RCEA Network

Rimini Centre for Economic Analysis
Working Paper Series

wp 22-15

On the identification of the oil-stock market relationship

Ioannis Arampatzidis
Theodore Panagiotidis

On the identification of the oil-stock market relationship

Ioannis Arampatzidis*† Theodore Panagiotidis‡

Abstract

The alternative identification techniques for oil market shocks could be responsible for the mixed results in the oil-stock market literature. This study employs a Bayesian Structural Vector Autoregression (SVAR) to compare the implications of traditional identification approaches (SVAR with zero/sign restrictions) with those from the baseline model (Bayesian SVAR) for the case of the US. We find that the baseline model implies more plausible posterior price elasticities of oil supply and demand and a more profound effect of oil supply shocks on oil prices. Nonetheless, all models provide qualitatively similar conclusions for the effects of oil market shocks on the US stock market, with shocks coming from the demand side playing a more important role than oil supply shocks. Overall, this study reveals that traditional identification schemes remain a good approximation in practice for the oil-stock market relationship.

Keywords: Bayesian SVAR; Identification; Oil market shocks; Stock market; US industries

JEL Classification: C32, Q43, G15

---

*We gratefully appreciate the very helpful comments and suggestions by the Editor Sushanta Mallick, an anonymous guest editor and two anonymous referees that significantly improved our paper. We would also like to thank Volker Clausen, the participants of the memorial event for Ansgar Belke, RGS Jamboree workshop and VfS annual conference for their helpful comments and suggestions. We dedicate this paper to the memory of Ansgar Belke.

†Department of Economics, University of Duisburg-Essen, Essen, Germany, & Ruhr Graduate School in Economics, Germany, email: Ioannis.Arampatzidis@vwl.uni-due.de

‡Department of Economics, University of Macedonia, Thessaloniki, Greece, email: tpanag@uom.edu.gr

1. Introduction

The seminal work of Kilian (2009) influenced the way the literature studies oil market shocks and their effects on commodity and financial markets. Since then, it has been widely accepted that different causes of oil market shocks have different effects on oil prices and thus on the economy. However, the model proposed by Kilian (2009), namely a SVAR with zero restrictions, has been criticized for its restrictive theoretical assumptions. Alternative approaches that try to relax the restrictive features of that model were developed. Baumeister and Peersman (2013) recommend replacing zero with sign restrictions, whereas Kilian and Murphy (2012, 2014) argue that sign restrictions alone are not enough, but must be further combined with specific bounds on the elasticities of demand and supply.

Another methodology has been introduced in the literature recently, namely a Bayesian SVAR model in which the identification is based on prior information in the form of prior distributions (Baumeister and Hamilton, 2019). The introduction of this model, combined with specific prior beliefs, sparked a debate in the literature. Baumeister and Hamilton (2019) criticize all previous attempts to model oil market shocks based either on zero or sign restrictions. They argue that such methodologies involve an *"all-or-nothing"* approach in the use of prior information since they treat some parameters as known with certainty and others as completely unknown. Their conclusions contradict some earlier findings, such as the posterior short-run price elasticities of oil supply and demand. In response, Kilian and Zhou (2018) and Kilian (2019) criticize many modelling choices in Baumeister and Hamilton (2019). To name only a few, these include the imposition of an unrealistically large value on the prior for the oil supply elasticity, the usage of fewer autoregressive lags than recommended in the literature, the use of pre-1973 oil market data and a controversial measurement error for oil inventories.

The different approaches to the oil market shock identification are reflected in the oil-stock market literature[1]. The majority of studies uses the SVAR model with zero restrictions. Kilian and Park (2009) are the first to employ this identification scheme to examine the effects of oil market shocks on the US stock market, both at

[1] Other applications include the relationship between oil market shocks and economic uncertainty (Degiannakis et al., 2018) and their effects on the macroeconomy in general (De et al., 2022).

aggregate and disaggregate level. They find that each identified oil market shock has a distinct impact on the stock market, with shocks from the demand side playing a more profound role than oil supply shocks. Many subsequent studies that consider different samples and examine both oil-importing and oil-exporting countries confirm the main findings of Kilian and Park (2009) (see among others, Wang et al., 2013; Güntner, 2014). There are also studies that draw conclusions which contradict some of the previous findings. For instance, Apergis and Miller (2009) conclude that stock returns show no reaction to oil market shocks, no matter whether they originate from the demand or the supply side. Arampatzidis et al. (2021) find that although aggregate and oil-specific demand shocks have in general a larger impact than oil supply shocks on 49 US industries, any statistically significant responses of the stock market are sporadic and time-dependent.

There are considerably fewer studies that employ SVAR models identified with sign restrictions. Gupta and Modise (2013) find that aggregate demand shocks have a positive but not persistent effect on stock prices, whereas oil supply shocks have a negative and larger impact than implied by the majority of models identified with zero restrictions. Koh (2017) reaches similar conclusions about oil supply shocks by examining the oil-stock market relationship in various Asian stock markets. Basher et al. (2018) further confirm the importance of oil supply shocks in some major oil-exporting countries. Overall, it seems that one difference between the two competing identification schemes is that oil supply shocks are viewed as more important in models identified with sign restrictions.

To the best of our knowledge, there are no studies so far that employ the Bayesian SVAR to examine the oil-stock market relationship in a systematic manner, with the exception of Güntner and Öhlinger (2021). However, their goal is different than ours, as they examine the comovement of oil prices and stock returns in response to structural oil market shocks. They use the model proposed in Baumeister and Hamilton (2019) with the inclusion of stock prices and they focus on the airline industry. In contrast, our study modifies the original model in the spirit of the recent critiques (Kilian and Zhou, 2018; Kilian, 2019; Herrera and Rangaraju, 2020), analyses the responses of stock returns to oil market shocks and examines a wider spectrum of the US stock market.

Given the ongoing debate in the oil market literature and since the Bayesian

SVAR has not been adequately tested yet, this study seeks to answer the following question: How important are the different identification schemes for the oil-stock market relationship in the US? We try to answer this question by comparing the main models that exist in the oil market literature for the identification of oil market shocks, namely (i) SVARs with zero restrictions (Kilian, 2009); (ii) SVARs with sign restrictions (Kilian and Murphy, 2012); and (iii) the Bayesian SVAR (Baumeister and Hamilton, 2019). One convenient way to carry out the empirical analysis, which also offers a direct comparison between the models in a unified framework, is to work with the Bayesian SVAR. Therefore, this Bayesian SVAR is both our baseline model and as Baumeister and Hamilton (2019) show, it also provides a general framework that allows the comparison of alternative models.

Our baseline model is based on the Bayesian SVAR model proposed by Baumeister and Hamilton (2019) and is modified according to economic theory and the recent critiques (Kilian and Zhou, 2018; Kilian, 2019; Herrera and Rangaraju, 2020). More specifically, (i) we introduce an alternative prior distribution for the oil supply elasticity; (ii) we propose modifications to the price and income elasticities of oil demand; (iii) we make a more conservative choice for the lag structure; (iv) we do not consider pre-1973 data as they are deemed questionable; and (v) we do not model the measurement error in global oil inventories as it tends to create more problems than it actually solves[2].

The contribution of this paper is threefold. First, to the best of our knowledge, this is the first time that a Bayesian SVAR is used in order to examine the oil-stock market relationship, both at aggregate and disaggregate level. Second, we propose the use of an exponential prior distribution for the oil supply elasticity. Its advantages compared to the student t-distribution (Baumeister and Hamilton, 2019) and the upper bound in SVAR models identified with sign restrictions (Kilian and Murphy, 2012) are discussed in Sections 3.3.1 and 4.1. Third, the generality of this model allows us to examine different identification techniques and their implications for the US stock market. Our analysis shows that all models yield qualitatively similar conclusions for the oil-stock market relationship. This implies that traditional methods remain a good approximation in this research area.

Our study is in the spirit of Herrera and Rangaraju (2020), but it differs along

[2]For more information see our discussion in Section 3.3.4.

three dimensions. Herrera and Rangaraju (2020) focus on the effects of oil market shocks on US GDP by testing different identification schemes, whereas we examine the oil-stock market relationship. Another important difference is that, since GDP is available only at quarterly frequency, their methodology involves a two-stage process. In the first stage they use each SVAR in order to identify the structural oil market shocks, while in the second stage they examine their effects on GDP using OLS regressions. In contrast, our study involves a single stage, as the stock market variable is included in the Bayesian SVAR from the start of the analysis. Finally, they only analyze the impact of oil supply shocks on US GDP, while we discuss the different effects of all identified oil market shocks on the US stock market.

The remainder of the paper is organized as follows: Section 2 presents the data and Section 3 describes the methodology, the technical and theoretical differences between the models as well as the reasons behind the modifications we propose to the original model of Baumeister and Hamilton (2019). Section 4 presents the results, while Section 5 concludes.

2. Data

Our dataset consists of four global oil market variables as well as stock prices for the aggregate US stock market and selected US industries for the period 1973:01-2019:12[3]. Figure 1 depicts the raw data of the oil market variables, whereas Table B.1 in the online Appendix B provides descriptive statistics for the stock returns.

As a proxy for global oil supply (q_t^s), we employ monthly world oil production data measured in million barrels of oil pumped per day. To proxy global real economic activity (y_t), we use the industrial production index constructed by Baumeister and Hamilton (2019)[4]. To get the global real price of crude oil (p_t), we deflate the US refiner's acquisition cost of imported crude oil with the US CPI. Our estimate of global oil inventories is obtained by multiplying the US crude oil inventories

[3] We use all variables in first logarithmic differences. A list of data sources is given in the online Appendix A.

[4] We prefer this conventional measure of industrial production over the widely used real economic activity index constructed by Kilian (2009) as it allows us to use previous empirical evidence on the income elasticity of oil demand. For more information, see the online Appendix in Baumeister and Hamilton (2019), available at `https://www.aeaweb.org/content/file?id=9558`.

by the ratio of OECD to US inventories of petroleum and petroleum products, as in Baumeister and Hamilton (2019). We then take the change in OECD oil inventories as a fraction of last period's oil production. This provides our final proxy, namely the change in global oil inventories (ΔI_t).

Figure 1: Global oil market variables

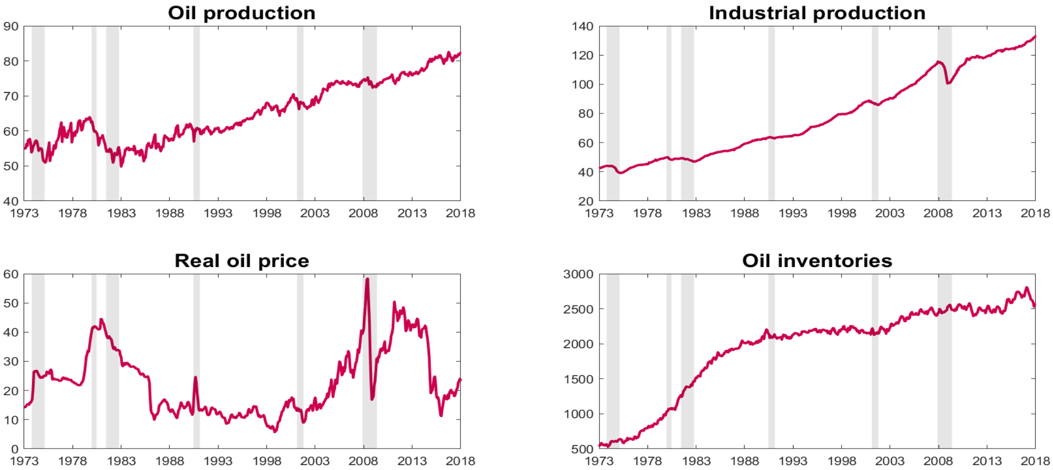

Note: (i) The shaded areas denote the NBER defined recession periods; (ii) oil production is measured in million of barrels/day; industrial production is an index; real oil price is the ratio of the nominal oil price (measured in dollars per barrel) and the US CPI; oil inventories are measured in million of barrels.

Finally, we use an US stock market index deflated by the US CPI to proxy the aggregate US stock market (ret_t). We further examine the effects of oil market shocks on four US industries that are expected to be more vulnerable to oil price changes. We start with the Petroleum & Natural Gas industry since oil plays a vital role for this sector. We further examine the Automobiles & Trucks industry, as its activity is directly affected by the level of oil prices. In addition, we consider the Precious Metals industry because of the widely held view that in times of uncertainty investors resort to precious metals, such as gold, which puts upward pressure on their share prices. Finally, we also include the Retail sector, given the view that higher oil prices hurt this sector as they lower the real disposable income of households.

3. Econometric Methodology

To model oil market shocks, the following equations are usually taken into account:

$$q_t^s = \alpha_{qy} y_t + \alpha_{qp} p_t + b_1^{'} x_t + u_{1t} \tag{1}$$

$$y_t = \alpha_{yq} q_t^s + \alpha_{yp} p_t + b_2^{'} x_t + u_{2t} \tag{2}$$

$$p_t = \alpha_{pq} q_t^d + \alpha_{py} y_t + b_3^{'} x_t + u_{3t} \tag{3}$$

Eq. (1) is the oil supply curve, Eq. (2) gives economic activity as a function of oil production and oil prices, Eq. (3) is the oil demand curve, written here in inverse form. All α_{ij} parameters (for $i, j = q, y, p$) capture the contemporaneous relationships between the variables, whereas b_j (for $j = 1, 2, 3$) give the corresponding effects with a lag. Note that the quantities of oil produced (q_t^s) and consumed (q_t^d) at time t are taken to be identical ($q_t^s \equiv q_t^d \equiv q_t$), an assumption that is relaxed in the baseline model (see Section 3.3).

