Crawford School of Public Policy

CAMA

Centre for Applied Macroeconomic Analysis

How do oil shocks transmit through the U.S. economy? Evidence from a large BVAR model with stochastic volatility

CAMA Working Paper 13/2021
January 2021

Renée Fry-McKibbin
Centre for Applied Macroeconomic Analysis, ANU

Beili Zhu
Centre for Applied Macroeconomic Analysis, ANU

Abstract

This paper employs a large BVAR model with common stochastic volatility to examine the effects of oil supply shocks, global oil demand shocks and precautionary oil shocks on 17 U.S. macroeconomic and financial market variables from 1986Q1 to 2019Q2. Generalized impulse response functions calculated using stochastic volatility provide a time-varying account of the impacts of the shocks occurring in each quarter. We also compute standard impulse response functions for shocks of the sizes evident in 2019Q2 and 2008Q4. The magnitudes of the generalized impulse response functions vary over time, but the fluctuations are not particularly different except during the global financial crisis. All oil shocks have permanent inflationary effects; there is evidence of long-run adverse effects on several macroeconomic variables because of global oil demand shocks despite rising GDP, and all oil shocks negatively affect the U.S. stock and currency markets in the long term, but the effects on the bond market differ.

Keywords

Generalised impulse responses, Sign restrictions

JEL Classification

C32, E31, E32

Address for correspondence:

(E) cama.admin@anu.edu.au

ISSN 2206-0332

How do oil shocks transmit through the U.S. economy? Evidence from a large BVAR model with stochastic volatility*

Renée Fry-McKibbin and Beili Zhu

CAMA, The Australian National University

Abstract

This paper employs a large BVAR model with common stochastic volatility to examine the effects of oil supply shocks, global oil demand shocks and precautionary oil shocks on 17 U.S. macroeconomic and financial market variables from 1986Q1 to 2019Q2. Generalized impulse response functions calculated using stochastic volatility provide a time-varying account of the impacts of the shocks occurring in each quarter. We also compute standard impulse response functions for shocks of the sizes evident in 2019Q2 and 2008Q4. The magnitudes of the generalized impulse response functions vary over time, but the fluctuations are not particularly different except during the global financial crisis. All oil shocks have permanent inflationary effects; there is evidence of long-run adverse effects on several macroeconomic variables because of global oil demand shocks despite rising GDP, and all oil shocks negatively affect the U.S. stock and currency markets in the long term, but the effects on the bond market differ.

Keywords: Generalised impulse responses, Sign restrictions
JEL classifications: C32, E31, E32

*We gratefully acknowledge ARC Discovery Project funding DP120103443 and thank Joshua Chan, Warwick McKibbin, Tatsuyoshi Okimoto, Shaun Vahey, Joaquin Vespignani and Benjamin Wong for useful comments. Author email addresses are renee.mckibbin@anu.edu.au and beilizhu.anu@gmail.com. The corresponding author is Renée Fry-McKibbin.

1 Introduction

A frequently used approach to estimate the impacts of oil price shocks on macroeconomic performance is to estimate structural vector autoregression (SVAR) models (Filippidis, Filis and Kizys, 2020; Herrera and Rangaraju, 2019; Kim, Hammoudeh, Hyun and Gupta, 2017; Kilian, 2008; Kilian 2009; Lardic and Mignon, 2008; Riggi and Venditti, 2015). This paper extends this previous literature by providing a quantitative analysis of the impacts of oil sector shocks using a large Bayesian VAR (BVAR) with common stochastic volatility on 17 U.S. macroeconomic and financial market aggregates. The paper explores the impacts of oil supply shocks, global oil demand shocks, and precautionary oil shocks from 1986Q1 to 2019Q2.

The recent literature suggests two critical components for specifying a good VAR model for structural analysis and forecasting of macroeconomic series. The first is a large set of macroeconomic variables. Carriero, Clark and Marcellino (2016), Chan (2020) and Koop (2013) show that VAR models with 15-20 variables perform better for structural analysis and forecasting than models with a small number of variables. Models with a large dataset more likely incorporate the information typically monitored by financial market participants and policymakers, ensuring better identification of financial market and monetary policy responses (Aastveit, 2014). In the energy literature, several papers examine the effects of energy price shocks using large macroeconomic models. For example, Stock and Watson (2016) use a 207-variable structural dynamic factor model to calculate the impacts of oil supply shocks on the U.S. economy, while Kapetanios, Marcellino and Venditti (2019) examine the time-varying effects of oil sector shocks on sectoral U.S. industrial output using a 28-variable non-parametric time-varying parameter VAR model. However, VAR models with large information sets involve many parameters and computationally-intensive estimation processes. Bayesian methods using prior information allowing the shrinkage of VAR coefficients is a popular solution for speeding up the computational process (Bańbura, Giannone and Reichlin, 2010; Carriero, Clark and Marcellino, 2016; Chan, 2020).

The second critical component includes time variation in the error terms to capture the time-varying volatility feature of macroeconomic data (Cogley and Sargent, 2005; Primiceri, 2005). Baumeister and Peersman (2013b) show that including time-varying volatility in the VAR specification is essential given observed fluctuations in oil price volatility and growth in macroeconomic volatility in recent years. Given the importance of including time-varying volatility in VAR models for oil shock analysis, some researchers have started to include this feature. Baumeister and Peersman (2013b) use a 4-variable VAR model, including time-varying parameters and stochastic volatility, to analyze the dynamic impacts of an oil supply shock on the U.S. macroeconomic performance. Riggi and Venditti (2015) identify foreign productivity and oil supply shocks in a small VAR model including drifting parameters and stochastic volatility to investigate the nexus between the real oil price and real exports in the eurozone.

Despite the demonstrated improvements for macroeconomic analysis by using large datasets and time variation in volatility in the specification of VAR models, few papers combine both of these features (see Carriero, Clark and Macellino, 2016; Carriero Clark and Macellino, 2019; and Chan, 2020 for exceptions). The difficulty is in incorporating stochastic volatility in VAR error terms if the model is of high dimension. The model loses symmetry due to the addition of time-varying volatility, which quickly makes the estimation process unmanageable as the number of variables increases. As a solution, Carriero, Clark and Macellino (2016) introduce a Kronecker structure of the likelihood to expedite the sampling process. Chan (2020) further improves the algorithm by using a fast band matrix routine and vectoring the operations. There is no analysis of the effects of oil price shocks on the macroeconomy in the oil shock literature combining both features. This paper builds on Carriero, Clark and Macellino (2016) and Chan (2020) by using the Kronecker structure of the likelihood and fast band matrix and vectoring the operations in the estimation algorithm to estimate a large BVAR model integrating common stochastic volatility to examine the impacts of oil sector shocks on the U.S. economy. The empirical results illustrate that the model with common stochastic volatility is statistically superior to the model excluding stochastic volatility using the marginal likelihood as the Bayesian model comparison procedure (Chan, 2020).

Few recent papers study oil price shocks using models that integrate time-varying parameters and large datasets. For example, Kapetanios, Marcellino and Venditti (2019) study the dynamic impacts of oil sector shocks on U.S. economic performance, applying a large time-varying coefficients VAR model with 28 variables. They find that using the Kalman filter and forgetting factors to estimate parameters dramatically improves the computational efficacy. Koop and Korobilis (2013) show that using forgetting factors means that the MCMC algorithm is not required to estimate the parameters, thereby reducing the computational burden and increasing estimation speed. Although time-varying coefficient VAR models have been applied in the oil studies to analyse oil price shocks, there is a long-standing debate about whether to include time-varying parameters in oil market models. The advocates state that the inclusion of time variation in the parameters of the oil market models can smooth structural change (Baumeister and Peerman, 2013a,b; Kapetanios, Marcellino and Venditti, 2019).

On the other hand, Herrera and Rangaraju (2019) show that nonlinear impulse responses in the time-varying coefficient models have wide error bands. Kilian and Zhou (2020) point out that the estimated coefficients in a time-varying parameters VAR model can change over time even if there is no time variation in the real world because the time variation in estimated coefficients derives from overfitting instead of time variation. Because of the drawbacks of the inclusion of drift parameters in VAR models and the computational difficulties of combining time-varying parameters with time-varying volatility in large VAR models, this paper focuses on the stochastic volatility measure of time variation to examines the effects of oil sector shocks on the U.S. economy. The concurrent integration of time variation in parameters and volatility in large VAR models is left for future work.

The literature proposes several identification schemes to identify oil shocks. For example, Kilian (2009) utilizes a recursive identification method and assumes that the short-run crude oil supply curve is vertical and that a precautionary oil shock does not have any contemporaneous impact on real economic activity. Peersman and van Robays (2009) attain identification via sign restrictions. An extension of the sign restrictions approach is to impose boundary restrictions on the sizes of the price elasticities of oil demand and oil supply in the short-run to reduce the structural models compatible with the sign restrictions (Baumeister and Peersman, 2013a; Kilian and Murphy, 2012). This study adopts the Baumeister and Peersman (2013a) identification scheme to distinguish the oil shocks.

The impulse response functions reported are the Fry and Pagan (2011) median target impulses along with the point-wise medians and associated percentiles of the posterior distribution of the impulses often reported in the literature. Fry and Pagan (2011) show that the latter approach misleads as the median responses do not have a structural interpretation as the reported medians come from different models. To solve this problem, Fry and Pagan develop the median target method, which selects a model with the impulse response functions most similar to the median values of the distribution of the impulses, rather than the median impulse response functions which come from different models. Robustness to other methods of impulse selection such as using the largest response of the real oil price to an oil supply shock or using the model with an oil demand elasticity most similar to the posterior medians of the elasticity among admissible models is also conducted (Kilian and Murphy, 2012; 2014).

An advantage of modeling common stochastic volatility is that impulse response functions can be computed explicitly for the shocks occurring in each quarter. This paper computes impulse response functions to shocks in two representative quarters (2019Q2 and 2008Q4). The generalized impulse response functions of Koop, Pesaran and Potter (1996) are also presented over the period 1986Q1 to 2019Q2. The generalized impulse response functions plot the reactions of the oil sector variables to the shocks contemporaneously and the U.S. macroeconomic and financial market variables one year after the shock at each point in time to show the time-varying impacts of oil price shocks on the U.S. macroeconomy and financial markets. Time-varying VAR models of oil sector shocks often use the generalized impulse response functions to analyze their macroeconomic impacts (Baumeister and Peersman, 2013a,b; Jebabli, Arouri and Teulon, 2014).

The framework yields several important findings. The generalized impulse response functions show that the magnitudes of the reactions of U.S. variables to oil shocks change over time, culminating during the global financial crisis. The impulse responses functions to a representative shock taken to be the last data point in 2019Q2 show that the effects of the oil sector shocks on the macroeconomic and financial market aggregates are strikingly dissimilar. However, all shocks generate permanent inflationary effects. There is evidence of some long-run adverse effects on several macroeconomic variables because of global oil demand shocks despite GDP rising permanently. All of the oil shocks negatively affect the U.S. stock market and currency market in the long run, while the effects on the U.S. bond market vary. Our paper calculates the impulse response functions for the shock in 2008Q4 for comparison. The qualitative results hold—however, the shock magnitudes and hence impulse magnitudes differ.

The paper proceeds as follows: Section 2 outlines the econometric methodology, comprising model specification, model comparison methods and the strategy for identification. Section 3 provides a descriptions of the data set. Section 4 shows the empirical results and sensitivity analysis. The conclusion is provided in Section 5.

2 Empirical methodology

This Section sets out the large Bayesian VAR model with common stochastic volatility (BVAR-CSV) used to analyze the structural oil sector shocks on the U.S. macroeconomic and financial market variables. Section 2.1 sets out the econometric model. Section 2.2 provides a model comparison criterion to compare the BVAR-CSV to its counterpart with constant volatility. Section 2.3 outlines the strategy to identify the structural shocks and Section 2.4 describes the generalized impulse response functions.

2.1 The BVAR-CSV model

We follow the approach set out in Carriero, Clark, and Marcellino (2016) and Chan (2020) to estimate a large BVAR-CSV by allowing a Kronecker structure on the posterior variance of the VAR coefficients. This structure accelerates the sampling process, given a large number of coefficients in the BVAR (Carriero, Clark, and Marcellino, 2016). Chan (2020) extends this work using fast band matrix and vectorization operations to reduce the computational intensity of the method further. This Section uses the Chan (2020) algorithm to estimate the BVAR-CSV model.

The model has the following generic representation:

$$\mathbf{y}_t = \mathbf{c_0} + \mathbf{c_1}t + \mathbf{B_1}\mathbf{y}_{t-\mathbf{1}} + ... + \mathbf{B_p}\mathbf{y}_{t-\mathbf{p}} + \mathbf{u}_t, \quad \mathbf{u}_t \sim \mathcal{N}(\mathbf{0}, e^{h_t}\mathbf{\Sigma}) \tag{1}$$

$$h_t = \rho h_{t-1} + \xi_t^h \qquad \xi_t^h \sim \mathcal{N}(0, \delta_h^2), \tag{2}$$

with an $n \times 1$ vector of variables $\mathbf{y}_t$ over the period $t = 1, 2, \cdots, T$, $\mathbf{c_0}$ is an $n \times 1$ constant vector, $\mathbf{c_1}$ is an $n \times 1$ vector, and $\mathbf{B_1}, \cdots, \mathbf{B_p}$ are $n \times n$ coefficient matrices. The heteroscedastic errors $\mathbf{u}_1 ... \mathbf{u}_T$ are distributed as $\mathcal{N}(\mathbf{0}, \mathbf{e}^{h_t \mathbf{\Sigma}})$ and the log volatility h_t is a stationary AR(1) process with $|\rho| < 1$.

To reduce computational intensity, $\mathbf{x}_t' = (1, t, \mathbf{y}_{t-1}', \cdots, \mathbf{y}_{t-p}')$ is set to contain a $1 \times k$ vector of a constant, a trend and lags with $k = 2 + np$. Equation (1) can be stacked over $t = 1, ..., T$ as

$$\mathbf{Y} = \mathbf{X}\mathbf{B} + \mathbf{\Theta}, \tag{3}$$

where $\mathbf{Y} = (\mathbf{y}_1, \mathbf{y}_2, \cdots, \mathbf{y}_T)'$ is a matrix of dimension $T \times n$, $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \cdots, \mathbf{x}_T)'$ is a matrix of dimension $T \times k$, and $\mathbf{B} = (\mathbf{c_0}, \mathbf{c_1}, \mathbf{B_1}, \cdots, \mathbf{B_p})'$ is a matrix of dimension $k \times n$. The error term Θ is a $T \times n$ matrix that can be succinctly written as $\text{vec}(\mathbf{\Theta}) \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma} \otimes \mathbf{\Psi})$, where variance-covariance matrix $\mathbf{\Sigma}$ is of dimension $n \times n$, $\otimes$ is the Kronecker product and vec$(\cdot)$ transforms the $T \times n$ matrix into a $Tn \times 1$ column vector by stacking the columns. $\mathbf{\Psi}$ is a diagonal covariance matrix of $T \times T$ dimension

$$\mathbf{\Psi} = \begin{pmatrix} e^{h_1} & 0 & 0 & \cdots & 0 \\ 0 & e^{h_2} & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & & \\ 0 & 0 & 0 & \cdots & e^{h_T} \end{pmatrix}.$$

The cross-sectional covariance and the serial covariance of $\mathbf{Y}$ are separately modelled respectively in $\mathbf{\Sigma}$ and $\mathbf{\Psi}$. Specifically, stochastic volatility is modelled by the matrix $\mathbf{\Psi} = \text{diag}\ (e_1^h, \ldots, e_T^h)$. The Appendix contains descriptions of the prior specifications and estimation details.

2.2 Model comparison criterion

This Section sets out the model comparison criterion to test whether allowing for common stochastic volatility in a large BVAR model fits the data better than its counterpart with constant volatility (BVAR-CV). The standard BVAR-CV is nested within the BVAR-CSV model. Specifically, in the BVAR-CV model with independent and identical (iid) Gaussian innovations vec$(\Theta) \sim \mathcal{N}(0, \Sigma \otimes I_T)$, where Σ is a $n \times n$ covariance matrix and I_T is the identity matrix of dimension T. The serial covariance matrix Ψ is an identity matrix with $\Psi = I_T$ in the large BVAR-CV model, while the serial covariance matrix Ψ is a diagonal matrix with $\Psi = diag(e_1^h, \ldots, e_T^h)$ in the BVAR-CSV model. To simplify comparison, the parameters common across the models have the same priors.

