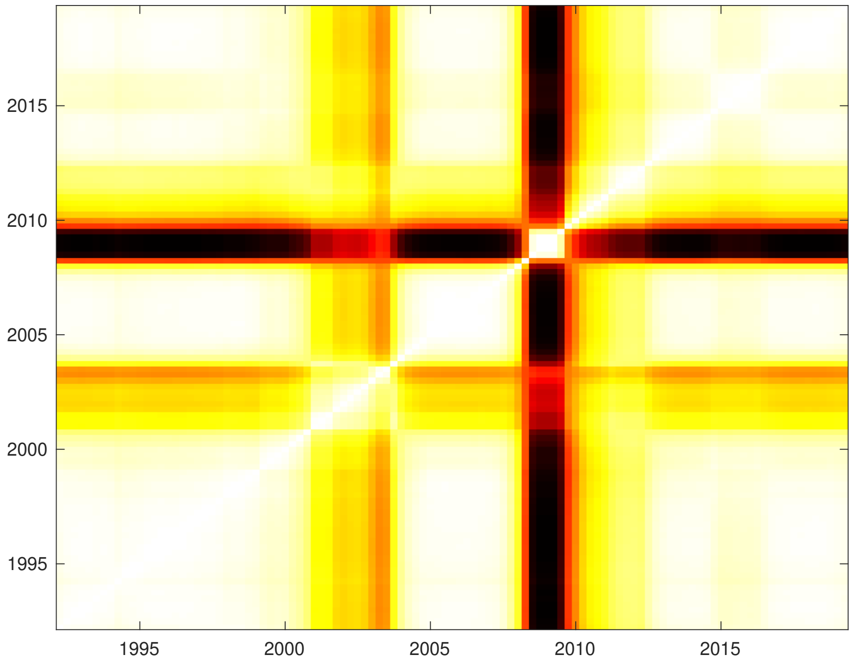

Figure 3: Heat map

Note: The figure plots the heat map for regime clustering based on the estimation of $P(s_i = s_j | y_{1:T})$.

3.3 Identification

To identify the structural VAR model we adopt the same set of agnostic sign restrictions used in Cross and Nguyen (2017) and restrictions on the demand and supply elasticities used in Kilian and Murphy (2012, 2014).[4] The sign restrictions summarized in Table 3 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 directional signs for these restrictions of

[4]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).

the impact matrix are summarized in Table 3 and are implemented with the algorithm in Rubio-Ramirez et al. (2010).

In the first instance, a *supply shock*, denoted by $\epsilon_{Q,t}$, represents an exogenous disruption of global oil production. Such disruptions may be caused by, for example, geopolitical turmoil in the middle east (Hamilton, 1983). 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 tends to generate higher oil prices, and are therefore called *rest of world (ROW) aggregate demand shocks*, $\epsilon_{Y_W,t}$. The third type of shocks originates from specific factor generated demand and are therefore called *oil specific demand shocks* or *residual shocks*, $\epsilon_{P,t}$. 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 economic activity 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 GDP 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 aggregate 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.

Table 3: Sign restrictions

	$\epsilon_{Q,t}$	$\epsilon_{Y_W,t}$	$\epsilon_{P,t}$	$\epsilon_{Y_C,t}$
Oil production	$+$	$+$	$+$	$\times$
World economic activity	$+$	$+$	$-$	$\times$
Real oil price	$-$	$+$	$+$	$\times$
China economic activity	$\times$	$\times$	$\times$	$+$

Note: $+$ and $-$ respectively indicate positive and negative responses, while $\times$ leaves the effect unrestricted. To ensure that the China shock is well identified we also impose a magnitude restriction in which a China specific aggregate demand shock is assumed to yield a greater contemporaneous response in China's economic growth compared to alternative shocks.

The cost of providing such an agnostic identification procedure is that the structural shocks may be not be identified. For instance, if a China specific demand shock elicits 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 contemporaneous responsive than alternative shocks. E.g. to identify a China specific demand shock we impose that the response of China's economic growth is greater (in absolute value) than world economic activity. Similar assumptions have been adopted in a range of papers, e.g. Peersman (2005); Aastveit et al. (2015) and Cross and Nguyen (2017). In Section 5 we provide estimation results with no such restrictions and find that they are robust.

