CENTRE FOR APPLIED MACRO – AND PETROLEUM ECONOMICS (CAMP)

CAMP Working Paper Series
No 8/2018

The Shale Oil Boom and the U.S. Economy: Spillovers and Time-Varying Effects

Hilde C. Bjørnland and Julia Skretting

BI NORWEGIAN BUSINESS SCHOOL

The Shale Oil Boom and the U.S. Economy: Spillovers and Time-Varying Effects*

Hilde C. Bjørnland[†] Julia Skretting[‡]

February 28, 2024

We provide new evidence that the transmission of oil prices to the U.S. economy has changed with the shale oil boom. To show this, we develop a time-varying parameter factor-augmented vector autoregressive (VAR) model with a large data environment of state-level, industry and aggregate U.S. data. The model effectively captures potential spillovers between oil and non-oil industries, as well as variation over time. Specified in this way, we find that investment, income, industrial production and (non-oil) employment in most oil-producing and some manufacturing-intensive U.S. states to increase following an oil-specific shock - effects that were not present before the shale oil boom.

JEL-codes: C11, C55, E32, E42, Q43

Keywords: Shale oil boom, Oil prices, Time-varying factor-augmented VAR model, Spillovers, Geographical dispersion

*We gratefully acknowledge comments and suggestions from the Editor Barbara Rossi, three anonymous referees, Rabah Arezki, Fabio Canova, Ana Maria Herrera, Michael McMahon, James Morley, Francesco Ravazzolo, Leif Anders Thorsrud, Ragnar Torvik, Rob Vigfusson, Jing Cynhia Wu and participants at the "Economics of Oil" conference in Rio de Janeiro, the 26th SNDE Symposium in Tokyo, the 2019 IJCB Annual Conference in Oslo, the 2019 North American Meeting in Econometrics in Seattle and seminar participants in the World Bank and Norges Bank. This paper is part of the research activities at the Centre for Applied Macro and commodity Prices (CAMP) at the BI Norwegian Business School.

[†]Centre for Applied Macroeconomics and commodity Prices (CAMP), BI Norwegian Business School and Norges Bank. Email: `hilde.c.bjornland@bi.no`

[‡]Centre for Applied Macroeconomics and commodity Prices (CAMP), BI Norwegian Business School. Email: `julia.skretting@bi.no`

1 Introduction

The history of the petroleum industry in the United States goes back to the early 19th century. Petroleum became a major industry following the discovery of oil at Oil Creek, Pennsylvania, in 1859, and for much of the 19th and 20th centuries, the U.S. was the largest oil-producing country in the world. However, after production peaked in 1970, the U.S. experienced decades of production decline. Over time, the country became increasingly dependent on imports of oil, and in 1973, the U.S. government banned firms from exporting oil.

The empirical oil-macroeconomic literature which was sparked off by the seminal paper by Hamilton (1983), has typically analyzed the effects of adverse oil price shocks on the U.S. economy focusing on the period when the U.S. was a net oil importer. In line with this, scholars have found that the U.S. economy responds negatively to adverse shocks that increase the oil price, as both consumers and producers have to pay more for the imported energy products and for the complementary products to energy (see e.g., Bjørnland, 2022).

The shale revolution may have changed this relationship. The massive surge in the production of shale (i.e., unconventional) oil, which started in the early 2000s, has in a few years made the United States the world's largest oil producer. Such a transition, however, did not happen by itself. As shale oil is trapped in petroleum-bearing formations with low permeability, this requires the combination of horizontal drilling in conjunction with hydraulic fracturing to extract the oil. This process is technological challenging, requiring capital, new technology, labor, skills and "learning by doing" (LBD) over a prolonged period of time, and with potential spillovers to other industries.[1] To the extent that these spillovers also affect production and employment across U.S. states, the relationship between oil prices and aggregate U.S. activity could also have changed.

So far, empirical studies addressing the implication of the shale oil boom for the aggregate U.S. economy are lacking. Some studies have documented local effects of recent oil and gas booms, including shale, in resource-rich areas. In particular, using cross-section or panel data analysis, Weber (2012), Fetzer (2014), Gilje et al. (2016), Feyrer et al. (2017) and Allcott and Keniston (2018), among others, show that oil booms have indeed had positive spillovers on production and employment in oil rich regions in the U.S. However, none of these papers analyse effects outside the local areas. Our hypothesis is that these effects are not only important for the oil rich regions, but may also imply wider benefits to the U.S. economy, thereby changing the transmission mechanism of the

[1] The seeds to the shale oil boom were planted already in the 1970s when the U.S. government decided to fund R&D programs and provide tax credits to enterprises developing unconventional natural gas. Still, it was not before Mitchell Energy experimented with new techniques for drilling shale oil in the early 2000s that the process escalated, and the natural gas boom spread to oil (Wang and Krupnick, 2013).

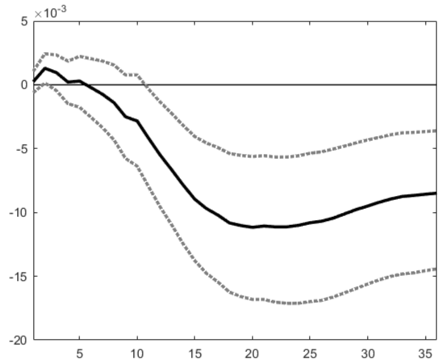

(a) IP Index: U.S. 1974-2009

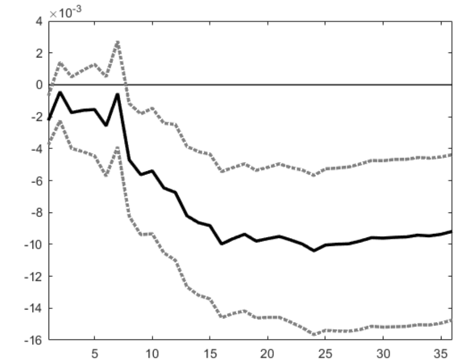

(b) IP Index: U.K. 1974-2009

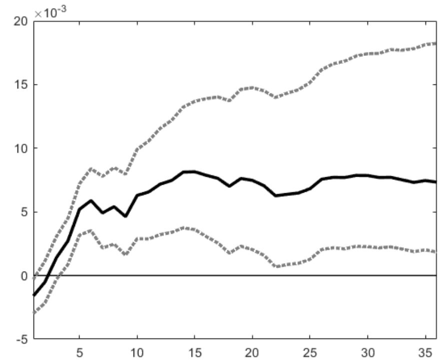

(c) IP Index: U.S. 2010-2018

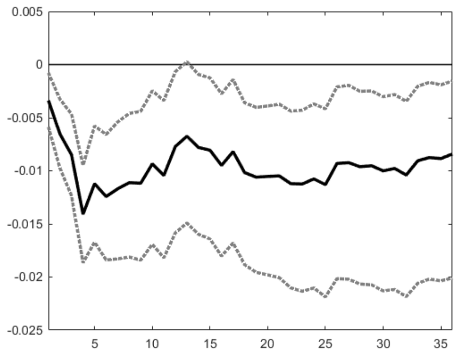

(d) IP Index: U.K. 2010-2018

Figure 1. The effect of an oil-specific shock (normalized to increase oil prices by 10%): Impulse responses from VAR models with crude global oil production (less U.S. production), global activity (GECON indicator), the real oil price and the country's industrial production index. Shocks are identified using recursive restrictions, (see Section 2 for details on data and identification of shocks). Upper row: sample period 1974:M1-2009:M12; lower row: sample period 2010:M1-2018:M12. Monthly frequency. Point estimate (solid line) and 84% confidence intervals (dashed lines). The confidence intervals were constructed using wild bootstrap. Calculations are based on 10000 iterations.

oil price shocks.

In this paper we set out to analyze if the effects of oil market shocks on the U.S. economy have changed with the shale oil boom, using a large data environment of disaggregated and aggregated data. In so doing, we take into account spillovers between the oil industry and non-oil industries across U.S. states, while also allowing the dynamics to vary over time. Figure 1 motivates our claim. It compares the effect of an oil-specific shock on industrial production in the U.S. and the U.K. To identify the oil market shocks, we estimate a small structural VAR model with oil production, global activity, real oil prices and industrial production, identified in a recursive manner as in Kilian (2009). The two countries are interesting to compare, as both are large, open and important oil producers. However, while the U.S. has seen oil production increase steadily since the early

2000s, oil production in the U.K. has declined in this period, and by the middle 2000s, the U.K. switched from being a net oil exporter to a net oil importer. In the figure, we compare results during two different subsamples: 1974-2009 and 2010-2018.[2] The graph suggests that the effects of an oil-specific shock on U.S. industrial production has changed over time, from being negative in the 1980-90s, to being positive from the mid 2000s. For the U.K., the responses are persistently negative, i.e., we find no evidence of time-varying changes. If anything, responses have become more negative recently.

These results are illustrative. In order to consistently analyze time-varying effects of oil market shocks across industries and geographical areas, we need a multivariate framework that incorporates the oil sector, accounts for heterogeneity in several dimensions and allows effects to change over time. Previous times series studies addressing this issue for the U.S. have typically been aggregate and focus on only a few macroeconomic variables. Furthermore, most often they rely on time-invariant regressions. Thus, their maintained assumption is that the effect of a shock to oil prices has not changed over time, and that the role of the oil sector is of little importance when analysing the dynamic effects of oil prices on the U.S. economy (see for instance Baumeister and Kilian (2016), who analyze the effects of the 2014-16 oil price decrease).[3] On the other hand, the recent cross-section or panel data analysis referred to above can account for dynamic effects at the local level in resource abundant states, but ignore spillovers outside the resource rich areas.