The criterion for choosing the best model given the data is the Bayes factor defined as

$$BF_{ab} = \frac{p(y|\ \mathbb{M}_a)}{p(y|\ \mathbb{M}_b)}, \tag{4}$$

where

$$p(y|\mathbb{M}_i) = \int p(y|\Theta_i, \mathbb{M}_i) P(\Theta_i|\mathbb{M}_i) d\Theta, \tag{5}$$

which is the marginal likelihood (ML) under model $\mathbb{M}_i, i = a, b$. Each model $\mathbb{M}_i$ consists of two parts: a likelihood function $p(y|\Theta_i, \mathbb{M}_i)$ and a prior density

$p(\Theta_i|\mathbb{M}_i)$. Θ_i is the model-specific parameter vector. The marginal likelihood measures the probability that a specific model captures the true data generating process. If the true data are likely under model $\mathbb{M}_a$, the associated ML of model $\mathbb{M}_a$ would be large and $BF_{ab} > 1$ meaning that the true data are more likely under model $\mathbb{M}_a$ than $\mathbb{M}_b$.

The marginal likelihood of the BVAR-CV model with natural conjugate priors is easily computed as it is available analytically (Karlsson, 2013). The computation of the marginal likelihood of the BVAR-CSV model is not trivial, given the high dimensionality of the time-varying model and needs to be estimated using Chibb's method. There are two ways to compute the marginal likelihood of latent variable models with stochastic volatility: one uses the conditional likelihood (Koop, Leon-Gonzalez and Strachan, 2009), and the other uses the integrated likelihood (Chan, 2020; Chan and Eisenstat, 2015; Chan and Grant, 2016). Chan and Grant (2016) show that the marginal likelihood predicated on the integrated likelihood is more accurate than those using the conditional likelihood. For discussion of the integrated likelihood computation, see Chan and Eisenstat (2015) and Chan and Grant (2016).

2.3 Identification of the oil sector shocks

The model comparison criterion discussed in the previous section can be used to compare reduced form models. However, the variance-covariance matrix $\mathbf{u_t} = \mathbf{e^{h_t}\Sigma}$ in equation (1) is non-diagonal and therefore the disturbances are correlated with each other. To give the impulse responses a structural interpretation, the orthogonalized structural shocks $\boldsymbol{\zeta_t}$ need to be recovered from the reduced form model in equation (1). Specifically, the structural shocks $\boldsymbol{\zeta_t}$ have the following representation:

$$\mathbf{u_t} = D_t\zeta_t, \tag{6}$$

where ζ_t is a $n \times 1$ vector of disturbances, whose covariance matrix K is a diagonal matrix. This means each shock in ζ_t is independent. D_t is a $n \times n$ matrix representing the contemporaneous interactions between variables with ones on the main diagonal.

The parameters in D_t need to be identified to have a structural interpretation. The identification method of sign restrictions, popular in the oil literature, restricts the choice of the candidate impulse response functions of the variables in the model to the oil sector shocks to those that satisfy the signs expected by economic theory (Baumeister and Hamilton, 2019; Kilian and Murphy, 2012; Peersman, 2005). This study also utilizes sign restrictions to disentangle oil price shocks. Sign restrictions overcome some of the limitations of alternative identification methods.

Not all oil sector shocks are identical. Various demand and supply factors can drive oil price fluctuations (Baumeister, Peersman and van Robays, 2009; Cross and Nguyen, 2017; Fang and You, 2014; Kilian, 2009; Peersman and van Robays, 2012). This paper follows the identification method applying sign restrictions proposed by Baumeister and Peersman (2013a).[1] Sign restrictions are applied to identify the three oil sector shocks. These are an oil supply shock, a global oil demand shock, and a precautionary oil shock.

The oil supply shock reflects that the shift in oil supply causes a shift in the oil price. Sources of such shocks may include oil quota cuts from oil-exporting countries, geopolitical tension and civil wars. A negative oil supply shock causes an increase in the oil price (Q^{oil}) but a reduction in global oil production (P^{oil}) and world economic activity (Y^w). The global demand shock reflects that an economic boom, such as that evident recently in China and India, with high commodity demand tends to stimulate commodity prices, leading to higher demand for oil and therefore the oil price. The positive global demand shock comes with the growth in global oil production (P^{oil}), the price of oil (Q^{oil}), and world economic activity (Y^w). The precautionary oil shock is similar to the global demand shock as it also causes a positive co-movement between the oil price and output. However, the unfavorable precautionary shock leads to a surge in speculative or precautionary oil demand because of a rise in uncertainty about future oil production and world economic activity. The positive precautionary shock reduces world economic activity. Table 1 contains a description of the sign restrictions. The $-$, and $+$ denote sign of the response relative to the price of oil. The imposed sign restrictions hold for one year after the shocks, common in literature on oil price dynamics (Baumeister and Peersman, 2013b; Hamilton 2003).

The sign restrictions are implemented on the global variables of Q^{oil}, P^{oil}, Y^w.

[1] Kilian (2009) proposes an alternative identification technique to identify three different oil sector shocks. These are an oil supply shock, a global demand shock, and a precautionary oil shock. Kilian assumes that the short-run oil supply curve is vertical so that an oil demand shock cannot instantaneously impact world oil production but responds with a month delay. However, this assumption is for models using monthly frequency data and is less appropriate, given the quarterly data of the application in this paper.

Table 1: Sign restrictions. The $-$ and $+$ denote a negative or positive response to the oil sector shocks relative to the oil price, respectively.

Oil sector shocks	Oil production (Q^{oil})	Oil price (P^{oil})	World economic activity (Y^{w})
Supply	$-$	$+$	$-$
Global demand	$+$	$+$	$+$
Precautionary	$+$	$+$	$-$

Hence, the variables in equation (1) can be divided into two parts. The first set of variables $\mathbf{x}_t$ includes the three global variables: global oil production (Q^{oil}), the real oil price (P^{oil}), and world economic activity (Y^w). The second set of variables $\mathbf{z}_t$ contains the U.S. macroeconomic and financial market variables. This means that equation (1) can be written as

$$\begin{pmatrix} \mathbf{x}_t \\ \mathbf{z}_t \end{pmatrix} = \mathbf{c_0} + \mathbf{c_1} + \mathbf{B_1} \begin{pmatrix} \mathbf{x}_{t-\mathbf{1}} \\ \mathbf{z}_{t-\mathbf{1}} \end{pmatrix} + \cdots + \mathbf{B_p} \begin{pmatrix} \mathbf{x}_{t-\mathbf{p}} \\ \mathbf{z}_{t-\mathbf{p}} \end{pmatrix} + D_t \begin{pmatrix} \zeta_t^x \\ \zeta_t^z \end{pmatrix}. \tag{7}$$

The independent disturbances (ζ_t^x) include the oil supply shocks, the global oil demand shocks and the precautionary oil shocks, and ζ_t^z contain shocks for the U.S. macroeconomy and financial markets.

To estimate the contemporaneous impacts D_t and to identify the structural shocks ζ_t^x and ζ_t^z, the imposition of some restrictions on D_t is necessary. Since the focus of our study is to analyze the effects of ζ_t^x, we assume that the three global variables Q^{oil}, P^{oil} and Y^w contemporaneously affect the oil sector variables, which means that the contemporaneous matrix D_t needs to be restricted to a block lower triangular matrix for the U.S. macroeconomic and financial market variables.

The identification proceeds by randomly drawing from the possible orthogonal shocks of the model until a set of 1.5 million admissible impulse response functions that fulfill the sign restrictions are obtained. Usually, some version of the median of the impulses that satisfy the sign restrictions is reported as the estimated impulse response function. However, the sign restrictions method suffers from the model identification problem, where many models can satisfy the specified sign restrictions (Fry and Pagan, 2011).

Several methods have been suggested in the literature to select the impulse responses to report from the credible set in empirical analyses. One way is to report point-wise medians and associated percentiles of the posterior distribution of the impulse responses. However, Fry and Pagan (2011) show that this approach is misleading as the median responses do not have a structural interpretation since the medians reported come from different models. They propose the median target method (MT) where the reported impulses come from a *single* model that is as close as possible to the medians of the impulse response functions of the models that fulfill the sign restrictions. The impulse response function plots in Section 4 report the MT and the point-wise median and associated percentiles of the admissible impulse response functions, with our preferred impulses being those generated using the MT method. Sensitivity to alternative methods of selecting the impulse response functions from the admissible set and sensitivity to alternative models with constant volatility is explored in Section 4.4.

Kilian and Murphy (2012) show that it is insufficient to use sign restrictions to infer economically meaningful responses to the oil sector shocks. The literature imposes additional restrictions to eliminate the set of admissible structural models with implausible impulses to solve this problem. For example, Antolín-Díaz and Rubio-Ramírez (2018) propose narrative sign restrictions ensuring that the structural shocks and historical decompositions are consistent with the established narrative around chosen historical events to constrain the structural parameters. Additional boundary restrictions on the sizes of short-run oil demand and oil supply elasticities are also imposed in recent literature to exclude those models that come with implausibly high elasticities (Baumeister and Peersman, 2013a, 2013b; Kilian and Murphy 2012, 2014). Baumeister and Peersman (2013a) set a bound of $(-0.8, 0)$ for the short-run oil demand elasticity as the short-run elasticity cannot surpass the long-run elasticity, calculated to be around -0.8 in the Hausman and Newey (1995) U.S. household surveys. The bound for short-run price elasticity of the oil supply is between 0 to 0.6. This paper follows Baumeister and Peersman (2013a), who impose additional boundary restrictions on the sizes of the short-run oil demand and oil supply elasticities.[2]

[2] Kilian and Murphy (2012) set the upper bound of the short-run price elasticity of oil supply to be 0.0258, which is much less than the upper bound imposed by Baumeister and Peersman (2013a). The difference comes from data frequency as Kilian and Murphy use monthly data while Baumeister and Peersman (2013a) use quarterly data. Because our paper focuses on the influences of the oil sector shocks on the macroeconomic conditions,

2.4 Generalized impulse response functions

The main obstacle in constructing impulse response functions in drifting volatility models is that these models have non-linearities. In a VAR model with non-homogenous time-varying shocks, the scale of the impulse response functions will change over time. Authors have approached the issue of what impulse response function to report in different ways. Nakajima, Kasuya and Watanabe (2011) use the average of the time series of stochastic volatility over the sample period as the size of the shock, while Chan and Eisenstat (2015) use the shocks identified in the final period of the dataset.

This paper follows the Chan and Eisenstat (2015) approach and uses the impulse response functions to the shocks in the final period of the sample as the reported impulse responses. For comparison, we also present the impacts of a shock from the 2007-2008 global financial crisis. These impulse response functions demonstrate the short and long-run effects of the oil price shocks from these specific periods on the U.S macroeconomy and financial market variables. For VAR models with time-varying volatility, time-varying impulse responses capture the dynamic impacts of time-varying oil price shocks at each point in time. Baumeister and Peersman (2013a) compute the generalized impulse response functions as the drifting impulse responses and depict the dynamic responses of the U.S. macroeconomic and financial market variables one year after each shock. These responses better reflect what occurred in each quarter, given the historical shocks.

The generalized impulse response functions developed by Koop, Pesaran and Potter (1996) can be defined as:

$$GIRF_{t+h} = \mathbb{E}[y_{t+h}|\Gamma_t, \varepsilon_t] - \mathbb{E}[y_{t+h}|\Gamma_t]. \tag{8}$$

The variable y_{t+h} is the forecast of the endogenous variables h periods ahead, Γ_t is the current information set and ε_t the structural shock. The generalized impulse response function represents the discrepancy between conditional expectations, including and excluding the shock ε_t. The information set Γ_t includes the actual values of the lagged endogenous variables, and knowledge of the model parameters, hyper-parameters and structural shocks up to period t.

we estimate the model using quarterly data as much of the macroeconomic data series are of quarterly frequency. Given that our model uses quarterly data, we confine the short-term oil supply price elasticity to be between 0.0 and 0.6).

3 Data

The paper studies the effects of the oil sector shocks on 17 U.S. macroeconomic and financial market variables over 1986Q1 to 2019Q2 for 20 variables. The starting date of 1986Q1 is consistent with similar studies finding a sizeable break in the nexus between the oil market and the real economy at this time (Baumeister and Peersman, 2013b; Peersman and van Robays, 2009). This starting date corresponds to the OPEC collapse and the beginning of the Great Moderation.[3] With the collapse of OPEC, oil prices are more likely to respond endogenously to U.S. macroeconomic conditions.

The oil sector consists of oil price and production variables, and the model includes a world economic activity variable to distinguish the oil demand, oil supply and precautionary oil shocks. The oil price variable (P^{oil}) is the real U.S. refiner's acquisition cost of imported crude oil (IRAC) following the finding of Baumeister, Peersman and van Robays (2009) that the IRAC is a superior proxy for the free global price of oil. This price is commonly applied in the literature as a proxy for the free market of imported crude oil (Baumeister and Peersman, 2013b; Cross and Nguyen, 2017; Kilian, 2009).

Global oil production (Q^{oil}) is the world crude oil supply, and world economic activity (Y^w) is proxied by global industrial production (Baumeister and Peersman, 2013b; Melolinna, 2014). Baumeister and Peersman (2013b) show that global industrial production is the best index at the quarterly frequency to measure the world economy.[4] The oil price, oil production and

[3]In the 1970s to early 1980s OPEC operated as a cartel to set or influence the oil price. Saudi Arabia in late 1985 failed to bolster the oil price by reducing its oil provision. As a result, Saudi Arabia reversed its policy of restricting oil production, causing OPEC's collapse (Baumeister and Kilian, 2016). Although OPEC attempted to reconvene and curtail production in the following year, there is no evidence that these attempts are unsuccessful after that (Kilian and Murphy, 2014).

[4]The oil price is from the Energy Information Administration (EIA)'s *Monthly Energy Review*. World oil production data is the seasonally adjusted crude oil production in thousand barrels daily from the EIA. The global industrial production index is from the United Nations *Monthly Bulletin of Statistics* computed as the weighted average of the industrial activity of a large number of countries. Refer to Peersman and van Robays (2012) for a full description of the index. We splice the series with world industrial production data retrieved from the Netherlands Bureau for Economic Policy Analysis because of the discontinuation of the United Nations series in 2008Q3. This series started in 1991, and the index covers 97 percent of world industrial production. An alternative is to use proxy world economic activity using the dry cargo single voyage ocean freight rate (Cross and Nguyen, 2017; Kilian, 2009; Ratti and Vespignani, 2013). An alternative

Table 2: Summary of the U.S. macroeconomic and financial market data and transformations, 1986Q1 to 2019Q2.

Variables	Transformation
Real gross domestic product	log-difference
Consumer price index	log-difference
Effective federal funds rate	level
Personal income	log-difference
Real personal consumption expenditure	log-difference
Industrial production index	log-difference
Civilian unemployment rate	level
Housing starts	log-difference
Producer price index	log-difference
Average hourly earnings: manufacturing	log-difference
Real gross private domestic investment	log-difference
Capacity utilization: manufacturing	log-difference
M1 money stock	log-difference
10-Year Treasury constant maturity rate	level
Corporate bond yield[1]	level
S&P 500 index	log-difference
Nominal exchange rate	log-difference

[1] The corporate bondy yield is Moody's seasoned Baa corporate bond Yield relative to yield on 10-Year Treasury constant maturity.

world economic activity variables are converted into growth rates by taking the first difference of the natural logarithm of each.

Table 2 lists the macroeconomic and financial market variables in the U.S. component of the model.[5] These U.S. macroeconomic aggregates are commonly used in the macroeconomic models such as in Chan (2020) and

indicator of world economic conditions is that recently proposed by Baumeister, Korobilis and Lee (2020), covering multiple dimensions of the global economy by including the information of 16 variables through the first principal component while previously existing indicators are constructed based on a single category of the global economy.