As a sensitivity analysis, following Kilian and Murphy (2012, 2014), we combine these sign restrictions with empirically plausible bounds on the magnitude of the short-run oil demand and supply elasticities. In particular, we assume upper bounds on the monthly elasticities of the demand and supply curves to respectively be -0.8 and 0.0258. In our model, these monthly values are normalized as quarterly elasticities to be in line with our data frequency.

4 Estimation Results

For expository purposes we focus on the oil price response to the structural shocks and evaluate the China factor. The supply shock is normalized as a negative shock implying a reduction in world oil production that would lead to increase the oil price. We then examine the impact of oil market shock on China's GDP. In each case, the shock size is one standard deviation. It is worth mentioning that our MS-VAR model allows for regime recurrence, which is distinguished from other non-linear VAR models, and hence the two regimes are frequently observed over the sample period. To examine the characteristics of these regimes and the corresponding impulse responses, we therefore report median impulse responses at two selected dates. One impulse response represents for periods that the oil market experiences low volatility and the other associates with periods of high volatility.

4.1 Oil Price Responses and the China Factor

In this section we quantify the overall sensitivity of the real price and production to the structural shocks and access the impact of the China factor. We first focus on the response of the oil price displayed in Figure 4. The figure plots the cumulative impulse response of the oil price and oil quantity to different types of oil market shocks in the two identified regimes. Consistent with the empirical evidence found in the oil literature, the oil price increases immediately after an unexpected oil supply disruption. The results further show clear evidence that the price reacts differently across the two states. We observe the oil price becomes more sensitive to shocks to oil supply when the oil market in the state of low volatility. Quantitatively, the magnitude of the price response in the low volatility regime is consistently larger than that of the corresponding response in the high volatility regime. A similar impact pattern is also found when examining the reaction of oil price to the residual shock although the degree of market volatility does not matter much to this type of shock.