We combine both the times series and cross section dimensions, while also allowing these effects to change over time. In so doing, we investigate whether the effects of a shock to oil prices have changed during the last two decades in the U.S. For this purpose we specify and estimate a time-varying parameter (TVP) factor-augmented VAR (FAVAR) model with stochastic volatility, building on the seminal contributions used to analyze the effects of monetary policy shocks. examples include Bernanke et al. (2005) and Stock and Watson (2005) regarding FAVAR models, Primiceri (2005) for time-varying VARs, and Korobilis (2013), Baumeister et al. (2013) and Ellis et al. (2014) for time-varying FAVAR models. Our model also allows us to estimate and control for other factors that may have changed the oil-macro relationships, such as increased global demand, changes in oil production or lower interest rates.

Doing so, we find substantial changes in the way oil prices are transmitted to the U.S. economy. In contrast to previous studies, our analysis suggests that an increase

[2] The split date reflects the period when production of tight oil in the U.S. took off. However, results are robust to splitting the sample a bit earlier or later, and results can be obtained at request.

[3] Baumeister and Kilian (2016) show that while real oil investments did decline following the oil price collapse in 2014, private real consumption and non-oil related business investments were positively stimulated, offsetting the negative setback from the oil sector. They therefore conclude that nothing has really changed: the U.S. still responds like a typical net oil importer.

in the oil price due to an oil-specific shock has positive spillovers to many parts of the U.S. economy, effects that were not present before the shale oil boom. In particular, we find that nonresidential business investment, industrial production and (non-oil) employment in both oil-producing and some manufacturing-intensive states increase following oil-specific shocks that increase oil prices. What is more, real personal income also increases temporarily. The reason is simply that the U.S. has increased its reliance of oil, not as a consumer, but by becoming the world's largest oil producer. On the other hand, we find that several sectors respond negatively as before. In particular, energy intensive industries such as motor vehicles, still respond negatively to an oil-specific shock, as do average consumption, most likely due to higher costs. Going forward, policymakers need to take into account that the transmission of oil price shocks in the U.S. has changed with the shale oil boom, and that there are heterogeneous effects across industries and U.S. states.

Our paper relates to and combines several approaches already developed in the literature, but in a separate manner. First, we relate to the above mentioned literature analysing local effects of oil booms in the U.S. (cf. Weber, 2012; Allcott and Keniston, 2018; Fetzer, 2014; Feyrer et al., 2017; Gilje et al., 2016). However we also relate to studies documenting heterogeneous responses to oil price shocks across net oil importers and exporters (cf. Bjørnland, 2000; Jiménez-Rodríguez and Sánchez, 2005; Peersman and Robays, 2012; Aastveit et al., 2015; Herrera et al., 2015; Guerrero-Escobar et al., 2019; De Michelis et al., 2020). These studies focus on various small and large open economies, being either classified as oil exporters or importers. However, this paper is the first comprehensive examination of how the shift from being an importer to an exporter influences a country's reaction to oil price shocks.[4]

Second, we relate to the large literature that analyses the effect of oil price shocks, emphasising different sources of shocks and identification methods (cf. Bjørnland, 2000; Hamilton, 2009; Kilian, 2009; Kilian and Murphy, 2012, 2014; Kilian and Vigfusson, 2011; Lippi and Nobili, 2012; Peersman and Robays, 2012; Cashin et al., 2014; Aastveit, 2014; Aastveit et al., 2015; Stock and Watson, 2016; Baumeister and Hamilton, 2019; Känzig, 2021). However, in contrast to these papers, we allow for changing dynamics. Furthermore, while these studies typically focus on aggregate macroeconomic variables, we explicitly include data for the oil sector, disaggregate production and state-level employment into the analysis to account for potential spillovers and resource movement due to

[4]Although the U.S. is not yet a net oil exporter, the massive growth in oil production (and exports) over a prolonged period supports our claims. In particular, the proceeds from oil have increased rapidly, making the net present value of the shale boom close to 9 percent of U.S. GDP in 2015 (see Figure 10 in Section A in the Online Appendix). This makes the shale oil boom comparable in relative size to some of the larges oil discoveries in the world (see Arezki et al., 2017).

the shale oil boom. For this purpose, we use a FAVAR model with a large data set and time-varying parameters.

Third, our TVP framework builds on a growing literature allowing for time-varying parameters and stochastic volatility when analysing the effect of oil price shocks (i) on the U.S. macroeconomy (e.g. Baumeister and Peersman, 2013b; Bjørnland et al., 2018), (ii) on the inflation passthrough (e.g. Clark and Terry, 2010), (iii) on the U.S. stock market (e.g. Kang et al., 2015; Foroni et al., 2017) and (iv) on the oil market (e.g. Baumeister and Peersman, 2013a). However, this is the first study to examine the time-varying changes when a country is transforming from being an importer to an exporter of oil.

Finally, we relate to a branch of the literature that has documented important heterogeneous effects in the transmission channels of oil price shocks to disaggregate industries (e.g. Bresnahan and Ramey, 1993; Davis and Haltiwanger, 2001; Lee and Ni, 2002; Herrera and Karaki, 2015; Herrera et al., 2017). However, while these papers have primarily studied how the negative effects of an oil price shock are transmitted to industries when the U.S. was an oil importer, our focus is to unravel potential heterogeneous effects at the industry level and across U.S. states, following the shale oil boom. To the best of our knowledge, this is the first paper that models the interaction between the oil market and the U.S. economy in a large data environment, allowing also for time-varying changes during the fracking revolution.

The remainder of the paper is structured as follows: Section 2 introduces the model framework and the dataset. In Section 3 we discuss empirical results, focusing on the effects of an oil-specific shock on various industries, the overall macroeconomy, and the geographical dispersion of shocks at the state level in terms of employment. Section 4 demonstrates that our findings cannot be attributed to alternative hypotheses, confirming the robustness of our results to variations in model specification and identification framework. Section 5 provides concluding remarks.

2 Empirical modeling framework

While most of the reserves of shale oil in the United States have been known for decades, it was long thought to be too costly and technologically impossible to extract. The breakthrough in technological innovation in the early-2000s, combining horizontal drilling with hydraulic fracturing, allowed oil to be extracted from shale formations on a large scale. From the middle of the 2000s, imports of crude oil to the U.S. plummeted as the shale oil boom sparked a strong recovery in the domestic production of crude oil. By 2015 the U.S. had surpassed Russia and Saudi Arabia to become the world's biggest producer of crude oil and natural gas, and by the end of 2015, the export ban on crude oil was lifted.

This extraordinary boom motivates our question: namely, to what extent has the transmission of oil market shocks in the U.S. changed with the shale oil boom? To analyze this question, we specify an empirical model that can account for: (i) heterogeneous responses to oil price shocks across the U.S.; (ii) spillovers between oil and non-oil industries; and (iii) time-varying responses. Many recent papers, as cited in the introduction, use SVAR models to study the effects of oil price shocks on the U.S. economy. As we aim to assess the oil industry's role in dispersing oil market shocks to economic activity, we enhance a standard aggregate VAR model by incorporating estimated factors that reflect dynamics from both oil and non-oil sectors.

To that end, we specify a factor-augmented vector autoregressive (FAVAR) model that includes both observable and unobservable factors. The observable factors will be driven by shocks that have the potential to affect all sectors of the U.S. economy. To also take into account the fact that there may be heterogeneous responses across U.S. industries, we estimate separate latent factors for the U.S. economy that explain a sufficient amount of variation in the data. The inclusion of latent factors also enables us to simultaneously estimate spillovers between different industries and states in the U.S. The simultaneous spillovers between different sectors at different geographical levels can not be captured by including only observable variables in a small panel of data and have therefore not been taken into account in previous studies.

We utilize factors in a time-varying parameter (TVP) Vector Autoregressive model, which features both time-varying coefficients and a time-varying variance-covariance matrix of innovations. This approach allows us to account for potential non-linearities or time-dependent variations between oil prices and the U.S. economy. We also address possible heteroscedasticity of structural shocks and nonlinearities in the simultaneous relations among variables by incorporating multivariate stochastic volatility. Our decision to adopt a TVP approach stems from our belief that the transition of the U.S. from a net oil importer to a major oil producer occurs gradually. This gradual transition is better captured through the TVP model, which allows for smooth changes in the shock transmission mechanism, rather than a model with discrete breaks. Altogether, this framework enables us to thoroughly investigate whether the transmission of oil market shocks to the U.S. economy has undergone significant changes over time.

2.1 The time-varying FAVAR model

Our framework builds on the FAVAR model, first proposed by Stock and Watson (2005) and Bernanke et al. (2005). Technically, the developed and employed model is most closely related to the setup used in Korobilis (2013). In particular, we use a two-step estimator and replace the factors with the first principal components obtained from the singular value decomposition of the data matrix, and consequently treat them as observables. These factors are then used in a time-varying VAR model with both time-varying coefficients and time-varying variance covariance matrix of innovations (see Primiceri, 2005).[5]

Still, we deviate from Korobilis (2013) in several important ways. First, while Korobilis (2013) uses a framework based on Bernanke et al. (2005) and Belviso and Milani (2006) to identify the factors, we follow Boivin and Giannoni (2007) and Boivin et al. (2009).[6] Doing so, we impose the constraint that the three observable variables are equivalent to the three factors in the first-step estimation, guaranteeing that the estimated latent factors identify dynamics not already captured by the three observable variables. Second, to keep our model as parsimonious as possible, we do not allow for stochastic volatility in the factor analysis regression. Finally, we stick to the standard convention in the literature and model the random walk evolution of the VAR parameters as in Primiceri (2005).