As our goal is to examine the effects of oil market shocks on the US stock market, an additional equation is necessary:

$$ret_t = \chi_1 q_t^s + \chi_2 y_t + \chi_3 p_t + b_4^{'} x_t + u_{4t} \tag{4}$$

Note that, although we allow the oil market variables to contemporaneously affect stock returns, we assume that the inverse does not hold. This is a standard assumption in the literature (see e.g. Kilian and Park, 2009; Wang et al., 2013) that takes the global oil market as predetermined.

The above equations can be written in compact form:

$$\boldsymbol{A} z_t = \boldsymbol{B} x_t + u_t \tag{5}$$

where $z_t = (q_t^s, y_t, p_t, ret_t)^{'}$ is a $(n \times 1)$ vector of the observed variables, x_t is a $(k \times 1)$ vector of the lagged variables and a constant $(x_t^{'} = (z_{t-1}^{'}, ..., z_{t-m}^{'}, 1)^{'}, k = mn + 1)$, u_t is a $(n \times 1)$ vector of structural disturbances with variance-covariance matrix ($\boldsymbol{D}$) assumed to be diagonal, $\boldsymbol{B}$ is a $(n \times k)$ matrix of autoregressive parameters, and $\boldsymbol{A}$ is a $(n \times n)$ matrix of the contemporaneous coefficients.

3.1. SVAR with zero restrictions

Kilian (2009) is the first to employ a SVAR model to disentangle the causes of oil price changes, whereas Kilian and Park (2009) augment that model to further examine the effects of oil market shocks on the stock market. For identification, both studies rely on the familiar Cholesky decomposition, which implies:

$$\boldsymbol{A} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ -\alpha_{yq} & 1 & 0 & 0 \\ -\alpha_{pq} & -\alpha_{py} & 1 & 0 \\ -\chi_1 & -\chi_2 & -\chi_3 & 1 \end{bmatrix} \tag{6}$$

This further implies that we act as if we know with certainty that the contemporaneous effect of stock returns on oil market variables (4th column) is zero as well as that $\alpha_{qy} = \alpha_{qp} = \alpha_{yp} = 0$. At the same time we assume that we know nothing about the remaining parameters in Eq. (6). Baumeister and Hamilton (2019) show that we can give a Bayesian interpretation to this traditional approach by assuming extremely flat priors for the non-zero parameters in Eq. (6), as shown in Table 1.

Table 1: Prior distributions for parameters in matrix $\boldsymbol{A}$ - SVAR with **zero restrictions**

Parameter	Meaning	Location	Scale	Degrees of freedom	Skew	Sign restriction
		Student t-distribution				
α_{yq}	Effect of q on economic activity	0	100	3	-	none
α_{pq}	Reciprocal of the price elasticity of oil demand	0	100	3	-	none
α_{py}	Reciprocal of the income elasticity of oil demand	0	100	3	-	none
χ_1	Effect of q on stock returns	0	100	3	-	none
χ_2	Effect of y on stock returns	0	100	3	-	none
χ_3	Effect of p on stock returns	0	100	3	-	none

3.2. SVAR with sign restrictions

The recursive identification scheme employed in Kilian and Park (2009) might be viewed as restrictive. Kilian and Murphy (2012) and Baumeister and Peersman (2013) replace the zero restrictions with sign restrictions on the impact multiplier matrix:

$$\boldsymbol{H} = \boldsymbol{A}^{-1} = \begin{bmatrix} + & + & + & 0 \\ + & + & - & 0 \\ - & + & + & 0 \\ ? & ? & ? & + \end{bmatrix} \tag{7}$$

These restrictions are based on the following assumptions: (i) a positive oil supply shock raises oil production and economic activity, but decreases oil prices (1st column); (ii) a positive economic activity shock leads to higher oil production, economic activity and oil prices (2nd column); (iii) an increase in oil-specific demand has a positive impact on oil production and oil prices, but depresses economic activity (3rd column); (iv) stock market shocks do not have an immediate impact on oil market variables (4th column), whereas the response of the stock market to oil market shocks on impact is ambiguous (4th row).

Baumeister and Hamilton (2019) show that in order to satisfy those sign restrictions we need the theoretically plausible assumptions that: i) the slope of the oil supply curve is positive ($\alpha_{qp} > 0$); ii) the slope of the oil demand curve is negative ($\alpha_{pq} < 0$); iii) a higher oil price decreases economic activity ($\alpha_{yp} < 0$); iv) higher income increases oil demand and in turn the price of oil ($\alpha_{py} > 0$). In addition, we need to assume that $\alpha_{qy} = \alpha_{yq} = 0$, which means that matrix $\boldsymbol{A}$ takes the form:

$$\boldsymbol{A} = \begin{bmatrix} 1 & 0 & -\alpha_{qp} & 0 \\ 0 & 1 & -\alpha_{yp} & 0 \\ -\alpha_{pq} & -\alpha_{py} & 1 & 0 \\ -\chi_1 & -\chi_2 & -\chi_3 & 1 \end{bmatrix} \tag{8}$$

Kilian and Murphy (2012) argue that these restrictions alone are not sufficient since the admitted models might imply unrealistic values for the price elasticity of oil supply (α_{qp}) and the equilibrium effect of oil-specific demand shocks on economic activity ($h_{23} = \boldsymbol{H}(2,3)$). To circumvent this, they propose specific bounds for these parameters based on historical evidence. Those assumptions can again be given a Bayesian interpretation once we use priors specified as in Table 2.

Table 2: Prior distributions for parameters in matrix $\boldsymbol{A}$ - SVAR with **sign restrictions**

Parameter	Meaning	Bounds				Sign restriction
		Uniform distribution				
α_{qp}	Oil supply elasticity	[0, 0.0258]				positive
h_{23}	Effect of oil-specific demand shock on economic activity	[-1.5, 0]				negative
Parameter	**Meaning**	**Location**	**Scale**	**Degrees of freedom**	**Skew**	**Sign restriction**
		Student t-distribution				
α_{yp}	Effect of p on economic activity	0	100	3	-	negative
α_{pq}	Reciprocal of the price elasticity of oil demand	0	100	3	-	negative
α_{py}	Reciprocal of the income elasticity of oil demand	0	100	3	-	positive
χ_1	Effect of q on stock returns	0	100	3	-	none
χ_2	Effect of y on stock returns	0	100	3	-	none
χ_3	Effect of p on stock returns	0	100	3	-	none

3.3. Bayesian SVAR

This section follows Baumeister and Hamilton (2019) and builds a model that uses prior information in a more intuitive way. The first two equations in our baseline model are identical to those used in traditional approaches, namely Eq. (1) and (2). Instead of using the oil demand curve in its inverse form (Eq. (3)), we employ it in its standard form:

$$q_t^d = \beta_{qy} y_t + \beta_{qp} p_t + b_3^{'} x_t + u_{3t} \tag{9}$$

This allows us to use prior information on the income (β_{qy}) and price (β_{qp}) elasticities of oil demand.

Kilian and Murphy (2014) were the first to notice that an additional factor is missing from the system. More specifically, there might be a difference between the quantity of oil produced (q_t^s) and consumed (q_t^d) at time t. This surplus/deficit in the amount of oil is then captured by changes in oil inventories, denoted by ΔI_t:

$$q_t^s - q_t^d = \Delta I_t \tag{10}$$

We thus obtain a refined version of the oil demand curve augmented with oil inventories, by plugging Eq. (10) into Eq. (9):

$$q_t^s = \beta_{qy} y_t + \beta_{qp} p_t + \Delta I_t + b_3^{'} x_t + u_{3t} \tag{11}$$

In addition, Eq. (12) shows how each variable in the system affects oil inventories:

$$\Delta I_t = \psi_1 q_t^s + \psi_2 y_t + \psi_3 p_t + b_4^{'} x_t + u_{4t} \tag{12}$$

This allows one to identify an additional oil market shock (u_{4t}), which is often described as speculative or inventory demand shock.

Finally, the last equation in our system captures the impact of oil market variables on the stock market:

$$ret_t = \chi_1 q_t^s + \chi_2 y_t + \chi_3 p_t + \chi_4 \Delta I_t + b_5^{'} x_t + u_{5t} \tag{13}$$

This completes the description of the baseline model, which consists of Eq. (1), (2), (11), (12) and (13), such that matrix $\boldsymbol{A}$ takes the form:

$$\boldsymbol{A} = \begin{bmatrix} 1 & -\alpha_{qy} & -\alpha_{qp} & 0 & 0 \\ -\alpha_{yq} & 1 & -\alpha_{yp} & 0 & 0 \\ 1 & -\beta_{qy} & -\beta_{qp} & -1 & 0 \\ -\psi_1 & -\psi_2 & -\psi_3 & 1 & 0 \\ -\chi_1 & -\chi_2 & -\chi_3 & -\chi_4 & 1 \end{bmatrix} \tag{14}$$

As in models identified with sign restrictions, we are also interested in the signs in $\boldsymbol{H} = \boldsymbol{A}^{-1}$. To be in line with economic theory, it is necessary that the elements in $\boldsymbol{H}$ have signs as in Eq. (15):

$$\boldsymbol{H} = \boldsymbol{A}^{-1} = \begin{bmatrix} + & + & + & + & 0 \\ + & + & - & - & 0 \\ - & + & + & + & 0 \\ ? & ? & ? & + & 0 \\ ? & ? & ? & ? & + \end{bmatrix} \tag{15}$$

To satisfy those sign restrictions we follow Baumeister and Hamilton (2019) and set $\alpha_{qy} = \alpha_{yq} = \psi_2 = 0$. Thus, matrix $\boldsymbol{A}$ takes the new form[5]:

$$\boldsymbol{A} = \begin{bmatrix} 1 & 0 & -\alpha_{qp} & 0 & 0 \\ 0 & 1 & -\alpha_{yp} & 0 & 0 \\ 1 & -\beta_{qy} & -\beta_{qp} & -1 & 0 \\ -\psi_1 & 0 & -\psi_3 & 1 & 0 \\ -\chi_1 & -\chi_2 & -\chi_3 & -\chi_4 & 1 \end{bmatrix} \tag{16}$$

3.3.1. Priors for matrix A

Table 3 summarizes the priors used in our baseline model. We start with the oil supply elasticity (α_{qp}), the coefficient that sparked a debate in the literature. The traditional approach is to assign a zero value to this elasticity (see Section 3.1),

[5]One might argue that the system in Eq. (16) is under-identified and that this poses a threat to the identification of oil market shocks and the estimation of the model in general. This is not an issue in this Bayesian SVAR (for more details see the discussion in Baumeister and Hamilton, 2015), but rather an advantage compared to traditional models that must be just-identified.

which implies a vertical short-run oil supply curve, usually justified by the considerable short-run adjustment costs in the production of oil. Nevertheless, subsequent research tried to relax this restriction by imposing bounds on this elasticity (see Section 3.2). Baumeister and Hamilton (2019) criticize both approaches by arguing that prior information is used in unappealing ways. They instead propose a student t prior distribution, which implies a posterior median of 0.15, about six times larger than the upper bound used in Kilian and Murphy (2012) (see Table 2).

The estimates of this elasticity vary in the literature, with most studies suggesting values close to zero. Anderson et al. (2018) show that oil production does not respond to oil prices in the short run, implying an oil supply elasticity that is essentially zero. Kilian and Murphy (2012) arrive at an upper bound of 0.0258 by focusing on the period during the invasion of Iraq in Kuwait. Using the same data, but accounting for potential confounding events, Caldara et al. (2019) find a higher upper bound, namely 0.043. Furthermore, by analyzing other historical events, they find an oil supply elasticity of 0.077, with a 0.037 standard error. Given previous empirical evidence and economic theory, we propose an exponential prior distribution for α_{qp}, with rate equal to 50 (i.e. mean equal to 0.02). The mass of this distribution is placed close to zero, with 72.5% of the points being below 0.0258, i.e. the upper bound proposed by Kilian and Murphy (2012). At the same time though, our prior distribution allows α_{qp} to take higher values with decreasing probability.

For the remaining oil market parameters, we follow Baumeister and Hamilton (2019) with some modifications. We assume a truncated student t-distribution for α_{yp}, with mode at -0.05, scale 0.1 and 3 degrees of freedom. The reason is that the dollar share of crude oil expenditures is small compared to total GDP, thus the effect on economic activity is expected to be small in the short run. Although small, it is expected to be negative given the view that rising oil prices are associated with decreasing economic activity (see e.g. Bernanke et al., 1997; Hamilton, 2011).