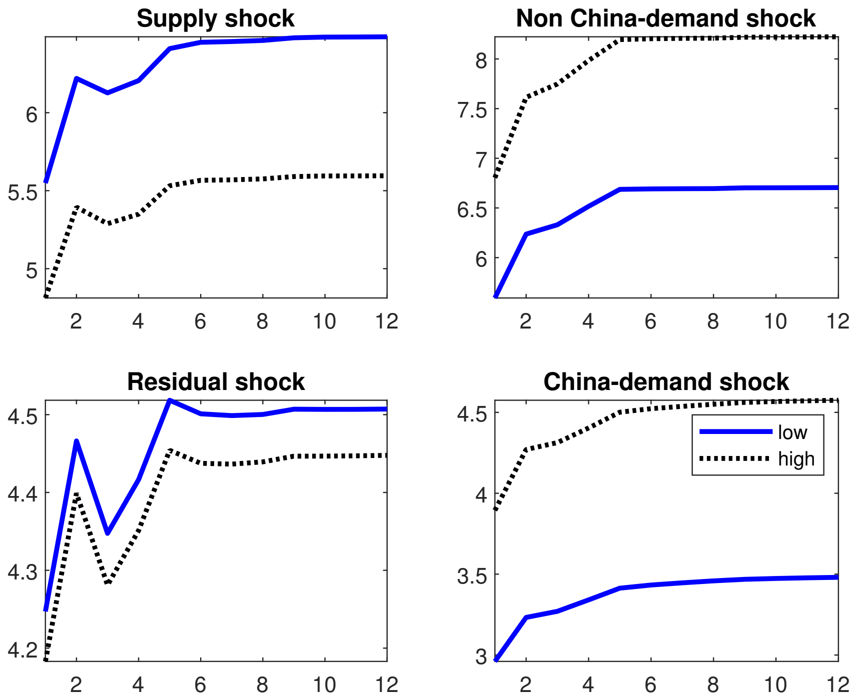

Figure 4: Responses of the oil price to the structural shocks

Turing to responses of the oil price to shocks in oil demand from China and the rest of the world, or non-China demand. As expected, the two demand shocks exert increasing pressure on the oil price. More importantly, the results strongly indicate that the responses depend on the level of oil volatility. We find that under a high-volatility environment, the impact of the two demand shocks is relatively larger than that of the same shocks hitting in times of low volatility. With respect to our research question, the responses show that the China factor has a significant positive impact on the oil price, however this impact is relatively smaller than the impact of a demand shock from the rest of the world. The influence of oil demand shocks from China is also found to be more pronounced during periods when the oil market experiences high volatility. Our finding is new to the literature and therefore contrasts those in Mu and Ye (2011); Wu and Zhang (2014) and Cross and Nguyen (2017) who each find little evidence to support the hypothesis that China's demand for oil has impacted the world price. Part of the explanation could be that, different from non-switching regime models, our model are able to gauge the state-dependent impulse responses, which explicitly distinguishes high and low volatility regime.

4.2 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 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. In what follows, we discuss two broad conclusions.

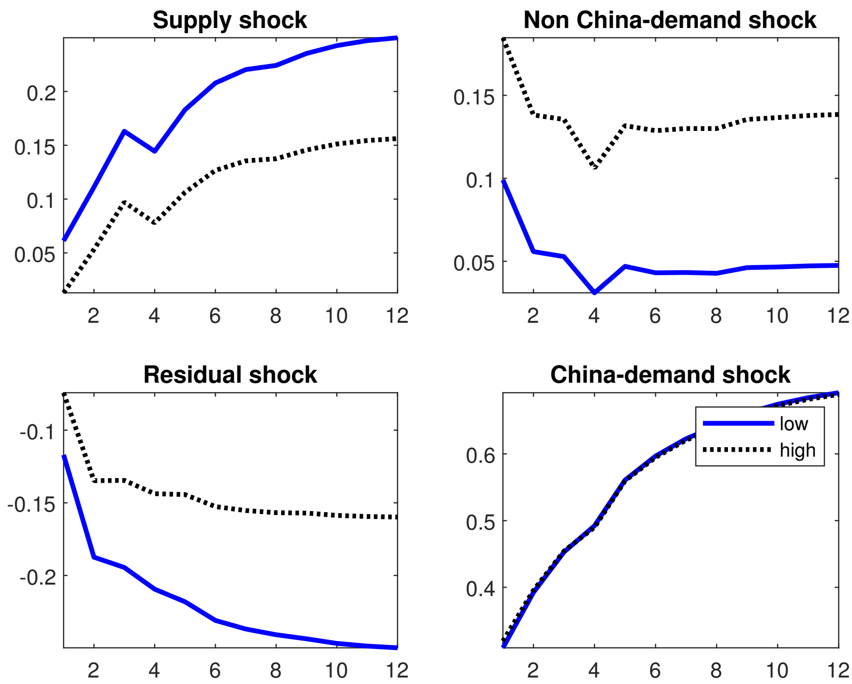

Figure 5: China's GDP response to the oil market shocks

First, we observe that negative oil supply and non-China demand shocks have positive effects, while positive oil-specific demand shocks have negative effects. More importantly, the magnitude of these responses are different across the two identified regimes and China's GDP is likely less sensitive to the oil market shocks during the times of high volatility. For example, after controlling demand and supply shocks, an unexpected increase in the oil price (a residual shock) lead to China's GDP to fall about 0.15 percent after a year if the shock hits the economy in high volatility regime. In normal times, the same shock elicits a relative larger effect, which is about 0.3 percent. The same impact pattern is also found for the supply and non-China demand shock.

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. We also find that such shocks elicit an almost zero response on China's real GDP growth during the 2000 recession when the oil market experienced greatly volatility. 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, although the responses of China's GDP to the oil market shocks are found to be regime dependent, the magnitude response these 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 Robustness Checks

In this section we discuss the results from three robustness checks that were performed in addition to our main analysis. To conserved space, we defer the figures and table to Appendix A.