Let F_t be a $m \times 1$ vector of common factors assumed to drive the dynamics of the economy. In our application, F_t contains both observable factors y_t of dimension $l \times 1$, and unobservable latent factors, f_t of dimension $k \times 1$, such that $F_t = \begin{pmatrix} y_t \\ f_t \end{pmatrix}$ and $l + k = m$. The latent factors are extracted from a larger dataset X_t of dimension $n \times 1$, and assumed to summarize additional information not captured by the observable factors. We assume that X_t can be described by an approximate dynamic factor model given by

$$X_t = \Lambda F_t + e_t, \tag{1}$$

where Λ is $n \times m$ matrix of factor loadings and $e_t \sim \mathcal{N}(0, R)$, is $n \times 1$ vector of errors assumed to be uncorrelated with the factors F_t and mutually uncorrelated.

The joint dynamics of the factors F_t are given by the following transition equation:

$$F_t = c_t + b_{1t} F_{t-1} + ... + b_{pt} F_{t-p} + u_t, \tag{2}$$

where c_t is an $m \times 1$ vector of time-varying intercepts; b_{jt} are $m \times m$ matrices for $j = 1, .., p$ of time-varying coefficients; u_t is an unconditionally heteroskedastic disturbance term that is normally distributed with zero mean and time-varying covariance matrix Ω_t. According to the literature on efficiently parametrizing large covariance matrices, Primiceri (2005), we decompose Ω_t in the following way:

$$\Omega_t = A_t^{-1} \Sigma_t \Sigma_t' (A_t^{-1})', \tag{3}$$

[5]However, when estimating the model we modify the algorithm of Primiceri (2005) to reflect the correction detailed in Del Negro and Primiceri (2015) (see Section B in the Online Appendix for details).

[6]Bernanke et al. (2005) and Belviso and Milani (2006) perform a transformation of the principal components, exploiting the distinct behaviors of "slow-moving" and "fast-moving variables", this approach, while effective for monthly data, is not suitable for quarterly data.

where Σ_t is a diagonal matrix that contains the stochastic volatilities and A_t is a unit lower triangular matrix with ones on the main diagonal that models the contemporaneous interactions among the variables in (2):

$$A_t = \begin{bmatrix} 1 & 0 & \cdots & 0 \\ a_{21,t} & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ a_{m1,t} & \cdots & a_{m(m-1),t} & 1 \end{bmatrix}, \Sigma_t = \begin{bmatrix} \sigma_{1,t} & 0 & \cdots & 0 \\ 0 & \sigma_{2,t} & \ddots & 0 \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & \sigma_{m,t} \end{bmatrix}. \tag{4}$$

It follows that

$$F_t = c_t + b_{1t}F_{t-1} + ... + b_{pt}F_{t-p} + A_t^{-1}\Sigma_t\varepsilon_t, \tag{5}$$

where $\varepsilon_t \sim \mathcal{N}(0, I_m)$; I_m is an m-dimensional identity matrix.

Our model setup captures time variation by allowing the coefficients and the error covariance matrix in the transition equation to evolve over time, with factor loadings in the factor equation remaining constant. This approach aligns with our objective to explore broad, systemic changes in the economic structure from aggregate shocks in the oil market, rather than the detailed interactions between specific variables and factors. An alternative method would involve time variation in both factor loadings and transition equation coefficients (see Koop and Korobilis, 2014). However, we opt to focus on the spillover effects and their impact on the U.S. economy through changes in the transition equations, which simplifies the model and reduces computational complexity and identification challenges.

We follow the standard convention and assume that model's time-varying parameters and stochastic volatilities follow random walk processes. Let $B_t = (vec(c_t)', vec(b_{1t})', ..., vec(b_{pt})')'$ be the vector of all R.H.S. coefficients in (5), $\alpha_t = (a'_{j1,t}, ..., a'_{j(j-1),t})'$ for $j = 1, ..., m$ be the vector of nonzero and nonunity elements of the matrix A_t, and $\sigma_t = (\sigma'_{1,t}, ..., \sigma'_{m,t})'$ be the vector containing the diagonal elements of Σ_t .The dynamics of the three processes are specified as follows:

$$B_t = B_{t-1} + \eta_t^B \tag{6}$$

$$\alpha_t = \alpha_{t-1} + \eta_t^\alpha \tag{7}$$

$$log\sigma_t = log\sigma_{t-1} + \eta_t^\sigma \tag{8}$$

While the elements of the vector B_t and the free elements of the matrix A_t are modeled as random walks, the standard deviations (σ_t) are assumed to evolve as a geometric random walks. As discussed by Primiceri (2005) and Bianchi et al. (2017), the random walk assumption, despite its potential drawbacks, offers the benefit of emphasizing permanent shifts and minimizing the number of parameters to be estimated.

We assume that innovations in the model are jointly normally distributed with the following assumptions on the variance covariance matrices:

$$Var\left(\begin{bmatrix} e_t \\ \varepsilon_t \\ \eta_t^B \\ \eta_t^\alpha \\ \eta_t^\sigma \end{bmatrix}\right) = \begin{bmatrix} R & 0 & 0 & 0 & 0 \\ 0 & I_m & 0 & 0 & 0 \\ 0 & 0 & Q & 0 & 0 \\ 0 & 0 & 0 & S & 0 \\ 0 & 0 & 0 & 0 & W \end{bmatrix} \tag{9}$$

Following Primiceri (2005), we postulate a block-diagonal structure for S, with blocks corresponding to parameters belonging to separate equations. Thus, the shocks to the coefficients of the contemporaneous relations among variables in (5) are assumed to be correlated within equations, but uncorrelated across equations. While less restrictive assumptions concerning Equation 9 might be justifiable (see Primiceri, 2005) for a more comprehensive discussion, they would also increase computational complexity and burden of already heavily parameterized model. In our setting we therefore utilize the SVAR framework, where we assume that the structural shock ε_t are independent of other disturbances in the model.

2.2 Data and transformations

To accommodate the effects of oil price shocks on the U.S. economy we use a large panel of domestic and international quarterly data series. A full description of the data can be found in Section A in the Online Appendix. In short, for our benchmark model, we include a broad range of domestic macroeconomic indicators as observable variables. Among others, we include consumer and producer prices, real investment series, stock prices, real personal income, real consumption, the short term interest rate and various disaggregate industrial production (IP) series. In addition we include relevant variables for the oil sector, including mining, mining investment and production of conventional and unconventional oil. For unconventional oil we have data for tight oil, which refers to all unconventional resources, reserves, and production associated with low-permeability formations that produce oil, including shale formations (see the U.S. Energy Information Administration (EIA)). In total we use 42 domestic series in our benchmark model (i.e., $n = 42$). These are used to extract unobservable latent factor.

In addition to the domestic series, we include three observable "foreign" factors ($l = 3$) in the benchmark model; world crude oil production, excluding U.S. crude oil production, global activity and the real oil price. World crude production net U.S. production is used as a measure of foreign supply that are not U.S.-specific. We include the global economic conditions indicator (GECON) as a measure of global activity as proposed by

Baumeister et al. (2020). For the real oil price, we follow Lee and Ni (2002) and Herrera (2018), among many others, and use the U.S. refineries acquisition cost deflated by the U.S. CPI. However, we also analyze extensive robustness to our choice of variables in Section 4, by among other things replacing the chosen global activity variable with an estimate of industrial production for the OECD plus other major countries published by OECD Main Economic Indicators, and extended from November 2011 by Baumeister and Hamilton (2019). In addition, we also replace the chosen oil price with Western Texas Intermediate (WTI). In sum, this gives a panel of 47 international and domestic data series, covering a sample period from 1974Q1 to 2018Q4. With quarterly data, we use four lags in the estimation, consistent with Hamilton and Herrera (2004).[7] Correcting for the number of lags ($p = 4$) and the sample used to estimate the priors[8], leaves us with 111 observations that we use for estimation, covering the sample 1991:Q1–2018:Q4.[9]

Finally, to account for local effects at the state level, we re-estimate our model using an extended dataset. More specifically, instead of using the aggregate employment for the U.S., we include employment series for all U.S. states. Due to availability of disaggregated employment series, the sample period starts in 1990Q1. With the additional disaggregate series, we re-estimate the domestic factors. For oil-producing states, we subtract employment in the mining sector from total non-farm employment, and include these series separately. There are 16 significant oil-producing states in our sample. The most important shale producing states are (in alphabetical order) California, Colorado, Kansas, Louisiana, New Mexico, North Dakota, Ohio, Oklahoma, and Texas (cf. Feyrer et al., 2017). In addition, we include five states that have significant production of conventional oil, defined to be on average more than 10 million barrels during 2010-2018. The five states are (in alphabetical order) Alaska, Mississippi, Montana, Utah and Wyoming. Finally, we also include two states where oil production increased significantly during the last part of our sample. These states are Pennsylvania and West Virginia. The remaining states have trivial production relative to these states (see Figure 11 in Section A in the Online Appendix).

In total, this extended model gives us a domestic data block with 109 series ($n = 109$), from which we extract the latent factors. In sum, we have a panel of 113 international and domestic data series, covering a sample period from 1990Q1 to 2018Q4.[10]

[7] Hamilton and Herrera (2004) show that a too restrictive lag length can produce misleading results, while increasing the lag length above one year has negligible effects. In Section C in the Online Appendix we conduct a quasi-real-time forecasting experiment varying the number of lags. The findings indicate that our benchmark model using four lags outperforms some alternative lags specifications.

[8] We use the first 16 years as a training period to estimate priors (see Section D in the Online Appendix).

[9] In Section 4 we make several extensions and analyse robustness to our modelling choices.

[10] We use the same number of lags as in the baseline model ($p = 4$) and the first 10 years of the sample are used as a training period to estimate priors. The model is estimated over the sample 2000:Q3–2018:Q4.