[5]The U.S. data comes from the FRED database except for the S&P 500 index and the nominal exchange rate. The S&P 500 index is from Datastream, and the nominal exchange rate measured by the effective exchange rate narrow index is from the Bank for International Settlements.

Koop (2013). The non-stationary variables are converted into growth rates by taking the first difference of the natural logarithm of each variable, while the stationary variables remain in levels. The table summarises the transformations applied to each variable.

Carriero, Clark and Marcellino (2016) and Chan (2020) use quarterly data for the U.S. in developing their large BVAR-CSV models. We follow their approach and use quarterly frequency data given our focus on macroeconomic data where many of the series of interest are of quarterly frequency. Quarterly frequency data also reduces the number of parameters that need to be computed, given the large data set and time-varying volatility and significantly reduces the intensity of the computational process. The monthly oil price and oil production data are transformed to quarterly frequency by calculating the average of the monthly observations in each quarter. This paper also follows Baumeister and Peersman (2013b) and Hamilton and Herrera (2004) and sets the lag length to $p = 4$.

4 Empirical results

This Section begins by comparison of the BVAR-CSV and BVAR-CV models 4.1. The empirical results are based on the estimation of the large BVAR-CSV model as the model comparison criterion favoured the model allowing for stochastic volatility within the data.

Section 4.2 presents the generalized impulse response functions to examine the impacts of oil sector shocks on the U.S. macroeconomy and financial markets for the sample duration. Section 4.3 presents the impulse response functions in response to the oil sector shocks specific to 2019Q2. Section 4.4 examines the sensitivity of the impulse response functions to shocks occurring in 2008Q4, being the peak of the disruption induced by the global financial crisis and the sensitivity of the impulse response functions estimated by a large BVAR-CV model. This last Section also compares alternative methods of impulse selection of the sign restrictions algorithm.

4.1 Model comparison results

Table 3 compares the BVAR-CSV model with the BVAR-CV model using the Bayes factor as described in Section 2.2. Following Chan and Eisenstat (2015), the Bayes factors are expressed in natural logarithms. The large

BVAR-CV is the base model. The natural logarithm of the Bayes factors for the BVAR-CV model is equal to zero. If the natural logarithm of the Bayes factors of a BVAR-CSV model is greater than zero, then the BVAR-CSV model is the preferred model for the given data. The results underline the substance of allowing for stochastic volatility in modeling global oil and U.S. macro and financial data as the data strongly favors the BVAR-CSV model.

Table 3: Log Bayes factors for the large BVAR-CV and BVAR-CSV models.

	BVAR-CV	BVAR-CSV
Log Bayes factor	0	**279.3**

4.2 Generalized impulse response functions

This Section examines the time-varying effects of the oil sector shocks on the U.S. macroeconomy and financial markets using generalized impulse response functions. Figures 2 to 4 show the reactions to the oil supply, global oil demand and precautionary oil shocks. The plots of the responses of the oil sector variables are for the quarter in which the shock occurs. The plots of the U.S. macroeconomic and financial variables depict their reaction one year after the initial shock (Baumeister and Peersman, 2013b). The responses are cumulated and are in levels.

We draw several conclusions from these plots. First, the magnitudes of the responses of the U.S. macroeconomic and financial market variables to the shocks change over time because stochastic volatility matters for the evolution of the size of the shocks.[6] Second, the effects of all three of the

[6] Baumeister and Peersman (2013b) examine the time-varying effects of oil supply shocks on the U.S. macroeconomy using a time-varying coefficients VAR model with stochastic volatility in a small four-variate model. They show that the reactions of GDP and consumer prices to oil supply shocks change substantially over time. The exclusion of the drifting parameters in our model mutes the reactions to the oil price shocks compared to Baumeister and Peersman (2013b) as the inclusion of heteroskedasticity does not considerably perturb the VAR parameter estimates in our model. This paper trades off modeling the impact of the time-varying coefficients in favor of using a large dataset.

Figure 1: Posterior mean of the log volatility h_t under the BVAR-CSV

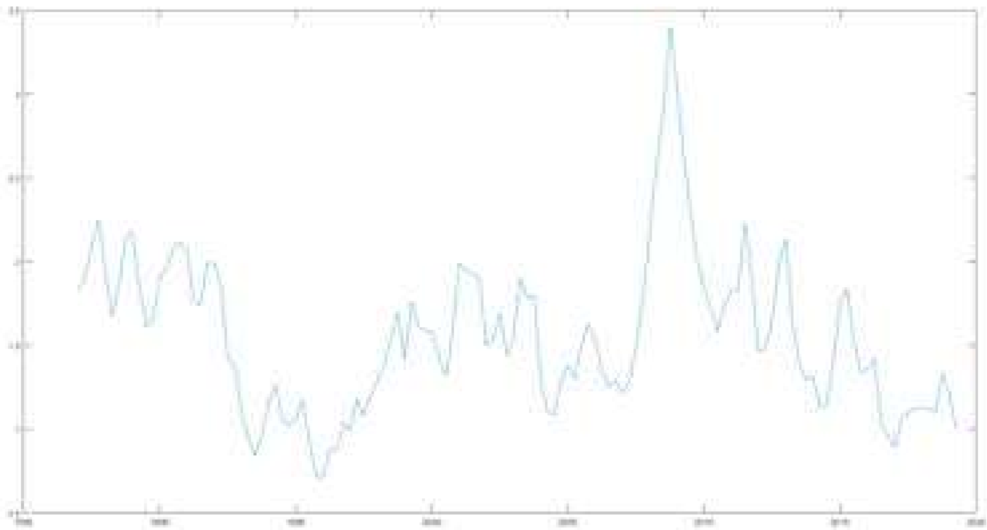

oil shocks peak during the global financial crisis and are often two to three times larger than the impacts of the equivalent shocks before and after the financial crisis. Correspondingly, the scales of the impulse response functions increase during the financial crisis.

Figure 1 shows the posterior mean of the common log volatility (h_t) through time. The log volatility increases and exceeds three during the 2008 financial crisis and hovers around a range between one and two in other periods. This result shows explicitly that the structural effects of the oil price shocks on the U.S. macroeconomy and financial markets vary over time, especially in the financial crisis, emphasizing the importance of including time-varying volatility in modeling the nexus between the U.S. macroeconomic performance and oil sector shocks.

Oil supply shock: Figure 2 illustrates the time-varying reactions of the oil sector and U.S. aggregates to the oil supply shocks. The results demonstrate that unfavorable oil supply shocks have adverse consequences for many U.S. macroeconomic and financial market variables such as GDP growth, housing starts, real gross private domestic investment and the exchange rate. As oil prices increase, consumer prices and the producer prices for all commodities rise, but average hourly earnings and personal income only slightly change, even in the financial crisis. As the Fed responds by stimulating the economy, the M1 money stock rises, and the federal funds rate falls. The magnitude of the effect of the monetary stimulus in 2008 is twice as large as those in other periods. For the bond market, surprisingly, the 10-year Treasury yield slightly rises in the short-run (Figure 2). One plausible reason is that it takes time for risk-off sentiment to build up and as a result there is a lag for investment to flood into the Treasury bond market.

Global oil demand shock: Figure 3 depicts the generalized impulse impacts of a global oil demand shock on the U.S. macroeconomy and financial markets. Overall the effects on the short-term reactions of the U.S. macroeconomic aggregates are positive. Similar to the responses to the oil supply shocks, the positive impacts of the global oil demand shocks increase dramatically in the global financial crisis in 2008. Surprisingly, average hourly earnings to global oil demand shocks are relatively small, while the positive responses of personal income are considerable. The stock price index declines in response to the positive oil demand shock due to rising global demand, although this effect is not significant. The rising appetite for risk of investors reduces the demand for 10-year Treasury bonds, and the demand for risky corporate bonds soars, increasing the 10-year Treasury bond yield and decreasing the corporate bond yield.

Precautionary oil shock: Figure 4 shows the generalized impulse responses of the U.S. variables to the precautionary oil shocks. The precautionary oil shocks negatively impact some U.S. macroeconomic variables, including GDP growth, real personal consumption, real gross private domestic investment and manufacturing capacity utilization. However, the adverse effects on other macroeconomic variables are not significant. For example, the responses of personal income and housing starts are small for all periods, including the global financial crisis, and the responses of average hourly earnings even increase after the precautionary oil shocks. In terms of U.S. stock and currency markets, the impacts of the precautionary oil shocks are also not significant for most periods. There is a small exception in reaction to the global financial crisis, but the influence of the shock vanishes after four quarters. Risk-off sentiment emerges in the bond market, and investors move funds to safer assets. The demand for Treasury bonds rises and declines for corporate bonds reflected in declines in the 10-year Treasury yield and growth in corporate bond yields.

Figure 2: Generalized impulse response functions of the oil sector and U.S. macroeconomic and financial variables to an oil supply shock, 1986Q1 to 2019Q2.

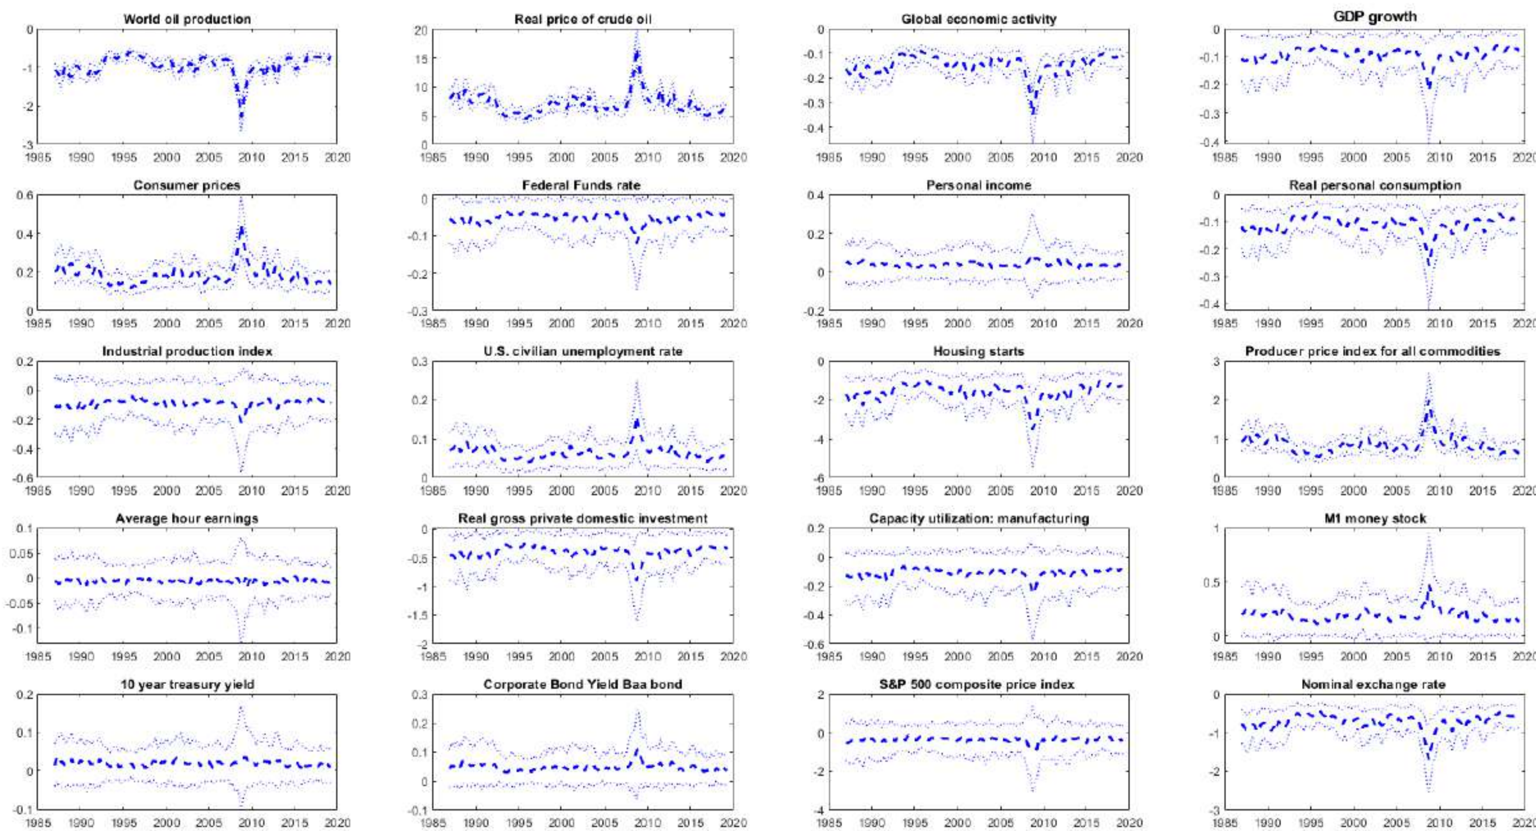

The responses for the oil sector variables are reported at $h = 0$; the responses for the macroeconomic and financial market variables are depicted at $h = 4$. The dotted lines indicate the 68% posterior credible set.

Figure 3: Generalized impulse response functions of the oil sector and U.S. macroeconomic and financial variables to an oil demand shock, 1986Q1 to 2019Q2.

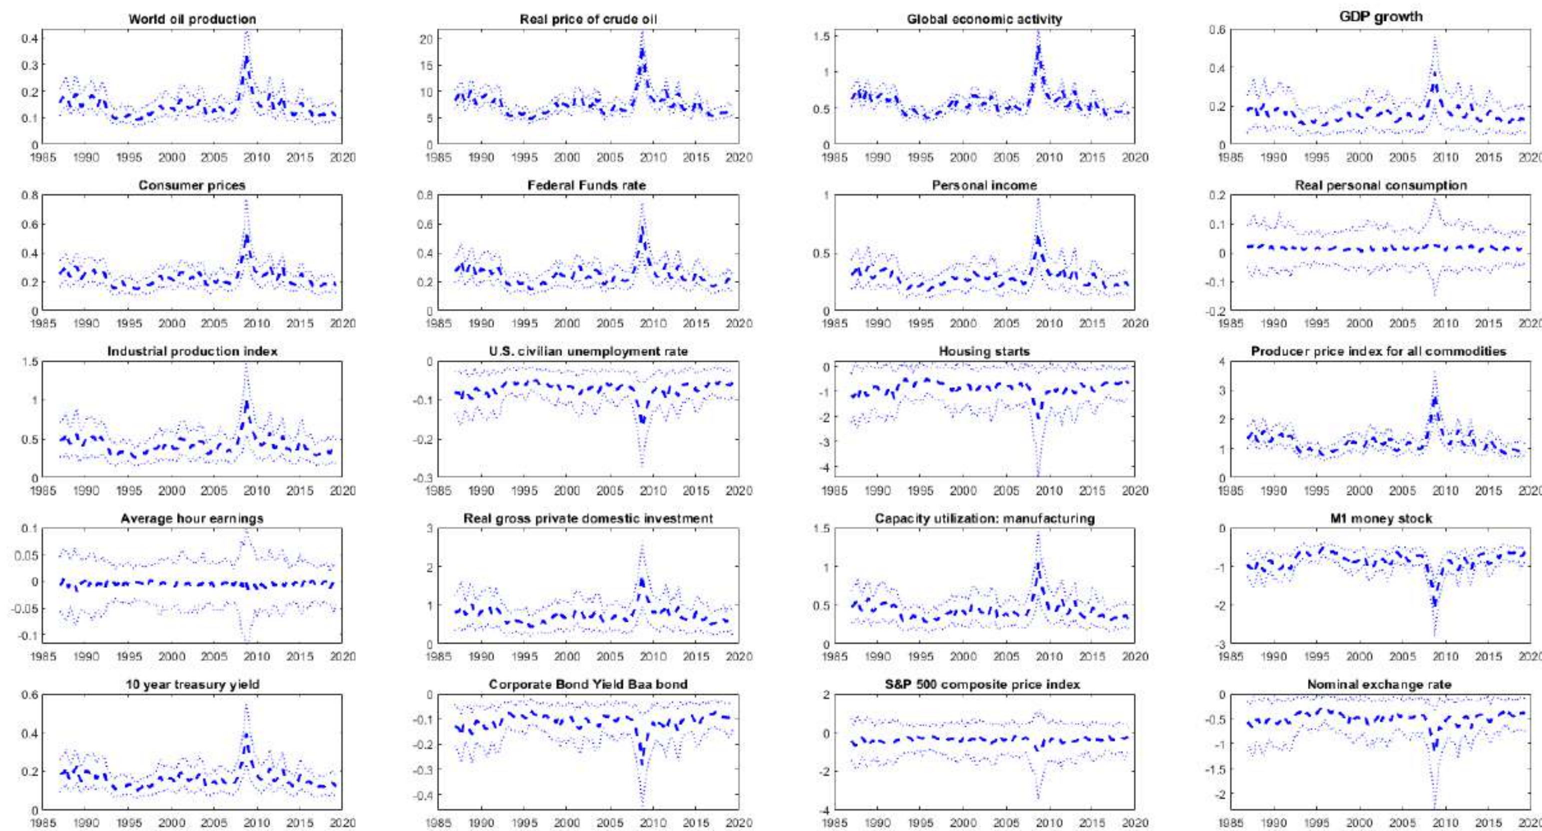

The responses for the oil sector variables are reported at $h = 0$; the responses for the macroeconomic and financial market variables are depicted at $h = 4$. The dotted lines indicate the 68% posterior credible set.