First, 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 A1—A2. 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.

Second, we test for robustness of the results when using either the WTI or Brent oil price. The results are provided in Figures A3—A6. 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.

Third, to examine the appropriateness of our MS-VAR methodology over the TVP-VAR-SV methodology employed in Cross and Nguyen (2017), we compute the log-marginal likelihood for their class of models using the importance sampling based algorithm proposed in Chan and Eisenstat (2018).[5] We find that the class of MS-VARs generally outperforms this class of TVP-VAR-SV models. This result suggests that using MS-VARs is critical to properly understanding the dynamics of the relationship between China and the world market for crude oil.

Table 4: Log marginal likelihoods for the class of TVP-VAR-SV models

C-VAR-SV	TVP-VAR-C	TVP-VAR-SV
-1037,9	-2069.50	-1148.90

Finally, we investigate whether our assumption of two regimes is a reasonable one. To this end, we compute the marginal likelihood for various C-VAR-MS models which differ in the number of possible regimes. The results contained in Table 5 show that the two regime C-VAR-MS is the best model.

Table 5: Log marginal likelihoods for C-VAR-MS with various regimes

No. of regimes	2	3	4	5
Log-ML	**-879.39**	-882.25	-881.38	-883.12

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

6 Conclusion

We have examined the relationship between China's macroeconomic growth and the world market for crude oil since the mid-1990s. In light of a potential non-linearity in this relationship, we began our analysis by showing that a class of flexible Markov-switching VARs is more appropriate than a constant VAR framework. We found that a Markov-switching VAR with a regime dependent covariance matrix (C-VAR-MS) was preferred to all other models. In the robustness section, we also found that this model outperformed

[5]The model set-up and priors used to estimate these these models is identical to those in Cross and Nguyen (2017).

a popular class of time-varying latent-parameter VARs that have been previously used in the literature.

The Markov-switching model identified key regime changes in oil market volatility, which a traditional constant parameter VAR model would be unable to detect. We found that the price of oil is more responsive to demand shocks from both China and the rest of the world in the high volatility state, while less responsive to supply shocks. In this sense, our results provide empirical support to the conjectured claims that China has been influential in driving oil price dynamics over the past two decades.

A similar state-dependent phenomenon was also observed for the impact of oil price shocks on China economic activity. In that case we observed that China's real GDP growth responds more to economic activity shocks from the rest of the world in the high volatility state, but is less responsive to supply shocks. That being said, the size of the responses was relatively small. Thus, despite China being a major player in international oil markets, we conclude that oil market shocks tend to have little impact on China's real GDP growth.

References

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

Basher, S. A., Haug, A. A., and Sadorsky, P. (2016). The impact of oil shocks on exchange rates: a markov-switching approach. *Energy Economics*, 54:11–23.

Baumeister, C. and Hamilton, J. D. (2019). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. *American Economic Review*, 109(5):1873–1910.

Baumeister, C., Korobilis, D., and Lee, T. K. (2020). Energy markets and global economic conditions. *Review of Economics and Statistics*, forthcoming.

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

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

Carriero, A., Clark, T. E., and Marcellino, M. (2019). Large Bayesian vector autoregressions with stochastic volatility and non-conjugate priors. *Journal of Econometrics*, 212(1):137–154.

Chan, J. C. and Eisenstat, E. (2018). Bayesian model comparison for time-varying parameter vars with stochastic volatility. *Journal of applied econometrics*, 33(4):509–532.

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

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

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

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

Doan, T., Litterman, R., and Sims, C. (1984). Forecasting and conditional projection using realistic prior distributions. *Econometric reviews*, 3(1):1–100.

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

Fong, W. M. and See, K. H. (2002). A markov switching model of the conditional volatility of crude oil futures prices. *Energy Economics*, 24(1):71–95.

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

Hamilton, J. D. (1983). Oil and the macroeconomy since World War II. *The Journal of Political Economy*, pages 228–248.