Baumeister and Hamilton (2019) assume a student t-distribution for β_{qy}, with location parameter 0.7, scale 0.2 and 3 degrees of freedom. Their choice is based on previous empirical evidence, with most estimates being around 0.7 (see e.g. Gately and Huntington, 2002; Csereklyei et al., 2016). However, this refers to studies that examine the long-run income elasticity of oil demand. Other studies that investigate the short-run elasticity find values in the range of 0.1 - 0.5. Huntington et al. (2019),

Table 3: Prior distributions for parameters in matrix $\boldsymbol{A}$ - **Baseline model**

Parameter	Meaning	Location	Scale	Degrees of freedom	Skew	Sign restriction
			Exponential distribution			
α_{qp}	Oil supply elasticity	0.02	-	-	-	positive
			Student t-distribution			
α_{yp}	Effect of p on economic activity	-0.05	0.1	3	-	negative
β_{qy}	Income elasticity of oil demand	0.3	0.15	3	-	positive
β_{qp}	Price elasticity of oil demand	-0.1	0.1	3	-	negative
ψ_1	Effect of q on oil inventories	0	0.5	3	-	none
ψ_3	Effect of p on oil inventories	0	0.5	3	-	none
χ_1	Effect of q on stock returns	0	100	3	-	none
χ_2	Effect of y on stock returns	0	100	3	-	none
χ_3	Effect of p on stock returns	0	100	3	-	none
χ_4	Effect of ΔI on stock returns	0	100	3	-	none
h_2	Effect of economic activity shock on y	0.8	0.2	3	-	none
			Asymmetric t-distribution			
h_1	Determinant of $\boldsymbol{A}$	0.2	0.9	3	2	none
			Gamma distribution			
d_{ii}^{-1}	Reciprocal of variance	$1/(a_i'\hat{\boldsymbol{S}}a_i)$	$1/(\sqrt{2}a_i'\hat{\boldsymbol{S}}a_i)$	-	-	positive
			Normal distribution			
b_i	Lagged autoregressive coefficients	0	M	-	-	none

based on a literature review, conclude that the average short-run income elasticity is equal to 0.39. Javan and Zahran (2015), using panel methods, find estimates between 0.1 and 0.47, depending on the country group. Agrawal (2015), using an ARDL model that allows for error correction, estimates a short-run elasticity around 0.1, with a standard error equal to 0.04. Since we are interested in the short-run rather than the long-run income elasticity of oil demand, our new prior has mode at 0.3, scale 0.15 and 3 degrees of freedom. This allows β_{qy} to take values within 0.1 - 0.5 with approximately 78% probability.

Numerous studies examine the price elasticity of oil demand (β_{qp}). Based on a literature survey, Hamilton (2009) finds an average short-run elasticity approximately equal to -0.06. Javan and Zahran (2015) estimate elasticities between -0.18 and 0, whereas Gelman et al. (2016) find an elasticity of -0.22, with a standard error of 0.05. Based on the rich empirical literature, Baumeister and Hamilton (2019) choose a student t prior with location parameter -0.1, scale 0.2 and 3 degrees of freedom. This distribution has approximately 71% of its mass below -0.3. Since most empirical estimates range between -0.3 and 0, we slightly modify their prior by using a scale parameter of 0.1. This puts 91% probability mass within this interval, while at the same time it allows for larger elasticities with smaller probability.

For the remaining coefficients of matrix $\boldsymbol{A}$, we use relatively uninformative priors. These include student t-distributions for ψ_1 and ψ_3, with 0 mode, 0.5 scale and 3 degrees of freedom, and for χ_1, χ_2, χ_3 and χ_4 0 mode, 100 scale and 3 degrees of freedom. Furthermore, in order to ensure that the sign restrictions in Eq. (15) hold, we need to make a couple of additional assumptions. More specifically, $\boldsymbol{H} = \boldsymbol{A}^{-1} = (1/det(\boldsymbol{A}))\boldsymbol{C}$, where $\boldsymbol{C}$ is the adjoint matrix of $\boldsymbol{A}$. No matter what signs the elements of $\boldsymbol{C}$ will have, it all comes to the sign of the determinant (h_1). If it is negative, everything can happen. Therefore, we must ensure that h_1 is positive. Following Baumeister and Hamilton (2019), we assume an asymmetric t-distribution, which allows us to put as much weight as we want on the probability that the sign of h_1 is positive. Our parameter choices (see Table 3) imply a roughly 90% probability that h_1 is larger than zero.

Finally, even if all our assumptions simultaneously hold, namely even if $\alpha_{qy} = \alpha_{yq} = \psi_2 = 0$ and $\alpha_{qp}, \beta_{qy}, h_1 > 0$, $\alpha_{yp}, \beta_{qp} < 0$, some of the equilibrium feedback effects in Eq. (15) might not have the desirable signs. This is the case for $h_2 = \boldsymbol{H}(2,2)$. In order to reflect our belief that this coefficient is positive, we assume a student t-distribution centered at 0.8, with scale 0.2 and 3 degrees of freedom. This implies an approximately 98% probability that this coefficient is positive.

3.3.2. Priors for matrices D and B

Table 3 also contains the priors for the matrices $\boldsymbol{D}$ and $\boldsymbol{B}$. Starting with matrix $\boldsymbol{D}$, it is common to specify a prior that reflects in part the scale of the data. We achieve this by assuming that the elements of this matrix, conditional on $\boldsymbol{A}$, follow a Gamma distribution, $d_{ii}^{-1}|\boldsymbol{A} \sim \Gamma(\kappa_i, \tau_i(\boldsymbol{A}))$. We set the mean and the scale of the prior equal to $1/(a_i^{'}\hat{\boldsymbol{S}}a_i)$ and $1/(\sqrt{\kappa_i}a_i^{'}\hat{\boldsymbol{S}}a_i)$, respectively[6]. $\hat{\boldsymbol{S}}$ denotes the variance-covariance matrix of the univariate residuals of an AR(m) fit to each variable in our model. Moreover, as in Baumeister and Hamilton (2019), we set $\kappa_i = 2$.

We assume that the autoregressive parameters in matrix $\boldsymbol{B}$ follow conditional

[6]Note that the mean of this Gamma distribution is $\kappa_i/\tau_i(\boldsymbol{A}) \overset{(C.2)}{=} \kappa_i/\kappa_i(a_i^{'}\hat{\boldsymbol{S}}a_i) = 1/(a_i^{'}\hat{\boldsymbol{S}}a_i)$, whereas the variance is given by $\kappa_i/[\tau_i(\boldsymbol{A})]^2 \overset{(C.2)}{=} 1/[\kappa_i(a_i^{'}\hat{\boldsymbol{S}}a_i)^2]$, thus the scale is $1/(\sqrt{\kappa_i}a_i^{'}\hat{\boldsymbol{S}}a_i)$.

Gaussian distributions, $b_i|\boldsymbol{A},\boldsymbol{D} \sim N(m_i, d_{ii}\boldsymbol{M}_i)$. Since we use all variables in first differences, we assume that all lagged coefficients are zero. We can put as much weight as we want on this prior belief. Following Doan et al. (1984), we have more confidence that the coefficients at higher lags are zero. In practice, we can impose this by using Eq. (C.4) - (C.7), as described in the online Appendix C.1. In the baseline model, we use $\lambda_0 = 0.5$, $\lambda_1 = 1$ and $\lambda_3 = 100$, which imply a relatively high confidence that the coefficients of higher lags are zero and a diffuse prior for the constant term[7].

3.3.3. Posterior Distributions

Upon observation of the data (Y_T), the joint prior distribution (Eq. (C.8)) takes the form of the joint posterior distribution:

$$p(\boldsymbol{A},\boldsymbol{D},\boldsymbol{B}|Y_T) = p(\boldsymbol{A}|Y_T)p(\boldsymbol{D}|\boldsymbol{A},Y_T)p(\boldsymbol{B}|\boldsymbol{A},\boldsymbol{D},Y_T) \tag{17}$$

which is the product of the individual posteriors. Eq. (17) summarizes all our uncertainty after observing the data. Since we use conjugate priors for matrices $\boldsymbol{D}$ and $\boldsymbol{B}$, the posterior distributions are known. The online Appendix C.2 provides the formulas.

Next, we can use a Metropolis-Hastings algorithm to generate one million draws (after an one million burn-in sample) from the joint posterior distribution (Eq. (17)). For more details about the algorithm, the reader is referred to the online Appendix in Baumeister and Hamilton (2019).

3.3.4. Additional Modeling Choices

This section explains a set of additional modeling choices. We choose 24 autoregressive lags (m) in all models, which differentiates us from Baumeister and Hamilton (2019) who use 12. Kilian and Zhou (2020) argue that we need at least two years of lags because of the existence of slowly building and declining cycles in global commodity markets. Hence, models with fewer lags tend to underestimate the importance of economic activity shocks. It is important to note though that the use

[7]For a more detailed description of the priors for matrices $\boldsymbol{D}$ and $\boldsymbol{B}$, the reader is referred to the online Appendix C.1.

of a relatively tight prior for matrix $\boldsymbol{B}$, necessary to achieve parsimony, somewhat weakens this argument.

We further deviate from Baumeister and Hamilton (2019) regarding the starting date of our sample and the inclusion of a measurement error in oil inventories. The authors use US oil market data over the period 1958:01-1972:12 to further inform their priors. Although there was a considerable structural break in that period, they argue that it is always optimal to use down-weighted pre-break data as prior information. However, it is widely accepted that in the pre-1973 era the US oil market was heavily regulated, which implies limited interdependence. Furthermore, Alquist et al. (2013) argue that pre-1973 oil price data do not allow for an autoregressive fit, which makes their use as a prior inappropriate. For this reason, we stick to the majority of empirical studies and choose January 1973 as our starting point.

Regarding the measurement error in oil inventories, Baumeister and Hamilton (2019) argue that the available data are an imperfect estimate of the true value[8]. Therefore, they propose the inclusion of a measurement error equation in the model. Kilian and Zhou (2018) argue that their approach is questionable. The main problem is that their error specification is time invariant. Thus, it cannot tackle the main source of the error, namely the rise in Chinese strategic oil inventories after 2010, which is not captured by the oil inventories proxy. Hence, it is not necessarily a good practice to include a questionable measurement error specification. For this reason, we again follow the majority of earlier studies that do not model measurement errors in oil inventories.

4. Empirical Results

4.1. Oil Market

Figure 2 presents the prior and posterior distributions of the oil supply elasticity and the price and income elasticities of oil demand of each model[9]. Panel A shows the oil supply elasticity. Since the SVAR model with zero restrictions assumes that the value of this elasticity is exactly zero, both the prior and posterior values are zero,

[8]For a detailed discussion of the measurement issues in oil inventories see Kilian and Lee (2014).

[9]Figures B.1 to B.3 in the online Appendix B present the remaining parameters of matrix $\boldsymbol{A}$.

thus there is no plot of the oil supply elasticity in Figure 2 for this model. On the other hand, models identified with sign restrictions bound this elasticity between 0 and 0.0258. Justifying the use of such a prior distribution (see Panel (A,2)) is difficult, since there is no valid reason to regard values below 0.0258 as perfectly plausible and values slightly larger than this upper bound as implausible.

Baumeister and Hamilton (2019) propose a truncated student t prior distribution (see Panel (A,3)). This specification goes to the other extreme though as it puts 94% probability mass above 0.0258, thereby allowing for unrealistically large values with high probability. On the one hand, we need a prior that relaxes the upper bound suggested by Kilian and Murphy (2012). On the other hand, the prior must be specified in a way that it places the mass close to zero and at the same time it allows for higher values with decreasing probability. This is exactly what our exponential prior (see Panel (A,1)) achieves, which marks one of the contributions of our study. The implied posterior median oil supply elasticity in our model is approximately 0.05, three times lower than the one in Baumeister and Hamilton (2019) and close to the recent estimates in Caldara et al. (2019).

Figure 2: Prior and posterior distributions of important elasticities

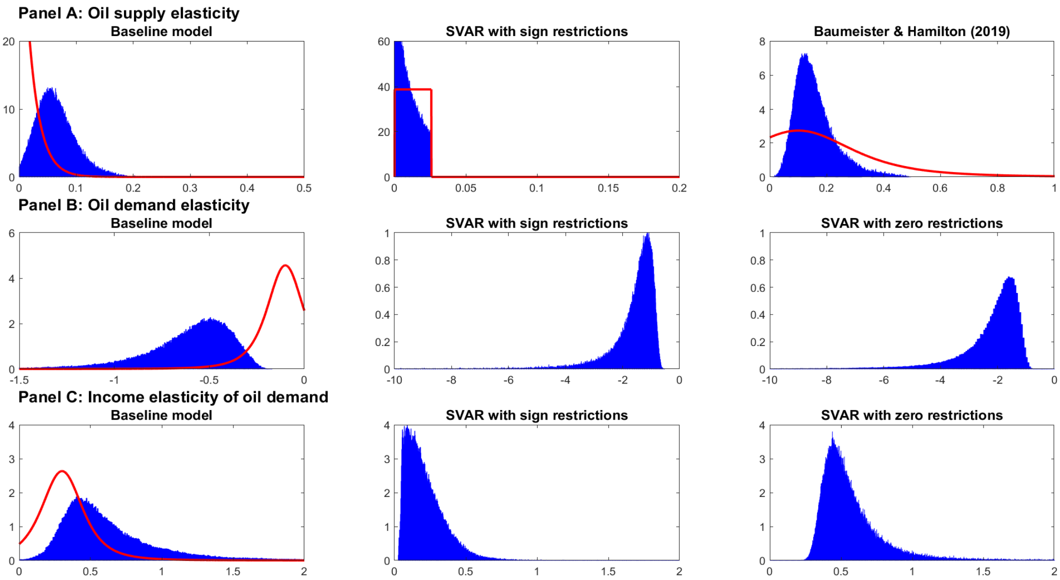

Note: Red lines: prior distributions; Blue histograms: posterior distributions.

Panel B depicts the price elasticity of oil demand. One undesirable feature of traditional SVAR models is the implication of a very elastic demand curve (Baumeister and Hamilton, 2019). More specifically, the implied posterior short-run price elasticities of oil demand (see Panels (B,2) and (B,3)) take values from -0.6 up to -8, with their mass being within -1 and -3. These values are significantly larger than most credible microeconometric estimates, which can be found in the range of -0.4 and 0[10]. Although the posterior median in our baseline model (see Panel (B,1)) is larger than our prior, it is considerably lower than the corresponding estimates in traditional approaches and much closer to the range suggested by previous microeconometric studies. This is another desirable feature of our baseline model.