Figure 4: Generalized impulse response functions of the oil sector and U.S. macroeconomic and financial variables to a precautionary oil shock, 1986Q1 to 2019Q2.

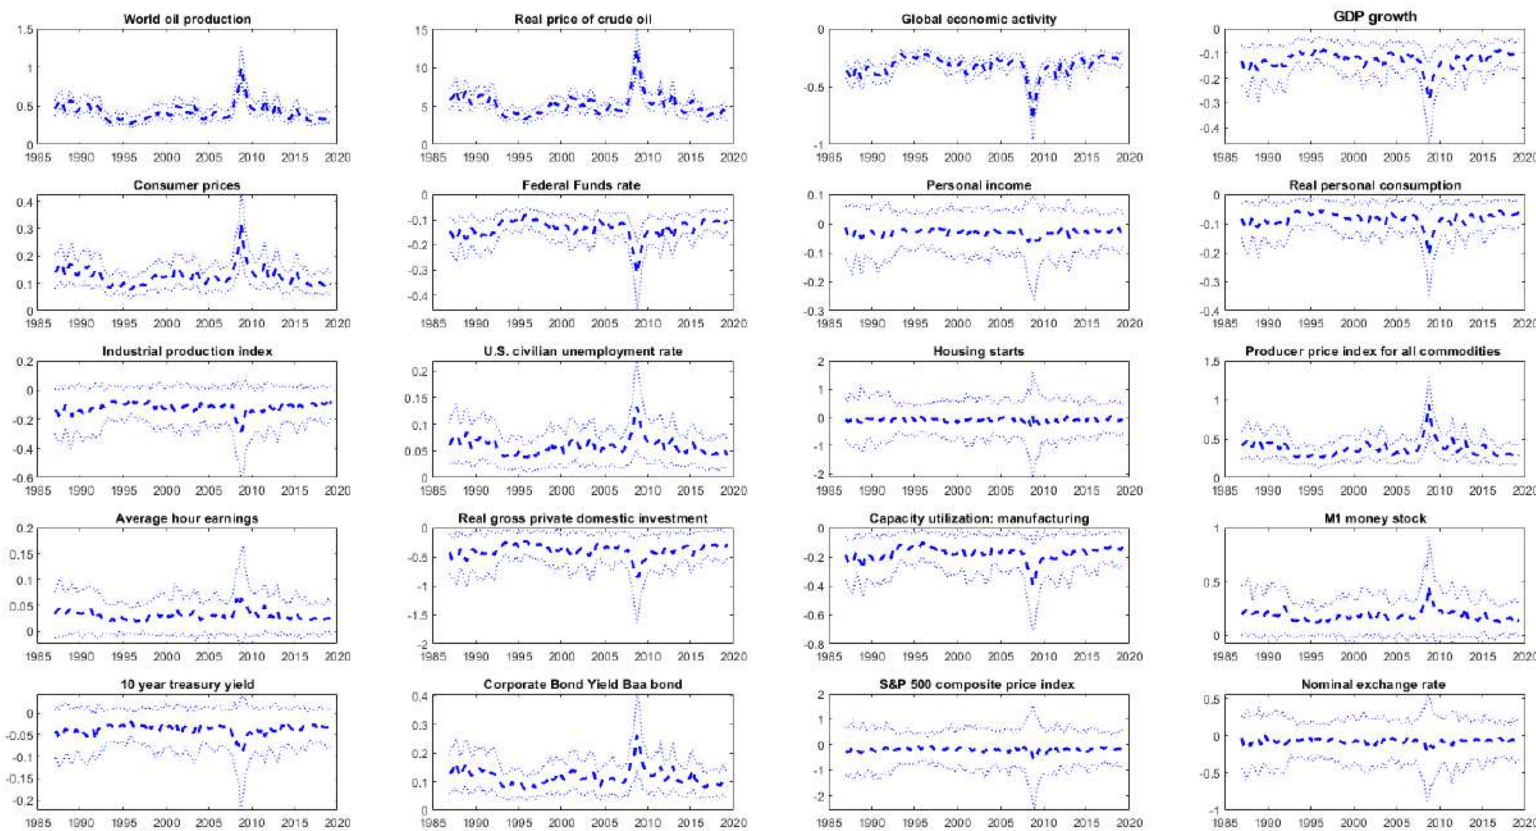

The responses for the oil sector variables are reported at $h = 0$; the responses for the macroeconomic and financial market variables are depicted at $h = 4$. The dotted lines indicate the 68% posterior credible set.

4.3 Impulse response functions

The generalized impulse response functions for the oil sector variables plot the shocks and reactions to the three oil shocks when the shocks occur. The generalized impulse response functions for the U.S. macroeconomic and financial market variables present plots of the responses of each variable four quarters after the oil sector shocks occur. However, a complete-time profile of the impulse response functions to the shocks showing their dynamic effects on the economy, comparable to standard impulse response functions estimated from models without stochastic time-varying volatility, are available. While standard impulse response functions are usually plots in response to one standard deviation shocks, the BVAR-CSV generalized impulse responses plot shocks specific to each quarter. The impulse response functions to the quarter specific shocks follow the same dynamic paths, but the magnitudes of the shocks differ. The magnitudes of the shocks and their contemporaneous effects are shown in the first three panels of the first rows of the generalized impulse response functions in Figures 2 to 4. This Section analyses the impulse response functions of the three oil sector shocks using the shocks occurring in 2019Q2 shown in Figures 5 to 7. The impulse responses are accumulated and shown in levels. The shaded areas indicate the 68 percent posterior credible set.

Oil supply shock: Figure 5 depicts the impulse responses to the oil supply shock in 2019Q2. The results demonstrate that the small unfavorable oil supply shock negatively influences many of the U.S. variables. The oil supply shock raises the unemployment rate in the U.S. by 0.1% in the long-term and causes a perpetual reduction in the level of GDP by 0.2%. Real personal consumption, the industrial production index, housing starts, and real gross private domestic investment experience a permanent decline in response to the adverse supply shock. The MT impulses for personal income and capacity utilization also reflect a reduction, although the credible set shows that an improvement in these variables is possible. Consumer prices in the U.S. rise by 0.1% in the long run for the 2019Q2 shock, and the producer price index for all commodities permanently increases by 0.4% associated with the permanent 3.6% growth in the oil price.

The lower federal funds rate and growth in the M1 money stock reflect the loosening of monetary policy to stimulate the economy. The negative oil supply shock has some adverse effects on the U.S. financial market variables. For example, the U.S. currency depreciates by around 0.5% in the short run and further deteriorates in the long run by 0.8%. Since the negative oil supply shock leads to lower GDP growth and higher inflation, economic uncertainty rises correspondingly.

The uncertainty encourages investors to move away from the relatively riskier stock and corporate bond markets. These movements cause a decline in stock prices, decreased corporate bond prices, and corporate bond yields rise. Surprisingly, the Treasury bond yield, which is a safe asset, increases in the short-term. The share market decline is relatively small in the short run but becomes more substantial in the long run. These movements coincide with the Treasury bond market movements, where yields start to decline around four quarters after the shock. The decline in yields may reflect investors' increasing uncertainty about the macroeconomic environment as at the onset of the negative oil supply shock; investors do not expect a permanent decline in the GDP level. However, as the loosening of monetary policy does not lift GDP growth or inflation, investors become more cautious about the economy and eventually reallocate funds from the stock market to the Treasury bond market.

Global oil demand shock: Figure 6 presents the MT selected impulse response functions and the median impulse response functions to the global oil demand shock for 2019Q2. Similar to the oil supply shock, the global oil demand shock is small and leads to the permanent increase in consumer prices of 0.3% and long-run growth in the producer price index for all commodities of 1.1%. However, the economic impacts of a global oil demand shock on other U.S. variables are different from the influences of the oil supply shock. The global oil demand shock raises U.S. GDP, personal income, industrial production, and manufacturing capacity utilization permanently. However, the 68% credible sets of impulses for U.S. GDP, capacity utilization, and the industrial production index shows that the signs of these effects are only certain for just under two years. The models consistent with the short-run sign restrictions in the short-term differ at the longer time horizon. For many macroeconomic variables, including real personal consumption, housing starts and real gross private domestic investment, the effects of global oil demand shock are not persistent. These variables rise temporarily in response to the global oil demand shock, and the unemployment rate falls for 16 quarters.

Monetary policy is tightened with a higher federal funds rate and a decrease in the M1 money stock to offset the inflationary consequences of the oil demand shock. Contractionary monetary policy sparks a trade-off between price stability and economic stabilization. The long-run impact on GDP is subdued, being less than 0.05%, while the maximum impact on GDP in the first few quarters is more than 0.1%.

In the U.S. financial markets, the nominal U.S. exchange rate experiences a long-run depreciation of 0.6%. Surprisingly, the global oil demand shock induced by the increased world economic conditions causes a permanent reduction in the stock price of 1%. For the bond market, the positive oil demand shock makes investment flow to the riskier corporate bonds and leads to a fall in corporate bond yields and growth in the 10-year Treasury yield.

Precautionary oil shock: Figure 7 contains the MT selected impulse responses to the precautionary oil shock. Although the precautionary shock in 2019Q2 leads to a permanent decrease in some variables such as GDP, personal income, real personal consumption, industrial production, real gross private domestic investment, manufacturing capacity utilization, and an increase in the long-run in the producer price index for all commodities and housing starts, the magnitudes of the impacts of precautionary oil shock are different to those of the oil supply shock. For example, the effects of a precautionary oil shock on personal income on average are larger than those to the oil supply shock and are negative in all horizons, suggesting that 2019Q2 reflects a precautionary oil shock rather than the oil supply or demand shock.

Consumer prices increase in the first year following the shock then declines slightly in the long run, while the point-wise median impulse response shows a permanent inflationary effect. The 68 % credible set of impulse response functions suggests that the response is positive and permanent. The impulse responses of average hourly earnings show a similar discrepancy between the MT selected response and the median response. The MT response shows that average hourly earnings rise slightly in the first few quarters and decrease by around 0.02% in the long run. However, the median response indicates that earnings per hour increase persistently following the shock.

Figure 5: Impulse response functions of the oil sector and U.S. macroeconomic and financial variables to an oil supply shock, 2019Q2.

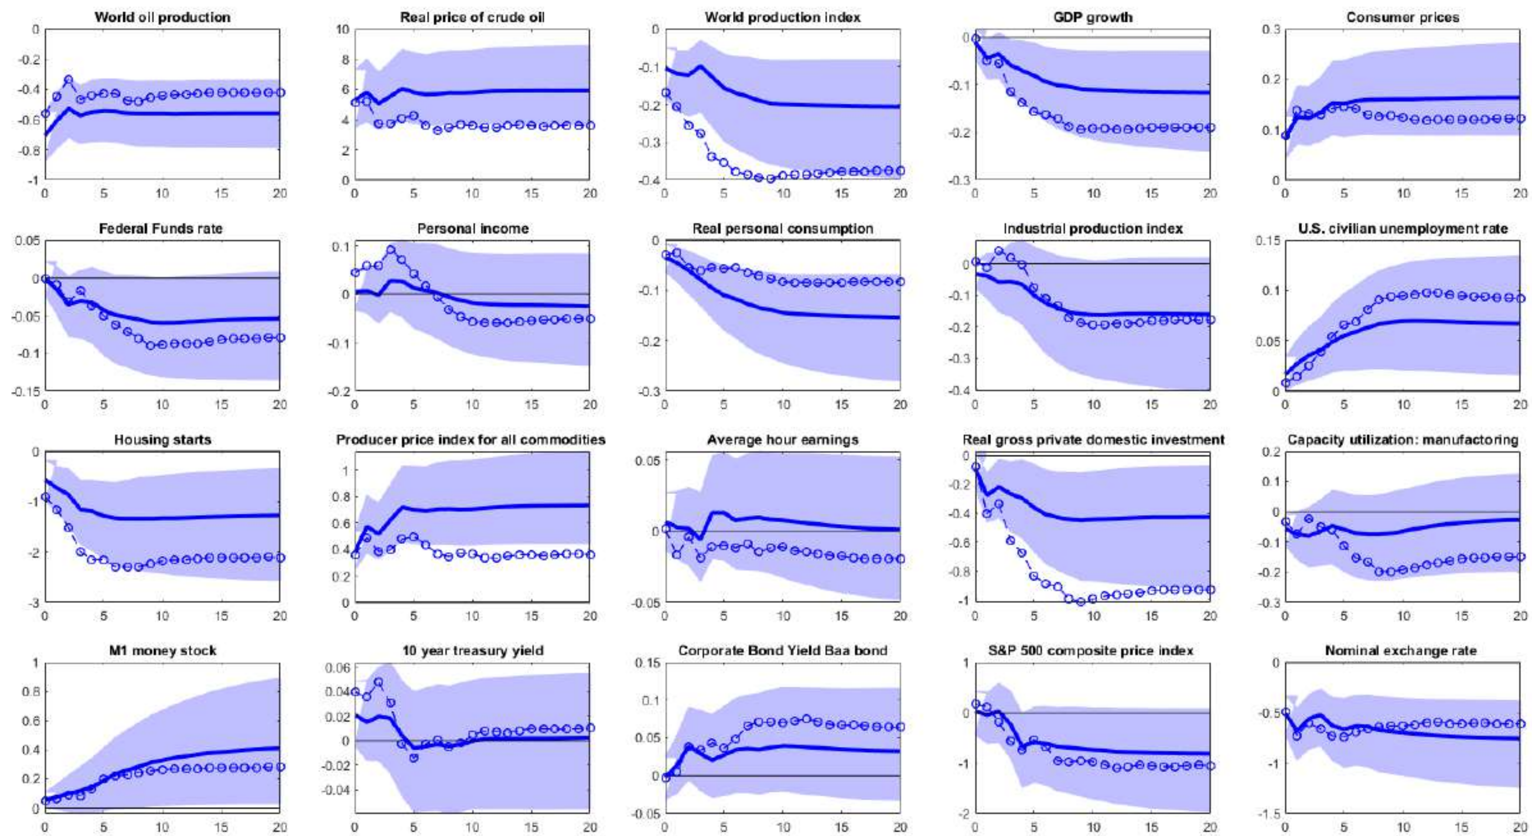

The solid blue lines are the median impulse responses. The blue dashed line with circle markers are the MT selected impulse responses. The shaded areas indicate the 68% posterior credible set.

Figure 6: Impulse response functions of the oil sector and U.S. macroeconomic and financial variables to an oil demand shock, 2019Q2.