Hamilton, J. D. (2009). Causes and consequences of the oil shock of 2007-08. Technical report, National Bureau of Economic Research.

Hamilton, J. D. (2019). Measuring global economic activity. *Journal of Applied Econometrics*.

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

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

Hou, C. (2017). Infinite hidden markov switching VARs with application to macroeconomic forecast. *International Journal of Forecasting*, 33(4):1025–1043.

Hou, C. and Nguyen, B. H. (2018). Understanding the US natural gas market: A Markov switching VAR approach. *Energy Economics*, 75:42–53.

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

Karlsson, S. (2013). Forecasting with Bayesian vector autoregression. In *Handbook of economic forecasting*, volume 2, pages 791–897. Elsevier.

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. (2013). Structural vector autoregressions. In Hashimzade, N. and Thornton, M. A., editors, *Handbook of Research Methods and Applications in Empirical Macroeconomics*, chapter 22, pages 515–554. Oxford University Press, Oxford.

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

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

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

Kilian, L. and Zhou, X. (2020). The econometrics of oil market var models. *CESifo Working Paper*.

Koop, G. and Korobilis, D. (2010). *Bayesian multivariate time series methods for empirical macroeconomics*. Now Publishers Inc.

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

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

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

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

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

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

Song, Y. (2014). Modelling regime switching and structural breaks with an infinite hidden Markov model. *Journal of Applied Econometrics*, 29(5):825–842.

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

Vo, M. T. (2009). Regime-switching stochastic volatility: Evidence from the crude oil market. *Energy Economics*, 31(5):779–788.

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

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

Appendix A Robustness Results

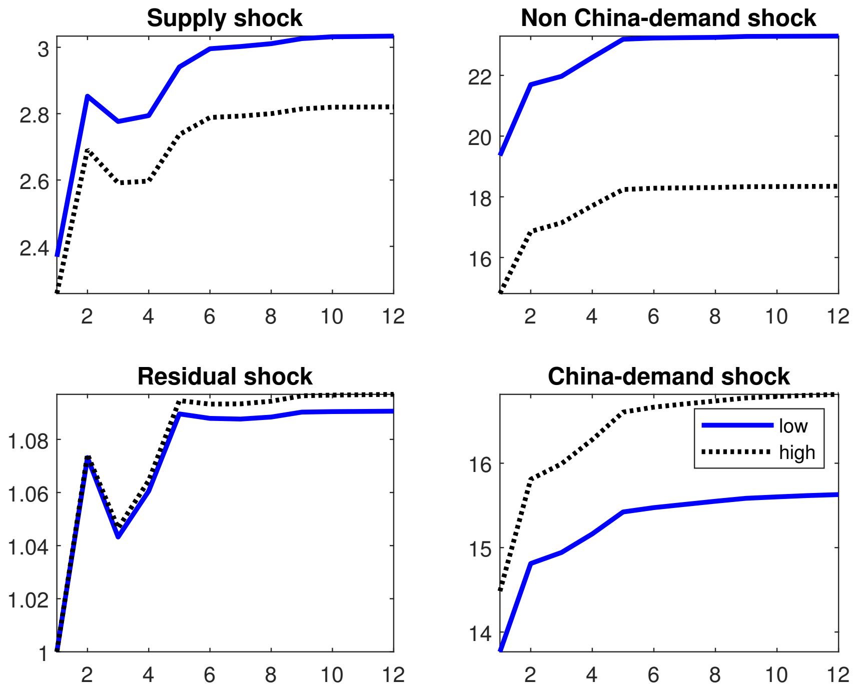

Figure A1: Robustness: Oil price response to supply and demand shocks with elasticity constraints.