Finally, all the series were initially transformed to log difference, except for GECON, that is already transformed and given in levels, and the interest rate series, where we take first differences. The data is then demeaned using the full sample. The domestic block of series used to extract factors is also standardized as is common practice in these type of analyses.

2.3 Identification

We estimate a model with both observable and unobservable factors, and with associated shocks that have the potential to affect all sectors of the U.S. economy. As described above, to be able to capture the oil price movements associated with unanticipated supply disruptions that are not U.S. specific, we include global oil production as the first foreign factor. We use global activity to capture global demand as the second foreign factor, and the real price of oil is the third foreign factor. These three factors are treated as observables. Further, we allow for two latent factors that capture different parts of the domestic activity in the U.S. and that are inferred from data.

Following the seminal paper of Kilian (2009), we identify three oil market shocks in a recursive manner; *flow supply* shocks, *flow demand* shocks and *oil-specific* shocks. First, as oil production is ordered first, it implies that crude oil supply can not respond to innovations to the global demand or other oil-specific price movements within the same quarter. Second, innovations to the global activity that are not captured by flow supply shock will be explained by the flow demand shock. Here we follow the usual assumption from the models of oil markets, and restrict global activity from responding to the oil-specific shocks at impact (see e.g., Hamilton, 2009). In turn, any unexpected news regarding oil production or global demand can affect oil prices contemporaneously.

While our identifying restrictions are consistent with those used in Kilian (2009) and many subsequent studies (i.e., Aastveit et al., 2015), our restrictions can also be questioned. In particular, as we are using quarterly data, the assumption that production of oil can only respond to oil-specific shocks with a lag, (i.e., the supply schedule is inelastic in the short run) is questionable. Such dogmatic (in this case exclusion) restrictions have also recently been criticized by Baumeister and Hamilton (2019) on the grounds that there is some uncertainty about the identifying restrictions themselves. In this case, the worry would be that if oil producers react to price signals within a quarter, we would underestimate the importance of oil supply shocks. So far, however, there is little evidence to suggest that conventional oil production is elastic in the short run (cf. Anderson et al., 2018). For shale oil producers, however, the situation is different. As pointed out by Bjørnland et al. (2021), Bornstein et al. (2022) and Aastveit et al. (2022), shale oil production responds more rapidly to oil price signals than conventional oil production, as shale wells have a more flexible production technology. Our model is consistent with these findings. In particular, while global oil production (primarily conventional oil), is restricted from responding to oil price signals within the quarter, U.S. oil production is permitted to respond on impact, as the series is included in the domestic data block (see below). Given that the U.S. series is dominated by shale oil, this makes sense. Still, in Section 4 we analyze robustness to the identification, by among other using monthly data.

Turning to the domestic economy, we assume shocks to the U.S. factors can not affect the three foreign factors on impact. Hence, the oil price is predetermined with respect to the U.S. variables, in line with findings of Kilian and Vega (2011). Finally, note that all observable variables in the vector X_t may respond to all shocks on impact inasmuch as they are contemporaneously related to the factors through the loading matrix, Λ.

2.4 Estimation

Following Stock and Watson (2005) and Korobilis (2013), we estimate our model using a conceptually and computationally simple two-step estimation method. In the first step, we estimate the space spanned by the factors using the approach advocated by Boivin and Giannoni (2007), to ensure that the estimated latent factors, f_t, will recover dimensions of the common dynamics not already captured by the observable variables, y_t. Once we have estimated the factors, we treat them as observables.[11] We then advance to the second step, where the model parameters are determined conditional on these factor estimates, employing Bayesian techniques for the estimation. To address the complexity of drawing from the joint posterior of our desired parameters, we employ Gibbs simulations. These simulations are a specific type of Markov chain Monte Carlo (MCMC) method, designed to sample a high-dimensional joint posterior by sequentially drawing from its lower-dimensional conditional posteriors. Given the data and the priors, Gibbs sampling is carried out in four blocks. In the first block, the parameters in the factor equation are sampled using standard arguments for linear regression models (see Koop, 2003). In blocks two to four, we draw the VAR model's unobserved state variables as well as the hyperparameters. A detailed description of the estimation procedure is given in Section B in the Online Appendix. In Section 4.2 we also provide a range of sensitivity analyses concerning prior specifications and model selection.

To generate the posterior draws, the simulations in this paper are all based on 90,000 iterations of the Gibbs sampler. The first 40,000 are discarded and only every tenth of the remaining are used for inference. In Section E in the Online Appendix we provide evidence of convergence of the Markov chain Monte Carlo Algorithm.

[11] As it is shown by Bai and Ng (2002a), when the number of series are large relative to the number of observations, the effect of estimated regressors can be completely ignored.

The system is estimated using three observable factors, with the number of unobservable factors set to two ($k = 2$). According to the criteria for static factors set forth by Bai and Ng (2002b), the suggested range for factors is between 1 and 8. The optimal factor counts, as determined by the ICP criteria tailored for factor estimation via principal components, are $ICP_1 = 3$, $ICP_2 = 2$, and $ICP_3 = 7$. In our benchmark configuration, five factors account for nearly half of the variance in X_t. The addition of an extra unobservable factor incrementally increases the explained variance by approximately 5 percent. However, employing more than two factors does not significantly alter our primary results and heightens computational requirements. Thus, we set the number of extracted unobservable factors in line with the most conservative recommendation provided by the ICP criteria.[12]

As discussed above, the five factors are included to capture different aspects of relevance to the U.S. economy. While the three observable factors are easily interpretable insofar as they capture oil supply (unrelated to the U.S.), global activity and the oil price, the two latent factors are unobservable, and not identified. Tables 3 and 4 in Section A in the Online Appendix report the correlation (above 0.4) with various variables and either the first or the second domestic factor, respectively. From Table 3, we observe that the first factor turns out to be a good proxy for real non-oil activity in the U.S. It shows a large correlation with non-oil variables and a weak negative correlation with U.S. oil production and exports. However, there is also a noticeable positive correlation with certain oil-related series, such as mining activity. In contrast, Table 4 shows that the second domestic factor closely mirrors changes in oil-related variables, suggesting its role as an oil activity indicator.

3 Empirical results

The aim of this paper is to analyze if the transmission of oil price shocks on the U.S. economy has changed with the shale oil boom. As we are analysing time-varying changes, we will report several types of impulse responses. First, as the impulse response may be different across time, we show median impulse responses at different selected dates: 1995:Q1, 2001:Q1, 2007:Q1, 2011:Q1, 2013:Q1, 2015:Q1, and 2017:Q1. The first two dates represent a period prior to the shale boom, 2007 is a year just before shale oil increased substantially, while the remaining years represent the shale boom period itself with two-year intervals. Beyond that, the dates are chosen arbitrarily and are not crucial for our

[12] For the extended model, all three ICP criteria concur that the optimal number of factors is two. Therefore, we specify that the extended model maintains the same factor count as the benchmark, explaining more than 50 percent of the variance in X_t.

conclusion. Second, we also report the impulse responses after two and four quarters across time with probability bands. In so doing we emphasize the maximum effect of an oil price shock, which typically occurs after about two to four quarters according to Hamilton (2008), Clark and Terry (2010), and Herrera (2018), among others. However, our conclusions are robust to alternative horizons. Third, we also display the difference in impulse responses over time, measured at the two- or four-quarter horizon, with 16-th and 84-th percentiles. For each draw in the posterior distribution, we compute the statistical difference between the current impulse responses and impulse responses in 1995:Q1.[13,14]

To ensure that we compare an equal sized shock over time, we normalize the dynamic effects of an oil-specific shock to a 10 percent increase in the oil price on impact (for all the calculated responses).[15] All the estimated responses are accumulated and shown in levels. Note that since the domestic data series are standardized, these represent changes relative to the base level, not the absolute original levels of the data. Consequently, the magnitude of the impulse responses should only be compared across variables.

Importantly, there are three type of shocks that can alter the oil price on impact; flow supply, flow demand and oil-specific shocks. Here we will focus on effects of the oil-specific shocks, as this is where we expect the time varying behavior (cf. Figure 1). However, in Section 4, we discuss results for the two other shocks. Importantly, both shocks increase oil prices, as expected. However, the effects on the U.S. variables are either negligible (flow supply shocks) or positive, but stable over time (flow demand shocks).

3.1 Real oil price and global activity

Before displaying results for the U.S., we examine how real oil prices and global activity respond to oil-specific shocks. Importantly, we want to learn if there are time varying changes also in other countries. This could be due to say, less energy dependence in production or consumption, or some other common global explanatory factors, such as the lower interest rates (zero lower bound) that may have damped the negative effect of the higher oil price (cf. Datta et al., 2021). If this is the case, we should expect the oil-specific shock to affect global economic more positively, and thereby also explain potential

[13]The base year 1995 has been chosen as the impulse responses in 1995 are close to the average in the pre-shale period. However, results are robust to other choices for base year.

[14]For each Gibbs draw and time point, we calculate the difference between the impulse response at each time and 1995 (base year). To detect significant divergences, we assess whether the zero mark lies outside the 68% confidence interval of these deviations. This is visualized as a plot of the median difference over time, with shaded areas indicating uncertainty. Non-overlapping shading signifies deviations from 1995.

[15]A common way to report impulse responses is to examine a one standard deviation shock. However, in the models where volatility changes over time, a one standard deviation shock corresponds to a different-sized shock at each point in time. Therefore, we normalize the impact effects of the shock over time.