Finally, Panel C presents the income elasticity of oil demand. We observe that although traditional approaches do not make any use of prior information (the priors for α_{py} in both models are flat, see Figures B.1 and B.2), their posterior elasticities are close to the one implied by our baseline model as well as previous empirical findings. Somewhat surprisingly, the mass of the posterior distribution of the model with sign restrictions is between 0.1 and 0.5, i.e. the range suggested by previous microeconometric estimates. On the other hand, the posterior distribution in our baseline model takes slightly larger values on average, which mimic the long-run estimates as well as the results in Baumeister and Hamilton (2019).

Figure 3 shows the impulse response functions of oil prices to oil market shocks in all models[11]. An oil supply disruption leads to an increase in oil prices in all cases. The responses are mostly insignificant in SVARs with zero and sign restrictions, with statistical significance found only for a short period of approximately six months. This observation is consistent with the previous findings (Kilian, 2009; Kilian and Murphy, 2012; Baumeister and Hamilton, 2019). In contrast, the response is larger and statistically significant at all forecast horizons in our baseline model. This shows that the effect of oil supply shocks on oil prices seems to be underestimated in traditional approaches. A potential explanation is related to the differences in the price elasticity of oil demand. More specifically, all models assume a very small oil supply elasticity, which implies a vertical oil supply curve in the model with zero restrictions ($\alpha_{qp} = 0$) and very steep curves in the other two models ($\alpha_{qp} < 0.1$). Under such circumstances, the response of oil price to an oil supply disruption (a shift of the oil supply curve to the left) depends to a large extent on the steepness of the oil demand curve. Our discussion in Figure 2 shows that traditional models

[10]See, among others, Javan and Zahran (2015) and Gelman et al. (2016).

[11]Figures B.4 to B.6 in the online Appendix B present the impulse response functions of the remaining oil market variables in all models.

imply a very elastic demand curve, whereas oil demand in our baseline model is price inelastic. As a consequence, the oil price is more responsive to oil supply shocks in our baseline model.

Figure 3: Impulse response functions of oil prices to oil market shocks

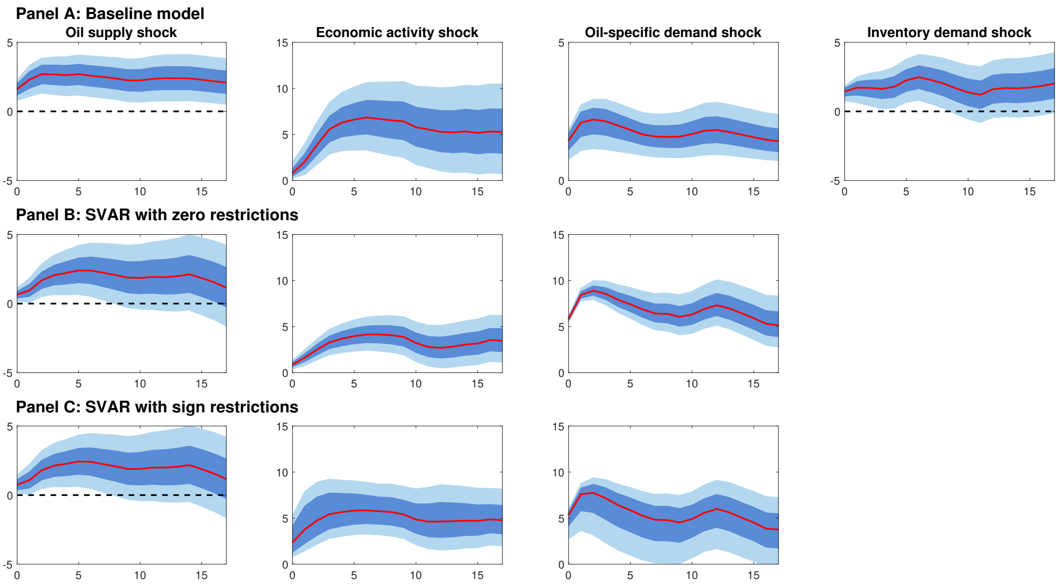

Note: (i) Red lines: posterior median; Dark blue shaded areas: 68% posterior credible sets; Light blue shaded areas: 95% posterior credible sets; (ii) all shocks are normalized such that to imply an increase in the price of oil; (iii) please note the different scaling of the graphs in the third column.

The second column in Figure 3 shows the response of oil prices to economic activity shocks. In general, we observe a similar pattern in all models, namely a positive, statistically significant and persistent reaction of oil prices. An oil-specific demand shock (third column) also causes a positive and persistent increase in the price of oil. The magnitude of the response though is lower in our baseline model compared to traditional approaches[12]. Finally, an inventory demand shock leads to a positive and statistically significant reaction of oil prices, which loses significance after the ninth month.

Overall, our analysis reveals both similarities and differences between the models. The most striking disagreements concern the implied posterior supply and demand elasticities. The discussion of Figure 2 reveals that these elasticities are more realistic and closer to previous microeconometric estimates in our baseline model,

[12] The reason is that in our baseline model we extract one factor from the residual oil-specific demand shock, namely the inventory demand shock.

which is also the main point of the critique of traditional approaches raised by Baumeister and Hamilton (2019). Another important difference concerns the effect of oil supply shocks on oil prices, which is found to be larger on average and much more persistent in our baseline model (see Figure 3). The next section examines whether those differences affect in turn the oil-stock market relationship.

4.2. Stock Market

Figures 4 to 6 show the impulse response functions of stock returns to oil market shocks in all models[13]. Figure 4 presents the impulse responses in the model identified with zero restrictions. Oil supply disruptions have no discernible impact on stock returns both at aggregate and disaggregate level. This is not a surprise since in such models the effect of oil supply shocks on oil prices is small (see Figure 3). This result is consistent with several earlier studies that employ the same model (see e.g. Kilian and Park, 2009; Güntner, 2014; Arampatzidis et al., 2021).

Figure 4: Impulse response functions of stock returns - **SVAR with zero restrictions**

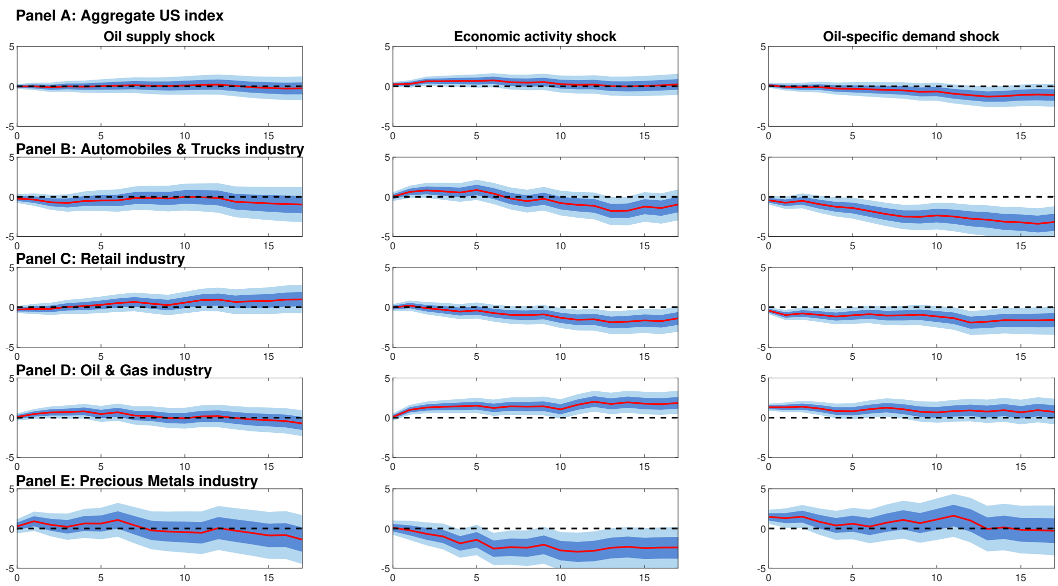

Note: (i) Red lines: posterior median; Dark blue shaded areas: 68% posterior credible sets; Light blue shaded areas: 95% posterior credible sets; (ii) all shocks are normalized such that to imply an increase in the price of oil.

[13]Figures B.7 to B.9 in the online Appendix B present the prior and posterior distributions of the stock market parameters in all models.

Consistent with the observation in Kilian (2009) that *"not all oil market shocks are alike"*, an economic activity shock has a different impact on the stock market. According to economic theory, a rise in global demand is viewed as good news for the domestic economy, and thus also for the stock market. On the other hand, increased aggregate demand means rising demand for oil, which puts upward pressure on oil prices. From the two opposing effects, the positive tends to prevail in the short to medium term, while in the longer run the negative effect becomes dominant.

The duration and persistence of those effects might vary by the industry. Indeed, the Petroleum & Natural Gas industry experiences a statistically significant and sustained rise in its share prices. Secondary industries, such as the Automobiles & Trucks and the Retail, experience only a rather short-term appreciation, if any. In such industries, the negative effect quickly becomes dominant, leading to even negative responses at longer horizons. Consistent with the view that the Precious Metals industry plays the role of a safe haven in periods of turmoil, a positive (negative) economic activity shock triggers a negative (positive) reaction of this industry already at very early horizons, which becomes statistically significant in the medium term. Finally, the response of the aggregate index follows the same pattern, but it is in general very close to zero. Similar conclusions can be found in Kilian and Park (2009) and Wang et al. (2013), among others.

An oil-specific demand shock causes a statistically significant and persistent increase in the price of oil (see Figure 3). The response of the Petroleum & Natural Gas industry is positive and statistically significant in the short run, which is reasonable since oil is the main output of this sector. Short-term benefits can also be identified in the Precious Metals industry, which can again be explained by the role of this sector as a safe haven in times of uncertainty. In contrast, the aggregate index as well as the Retail and the Automobiles & Trucks sectors respond negatively throughout the whole forecasting horizon, with the latter being affected the most. The strong, negative effect of oil-specific demand shocks for non-oil industries and the economy in general is known (see among others, Kilian and Park, 2009; Güntner, 2014). However, there are also some contradicting results (see e.g. Apergis and Miller, 2009).

Figure 5 shows the corresponding responses in the model identified with sign restrictions. An oil supply disruption causes almost identical responses as in the model with zero restrictions, which can again be explained by the minor impact of oil supply shocks on oil prices in the first place. This result partially contradicts the findings in earlier studies (Gupta and Modise, 2013; Koh, 2017; Basher et al., 2018), which assign a more important role to oil supply shocks. The response to economic activity shocks follows the same pattern as in Figure 4, with the positive effect prevailing in the short run and the negative effect dominating at longer horizons. This negative effect though appears slightly stronger in the model with sign restrictions, which changes the conclusions only for the Precious Metals industry. While in the model with zero restrictions the response of this particular industry is negative and statistically significant in the medium term, it remains negative in the model with sign restrictions, but it loses significance. Finally, an oil-specific demand shock leads in general to a smaller reaction of stock returns. Overall, although we observe some small discrepancies between the two traditional identification techniques, their main conclusions are quite similar.

Figure 5: Impulse response functions of stock returns - **SVAR with sign restrictions**

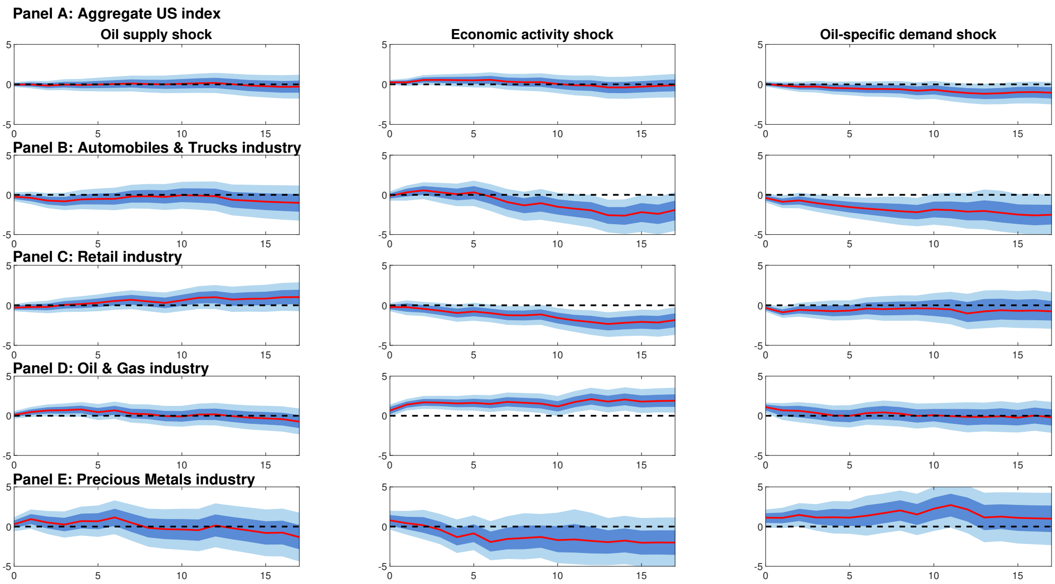

Note: (i) Red lines: posterior median; Dark blue shaded areas: 68% posterior credible sets; Light blue shaded areas: 95% posterior credible sets; (ii) all shocks are normalized such that to imply an increase in the price of oil.