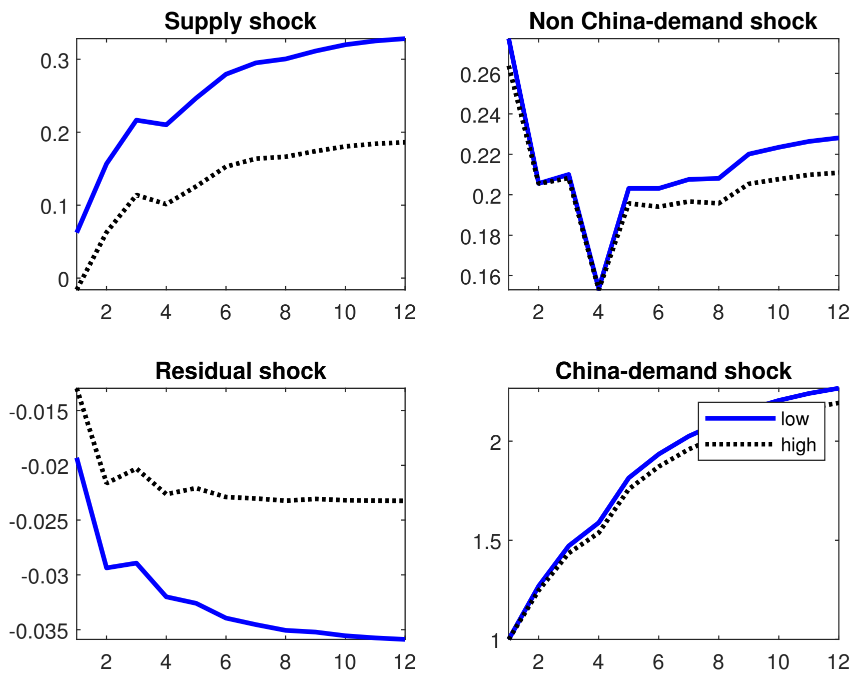

Figure A2: Robustness: China's GDP response to supply and demand shocks with elasticity constraints.

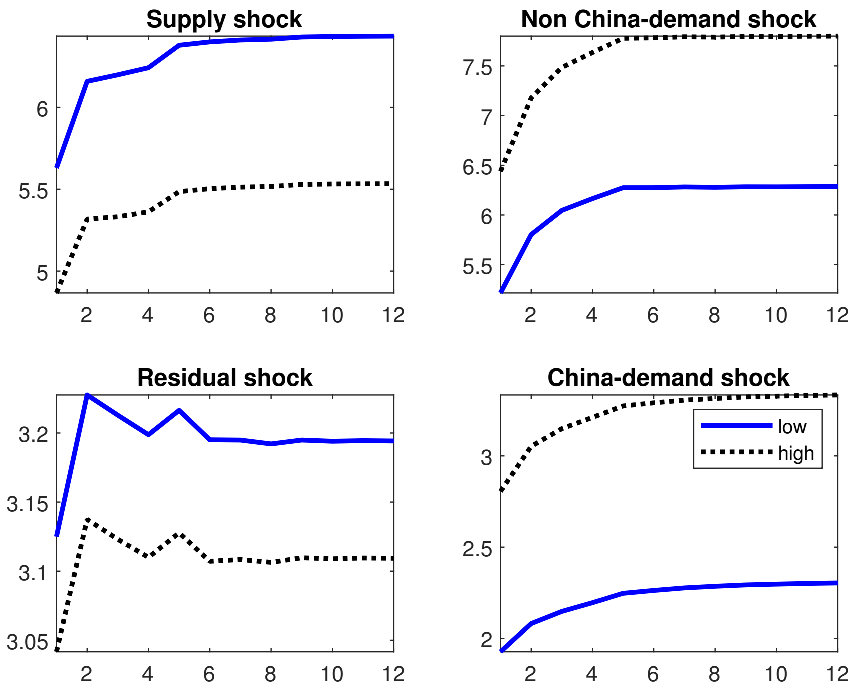

Figure A3: Robustness: Oil price response to supply and demand shocks with WTI oil price data.