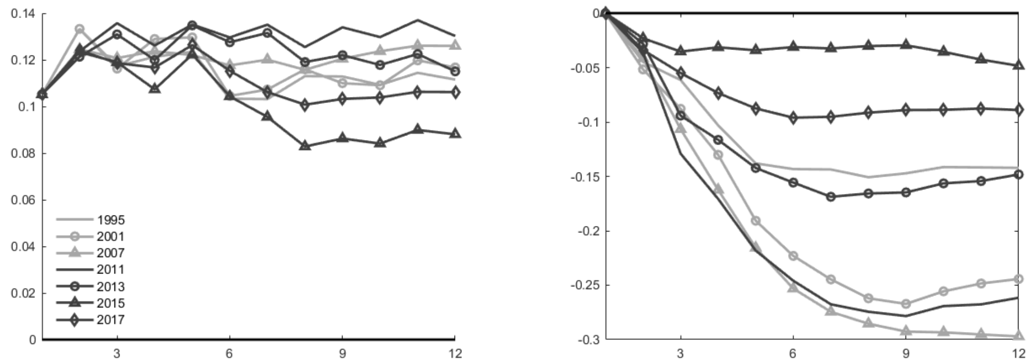

(a) Oil Price (median) **(b)** GECON (median)

Figure 2. The effect of an oil-specific shock: Impulse responses for oil price and global activity (GECON). The initial shock is normalized to increase oil prices by 10% . All responses are reported in levels. Posterior median of impulse responses at seven different periods in time.

time varying behavior in U.S. aggregates.

Figure 2 suggests that this is not the case. The figure shows the median impact of the oil-specific shock on oil prices and global activity at different time intervals in the left and right frame respectively. The figure confirms that an oil-specific shock that increases oil prices, has a persistent negative impact on global activity, in line with what Kilian (2009) and others have found. Furthermore, although the negative effect on global activity may have been slightly muted in recent years, we show in Section 4 that these changes are not statistically significant.

3.2 The U.S. oil sector

Having seen that global activity falls with a higher oil price, we now turn to examine detailed responses in the U.S., starting with the the oil sector. To that end, Figure 3 displays the impact of the oil-specific shock on mining investment and mining activity, focusing on median responses at different time intervals (left column) and time-varying responses after four quarters (right column) (c.f. the explanation above).

To the extent that higher oil prices also generate a resource boom[16] in the U.S. economy, we should expect to see investment and production in the oil sector to increase. And we do (cf. Figure 3). The responses of mining investment and production to an oil-specific shock are not only positive, but have also increased gradually, and in particular since 2008/09. The maximum effect seems to be in 2015, most likely as the industry was catching up after the oil decline the year before. The effects have come down slightly since

[16] A resource boom takes the form of either a new oil discovery, more productive oil fields or higher real oil prices (see Corden, 1984).

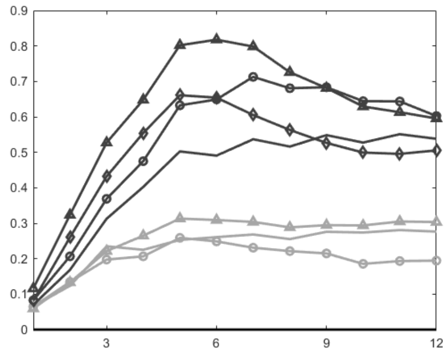

(a) Mining Investment (median)

(b) Mining Investment (after four quarters)

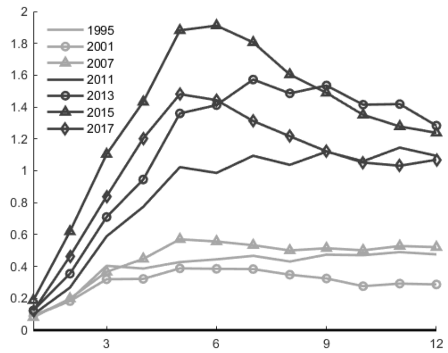

(c) Mining (median)

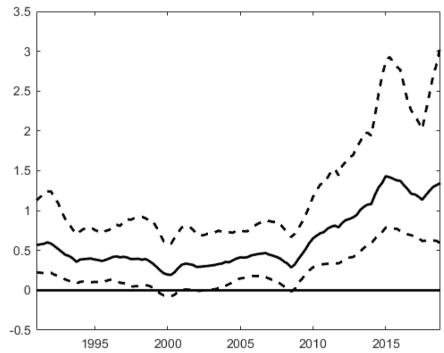

(d) Mining (after four quarters)

Figure 3. The effect of an oil-specific shock: Impulse responses for mining investment (top row) and production (bottom row). The initial shock is normalized to increase oil prices by 10%. All responses are reported in levels (of the standardized data). We display posterior median of impulse responses at seven different periods in time in the left column. In the right column we report responses across the sample four quarters after the shock. The dashed lines represent 68% posterior probability bands. The solid line is the median estimate.

then, but are still significant positive (see the right panels). These results are in line with our expectations: a higher oil price makes it more profitable for firms operating in the oil sector to produce oil, and therefore stimulates investment and subsequently production. Hence, for an equally sized increase in oil prices, investment and mining increase slightly more now than before the shale oil boom. We will delve deeper into the oil sector by separating effects for tight and conventional oil in Section 3.5 when we are analysing the responses at the state level. We conclude for now that the time-varying pattern observed in the mining sector seems closely related to the shale boom.

3.3 The aggregate U.S. economy

Having seen that oil investment and production respond more positively to an oil-specific shock, we turn to the non-oil sector to analyze indirect effects and spillovers. We start

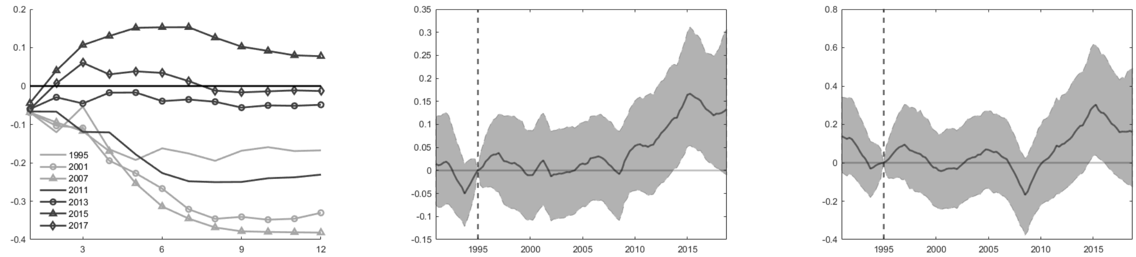

(a) Investment (median) **(b)** Investment (diff two quarters) **(c)** Investment (diff four quarters)

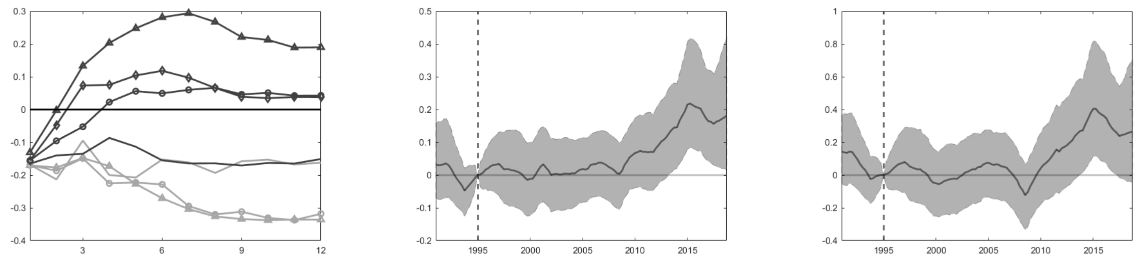

(d) Income (median) **(e)** Income (diff two quarters) **(f)** Income (diff four quarters)

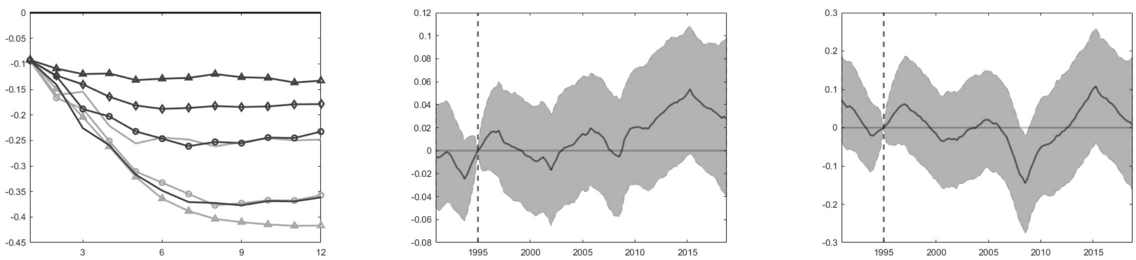

(g) Consumption (median) **(h)** Consumption (diff two quarters) **(i)** Consumption (diff four quarters)

Figure 4. The effect of an oil-specific shock: Impulse responses for non-residential, non-oil, investment, personal income and private consumption expenditure. The initial shock is normalized to increase oil prices by 10%. All responses are reported in levels. Left column: Posterior median of impulse responses at seven different periods in time. Middle and right columns: The difference between the responses in period 1991:Q1-2018:Q4 and the responses in 1995:Q1 after two and four quarters respectively. The solid line is the difference between the median estimates. The shaded area represents 68% posterior probability bands for the difference in impulse responses.

by examining real activity. Figure 4 presents results for (non-oil) real investment, income and consumption. We present the median response (left column) and the difference between impulse response after two and four quarters in the period 1991:Q1-2018:Q4 and 1995:Q1, with 16-th and 84-th percentiles (middle and right columns). The last two figures are included to analyze if the changes are statistical significant. The figure shows important time-variation in the median response around the period of the shale oil boom. In particular, in the first row, we find that non-residential (non-oil related) investment has responded systematically more positive to an oil-specific shock throughout the 2000s, and the effect is significant and positive from 2012/2013 (cf. Figure 4 (b) and (c)). Hence,