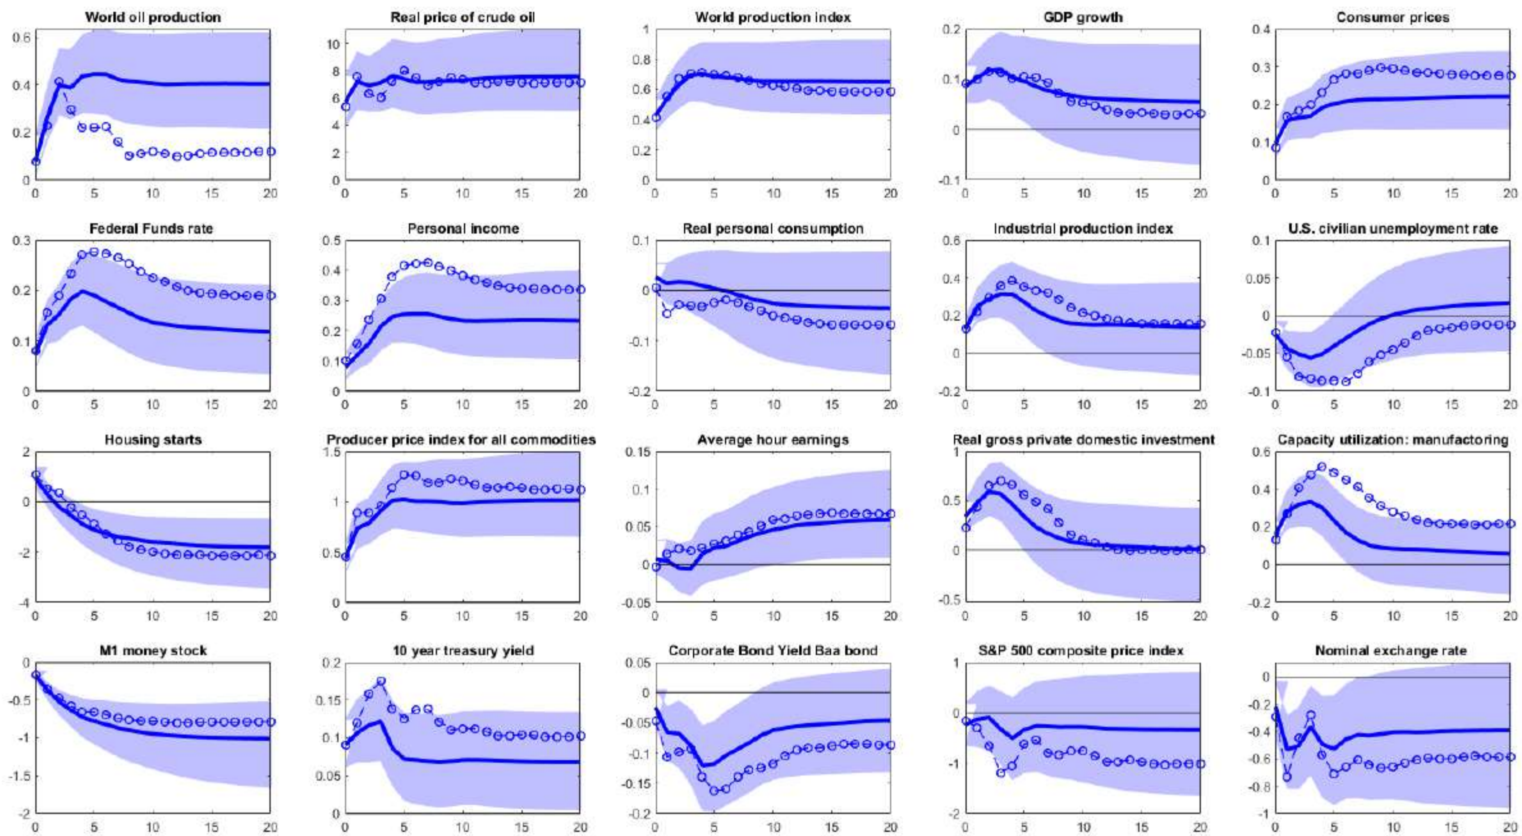

The solid blue lines are the median impulse responses. The blue dashed line with circle markers are the MT selected impulse responses. The shaded areas indicate the 68% posterior credible set.

Figure 7: Impulse response functions of the oil sector and U.S. macroeconomic and financial variables to a precautionary oil shock, 2019Q2.

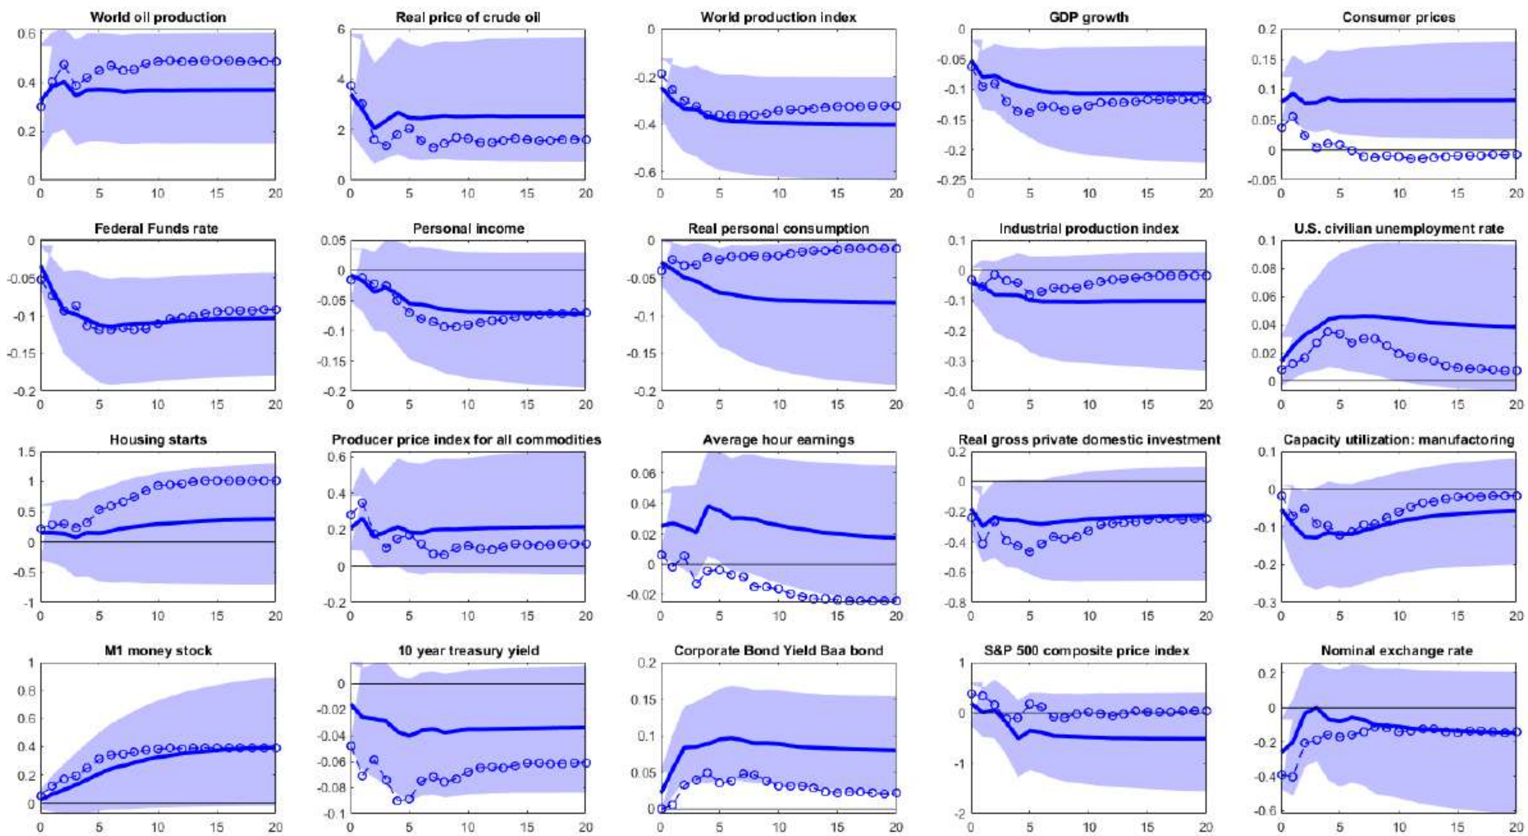

The solid blue lines are the median impulse responses. The blue dashed line with circle markers are the MT selected impulse responses. The shaded areas indicate the 68% posterior credible set.

4.4 Sensitivity analysis

Figures 8 to 10 show the median impulse responses and MT selected responses to the oil shocks in 2008Q4 to illuminate how the impulse response functions differ given the size of the shocks. The shocks in 2019Q2 are small, while the shocks in 2008Q4 are the largest in the sample corresponding to the global financial crisis as reflected by the log volatility of h_t in Figure 1. The parameter estimates governing the dynamics of the impulses are the same for each draw. The range of the credible set of impulses is often two to three times wider for the 2008Q4 shock than in 2019Q2. The differences in the dynamic paths of the impulse response functions that occur are a product of randomly drawing 1.5 million admissible impulse response functions that fulfill the sign restrictions. The discrepancies are most substantial for the response of average hourly earnings to the oil supply shock and the reaction of consumer prices to a precautionary oil shock.

Figures 11 to 13 display three alternative approaches to the median-based selection of the impulse response functions for the 2019Q2 shock. The approaches are to choose the model with an oil demand elasticity that is most similar to the posterior median of the elasticity among the set of admissible models (Kilian and Murphy, 2014) and to select the models with the largest response of the real oil price to an oil supply shock, and the smallest response of the real oil price to an oil supply shock. The patterns in the variable responses using the elasticity approach are more similar to those of the median responses than those of MT selected responses. However, some exceptions are where the impulses are dissimilar or even move in opposite directions after some shocks. For example, the elasticity-selected responses show that after a global oil demand shock, the nominal exchange rate declined in the first three quarters then rallies to a permanent 0.2% increase while the median impulse response shows that the U.S. currency depreciates more than 0.4% in the long run. Another discrepancy is that after the precautionary oil shock, the median impulse response for the CPI rises by around 0.08%. However, both the elasticity-selected response and the MT selected-response show that the CPI only rises in the first few quarters before declining in the long run.

Most of the impulse responses from the models with the largest and the smallest reactions of oil to the oil supply shock move in the same direction as their median counterparts. For some variables, the differences in the impulses taken from the models giving the largest and smallest responses are small. For example, the 10-year Treasury yield reactions to a global oil demand shock are similar, especially in the long run. Noteworthy is that the responses from the model with the smallest oil price response are not always smaller for the other variables in the model than those from the model with the largest response to the oil price. For example, after a precautionary oil shock, the sizes of the reactions of some aggregates such as the unemployment rate, the industrial production index, the 10-year Treasury yield from the small response model are more substantial than the equivalent selected from the model with the largest reaction of the oil price to the oil supply shock.

Although the model comparison measured in the previous Section shows that the BVAR-CV provides inferior model-fit for the sample data, this Section estimates the impulse responses to the shocks identified in the final period of the dataset by the BVAR-CV model as another robustness check. Figures 14 to 16 depict the median impulse and MT-selected responses to the oil price shocks in 2019Q2 using a BVAR-CV model to illustrate how the impulse response functions differ, estimated by a BVAR model with constant volatility. The results show some obvious dissimilarities between the magnitudes of the impulse responses estimated by the BVAR-CV model and the ones estimated by the BVAR-CSV model. For example, the responses estimated by the BVAR-CV model show that after an oil supply shock, S&P 500 index drops by around 2% in the long run, while the responses estimated using the BVAR-CSV model illustrate that the S&P 500 index only experiences a less than 1% permanent decline. The discrepancy between MT-selected responses and the median impulse responses, estimated using the BVAR-CV model, are superfluous to all three types of oil sector shocks.

Figure 8: Impulse response functions of the oil sector and U.S. macroeconomic and financial variables to an oil supply shock, 2008Q4.

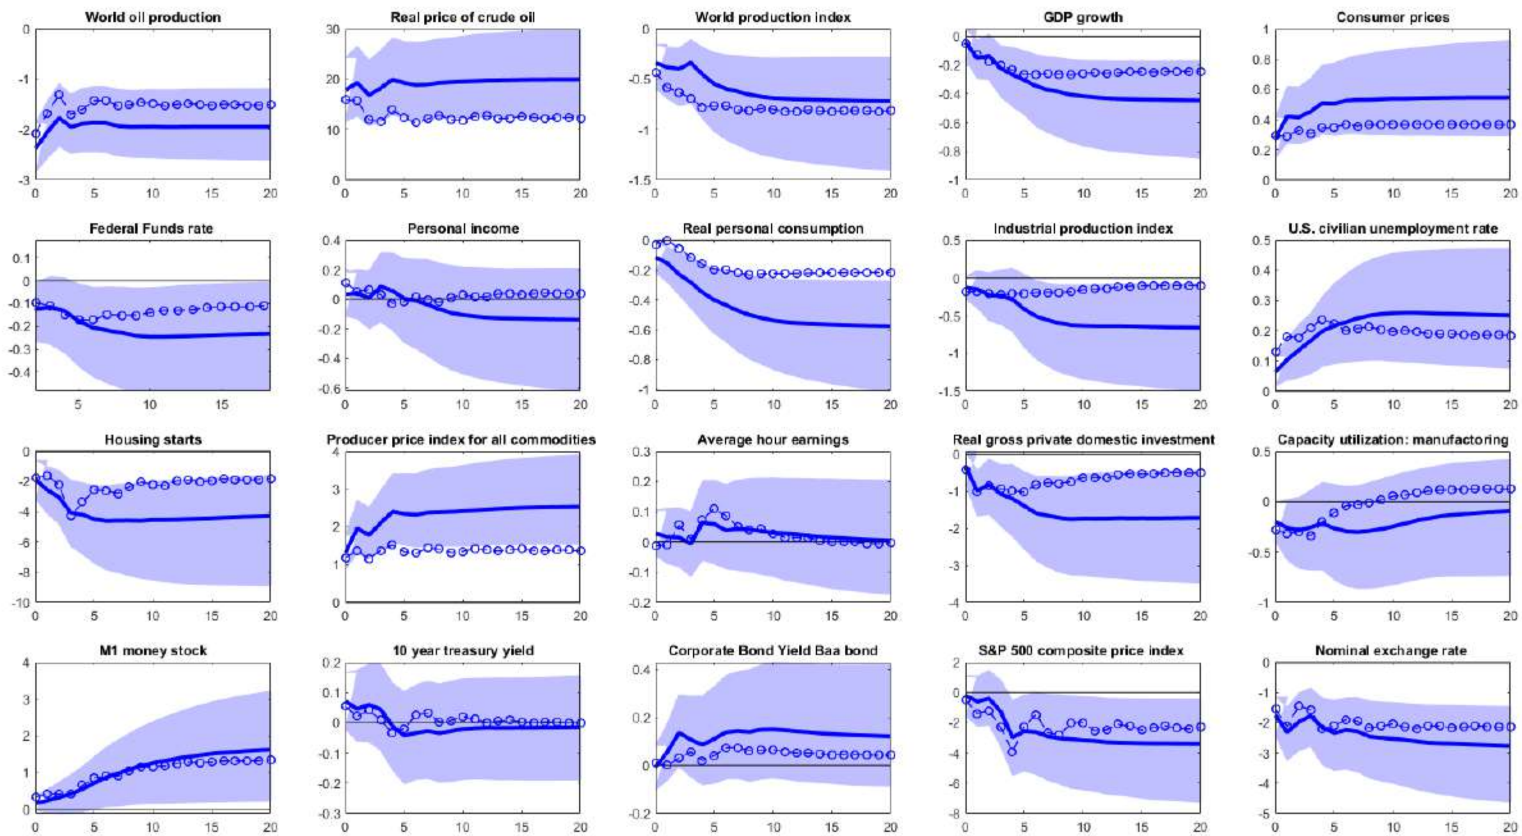

The solid blue lines are the median impulse responses. The blue dashed line with circle markers are the MT selected impulse responses. The shaded areas indicate the 68% posterior credible set.

Figure 9: Impulse response functions of the oil sector and U.S. macroeconomic and financial variables to an oil demand shock, 2008Q4.