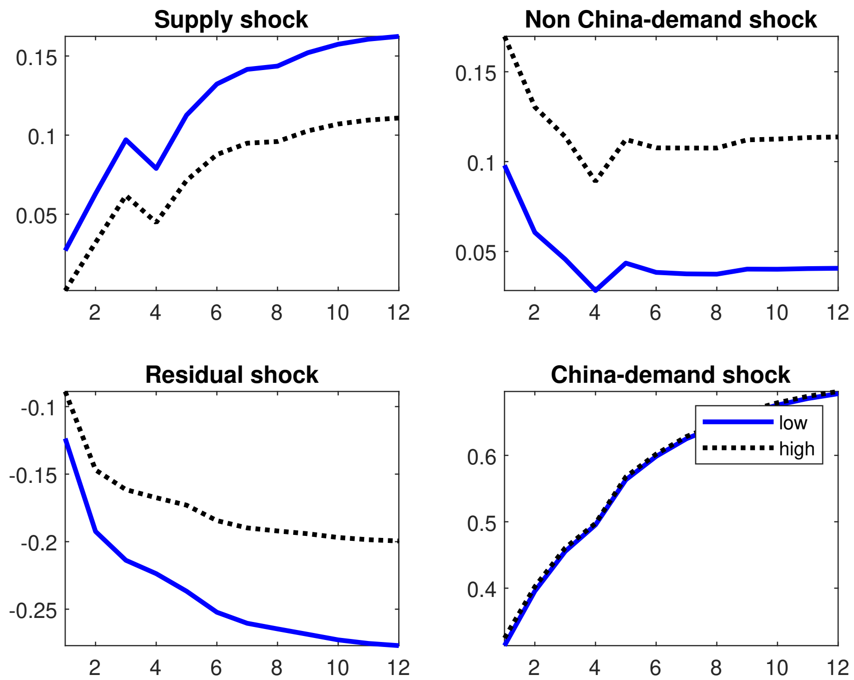

Figure A4: Robustness: China's GDP response to supply and demand shocks with WTI oil price data.

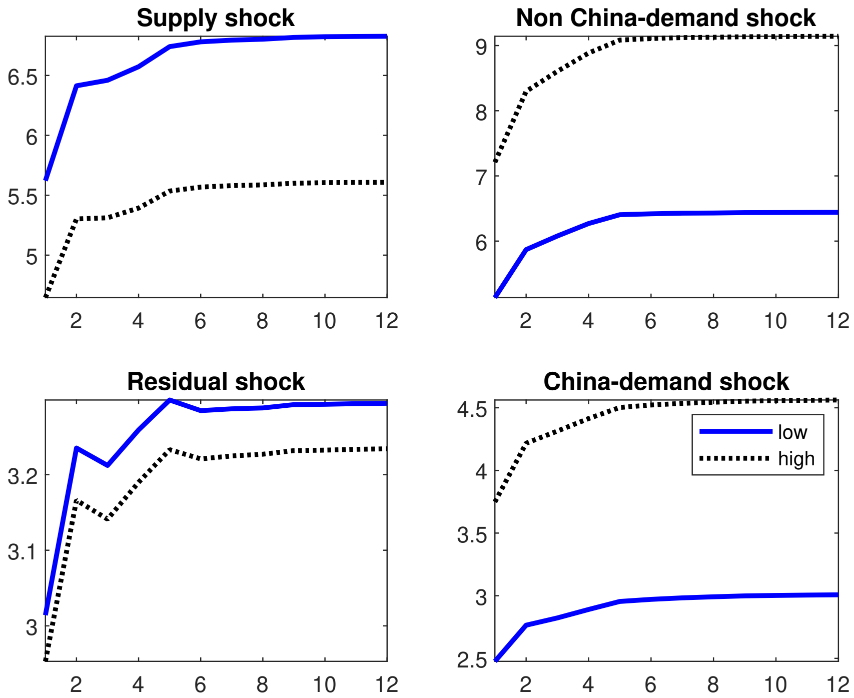

Figure A5: Robustness: Oil price response to supply and demand shocks with Brent oil price data.