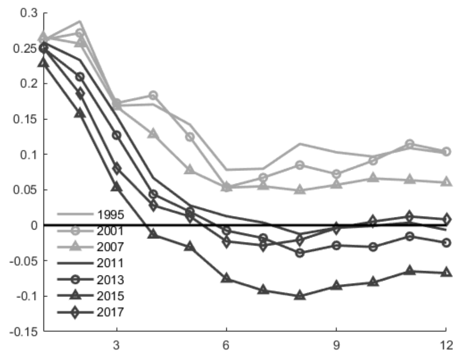

(a) CPI (median)

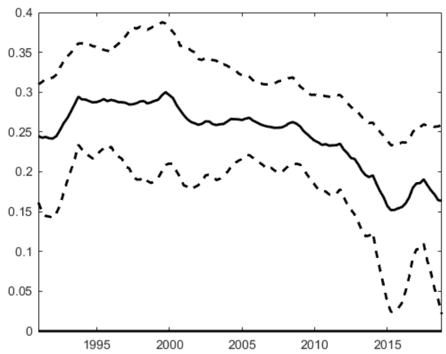

(b) CPI (after two quarters)

Figure 5. The effect of an oil-specific shock: Impulse responses for CPI (All Items). The initial shock is normalized to increase oil prices by 10%. The responses are reported in levels (of the standardized data). We display posterior median of impulse responses at seven different periods in time in the left column. In the right column we report responses across the sample two quarters after the shock. The dashed lines represent 68% posterior probability bands. The solid line is the median estimate.

we conclude that while non-oil investment in the U.S. economy previously fell when oil prices rose, it is now picking up. This is a new finding in the literature.

Second, for an oil importing country, a higher oil price is typically manifested in lower real income and purchasing power, as costs and prices increase. This was the case throughout the first part of the sample (see the impulse responses for real income in the middle row of Figure 4). However, from 2012 and onwards, real personal income drifts upwards following an oil-specific shock. This suggests that the increased investment may have had spillover effects on the U.S. economy, outside oil, by creating jobs, boosting industries, and finally contributing to higher real income levels. However, turning to consumption in the lower row, we see that the impulse responses are still consistently negative, although somewhat more muted over time. This is not surprising. Despite higher income, increased energy costs will lead to a decrease in consumer spending, as one has have to allocate more income to cover essential expenses like fuel and heating. Figure 5 confirms such a scenario. It shows that an oil-specific shock is still strongly associated with an increase in CPI, and it is in line with our expectations and previous findings in the literature (cf. Hamilton and Herrera, 2004): higher oil prices lead to higher costs for firms, hence prices rise. We note that there is no evidence of time variation in the way oil-specific shocks are transmitted to CPI, although there seem to have been a gradual decline in the magnitude of the effects, most likely as oil dependence in the U.S. has fallen in general (cf. Clark and Terry, 2010).

Taken together, these results are consistent with the U.S. becoming a major oil producer, where non-oil investment and income in the U.S. now increase following oil-specific shocks that push up oil prices. These effects suggest that there are spillovers from the oil boom to the aggregate economy. These are new results for the aggregate U.S. economy. Importantly, this suggests that higher oil prices may no longer be unambiguously negative for the U.S. economy. In particular, while consumers as still responding to higher oil (and gasoline) prices by lowering consumption, many parts of the U.S. economy are benefiting from the increased activity and income. Below we try to shed more light on these issues by first looking at disaggregate industry responses, before we look into the behaviour of employment across the different states.

3.4 Disaggregate industry effects

We analyze disaggregate responses at the industry level to understand better the mechanisms behind the time-varying changes. There are at least three hypotheses for changing coefficients; (i) through direct purchases of manufactured inputs; (ii) through indirect productivity and LBD spillovers from new knowledge and technology transfers created during the shale revolution; (iii) by creating energy independence and costs savings for U.S. companies that operate in energy-intensive sectors. We do not expect all industries to benefit equally from the shale oil boom. Some may also respond negatively as before, in particular if they are energy intensive in production. Figure 6 illustrates this. In line with Figure 1, it shows that industrial production (IP) has responded systematically more positively to an oil price shock throughout the 2000s, and the effect is significant from 2013/2014. Looking further at sub-groups of IP, we find strong upward drift for business equipment, while the effect for manufacturing production is still negative.

Table 1 analyzes this further by summarizing the response of 21 industry group (see Section A in the Online Appendix for details). Results are based on the impulse responses two quarters after the initial shock. The upper part of the table shows industries where there have been no time-varying changes (i.e., the responses have been either consistently negative (left column), insignificant (middle column) or positive (right column) throughout the sample). The lower part shows industries where the responses have changed significantly from negative to insignificant (left column) or to positive (right column). Overall, we find evidence of heterogeneous responses among different industries. First, for energy-intensive industries (i.e., food, beverage and tobacco products, motor vehicles, and consumer goods) responses are negative, as expected. Higher oil prices increase the cost of production and lead to a decline in demand, all else equal. Furthermore, there is no variation over time. Second, some industries respond significantly positively to an oil-specific shock over the whole period. These are industries that have had strong ties with the oil sector throughout the sample, such as primary metals, mining and defense and state equipment. Third, from the lower part of the table, we see that there are some

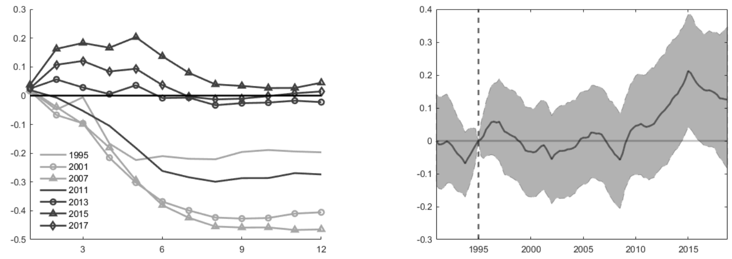

(a) IP Index (median) **(b)** IP Index (diff 2 quarters) **(c)** IP Index (diff 4 quarters)

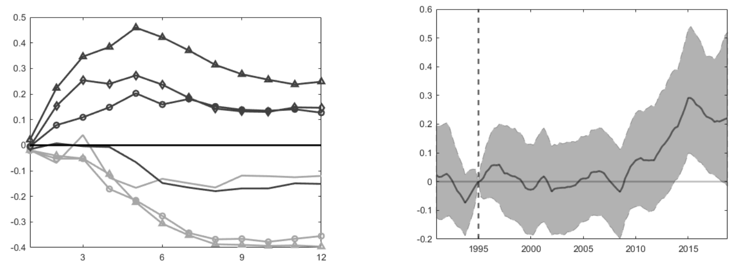

(d) Business Equipment (median) **(e)** Business Equipment (diff two quarters) **(f)** Business Equipment (diff four quarters)

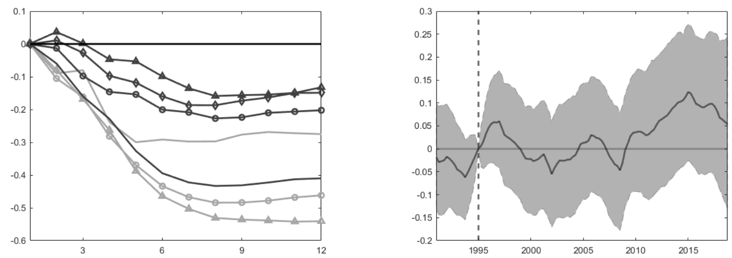

(g) Manufacturing (median) **(h)** Manufacturing (diff two quarters) **(i)** Manufacturing (diff four quarters)

Figure 6. The effect of an oil-specific shock: Impulse responses for industrial production series divided according to market groups. The initial shock is normalized to increase oil prices by 10%. All responses are reported in levels (of the standardized data). Left column: posterior median of impulse responses at seven different periods in time. Middle and right columns: the difference between the responses in period 1991:Q1-2018:Q4 and the responses in 1995:Q1 after two and four quarters respectively. The solid line is the difference between the median estimates. The shaded area represents 68% posterior probability bands for the difference in impulse responses.

industries that have observed changing coefficients. In particular, we find a systematically and significant more positive response for industries such as petroleum and coal products and fabricated metal products, that benefit directly from shale. In addition, technological advanced industries, such as aerospace, machinery, and to some extent electronic equipment and appliances and chemicals, are now also responding positively.[17]

[17]In Figure 12 in Section F in the Online Appendix, we show responses to subgroups of investment. Consistent with the responses reported above, we find investment in manufacturing, industrial and other equipment responds positively to the oil-specific shock after 2012/13, while investment in transportation equipment responds negatively all through the period.

Negative effects	Insignificant effects	Positive effects
Food, beverage, and tobacco	Computer electronic products	Primary metal
Furniture	Electronic and gas utilities	Mining
Motor vehicles	Miscellaneous	Defense and space equipment
Plastic and rubber products	Nonmetallic mineral products	
Printing related	Paper	
Wood products	Apparel and leather goods	
From negative to insignificant effects		**From negative to positive effects**
Chemicals		Petroleum and coal products
Electronical equipment appliances		Aerospace
		Fabricated metal products
		Machinery

Table 1. The effect of an oil-specific shock on industrial production (after two quarters). The first three columns in the top row show industries with stable responses: either negative, insignificant, or positive throughout the sample. The two columns in the lower row show industries where the responses have changed, either from negative to insignificant or to positive.

Thus, we suggest that a gradual shift has taken place for several industries. In line with Baumeister and Kilian (2016) we find that most of the energy intensive industries respond negatively as before, supporting the view that the cost channel (i.e., hypothesis (iii)) is of less importance. However, the industries that have direct and indirect relationship to the oil sector have benefited from higher activity and are now responding by increasing investment and activity when oil prices increase, (i.e., hypotheses (i) and (ii)).