Figure 6 shows the impulse responses of stock returns in our baseline model. An oil supply disruption leads in almost all cases to similar conclusions as in the traditional models. The only exception is the Petroleum & Natural Gas industry that shows a positive reaction in all models, which is marginally statistically significant in the short run only in our baseline model. An economic activity shock causes in general a reaction similar to the models with zero and sign restrictions. On the other hand, the effects of an oil-specific demand shock in our baseline model are closer to those observed in the model with sign restrictions. Nonetheless, the differences with the model identified with zero restrictions are negligible. Finally, an inventory demand shock has in all cases only a minor impact on stock returns. While the aggregate index as well as the Automobiles & Trucks industry and the Retail sector experience a negative effect in the medium to long term, it tends to be statistically significant in the medium run only in the latter.

Figure 6: Impulse response functions of stock returns - **Baseline model**

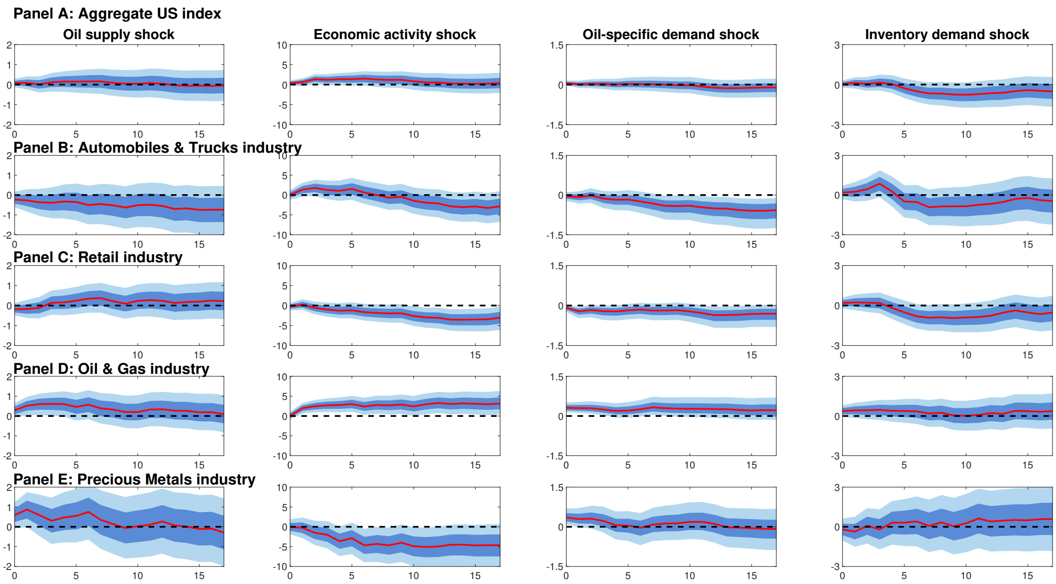

Note: (i) Red lines: posterior median; Dark blue shaded areas: 68% posterior credible sets; Light blue shaded areas: 95% posterior credible sets; (ii) all shocks are normalized such that to imply an increase in the price of oil.

Figure 7 offers a direct comparison of the impulse responses of stock returns in all models. Panel A presents the responses to oil market shocks at the 6th month, which corresponds to the short-term reaction of the stock market. Panel B in turn presents the impulse responses at the 18th month, as representative of the long-run reaction. In this way, this figure summarizes the main conclusions of this section, which in short are the following: (i) the effect of oil market shocks on the stock market depends on the underlying cause of the oil price change; (ii) it also depends on the industry; (iii) it does not depend though on the identification

Figure 7: Summary of impulse response functions

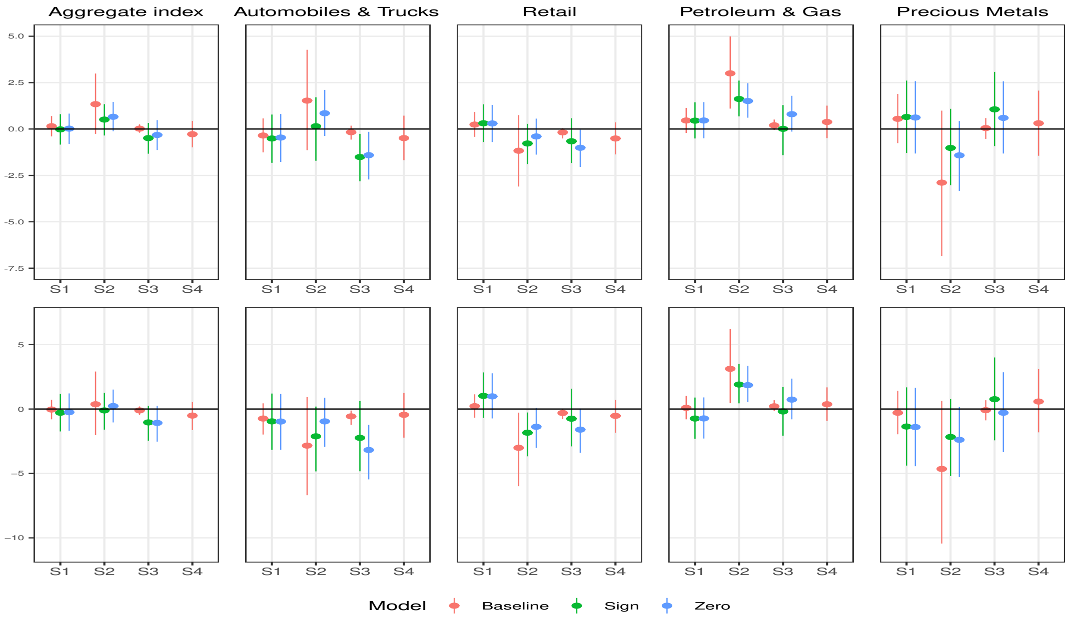

Note: This figure presents the median response of each stock return as well as the 95% credible set. **Panel A**: response at the 6th month; **Panel B**: response at the 18th month. **S1**: Oil supply shock; **S2**: Economic activity shock; **S3**: Oil-specific demand shock; **S4**: Inventory demand shock.

scheme. Put differently, besides some minor differences in the statistical significance, the magnitude and direction of the responses and thus the conclusions drawn from all models are qualitatively similar.

4.3. Sensitivity Analysis

In this section, we focus on our baseline model and assess the sensitivity of the results to modifications of one or more parameters at a time. Table 4 shows the posterior median and 68% credibility regions of the short-run price elasticities of oil supply and demand (Panels A-B) as well as the response of stock returns to each structural oil market shock at the 12th month (Panels C-F). The third row in Panels C-F refers to the correlation of the impulse responses of each alternative specification with those from the baseline model.

Column 1 presents the results for our baseline model. In the first alternative specification (Column 2), we use a slightly less tight prior for the price elasticity of oil demand (β_{qp}), i.e. we change the scale from 0.1 to 0.2, which corresponds to the prior used in Baumeister and Hamilton (2019). This modification brings only a small change in the posterior price elasticities of oil demand and supply (Columns 1 and 2, Panels A-B) and almost no change in the impulse responses (Columns 1 and 2, Panels C-F). Similarly, in the third column we modify the prior for the income elasticity of oil demand (β_{qy}) such that it uses the same prior information as in Baumeister and Hamilton (2019). More specifically, we change the location parameter from 0.3 to 0.7 and the scale from 0.15 to 0.2. Our conclusions for all parameters of interest (Columns 1 and 3, Panels A-F) do not change.

Table 4: Sensitivity of parameter inference

Baseline (1)	Price elasticity (2)	Income elasticity (3)	Both elasticities (4)	Minnesota prior (5)	Number of lags (6)	Lag structure (7)	All (8)
Panel A: Short-run oil supply elasticity (α_{qp})							
0.05	**0.04**	**0.05**	**0.04**	**0.06**	**0.05**	**0.06**	**0.04**
(0.03, 0.09)	(0.02, 0.07)	(0.02, 0.08)	(0.02, 0.07)	(0.03, 0.09)	(0.02, 0.09)	(0.03, 0.09)	(0.02, 0.07)
Panel B: Short-run price elasticity of oil demand (β_{qp})							
-0.57	**-0.64**	**-0.61**	**-0.66**	**-0.54**	**-0.57**	**-0.54**	**-0.62**
(-0.83, -0.41)	(-0.92, -0.46)	(-0.84, -0.45)	(-0.90, -0.50)	(-0.79, -0.39)	(-0.83, -0.41)	(-0.77, -0.38)	(-0.85, -0.47)
Panel C: Effect of oil supply shock on stock returns (12th month)							
0.07	**0.09**	**0.08**	**0.09**	**0.09**	**0.07**	**0.05**	**0.09**
(-0.26, 0.42)	(-0.25, 0.44)	(-0.26, 0.43)	(-0.25, 0.45)	(-0.29, 0.49)	(-0.29, 0.44)	(-0.35, 0.48)	(-0.32, 0.51)
100%	**99.6%**	**99.9%**	**99.4%**	**98.1%**	**94%**	**92.1%**	**91.4%**
Panel D: Effect of economic activity shock on stock returns (12th month)							
0.54	**0.54**	**0.52**	**0.52**	**0.54**	**-0.12**	**-0.33**	**-0.37**
(-0.51, 1.64)	(-0.51, 1.64)	(-0.54, 1.62)	(-0.53, 1.62)	(-0.58, 1.70)	(-1.26, 1.02)	(-1.52, 0.86)	(-1.57, 0.82)
100%	**99.9%**	**99.9%**	**99.9%**	**99.5%**	**93.6%**	**92.6%**	**92.1%**
Panel E: Effect of oil-specific demand shock on stock returns (12th month)							
-0.07	**-0.07**	**-0.07**	**-0.07**	**-0.10**	**-0.14**	**-0.17**	**-0.16**
(-0.21, 0.05)	(-0.20, 0.04)	(-0.20, 0.05)	(-0.19, 0.04)	(-0.26, 0.03)	(-0.29, 0.00)	(-0.35, 0.00)	(-0.32, 0.00)
100%	**99.9%**	**99.9%**	**99.9%**	**99.8%**	**97.3%**	**97.8%**	**97.4%**
Panel F: Effect of inventory demand shock on stock returns (12th month)							
-0.71	**-0.71**	**-0.71**	**-0.71**	**-1.14**	**-0.72**	**-0.87**	**-0.86**
(-1.19, -0.24)	(-1.20, -0.23)	(-1.20, -0.24)	(-1.20, -0.23)	(-1.68, -0.63)	(-1.22, -0.23)	(-1.41, -0.34)	(-1.40, -0.33)
100%	**99.9%**	**99.9%**	**99.9%**	**99.6%**	**98.6%**	**98.2%**	**98.1%**

Note: The table reports the posterior median (in bold), 68% credibility regions (in parentheses) and correlation of impulse responses of each alternative specification with those from the baseline model (in percentages). The baseline model uses the priors introduced in Table 3, whereas each of the alternative specifications relaxes one or more of those priors, as described in Section 4.3.

Kilian and Zhou (2018) criticize the robustness exercise in Baumeister and Hamilton (2019) for modifying only one parameter at a time. To take this into account, our sensitivity analysis considers also specifications that simultaneously modify more than one parameter of our baseline model (see Columns 4, 7 and 8). Starting with the fourth column, we apply the combination of prior information for the price and income elasticities used in Columns 2 and 3. The results for the impulse responses (Columns 1 and 4, Panels C-F) as well as the two posterior elasticities (Columns 1 and 4, Panels A-B) are almost identical to our baseline case.

Next, in Columns 5-7 we consider modifications that affect the lag structure of the model. Starting with Column 5, we follow Baumeister and Hamilton (2019) and lower the weight given to the Minnesota prior by changing the value of λ_0 from 0.5 to 1. This brings only a modest change (Columns 1 and 5, Panels A-F). The only exception seems to be the response to an inventory demand shock (Panel F), which equals -0.71 in our baseline model and -1.14 in the alternative specification. But this is merely a change in the magnitude, since a closer look reveals that the qualitative conclusions remain the same as the responses in both models are statistically significant and their evolution is nearly identical (99.6% correlation).

The next step is to consider the modification of a feature that is never relaxed in Baumeister and Hamilton (2019). Kilian and Zhou (2018) criticize their choice of relying only on the use of 12 lags without testing inference based on additional lags. In contrast, our baseline model follows the majority of empirical studies and uses 24 lags. Since there is a clear disagreement on the appropriate number of lags, Column 6 presents the results obtained after using $m = 12$ lags. In general, inference about most parameters does not change (Columns 1 and 6, Panels A-F). The only apparent exception is the response to an economic activity shock (Columns 1 and 6, Panel D). In this case, we observe a change in the magnitude and the sign one year after the shock. This comes in line with the argument in Kilian and Zhou (2018) that models with fewer lags tend to underestimate the effect of economic activity shocks. Nonetheless, the overall conclusions remain the same, as both responses are statistically insignificant and evolve in an almost identical fashion (the correlation is 93.6%). Finally, Column 7 considers a simultaneous change in the Minnesota prior and the number of lags. This again has only a modest effect on most results (Columns 1 and 7, Panels A-F).

As a final robustness exercise, in the last column we consider all changes in parameters discussed above at the same time. Although the prior information used in this alternative specification is considerably different from our baseline model, we observe that the inference on the short-run price elasticities of supply and demand does not change (Columns 1 and 8, Panels A-B). For the impulse responses (Columns 1 and 8, Panels C-F), we observe again a similar evolution (the correlation in all cases exceeds 91%) and identical conclusions about statistical significance. The only difference occurs in the magnitude of the response to an economic activity shock, which can again be justified as discussed above. Overall, our sensitivity analysis reveals that the results from our baseline model remain fairly robust to both small and large modifications in particular components of prior information.