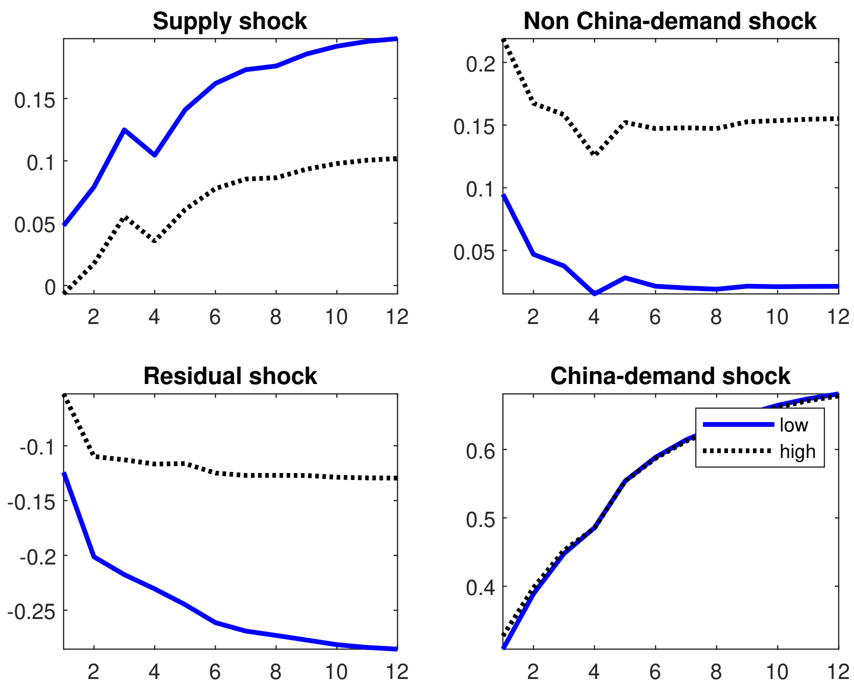

Figure A6: Robustness: China's GDP response to supply and demand shocks with Brent oil price data.

Appendix B Bayesian Estimation

We estimate the MS-VAR model using a Gibbs Sampling algorithm that successively samples from the following full conditional distributions:

1. $p(\mathbf{s}|\mathbf{\Theta}, \mathbf{y})$,

2. $p(\mathbf{\Theta}|\mathbf{s}, \mathbf{y})$,

3. $p(\mathbf{P}|\mathbf{s})$,

where $\mathbf{s} = (s_1, \ldots, s_T)'$ is a vector of regime indicators, $\mathbf{\Theta} = \{(\boldsymbol{\beta}_i, \mathbf{\Sigma}_i)\}_{i=1}^{M}$ denotes the collection of model parameters across the M regimes, and $\mathbf{P}$ be the $M \times M$ Markov transition matrix, i.e., $\mathbf{P}_{ij} = p_{ij}$. In our empirical analysis we obtain $15,000$ posterior draws, discarding the first $5,000$ as a burn-in. 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 $p(s_{t-1}|\mathbf{y}_{1:t-1}, \boldsymbol{\theta})$ we compute $p(s_t|\mathbf{y}_{1:t}, \boldsymbol{\theta})$ by

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

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 $\widehat{\mathbf{K}}_i = \mathbf{X}_i{}' \left(\mathbf{I}_{T_i} \otimes \boldsymbol{\Sigma}_i\right)^{-1} \mathbf{X}_i + \mathbf{V}_0^{-1}$, $\widehat{\boldsymbol{\beta}}_i = \widehat{\mathbf{K}}_i^{-1} \left(\mathbf{X}_i{}' \left(\mathbf{I}_{T_i} \otimes \boldsymbol{\Sigma}_i\right)^{-1} \mathbf{y}_i + \mathbf{V}_0^{-1} \boldsymbol{\beta}_0\right)$, $\widehat{\nu}_i = T_i + \nu_0$ 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 C Identification by sign restrictions

The identification procedure discussed in Section 3.3 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: $\boldsymbol{\Sigma}_{s_t}$, so that $\boldsymbol{\Sigma}_{s_t} = \mathbf{P}_{s_t} \mathbf{D}_{s_t} \mathbf{P}'_{s_t}$ 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 Section **??**. If yes, keep it. Otherwise, discard it and redraw $\mathbf{K}$.