3.5 Geographical dispersion at the state level

So far we have focused on aggregate and disaggregate responses for the U.S. as a whole. We now turn to investigate responses at the state level, focusing on employment, as an important part of an oil boom is the resource movement of labor (see e.g., Corden, 1984; Bjørnland et al., 2019). Doing so, we can also examine if proximity to the oil producing regions matter. To this end, we re-estimate the model including disaggregate employment series for all U.S. states. In addition, we also replace the aggregate U.S. oil production series with separate series for conventional and unconventional oil. The sample period for the extended model starts in 1990 due to the availability of disaggregated employment series.[18] As described in Section 2.2, we subtract employment in the mining sector from total non-farm employment for the oil producing states, and include these series separately. We consider 16 states to be significant oil producing states: nine unconventional oil producing states: California, Colorado, Kansas, Louisiana, New Mexico, North Dakota, Ohio,

[18] In Section F in the Online Appendix we show that the main results presented above are robust for the shorter sample, both in terms of the sign of the responses, and timing of the changes.

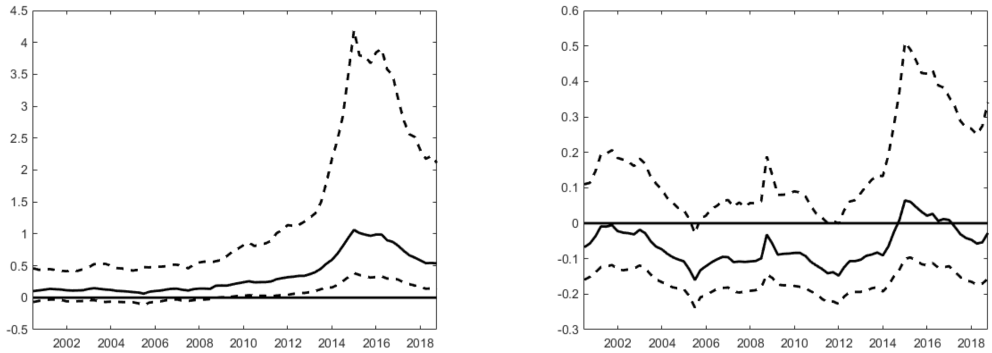

(a) Tight Oil (after 2 quarters) **(b)** Conventional Oil (after 2 quarters)

Figure 7. The effect of an oil-specific shock: Impulse responses for tight oil and conventional oil production. The initial shock is normalized to increase oil prices by 10%. All responses are reported in levels (of the standardized data), showing responses across the sample two quarters after the shock. The dashed lines represent 68% posterior probability bands. The solid line is the median estimate.

Oklahoma, and Texas,[19] and seven conventional oil producing states: Alaska, Mississippi, Montana, Pennsylvania, Utah, West Virginia and Wyoming.[20]

We start by examining responses in tight and conventional oil production to the oil-specific shocks in Figure 7. The results are striking. While conventional oil does not respond significantly to the oil-specific shock even after two quarters, unconventional (tight) oil responds strongly, and in particular from 2010. These results suggests that shale (unconventional) oil producers are more price responsive than conventional oil producers, in line with what other studies have found (cf. Bjørnland et al., 2021; Bornstein et al., 2022; Aastveit et al., 2022). Hence, we conclude that there is strong evidence that the time-varying pattern we observe is closely related to the shale oil boom.

Turning to employment, Figure 8 first shows results for oil employment and non-oil employment in three different oil producing states: California, North Dakota and Texas, before Figure 9 summarizes results for non-oil employment in all states by focusing on geographical dispersion.[21] For all three oil producing states, we find, as expected, that there is a significant positive effect from an oil-specific shock on mining employment (cf. the upper row in Figure 8). Not only does mining employment increase with higher oil prices, the effect has also increased over time, and North Dakota shows the largest response. This suggests that the resource boom relates to the the time varying parameters. Turning to the lower row in Figure 8, that graphs the effect of an oil-specific shock on non-oil employment in the same three oil producing states. We note that for North Dakota and

[19]Note that some states, like Texas, are important producers of both conventional and unconventional oil.

[20]The results are robust to inclusion of disaggregate series for all oil producing states, that is, also those where the average production is below 10 million barrels during the period 2010-2018.

[21]Detailed impulse responses for all states can be obtained on request.

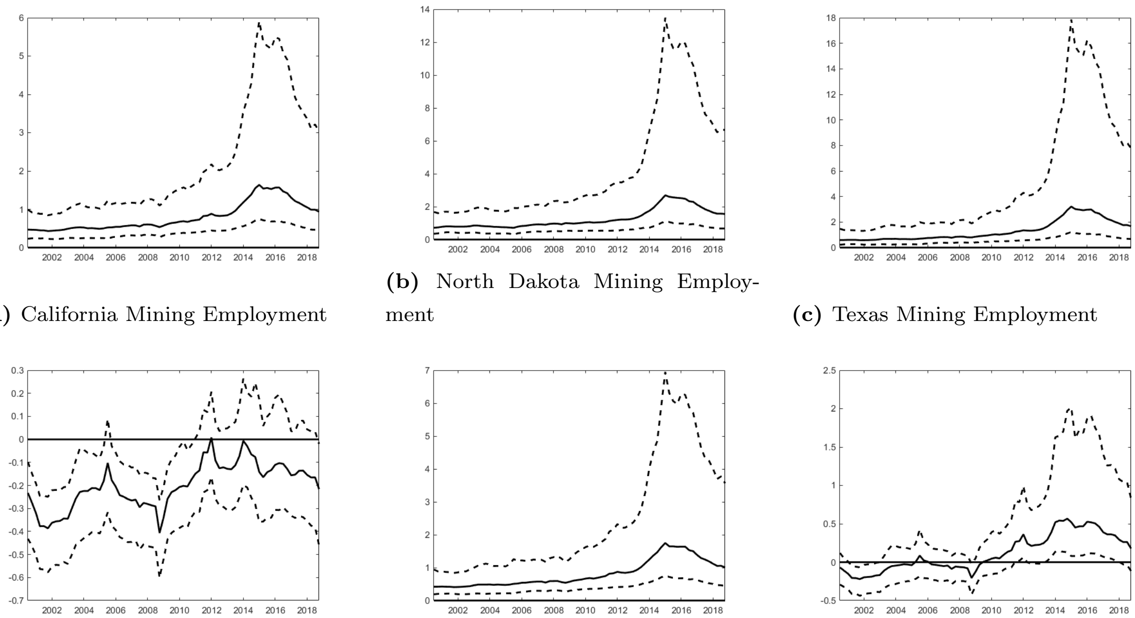

(a) California Mining Employment **(b)** North Dakota Mining Employment **(c)** Texas Mining Employment

(d) California Non-oil Employment **(e)** North Dakota Non-oil Eployment **(f)** Texas Non-oil Employment

Figure 8. The effect of an oil-specific shock: Impulse responses for mining and non-mining employment in three oil producing states. The initial shock is normalized to increase oil prices by 10%. All responses are reported in levels (of the standardized data), showing responses across the sample two quarters after the shock. The dashed lines represent 68% posterior probability bands. The solid line is the median estimate.

Texas, non-oil employment increases significantly following an oil-specific shock. However, whereas for North Dakota, the effect is significant positive over the whole sample (although even more so after 2010), for Texas non-mining employment is only affected positively from 2010/2011, whereas previously the effect was insignificant or negative. For California, however, where oil production is small relative to the aggregate activity, we do not find any positive responses or time-varying pattern for non-oil employment.

Figure 9 summarizes responses in non-oil employment to an oil-specific shock in all U.S. states, focusing on geographical dispersion. To this end, we let dark grey indicates states where employment responds significantly positive either all through the sample or in the final part. Light grey indicates states where employment responses change significantly from being negative to positive/insignificant in the final part of the sample, while white indicates states where there are no time varying changes, and responses are either negative or insignificant throughout the whole sample.

We first note that positive effects on non-oil employment are found in the majority of oil producing states; Alaska, Kansas, Louisiana, New Mexico, North Dakota, Oklahoma, Pensylvania, Texas and Wyoming. These results are consistent with the literature using cross-section data that find positive spillovers from oil activity on local employment, c.f.

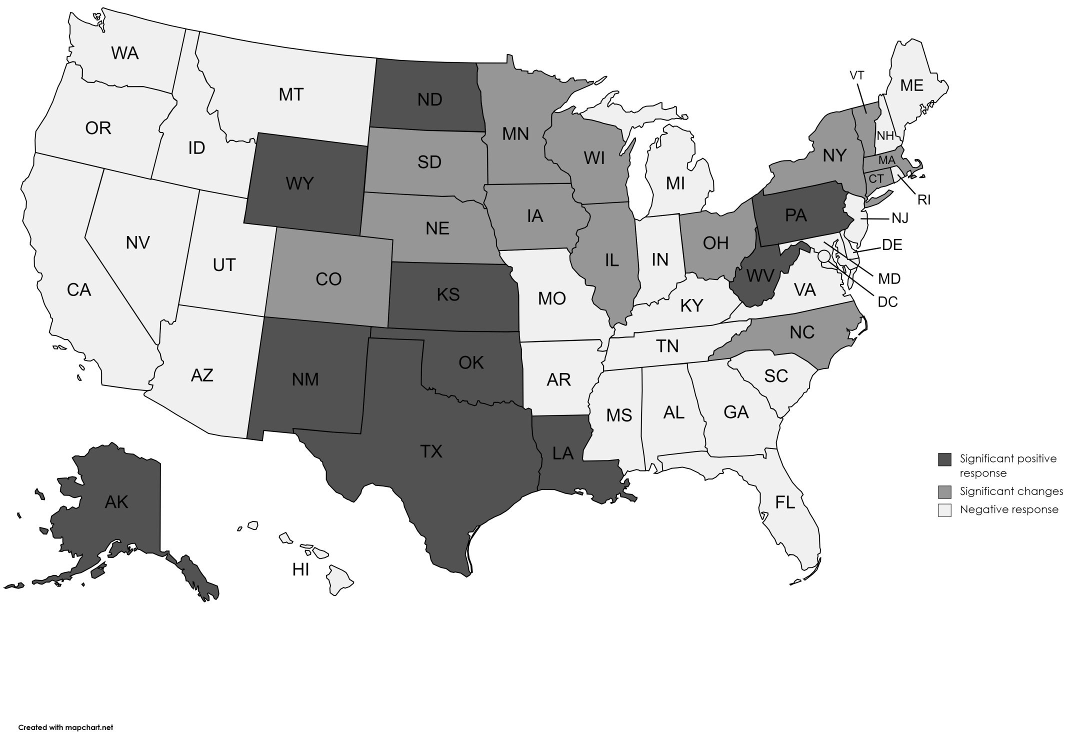

Figure 9. Geographical dispersion of oil-specific shocks to non-oil employment across U.S states. Dark grey indicates oil-rich states where employment responds significantly positive; light grey indicates states where employment responses change significantly from being negative to positive/insignificant in the final part of the sample,; white indicates states where there are no time varying changes, and responses are either negative or insignificant. See the main text for additional explanations.