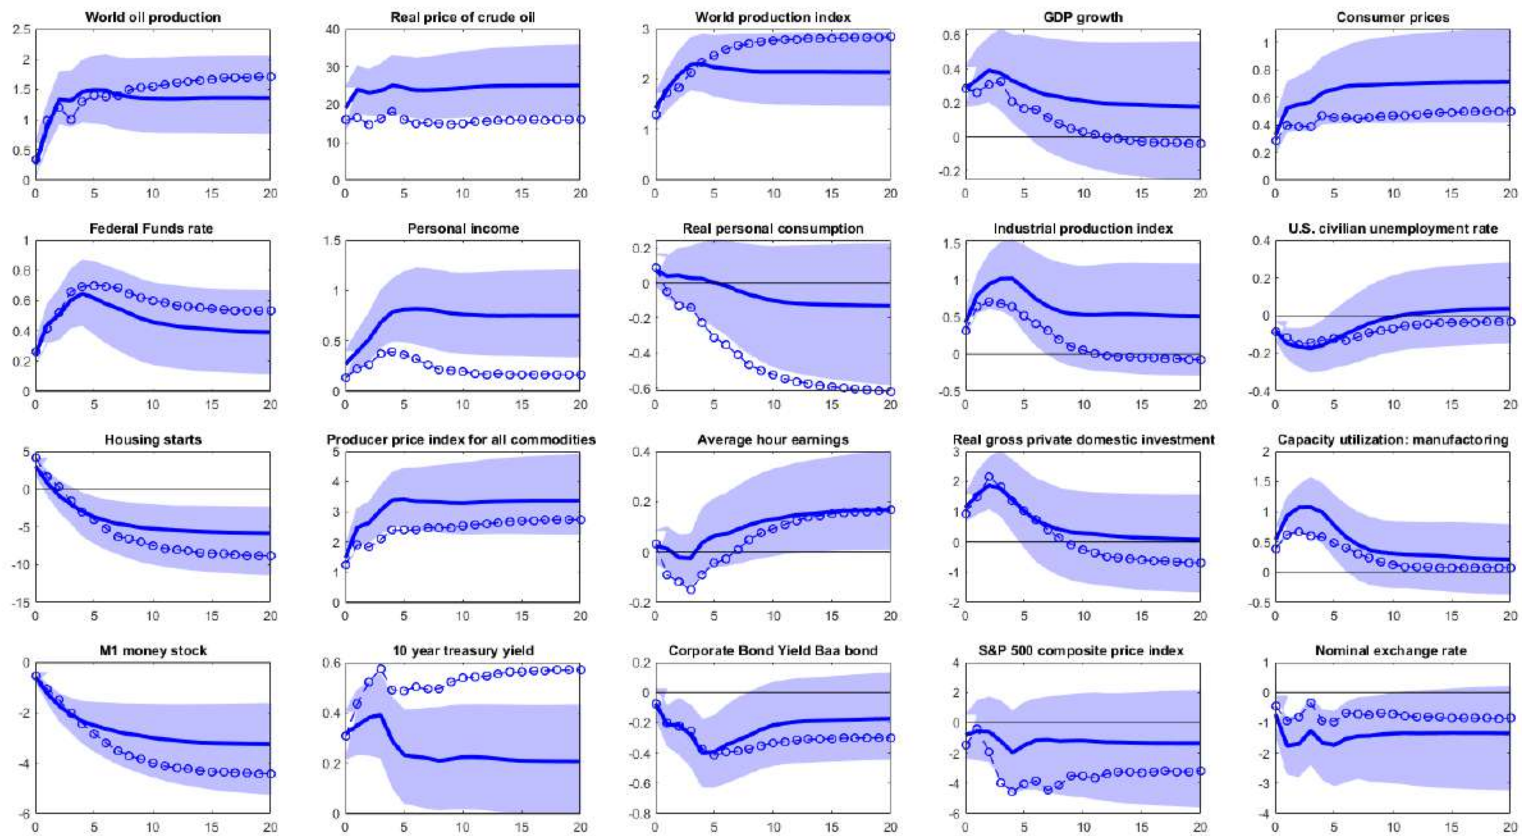

The solid blue lines are the median impulse responses. The blue dashed line with circle markers are the MT selected impulse responses. The shaded areas indicate the 68% posterior credible set.

Figure 10: Impulse response functions of the oil sector and U.S. macroeconomic and financial variables to a precautionary oil shock, 2008Q4

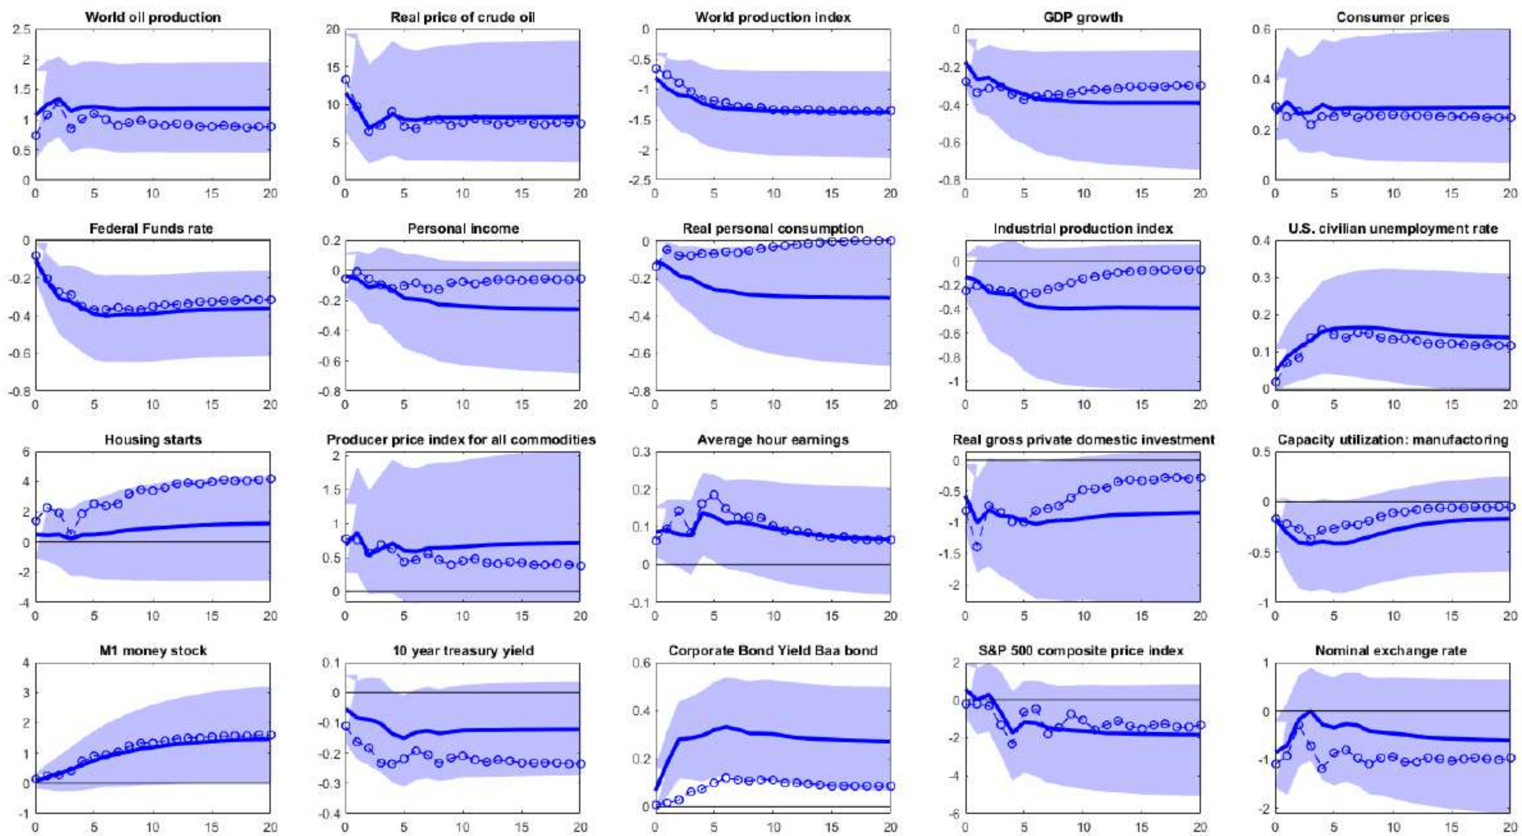

The solid blue lines are the median impulse responses. The blue dashed line with circle markers are the MT selected impulse responses. The shaded areas indicate the 68% posterior credible set.

Figure 11: Impulse response functions of the oil sector and U.S. macroeconomic and financial variables to an oil supply shock, 2019Q2: Alternative identification methods.

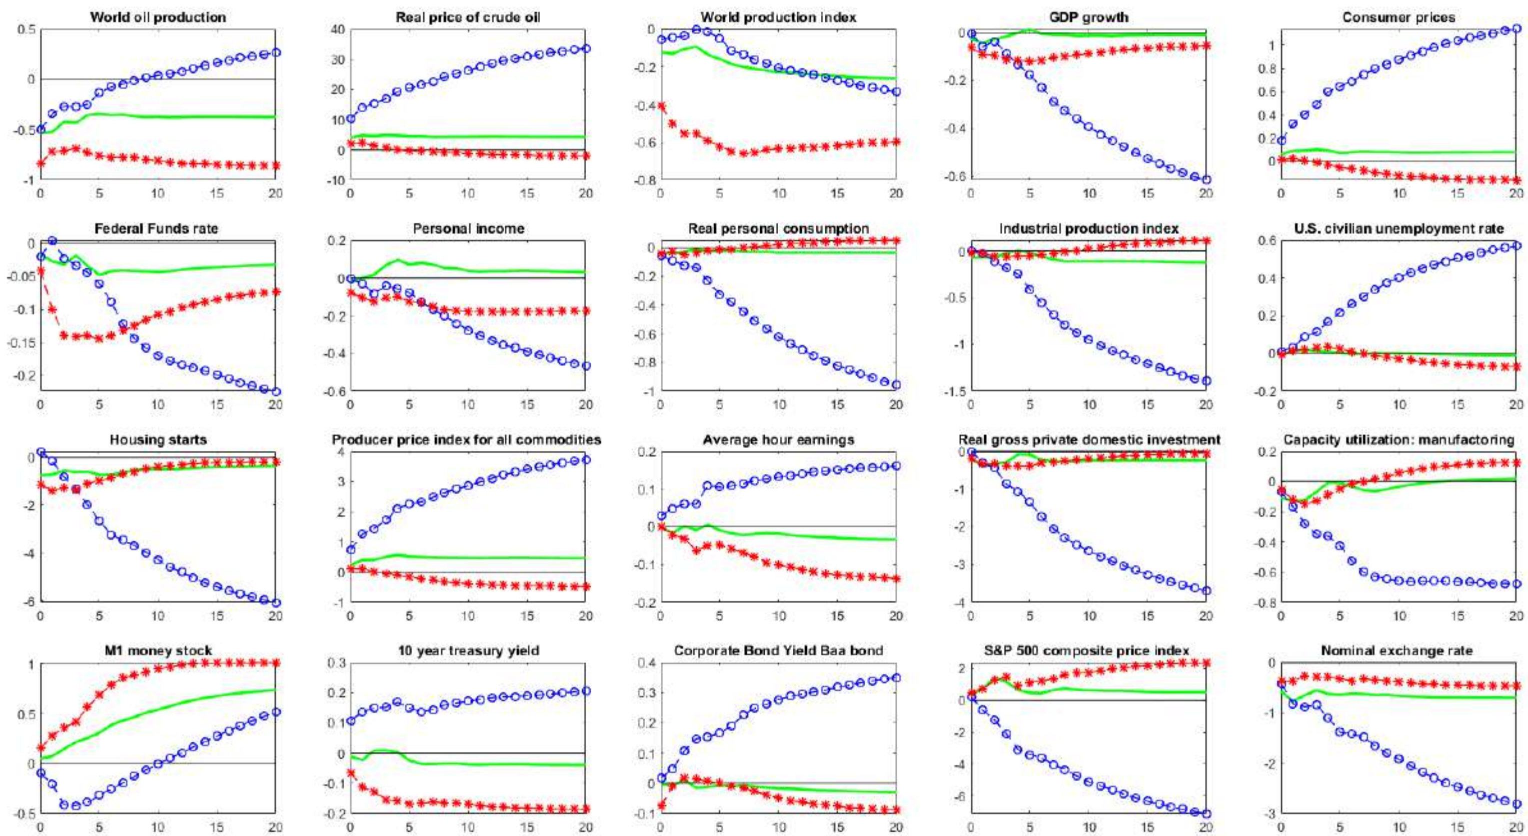

The solid green lines are the model with an oil demand elasticity closest to the posterior median of the elasticity of the admissible models. The blue dashed lines with circle markers are the model with the largest oil price response to an oil supply shock of the admissible models. The red dashed lines with star markers are the model with the smallest response of the oil price to an oil supply shock of the admissible models.

Figure 12: Impulse response functions of the oil sector and U.S. macroeconomic and financial variables to an oil demand shock, 2019Q2: Alternative identification methods.

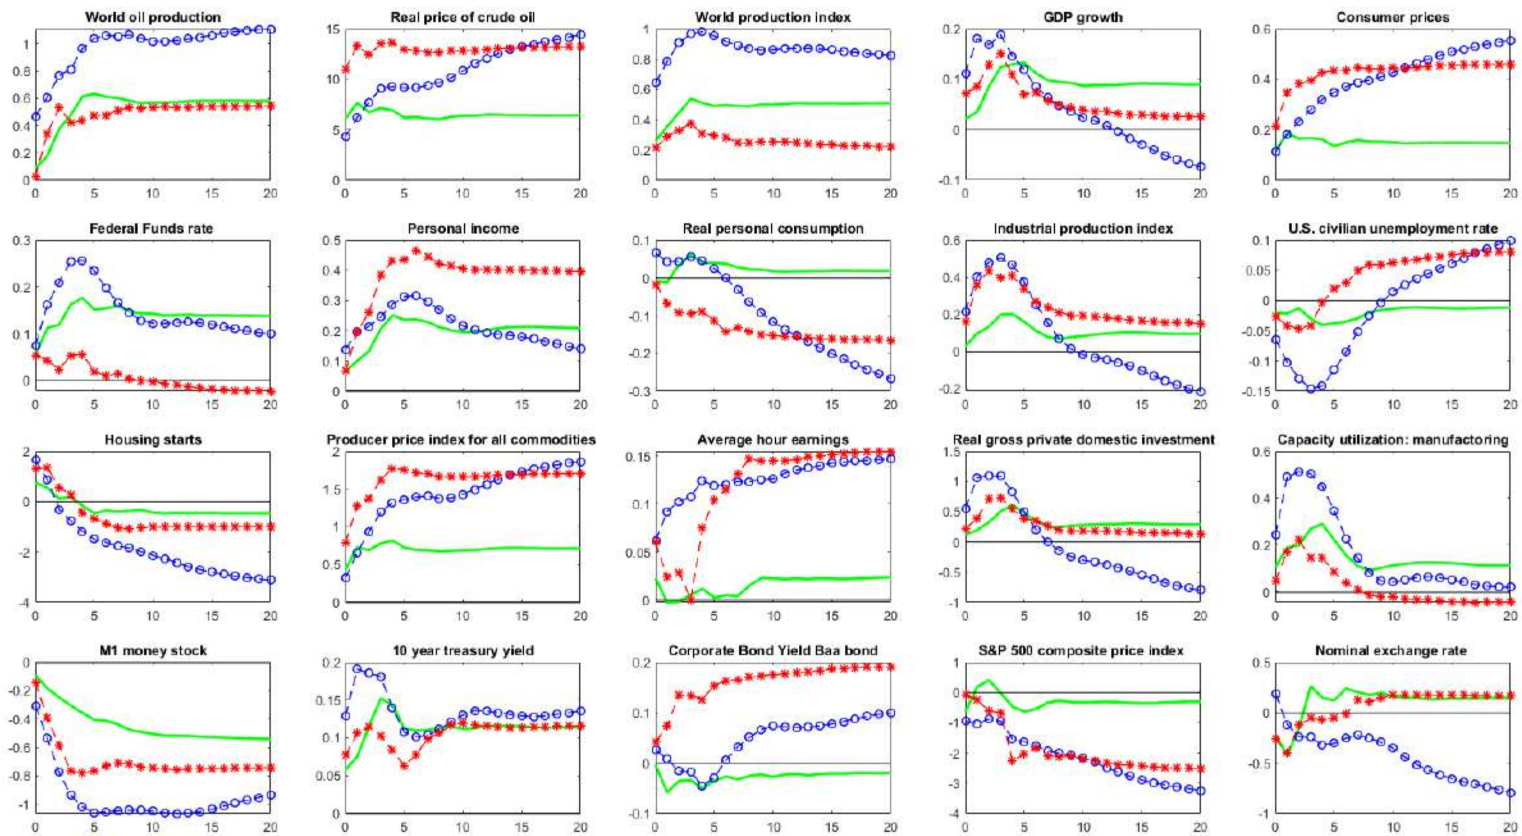

The solid green lines are the model with an oil demand elasticity closest to the posterior median of the elasticity of the admissible models. The blue dashed lines with circle markers are the model with the largest oil price response to an oil supply shock of the admissible models. The red dashed lines with star markers are the model with the smallest response of the oil price to an oil supply shock of the admissible models.