5. Conclusions

Withing a Bayesian SVAR that also serves as our baseline model, we examine the importance of different identification schemes (i.e. SVAR with zero restrictions, SVAR with sign restrictions, baseline model) for the relationship between oil prices and the US stock market. In this framework, we scrutinize the prior and posterior distributions of three important elasticities, namely the oil supply elasticity and the price and income elasticities of oil demand. We further examine the effects of oil market shocks on the price of oil in all models. Finally, we investigate whether the differences between the models identified in the oil market matter for the oil-stock market relationship.

Our findings suggest that there are important differences between our baseline model and the two traditional models regarding the oil market. First, the flexibility of the Bayesian SVAR allows us to specify an exponential prior distribution for the oil supply elasticity that uses prior information in a more intuitive way. Therefore, our model suggests more plausible posterior values for the oil supply elasticity compared to traditional models as well as the model in Baumeister and Hamilton (2019), adding in this way to the corresponding debate in the literature. Second, the implied posterior price elasticity of oil demand is unrealistically high in traditional models, which suggests a very elastic demand curve. In contrast, our baseline model implies a price inelastic oil demand, which is in line with empirical microeconometric studies. Third, in our Bayesian SVAR oil supply shocks have a larger and more persistent impact on oil prices.

The main finding of our study is that the effects of oil market shocks on the US stock market, both at aggregate and disaggregate level, remain fairly robust to the choice of identification scheme. More specifically, despite the theoretical and econometric differences between the models under consideration, the disparities regarding the oil-stock market relationship are small. Since there are numerous studies that employ SVARs with zero restrictions, and less frequently sign restrictions, for the investigation of the oil-stock market relationship, this observation has important implications. It suggests that the SVAR with zero restrictions, despite its limitations, remains a good approximation in practice for the oil-stock market relationship, as it yields results comparable to more sophisticated models, such as the Bayesian SVAR. This in turn means that it is difficult to refute the conclusions from previous empirical studies that employ this model to examine the oil-stock market relationship. Instead, any disagreement on their conclusions must be due to the use of different datasets, sample of countries/industries, or different sample periods, which suggests a time-varying relationship (Foroni et al., 2017; Arampatzidis et al., 2021).

What does this observation mean for the Bayesian SVAR model? The econometric appeal and flexibility of this model is thoroughly discussed in the previous sections. Nonetheless, there could be a trade-off between potential gains and additional complexity. Our analysis shows that: i) the modelling effort and time are considerably higher in the Bayesian SVAR; ii) the gains are important for the oil market, but negligible for the oil-stock market relationship. However, the Bayesian SVAR might be more useful in other applications. It depends on how realistic the identifying assumptions are in each particular framework. If zero/sign restrictions are not a good approximation, then the Bayesian SVAR could help in the identification process by allowing the use of more informative and intuitive priors. Our study shows that this is not the case though for the oil-stock market relationship, as zero restrictions alone seem to be a good approximation. In any case, as the Bayesian SVAR is relatively new in the literature, additional tests are necessary in order to assess its full potential.

There is still room for improvement also in the oil-stock market literature, if we exploit the additional capabilities of the Bayesian SVAR. Our study adopts a simple approach to allow for a direct comparison between the candidate models. There are still assumptions in the baseline model that could be further relaxed. One such example concerns the contemporaneous relationship between oil and stock markets.

In particular, the Bayesian SVAR allows us to relax the assumption that the US stock market does not affect the oil market at time t (see the last column of matrix $\boldsymbol{A}$, Eq. (16)). Although this assumption might be reasonable for the effect of a stock market shock on oil production, economic activity and oil inventories, it might be questionable for oil prices. Using a data-driven approach, Keweloh (2021) shows that there is indeed a contemporaneous effect of a stock market shock on oil prices. To take this into account, one can replace that specific zero restriction with a flat or even an informative prior based on previous empirical evidence. Another feature of our baseline model that can be improved is related to the use of non-informative priors for the stock market parameters (see the χ_i, for $i = 1, 2, 3, 4$, in Eq. (13)). Our choice of using completely uninformative priors for those parameters can be justified by the need to offer a direct comparison between all models. Future research could make broader use of the existing prior information in the oil-stock market literature by using more informative priors for the stock market.

Our findings have additional implications for investors and policy makers. No matter which SVAR model we employ, the need to distinguish between supply and demand shocks remains imperative. In general, shocks from the demand side were found to play a larger role for the determination of oil prices as well as for the oil-stock market relationship. The investigation of the effects of oil market shocks both at aggregate and disaggregate level is also important. Based on our analysis, it is clear that a study that focuses only on the aggregate stock market would not be able to uncover the idiosyncratic behavior of different industries. There are industries like the Automobiles & Trucks and Retail which respond negatively to an increase in the price of oil, whereas there are others, such as the Petroleum & Natural Gas and Precious Metals, which might experience considerable gains. Therefore, aggregating all individual industries to a single index might not be the most efficient way of examining the oil-stock market relationship from a policy maker or investor point of view. Our analysis though focuses only at the industry level. Although examining the behavior of individual firms will probably deepen our understanding further, we leave such an analysis for future research.

References

Agrawal, P. (2015). India's petroleum demand: Estimations and projections. *Applied Economics*, 47(12):1199–1212.

Alquist, R., Kilian, L., and Vigfusson, R. J. (2013). Chapter 8 - Forecasting the price of oil. In Elliott, G. and Timmermann, A., editors, *Handbook of Economic Forecasting*, volume 2 of *Handbook of Economic Forecasting*, pages 427–507. Elsevier.

Anderson, S. T., Kellogg, R., and Salant, S. W. (2018). Hotelling under pressure. *Journal of Political Economy*, 126(3):984–1026.

Apergis, N. and Miller, S. M. (2009). Do structural oil-market shocks affect stock prices? *Energy Economics*, 31(4):569–575.

Arampatzidis, I., Dergiades, T., Kaufmann, R. K., and Panagiotidis, T. (2021). Oil and the U.S. stock market: Implications for low carbon policies. *Energy Economics*, 103:105588.

Basher, S. A., Haug, A. A., and Sadorsky, P. (2018). The impact of oil-market shocks on stock returns in major oil-exporting countries. *Journal of International Money and Finance*, 86:264–280.

Baumeister, C. and Hamilton, J. D. (2015). Sign restrictions, structural vector autoregressions, and useful prior information. *Econometrica*, 83(5):1963–1999.

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. and Peersman, G. (2013). 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.

Bernanke, B. S., Gertler, M., Watson, M., Sims, C. A., and Friedman, B. M. (1997). Systematic monetary policy and the effects of oil price shocks. *Brookings Papers on Economic Activity*, 1997(1):91–157.

Caldara, D., Cavallo, M., and Iacoviello, M. (2019). Oil price elasticities and oil price fluctuations. *Journal of Monetary Economics*, 103:1–20.

Csereklyei, Z., Rubio-Varas, M., and Stern, D. I. (2016). Energy and economic growth: The stylized facts. *Energy Journal*, 37(2):223–256.

De, K., Compton, R. A., and Giedeman, D. C. (2022). Oil shocks and the U.S. economy in a data-rich model. *Economic Modelling*, 108:105755.

Degiannakis, S., Filis, G., and Panagiotakopoulou, S. (2018). Oil price shocks and uncertainty: How stable is their relationship over time? *Economic Modelling*, 72:42–53.

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

Foroni, C., Guérin, P., and Marcellino, M. (2017). Explaining the time-varying effects of oil market shocks on US stock returns. *Economics Letters*, 155:84–88.

Gately, D. and Huntington, H. G. (2002). The asymmetric effects of changes in price and income on energy and oil demand. *Energy Journal*, 23(1):19–56.

Gelman, M., Gorodnichenko, Y., Kariv, S., Koustas, D., Shapiro, M. D., Silverman, D., and Tadelis, S. (2016). The response of consumer spending to changes in gasoline prices. Working Paper 22969, National Bureau of Economic Research.

Gupta, R. and Modise, M. P. (2013). Does the source of oil price shocks matter for South African stock returns? A structural VAR approach. *Energy Economics*, 40:825–831.

Güntner, J. H. F. (2014). How do international stock markets respond to oil demand and supply shocks? *Macroeconomic Dynamics*, 18(8):1657–1682.

Güntner, J. H. F. and Öhlinger, P. (2021). Oil price shocks and the hedging benefit of Airline investments. Economics working papers 2021-14, Department of Economics, Johannes Kepler University Linz, Austria.

Hamilton, J. D. (2009). Causes and consequences of the oil shock of 2007-08. Working Paper 15002, National Bureau of Economic Research.

Hamilton, J. D. (2011). Nonlinearities and the macroeconomic effects of oil prices. *Macroeconomic Dynamics*, 15(S3):364–378.

Herrera, A. M. and Rangaraju, S. K. (2020). The effect of oil supply shocks on US economic activity: What have we learned? *Journal of Applied Econometrics*, 35(2):141–159.

Huntington, H. G., Barrios, J. J., and Arora, V. (2019). Review of key international demand elasticities for major industrializing economies. *Energy Policy*, 133:110878.

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

Keweloh, S. A. (2021). A generalized method of moments estimator for structural vector autoregressions based on higher moments. *Journal of Business & Economic Statistics*, 39(3):772–782.

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. (2019). Understanding the estimation of oil demand and oil supply elasticities. *CFS Working Paper, No. 649, 2020.*

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

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 Park, C. (2009). The impact of oil price shocks on the U.S. stock market. *International Economic Review*, 50(4):1267–1287.

Kilian, L. and Zhou, X. (2018). Structural interpretation of vector autoregressions with incomplete information: Revisiting the role of oil supply and demand shocks: Comment. *CEPR Discussion Paper No. DP13068*.

Kilian, L. and Zhou, X. (2020). The econometrics of oil market VAR models. *CESifo Working Paper No. 8153*.

Koh, W. C. (2017). How do oil supply and demand shocks affect Asian stock markets? *Macroeconomics and Finance in Emerging Market Economies*, 10(1):1–18.

Wang, Y., Wu, C., and Yang, L. (2013). Oil price shocks and stock market activities: Evidence from oil-importing and oil-exporting countries. *Journal of Comparative Economics*, 41(4):1220–1239.

Appendix A Data Sources

1. **World oil production**: measured in million of barrels/day; Source: EIA (available at https://www.eia.gov/opendata/qb.php?category=2134979&sdid=INTL.57-1-WORL-TBPD.M)

2. **Industrial production**: index; Source: Baumeister and Hamilton (2019) (available at https://sites.google.com/site/cjsbaumeister/research)

3. **US refiner's acquisition cost of imported crude oil**: measured in dollars per barrel; Source: EIA (available at https://www.eia.gov/totalenergy/data/browser/index.php?tbl=T09.01#/?f=M&start=197301&end=202109&charted=0-6)

4. **US CPI**: index, all items, 1982-1984=100; Source: FRED database (available at https://fred.stlouisfed.org/series/CPIAUCSL)

5. **US crude oil inventories**: measured in million of barrels; Source: EIA (available at https://www.eia.gov/opendata/qb.php?sdid=PET.MCRSTUS1.M)

6. **OECD inventories of petroleum and petroleum products**: measured in million of barrels; Source: EIA (available at https://www.eia.gov/opendata/qb.php?category=2134439&sdid=INTL.5-5-OECD-MBBL.M)

7. **US inventories of petroleum and petroleum products**: measured in million of barrels; Source: EIA (available at https://www.eia.gov/opendata/qb.php?sdid=PET.MTTSTUS1.M)

8. **Aggregate US stock market index**: index, 2015=100; Source: OECD database (available at https://stats.oecd.org/index.aspx?queryid=84)

9. **Industry nominal stock returns**: index; Source: Kenneth R. French database (available at http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html)

Appendix B Additional Data and Results

Table B.1: Main descriptive statistics of stock returns

Industries	Mean	Median	Max	Min	Std	Skewness	Kurtosis
Aggregate index	0.22	0.57	11.82	-24.60	3.69	-1.09	7.70
Automobiles & Trucks	0.55	0.64	49.46	-35.55	7.09	0.22	8.51
Retail	0.76	0.54	26.33	-29.38	5.58	-0.21	5.07
Petroleum & Natural Gas	0.68	0.80	23.60	-19.14	5.64	0.00	4.03
Precious Metals	0.62	0.01	80.02	-32.66	11.01	0.81	7.84

Figure B.1: Prior and posterior distributions of oil market parameters - **SVAR with zero restrictions**

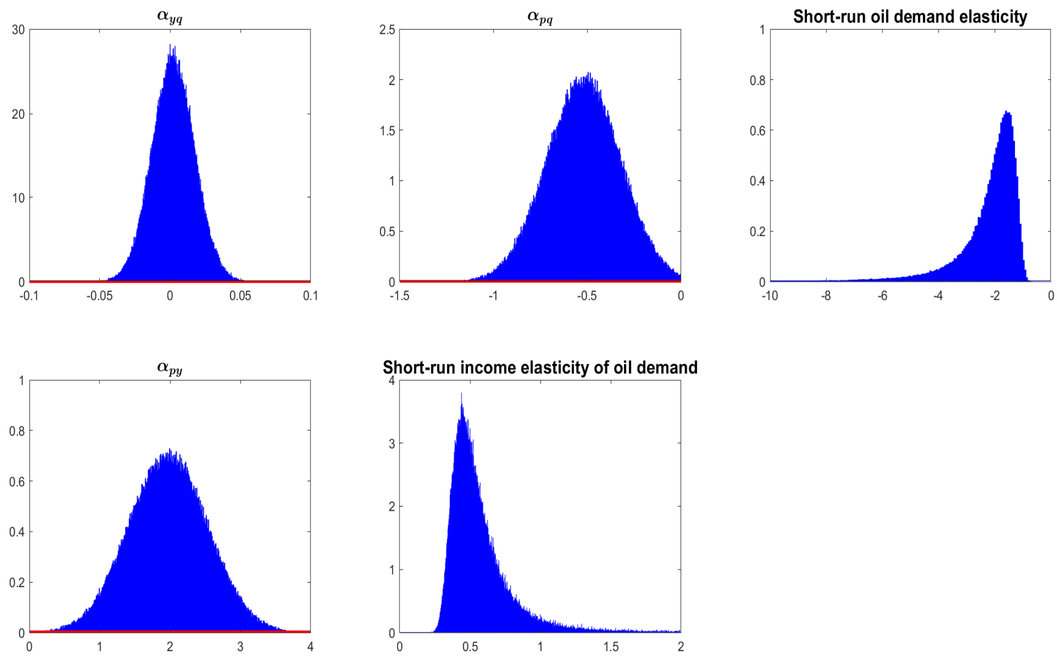

Note: Red lines: prior distributions; Blue histograms: posterior distributions.