Figure 13: Impulse response functions of the oil sector and U.S. macroeconomic and financial variables to a precautionary oil shock, 2019Q2: Alternative identification methods.

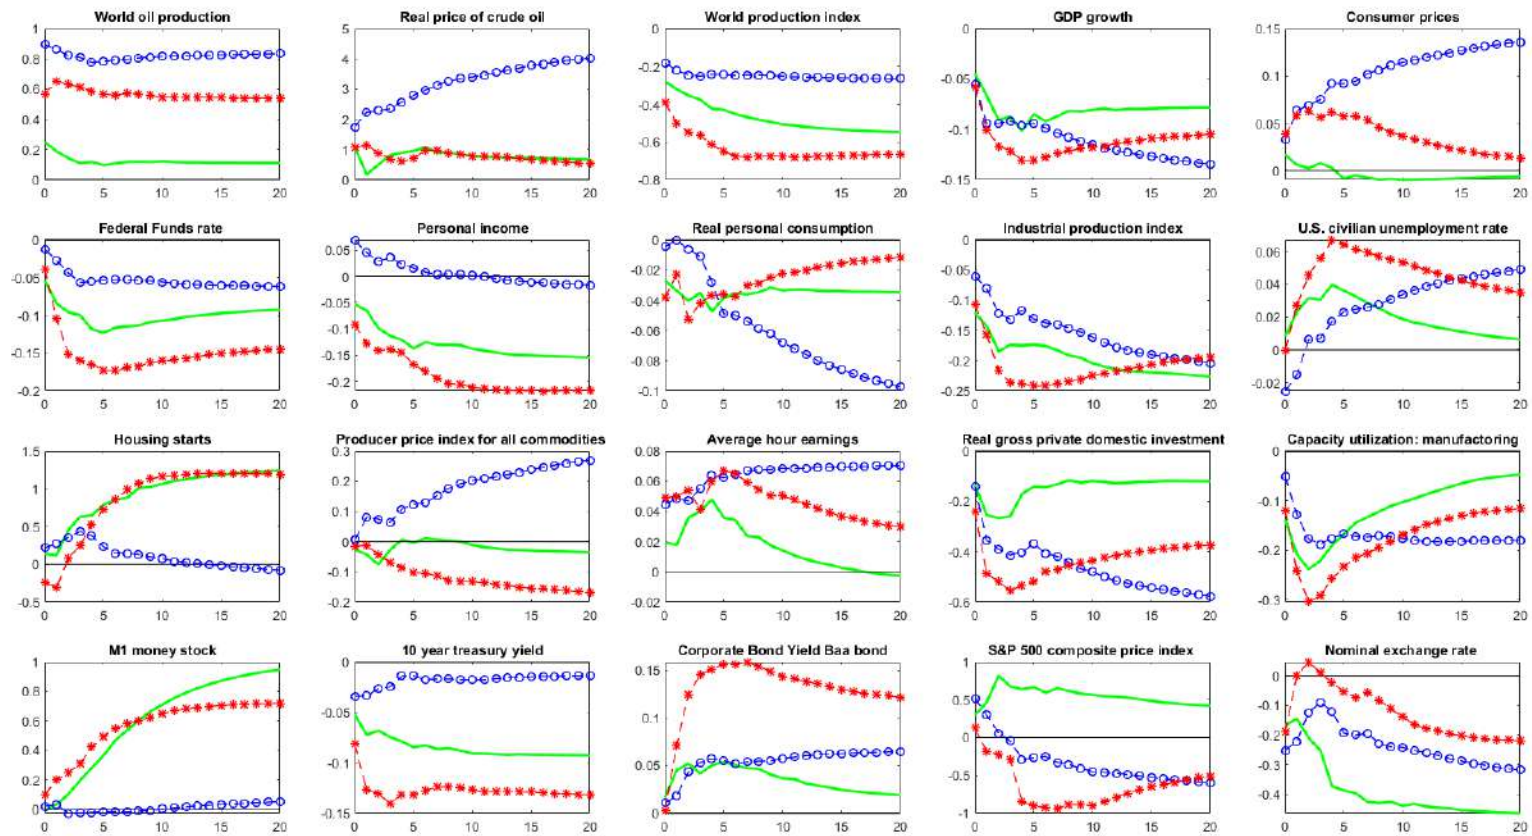

The solid green lines are the model with an oil demand elasticity closest to the posterior median of the elasticity of the admissible models. The blue dashed lines with circle markers are the model with the largest oil price response to an oil supply shock of the admissible models. The red dashed lines with star markers are the model with the smallest response of the oil price to an oil supply shock of the admissible models.

Figure 14: Impulse response functions of the oil sector and U.S. macroeconomic and financial variables to an oil supply shock, 2019Q2: Using the BVAR-CV model.

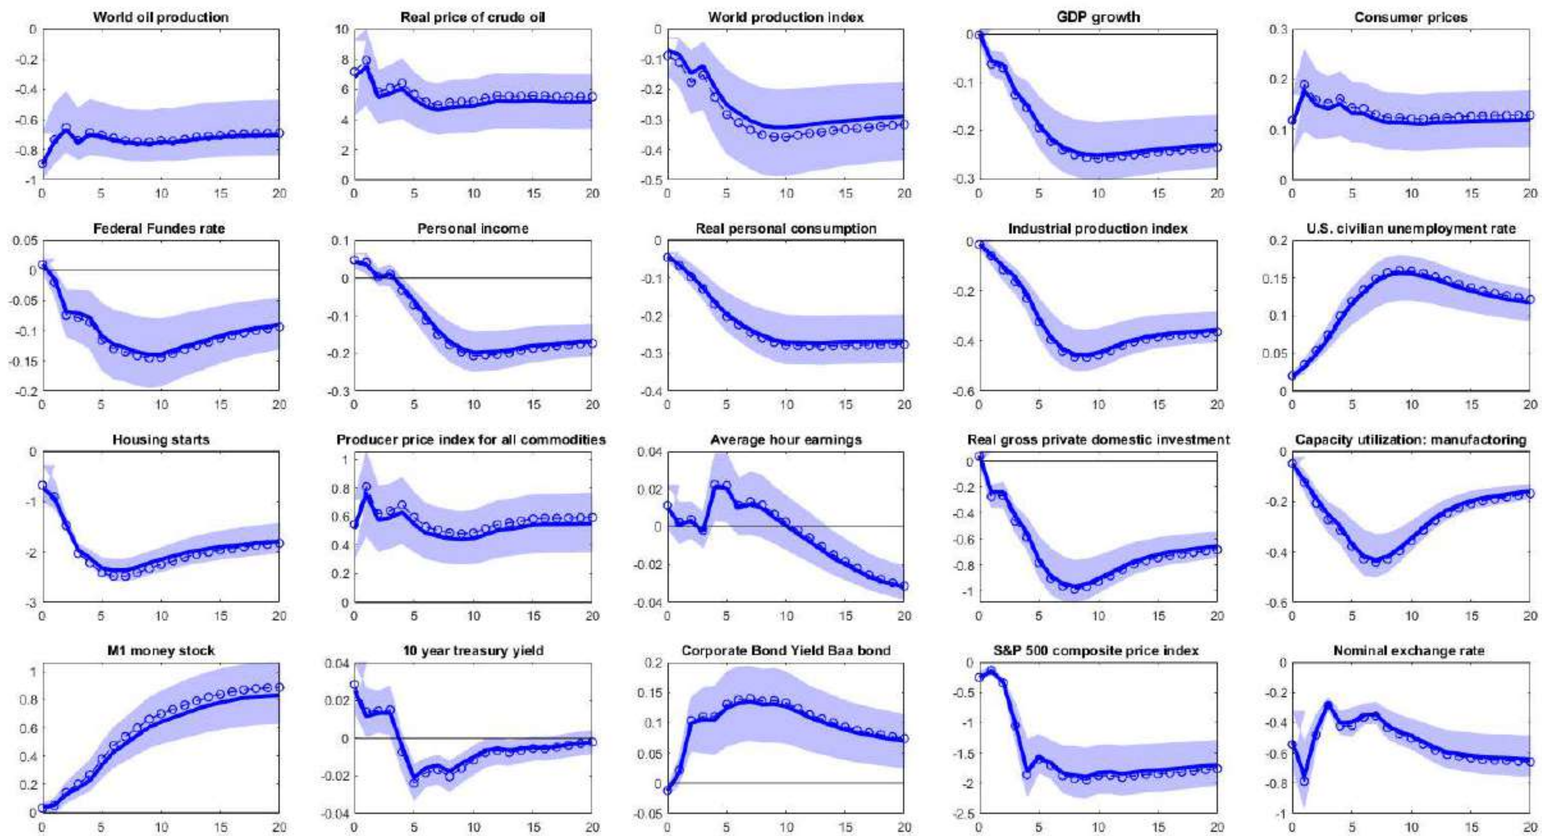

The solid blue lines are the median impulse responses. The blue dashed line with circle markers are the MT selected impulse responses. The shaded areas indicate the 68% posterior credible set.

Figure 15: Impulse response functions of the oil sector and U.S. macroeconomic and financial variables to an oil demand shock, 2019Q2: Using the BVAR-CV model.

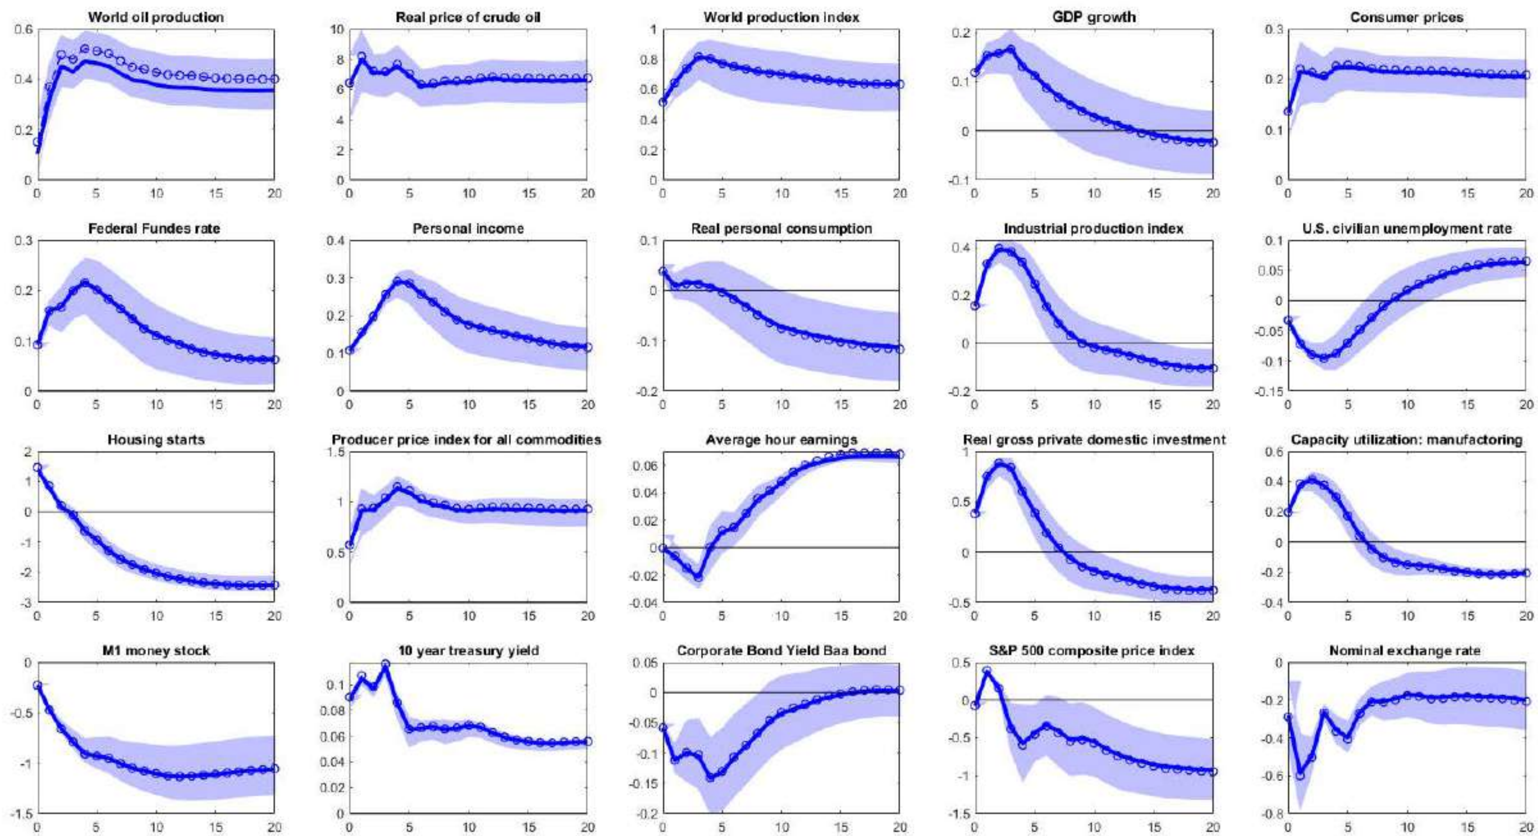

The solid blue lines are the median impulse responses. The blue dashed line with circle markers are the MT selected impulse responses. The shaded areas indicate the 68% posterior credible set.

Figure 16: Impulse response functions of the oil sector and U.S. macroeconomic and financial variables to a precautionary oil shock, 2019Q2: Using the BVAR-CV model.

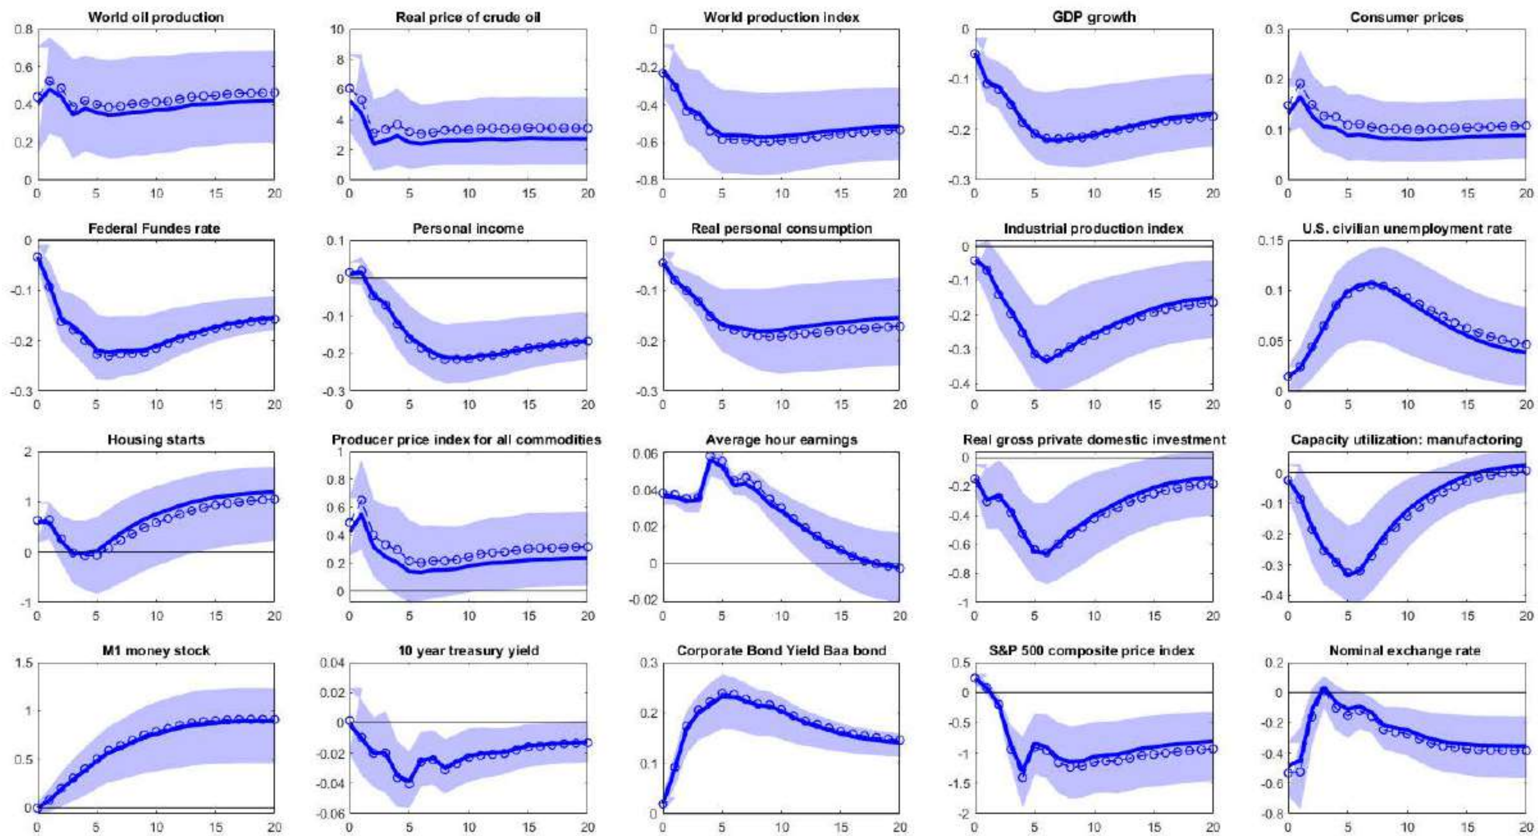

The solid blue lines are the median impulse responses. The blue dashed line with circle markers are the MT selected impulse responses. The shaded areas indicate the 68% posterior credible set.

5 Conclusion

This paper combines recent results in the Bayesian SVAR literature to estimate a large Bayesian VAR model including common stochastic volatility to provide empirical evidence on the impacts of oil sector shocks on U.S. macroeconomic and financial market variables. This recent work shows that BVAR models with a large number of variables outperform models with a small number of variables for structural analysis and forecasting, and time variation in the VAR error terms capture the time-varying volatility that is a feature of macroeconomic data. Formal Bayesian model comparison shows that the data prefer the large BVAR-CSV model to a counterpart with constant volatility. Sign and boundary restrictions are then applied to disentangle the three structural oil sector shocks of an oil supply shock, a global oil demand shock, and a precautionary oil shock.

The advantage of incorporating the time-varying volatility is that the magnitude of the oil sector shocks differs in each quarter. Generalized impulse response functions reveal the contemporaneous effects of the oil sector shocks and their effects on the U.S. variables after one year. The impulse response functions show the dynamic effect of the shocks. The analysis shows the dynamic impulse response functions in response to the oil sector shocks in 2019Q2 and, for comparison, in 2008Q4. The generalized impulse response functions show that while the magnitudes of the responses are not constant at each point in time, they are not considerably different, with the exception of the global financial crisis in 2008 when the size of the shocks and subsequent reactions of most of the U.S. variables to the three oil shocks often rise to be between two to three times larger than those in other periods.

The paper draws several conclusions. First, the impacts of the three oil sector shocks on the model variables are strikingly different. Second, there are permanent inflationary effects for all of the oil sector shocks. Third, although the global oil demand shock is associated with increased GDP, there are some adverse effects on several macroeconomic variables in the long run. Fourth, all oil sector shocks negatively affect the U.S. share market and currency market in the long run, while the influences on the U.S. bond market vary.

Appendix - estimation details

This Appendix sets out the method used to estimate the large BVAR-CSV described in Section 2 following Chan (2020).

The paper sets the prior $p(\mathbf{B}, \mathbf{\Sigma}, \mathbf{\Psi}) = p(\mathbf{B}, \mathbf{\Sigma})p(\mathbf{\Psi})$ with the parameter blocks $(\mathbf{B}, \mathbf{\Sigma})$ and *a priori* independent $\mathbf{\Psi}$. The prior is the standard normal-inverse-Wishart prior where $\mathbf{\Sigma} \sim \mathcal{IW}(\mathbf{K_0}, \delta_\mathbf{0})$ and $(\mathrm{vec}(\mathbf{B}) \mid \mathbf{\Sigma} \sim\mathcal{N}$ $(\mathrm{vec}(\mathbf{B_0}), \mathbf{\Sigma} \otimes \mathbf{V_B})$ (Kadiyala and Karlsson, 1997). $\mathbf{B_0} = \mathbf{0}$ and the covariance matrix $\mathbf{V_B}$ is chosen to be a diagonal matrix to induce shrinkage. The diagonal matrix $\mathbf{V_B}$ is assumed to have diagonal elements $\varsigma_{\mathbf{B},ii} = \tau_2$ for the intercept and $\varsigma_{\mathbf{B},ii} = \tau_1/(\iota^2 \widehat{s_r})$ for coefficients on the lags of variable r denoted ι, where $\widehat{s_r}$ is the sample variance of an AR(4) of r. The paper also sets $\delta_0 = n + 3, \mathbf{K_0} = \mathbf{I_n}, \tau_1 = 0.2^2$ and $\tau_2 = 10^2$. The intercepts are not shrunk to zero while the lag coefficients ι are shrunk to zero as the lag length increases (Koop and Korobilis, 2010). In terms of $\mathbf{\Psi} = diag(e^{h_1}, \ldots, e^{h_T})$ with $h_t = \rho h_{t-1} + \xi_t^h$; $\xi_t^h \sim \mathcal{N}(0, \delta_h^2)$, the prior is the independent inverse-gamma prior for δ_h^2: $\delta_h^2 \sim (\mathcal{IG})(\upsilon_h, K_h)$ and truncated normal prior for $\rho \sim \mathcal{N}(\rho_0, V_\rho)\mathbf{1}(|\rho| < 1)$. Finally, $\upsilon_h = 5, K_h = 0.04, \rho_0 = 0.9$ and $V_\rho = 0.2^2$.

To estimate the BVAR-CSV we sample jointly $\mathbf{B}$ and $\mathbf{\Sigma}$. Specifically, the natural conjugate priors for $\mathbf{B}, \mathbf{\Sigma}$ mean that the posterior draws can come from $p(\mathbf{B}, \mathbf{\Sigma} \mid \mathbf{Y}, \mathbf{\Psi})$ and $p(\mathbf{\Psi} \mid \mathbf{Y}, \mathbf{B}, \mathbf{\Sigma})$.

Combining the likelihood $p(\mathbf{Y}|\mathbf{B}, \mathbf{\Sigma}, \mathbf{\Psi})$ means that the joint prior density function $p(\mathbf{B}, \mathbf{\Sigma})$, the posterior density function $p(\mathbf{B}, \mathbf{\Sigma} \mid \mathbf{Y}, \mathbf{\Psi})$ is normal inverse-Wishart, making analytical results easy to obtain. There are two steps in sampling $(\mathbf{B}, \mathbf{\Sigma} \mid \mathbf{Y}, \mathbf{\Psi})$. First, sample $\mathbf{\Sigma}$ from $(\mathbf{\Sigma} \mid \mathbf{Y}, \mathbf{\Psi})$ from the inverse-Wishart distribution. Second, given $\mathbf{\Sigma}$, sample the $(\mathrm{vec}(\mathbf{B}) \mid \mathbf{Y}, \mathbf{\Sigma}, \mathbf{\Psi})$ from the normal distribution. Since the large BAVR model in this paper involves 20 variables, the high-dimension density makes the computation process time-consuming. To improve the computational process, given the large number of parameters, we apply a Kronecker structure to the likelihood and use fast band matrix methods following Carriero, Clark and Marcellino (2016) and Chan (2020).

The number of iterations to obtain the posterior draws is set to 25,000, with the first 5000 draws discarded. To reduce the autocorrelation among the draws, only every 10^{th} draw is selected. To compute the generalized impulse response functions the mean responses of the 20 endogenous variables are kept from 100 draws successfully satisfying the imposed sign restrictions. For each period, 500 current states of the economy are drawn, and the representative impulse responses are taken to be the median of this distribution. The impulse response functions to the shocks from the specific periods are calculated using 1.5 million random draws of the rotation matrix.

References

[1] Aastveit, K. A. (2014). Oil price shocks in a data-rich environment. *Energy Economics*, 45, 268-279.

[2] Antolín-Díaz, J., and Rubio-Ramírez, J. (2018). Narrative sign restrictions for SVARs. *American Economic Review*, 108(10), 2802-2829.

[3] Bańbura, M., Giannone, D., and Reichlin, L. (2010). Large Bayesian vector autoregressions. *Journal of Applied Econometrics*, 25(1), 71-92.

[4] 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.

[5] Baumeister, C., and Kilian, L. (2016). Forty years of oil price fluctuations: Why the price of oil may still surprise us. *Journal of Economic Perspectives*, 30(1),139-160.

[6] Baumeister, C., Korobilis, D., and Lee, T. K. (2020). Energy markets and global economic conditions. *The Review of Economics and Statistics*, Forthcoming, 2020.

[7] Baumeister, C., Peersman, G., and van Robays. I. (2009). *The economic consequences of oil shocks: Differences across countries and time* (D/2009/7012/81). Retrieved from University Gent website: http://wps−feb.ugent.be/Papers/ wp_ 09_ 630.pdf.

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

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

[10] Carriero, A., Clark, T. E., and Marcellino, M. (2016). Common drifting volatility in large Bayesian VARs. *Journal of Business & Economic Statistics*, 34(3), 375-390.

[11] 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.

[12] Chan, J. C. C. (2020). Large Bayesian VARs: A flexible Kronecker error covariance structure. *Journal of Business & Economic Statistics*, 38(1), 68-79.

[13] Chan, J. C. C., and Eisenstat, E. (2015). Marginal likelihood estimation with the cross-entropy method. *Econometric Reviews*, 34(3), 256-285.

[14] Chan, J. C. C., and Grant, A. L. (2016). Fast computation of the deviance information criterion for latent variable models. *Computational Statistics & Data Analysis*, 100, 847-859.

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

[16] Cogley, T., and Sargent, T. J. (2005). Drifts and volatilities: Monetary policies and outcomes in the post WWII US. *Review of Economic Dynamics*, 8, 262-302.

[17] Fang, C. R., and You, S. Y. (2014). The impact of oil price shocks on the large emerging countries' stock prices: Evidence from China, India and Russia. *International Review of Economics & Finance*, 29, 330-338.

[18] Filippidis, M., Filis, G., and Kizys, R. (2020). Oil price shocks and EMU sovereign yield spreads. *Energy Economics*, 86.

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

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

[21] Hamilton, J. D., and Herrera, A. M. (2004). Oil shocks and aggregate macroeconomic behaviour: The role of monetary policy: Comment. *Journal of Money, Credit and Banking*, 36(2), 265-286.

[22] Hausman, J. A., and Newey, W. K. (1995). Non-parametric estimation of exact consumer's surplus and deadweight loss. *Econometrica*, 63, 1445-1476.

[23] Herrera, A. M., and Rangaraju, S. K. (2019). The effect of oil supply shocks on U.S. economic activity: What have we learned? *Journal of Applied Econometrics*, 1-19. https://doi.org/10.1002/jae.2735.

[24] Jebabli, I., Arouri. M., and Teulon, F. (2014). On the effects of world stock market and oil price shocks on food prices: An empirical investigation based on TVP-VAR models with stochastic volatility. *Energy Economics*, 45, 66-98.

[25] Kadiyala, K., and Karlsson, S. (1997). Numerical methods for estimation and inference in Bayesian VAR-models. *Journal of Applied Econometrics*, 12(70), 99-132.

[26] Karlsson, S. (2013). Forecasting with Bayesian vector autoregressions, in G. Elliott, and A. Timmermann (Eds.), *Handbook of Economic Forecasting* (pp. 791–897). Retrieved from https://doi.org/10.1016/B978-0-444-62731-5.00015-4

[27] Kapetanios, G., Marcellino, M., and Venditti. F. (2019). Large time-varying parameter VARs: A non-parametric approach. *Journal of Applied Econometrics*, 34(7), 1027-1049.

[28] Kilian, L. (2008). Exogenous oil supply shocks: How big are they and how much do they matter for the U.S. economy? *The Reviews of Economics*, 90(2), 216-240.

[29] 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.

[30] 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.

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

[32] Kilian, L. and Zhou, X. Q. (2020). *The Econometrics of Oil Market VAR Models.* CESifo working paper (No.8153).

[33] Kim, W. J., Hammoudeh, S., Hyun, J. S., and Gupta, R. (2017). Oil price shocks and China's economy: Reactions of the monetary policy to oil price shocks. *Energy Economics*, 62, 61-69.

[34] Koop, G. (2013). Forecasting with medium and large Bayesian VARS. *Journal of Applied Econometrics*, 28(2), 177-203.

[35] Koop, G., and Korobilis, D. (2010). Bayesian multivariate time series methods for empirical macroeconomics. *Foundations and Trends(R) in Econometrics*, 3(4), 267-358.

[36] Koop, G., and Korobilis, D. (2013). Large time-varying parameter VARs. *Journal of Econometrics*, 177(2), 185-198.

[37] Koop. G., Leon-Gonzalez. R., and Strachan. R. W. (2009). On the evolution of the monetary policy transmission mechanism. *Journal of Economic Dynamics and Control*, 33(4), 997-1017.

[38] Koop. G., Pesaran. M. H., and Potter. S. M. (1996). Impulse response analysis in nonlinear multivariate models. *Journal of Econometrics*, 74(1), 119-447.

[39] Lardic, S., and Mignon, V. (2008). Oil prices and economic activity: An asymmetric cointegration approach. *Energy Economics*, 30(3), 847-855.

[40] Melolinna, M. (2014). *What is the role of emerging Asia in global oil prices.* BOFIT Discussion paper (No. 18/2014).

[41] Nakajima, J., Kasuya, M., and Watanabe, T. (2011). Bayesian analysis of time-varying parameter vector autoregressive model for the Japanese economy and monetary policy. *Journal of the Japanese and International Economies*, 25(3), 225-245.

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

[43] Peersman, G., and van Robays, I. (2009). Oil and the euro area economy. *Economic Policy*, 24(60), 603-651.

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

[45] Primiceri, G. E. (2005). Time varying structural vector autoregressions and monetary policy. *Review of Economic Studies*, 72, 821-852.

[46] Ratti, R. A., and Vespignani, J. L. (2013). Why are crude oil prices high when global activity is weak? *Economics Letters*, 121(1), 133-136.

[47] Riggi, M., and Venditti, F. (2015). The time varying effect of oil price shocks on euro area exports. *Journal of Economic Dynamics and Control*, 59, 75-94.

[48] Stock, J. H., and Watson, M. W. (2016). Dynamic factor models, factor-augmented vector autoregressions, and structural vector autoregressions in macroeconomics. In J. B. Taylor, and H. Uhlig (Eds.), *Handbook of Macroeconomics* (pp.415-525). Elsevier.