Figure B.2: Prior and posterior distributions of oil market parameters - **SVAR with sign restrictions**

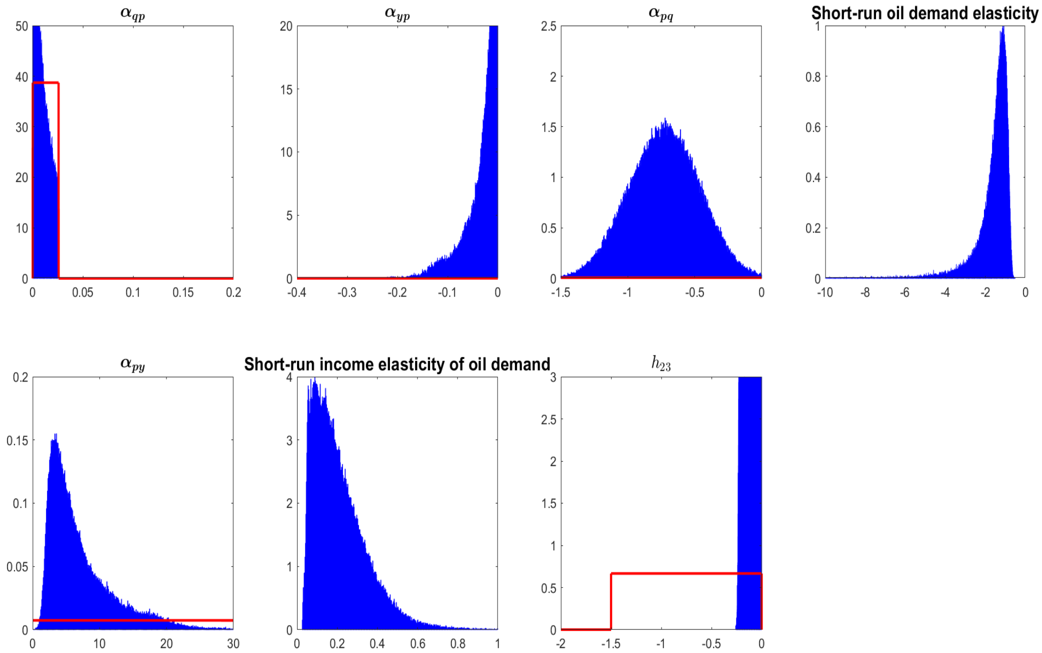

Note: Red lines: prior distributions; Blue histograms: posterior distributions.

Figure B.3: Prior and posterior distributions of oil market parameters - **Baseline model**

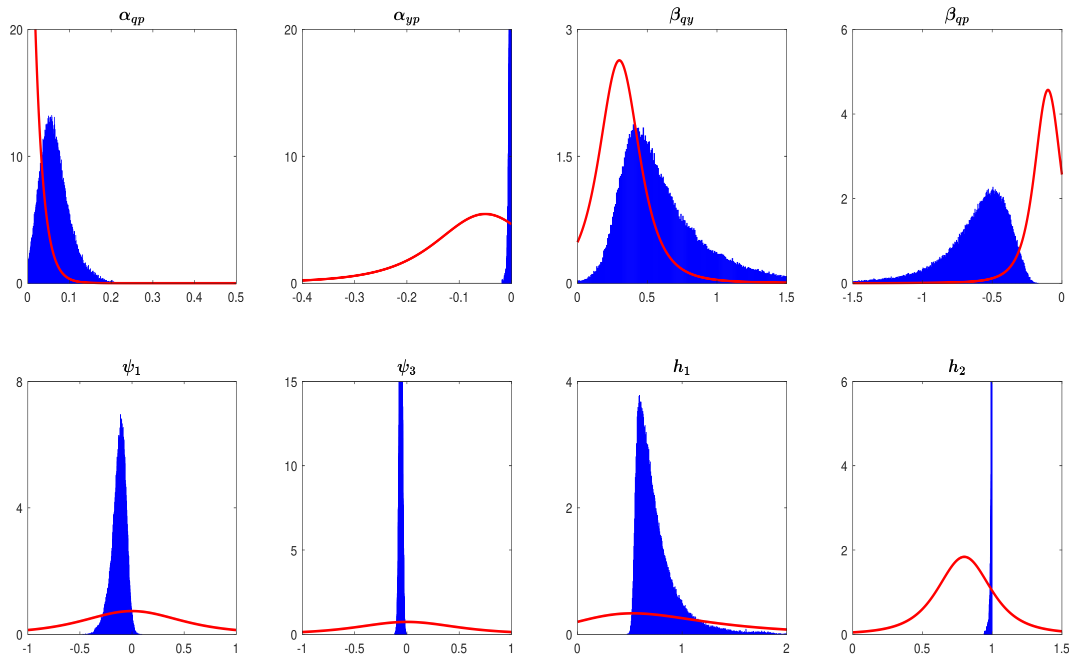

Note: Red lines: prior distributions; Blue histograms: posterior distributions.

Figure B.4: Impulse response functions of oil market variables - **SVAR with zero restrictions**

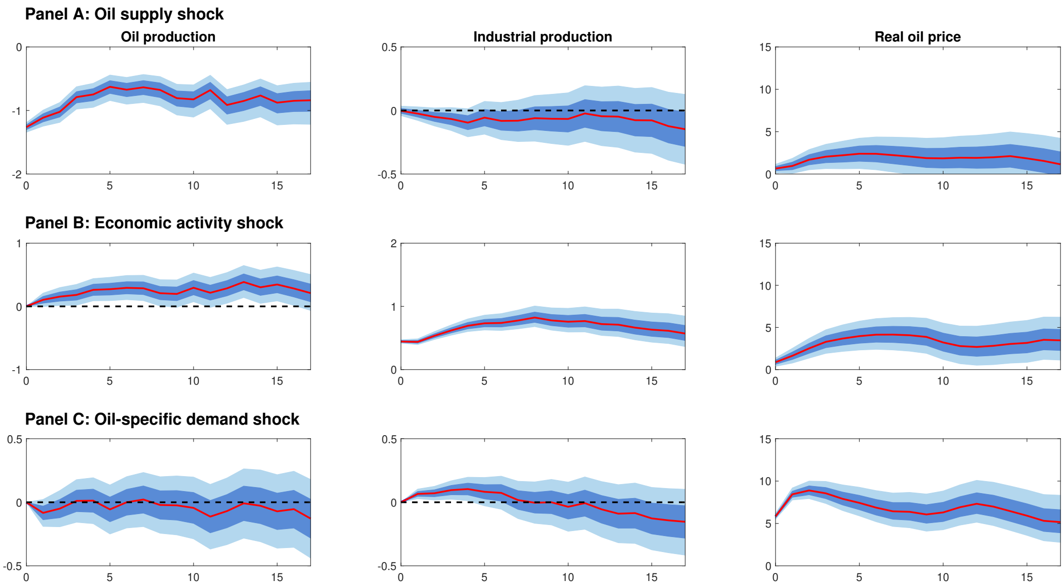

Note: (i) Red lines: posterior median; Dark blue shaded areas: 68% posterior credible sets; Light blue shaded areas: 95% posterior credible sets; (ii) all shocks are normalized such that to imply an increase in the price of oil.

Figure B.5: Impulse response functions of oil market variables - **SVAR with sign restrictions**

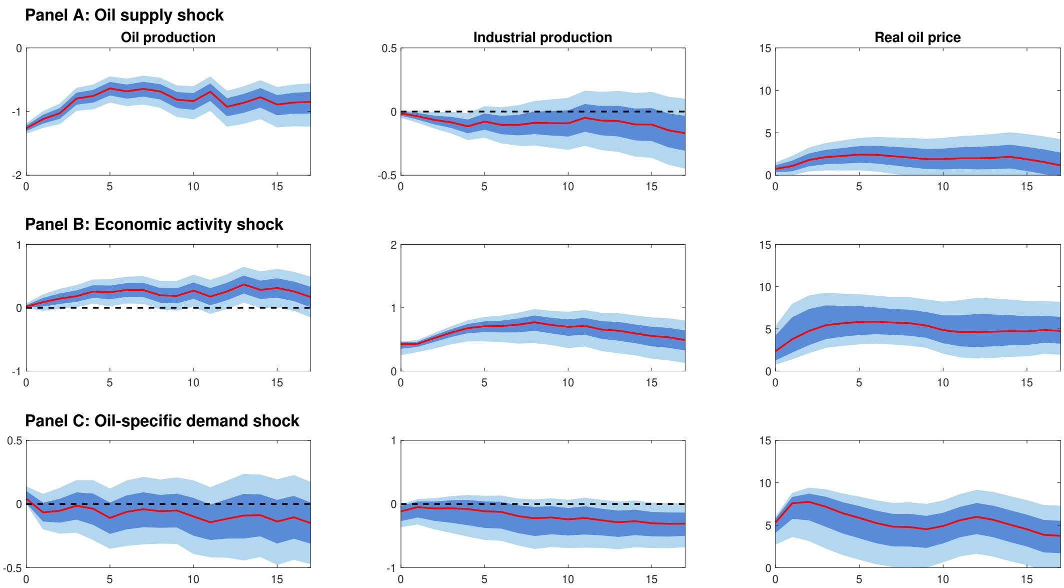

Note: (i) Red lines: posterior median; Dark blue shaded areas: 68% posterior credible sets; Light blue shaded areas: 95% posterior credible sets; (ii) all shocks are normalized such that to imply an increase in the price of oil.

Figure B.6: Impulse response functions of oil market variables - **Baseline model**

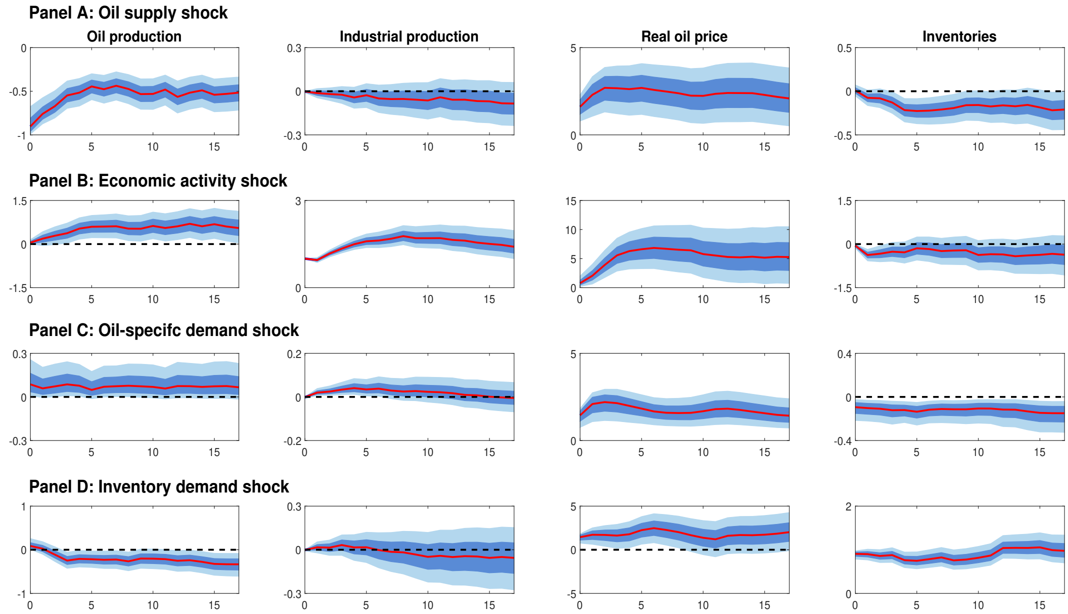

Note: (i) Red lines: posterior median; Dark blue shaded areas: 68% posterior credible sets; Light blue shaded areas: 95% posterior credible sets; (ii) all shocks are normalized such that to imply an increase in the price of oil.

Figure B.7: Prior and posterior distributions of stock market parameters - **SVAR with zero restrictions**

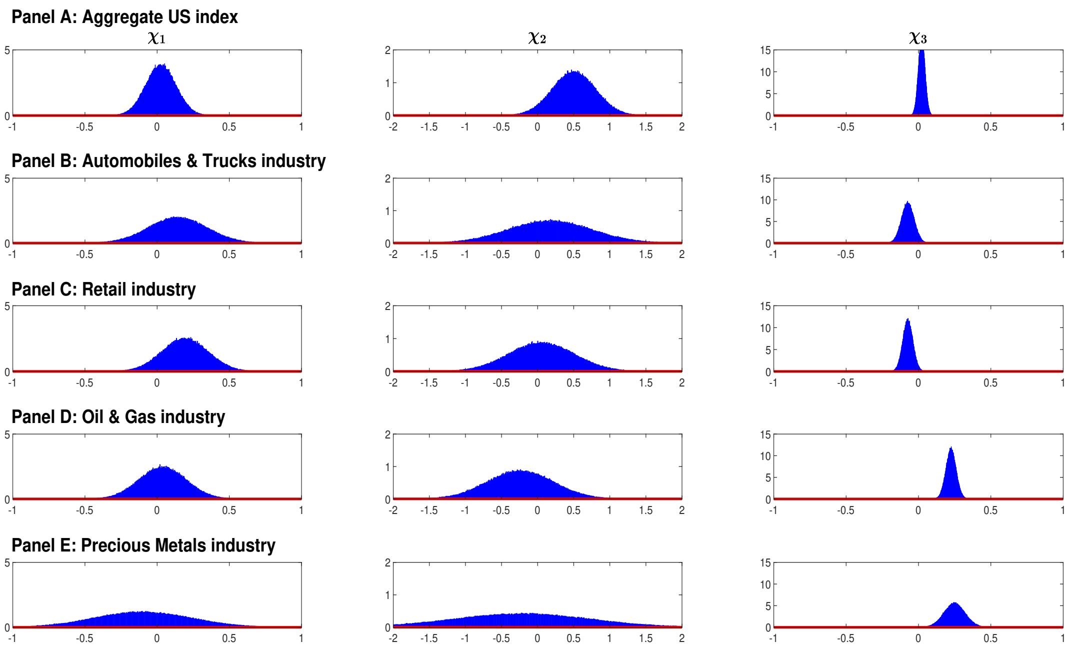

Note: Red lines: prior distributions; Blue histograms: posterior distributions.

Figure B.8: Prior and posterior distributions of stock market parameters - **SVAR with sign restrictions**

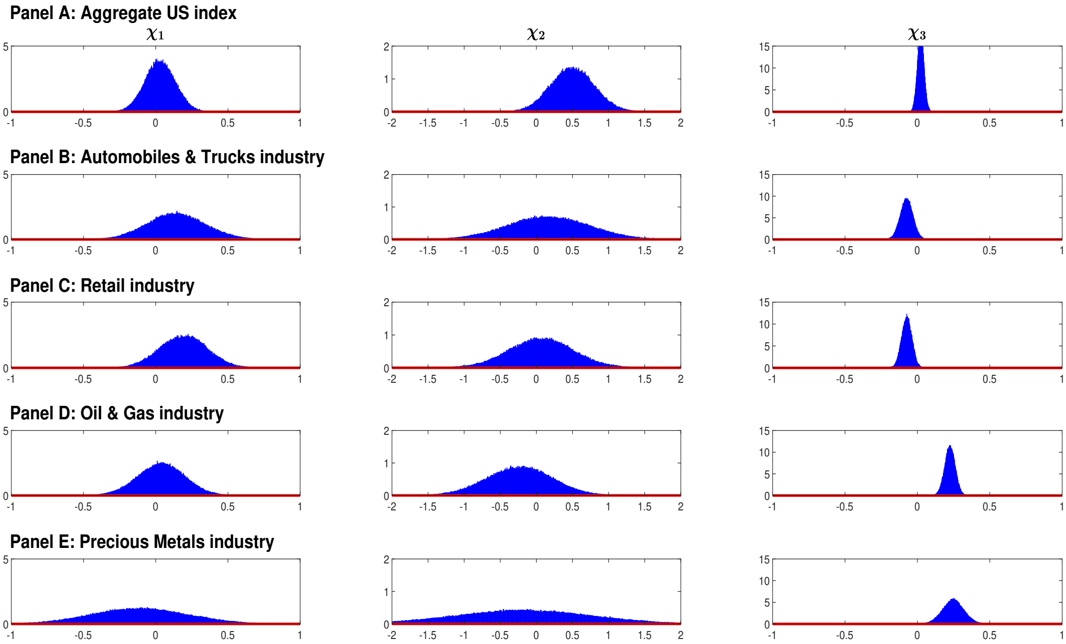

Note: Red lines: prior distributions; Blue histograms: posterior distributions.

Figure B.9: Prior and posterior distributions of stock market parameters - **Baseline model**

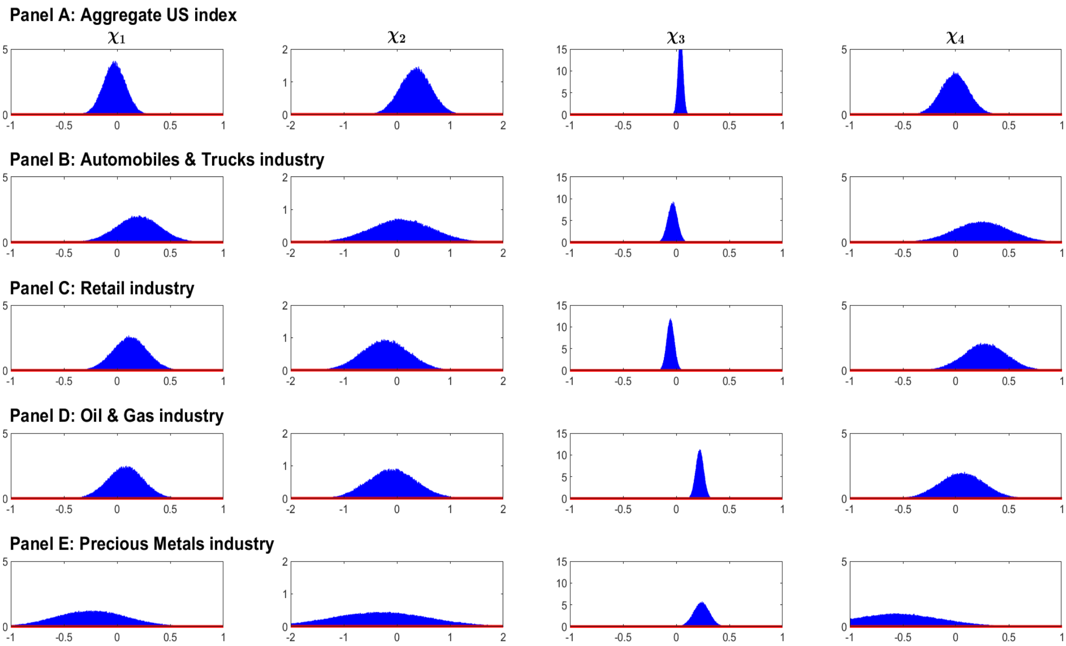

Note: Red lines: prior distributions; Blue histograms: posterior distributions.

Appendix C Prior and Posterior Distributions

C.1 Prior Distributions

Baumeister and Hamilton (2015) propose specific prior distributions for matrices $\boldsymbol{D}$ and $\boldsymbol{B}$ since closed-form analytic expressions are available for their posteriors. On the other hand, we are free to choose any prior distribution for $\boldsymbol{A}$ without any computational concerns. Therefore, we follow Baumeister and Hamilton (2019) and assume that the reciprocal elements of $\boldsymbol{D}$ are independent Gamma distributions:

$$d_{ii}^{-1}|\boldsymbol{A} \sim \Gamma(\kappa_i, \tau_i(\boldsymbol{A})) \tag{C.1}$$

where d_{ii} denotes the diagonal elements in $\boldsymbol{D}$, whereas κ_i and τ_i are parameters that control the weight of the prior. Baumeister and Hamilton (2015) set $\tau_i/\kappa_i = \boldsymbol{A}\hat{\boldsymbol{S}}\boldsymbol{A}'$, where τ_i/κ_i is the inverse of the prior mean of d_{ii}^{-1}. This implies that:

$$\tau_i(\boldsymbol{A}) = \kappa_i a_i' \hat{\boldsymbol{S}} a_i \tag{C.2}$$

We further assume that the elements of $\boldsymbol{B}$ are independent Normals:

$$b_i|\boldsymbol{A}, \boldsymbol{D} \sim N(m_i, d_{ii}\boldsymbol{M}_i) \tag{C.3}$$

where m_i denotes our prior beliefs on the values of the lagged coefficients, whereas $\boldsymbol{M}_i$ shows our confidence on the prior. Since all variables are used in first differences, it is intuitive to set $m_i = 0$, for all $i = 1, ..., n$. In addition, we place greater confidence on our beliefs that the coefficients at higher lags are zero. This can be achieved by using smaller values for the diagonal elements in $\boldsymbol{M}$ as the lags increase:

$$v_1' = (1/(1^{2\lambda_1}), 1/(2^{2\lambda_1}), ..., 1/(m^{2\lambda_1})) \tag{C.4}$$

$$v_2' = (s_{11}^{-1}, s_{22}^{-1}, ..., s_{nn}^{-1})' \tag{C.5}$$

$$v_3 = \lambda_0^2 \begin{bmatrix} v_1 \otimes v_2 \\ \lambda_3^2 \end{bmatrix} \tag{C.6}$$

$$\boldsymbol{M}_{i,jj} = v_{3j} \tag{C.7}$$

where λ_0 represents the overall weight given to the prior, with smaller values associated with higher confidence, λ_1 captures our confidence that the coefficients at higher lags are zero and λ_3 governs the tightness of the prior for the constant.

For the priors in $\boldsymbol{A}$ we provide a detailed discussion in Section 3.3.1. Taking everything together, our prior information can be summarized as:

$$p(\boldsymbol{A}, \boldsymbol{D}, \boldsymbol{B}) = p(\boldsymbol{A})p(\boldsymbol{D}|\boldsymbol{A})p(\boldsymbol{B}|\boldsymbol{A}, \boldsymbol{D}) \tag{C.8}$$

The goal is then to see how observation of the data causes us to revise our prior beliefs. Appendix C.2 discusses the form of the posteriors.

C.2 Posterior Distributions

Since the prior for $\boldsymbol{D}|\boldsymbol{A}$ follows a Gamma distribution, the posterior is:

$$d_{ii}^{-1}|\boldsymbol{A}, Y_T \sim \Gamma(\kappa_i^*, \tau_i^*(\boldsymbol{A})) \tag{C.9}$$

$$\kappa_i^* = \kappa_i + T/2 \tag{C.10}$$

$$\tau_i^*(\boldsymbol{A}) = \tau_i(\boldsymbol{A}) + (1/2)\zeta_i^*(\boldsymbol{A}) \tag{C.11}$$

The only unknown is $\zeta_i^*(\boldsymbol{A})$. Baumeister and Hamilton (2015) show that:

$$\zeta_i^*(\boldsymbol{A}) = (\tilde{Y}_i^{'}(\boldsymbol{A})\tilde{Y}_i(\boldsymbol{A})) - (\tilde{Y}_i^{'}(\boldsymbol{A})\tilde{X}_i)(\tilde{X}_i^{'}\tilde{X}_i)^{-1}(\tilde{X}_i^{'}\tilde{Y}_i(\boldsymbol{A})) \tag{C.12}$$

where

$$\tilde{Y}_i(\boldsymbol{A}) = [a_i^{'}y_1 \ \dots \ a_i^{'}y_T \ \ m_i(\boldsymbol{A})^{'}\boldsymbol{P}_i]^{'} \tag{C.13}$$

$$\tilde{X}_i = [x_0^{'} \ \dots \ x_{T-1}^{'} \ \ \boldsymbol{P}_i]^{'} \tag{C.14}$$

$$\boldsymbol{M}_i^{-1} = \boldsymbol{P}_i\boldsymbol{P}_i^{'} \tag{C.15}$$

Since the prior for $\boldsymbol{B}|\boldsymbol{A}, \boldsymbol{D}$ follows a Normal distribution, the posterior is:

$$b_i|\boldsymbol{A}, \boldsymbol{D}, Y_T \sim N(m_i^*, d_{ii}\boldsymbol{M}_i^*) \tag{C.16}$$

$$\boldsymbol{M}_i^* = (\tilde{X}_i^{'}\tilde{X}_i)^{-1} \tag{C.17}$$

$$m_i^*(\boldsymbol{A}) = \boldsymbol{M}_i^*(\tilde{X}_i^{'}\tilde{Y}_i(\boldsymbol{A})) \tag{C.18}$$

The posterior marginal distribution for $\boldsymbol{A}$ is given by:

$$p(\boldsymbol{A}|Y_T) = \frac{\kappa_T p(\boldsymbol{A})[det(\boldsymbol{A}\hat{\boldsymbol{\Omega}}_T\boldsymbol{A}^{'})]^{T/2}}{\prod_{i=1}^{n}[(2/T)\tau_i^*(\boldsymbol{A})]^{\kappa_i^*}} \prod_{i=1}^{n} \tau_i(\boldsymbol{A})^{\kappa_i} \tag{C.19}$$

Finally, $\hat{\boldsymbol{\Omega}}_{\boldsymbol{T}}$ is the sample variance matrix of the reduced-form residuals:

$$\hat{\boldsymbol{\Omega}}_T = T^{-1}\left\{\sum_{t=1}^{T} y_t y_t^{'} - \left(\sum_{t=1}^{T} y_t x_t^{'}\right)\left(\sum_{t=1}^{T} x_t x_t^{'}\right)^{-1}\left(\sum_{t=1}^{T} y_t x_t^{'}\right)\right\} \tag{C.20}$$