Feyrer et al. (2017) and Allcott and Keniston (2018). Here we show that the effect is positive also at the state level during the shale oil boom. Second, structural changes on non-oil employment are also observed in the oil producing state, Colorado, and in states located close to the oil producing states, South Dakota, Minnesota, Nebraska, Illinois, Iowa and Wisconsin, in addition to some states located in the Northeast: Ohio, New York, Connecticut, Vermont and Massachusetts, as well as North Carolina more in in the South. These are states that have industries that include some of machinery, fabricated metal products, aerospace equipment, computer and electronic products, and primary metals as one the three most important industries (cf. NAM, 2015). As discussed in Section 3.4 these industries are typically closely connected to the energy boom.

In the remaining states, including the four oil producing states California, Mississippi, Montana and Utah, the effects are either negative or insignificant. For the oil states, this is most likely due to oil being a relative small share of value added in the state. For the other states, the negative response could be due to the fact that some states are energy rich in production, such as states with large car manufacturers, i.e, Michigan, Indiana, Kentucky, and Tennessee, or are not benefiting (have no spillovers) from the oil boom.

We conclude that there are heterogeneous effects of an oil-specific shock across employment in U.S. states, as employment in many states is now behaving more in sync with the oil boom. This includes most oil producing states, many states with a large manufacturing sector, and some states that are closely located to the oil producing states and that benefit from various spillovers.

4 Alternative hypothesis and robustness

We have documented significant changes in the way an oil-specific shock is affecting the U.S. economy, and attribute this to the giant shale oil boom. Still, there could be other structural changes in the economy, unrelated to oil, that may explain the time-varying changes. We address this below by analyzing some alternative hypothesis in Section 4.1, before turning to the robustness tests in Section 4.2.

4.1 Alternative hypothesis

We investigate five alternative hypothesis. Below we comment briefly, while details can be found in Section G in the Online Appendix. First, one could expect that the other oil market shocks that increase oil prices, like the flow supply shock, could also explain the observed time varying changes. Section G.1 confirms that the flow supply shock has slight positive effects on employment, industrial production and manufacturing in the last part of the sample, but the time varying changes are not statistical significant, motivating our focus on the oil-specific shock. Second, we examine to what extent the time varying changes could relate to growth in the global economy rather than the shale boom, as global demand has pushed up oil prices in the 2000s (cf. Aastveit et al., 2015). To evaluate this, we examine the model implied effects of the identified (global) flow demand shock. The impulse responses displayed in Section G.2 confirm that the flow demand shock increases oil prices and activity in the U.S., as expected, but except for oil prices, there is no evidence of significant time-varying changes. For oil prices, however, the positive response escalates significantly throughout the 2000s, in line with the surge in demand for energy from Asia (cf. Aastveit et al., 2015). Third, could the observed time varying changes following an oil-specific shock be a common feature in many countries, not just related to the U.S.? In particular, less energy dependence in production or consumption, or changes in the transmission mechanism, could have dampened potential negative effects. In Section G.3 we show that this is not the case. Fourth, we analyse if the period of low interest rates has changed the relationship between oil prices and macroeconomic variables. As pointed out by Datta et al. (2021), the central bank can not respond to inflationary pressures induced by oil price changes at the zero lower bound. To account for this, we add a shadow rate defined by Wu and Xia (2016) to our dataset, in addition to inflation expectations and the real interest rate, and re-estimate the model. Doing so, we show in Section G.4 that the interest rates and inflation expectations respond as expected, but there are no evidence of time-varying changes. For the other variables, the time-varying results related to the oil-specific shock prevail. Fifth, we examine if results could be due to the fact that there are other commodity prices, say corn and coal prices, correlated with oil prices, that also show evidence of time varying behavior. Hence, we may have found significant effects on employment in states with high coal and corn production. Adding relevant variables, we show in Section G.5 that responses in coal or corn prices are significant, but stable over time. Hence, we conclude that while there are other hypothesis for structural changes that have been affecting the U.S. economy the last decades, these are not able to explain the changing relationship between the oil price and the U.S. economy documented here.

4.2 Robustness

We now analyze sensitivity to our choice of modelling framework. We comment briefly here, while details can be found in Section H in the Online Appendix. In particular, we analyze robustness in three dimensions. First we analyse our modelling choices with respect to time variation and choice of priors. In particular, we follow Bianchi et al. (2017) and illustrate how different time-varying components influence our main results. We show in section H.1 that time-varying coefficients and time-variation in simultaneous relationships are crucial for determining structural changes in the economy, while stochastic volatility is essential to differentiate between temporary and permanent shifts. This supports our model choices. Regarding choice of prior specification, various sensitivity analyses displayed in Section H.2 underscores the importance of balanced priors for time-varying parameters and demonstrate the robustness of our main results across alternative model specifications. Second, we analyze robustness to the use of recursive restrictions by re-estimating our model using monthly data, thereby restricting the delayed response to only be a month instead of a quarter. We confirm that various monthly series such as industrial production and business equipment show evidence of time variation following an oil-specific shock (see Section H.3). Third we show in Section H.4 that results are robust with respect to replacing the global activity (GECON) variable with an estimate of industrial production (see Baumeister and Hamilton, 2019; Hamilton, 2019). We further show in Section H.5 that results are robust to replacing U.S. refineries' acquisition cost with the WTI oil price.

5 Conclusion

It is widely acknowledged that the remarkable growth of U.S. shale oil has significantly bolstered oil investments since the year 2010. In this paper, we demonstrate that the shale oil boom has not only impacted oil investment, but has also changed the way oil specific shocks are transmitted to non-oil investment, income, employment and many manufacturing intensive industries across the U.S. To capture these effects, we use a time-varying parameter factor-augmented vector autoregressive (VAR) model with state-level, industry and aggregate macroeconomic data. Our framework allows us to study the effects of oil-specific shocks on a large number of U.S. macroeconomic variables and analyze the time variation in these effects. To the best of our knowledge this is the first paper that jointly models the interaction between the oil market and the U.S. economy in a large data environment, allowing also for time-varying changes during the fracking revolution.

In contrast to previous studies, we find substantial changes in the way oil prices are transmitted to the U.S. economy. In particular, we find that both oil and the non-oil nonresidential business investments, as well as production and employment in oil-producing and manufacturing-intensive states to pick up following an adverse oil-specific shock. What's more, there are positive spillovers to real personal income. Through extensive robustness we also show that while there are other structural changes in the economy, unrelated to oil, that have been affecting the U.S. economy the last decades, these are not able to explain the changing relationship between the oil price and the U.S. economy documented here. Furthermore, we show that the evidence of time varying changes is unique to the U.S. Going forward, economic policy needs to take into account that the transmission of an oil-specific shock has changed with the shale oil boom, and that there are heterogeneous effects across the U.S.

References

Aastveit, K. A. (2014). Oil Price Shocks in a Data-Rich Environment. *Energy Economics 45*, 268–279.

Aastveit, K. A., H. C. Bjørnland, and T. S. Gundersen (2022). The price responsiveness of shale producers: Evidence from micro data. Cama working papers 2022-70, CAMA.

Aastveit, K. A., H. C. Bjørnland, and L. A. Thorsrud (2015). What Drives Oil Prices? Emerging Versus Developed Economies. *Journal of Applied Econometrics 30*(7), 1013–1028.

Allcott, H. and D. Keniston (2018). Dutch Disease or Agglomeration? The Local Economic Effects of Natural Resource Booms in Modern America. *Review of Economic Studies 85*(2), 596–731.

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

Arezki, R., V. A. Ramey, and L. Sheng (2017). News Shocks in Open Economies: Evidence from Giant Oil Discoveries. *The Quarterly Journal of Economics 132*(1), 103–155.

Bai, J. and S. Ng (2002a). Determining the number of factors in approximate factor models. *Econometrica 70*(1), 191–221.

Bai, J. and S. Ng (2002b). Determining the Number of Factors in Approximate Factor Models. *Econometrica 70*(1), 191–221.

Baumeister, C. and J. D. Hamilton (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 L. Kilian (2016). Lower Oil Prices and the U.S. Economy: Is This Time Different? *Brookings Papers on Economic Activity*, 287–357.

Baumeister, C., D. Korobilis, and T. K. Lee (2020). Energy markets and global economic conditions. *Review of Economics and Statistics*, 1–45.

Baumeister, C., P. Liu, and H. Mumtaz (2013). Changes in the effects of monetary policy on disaggregate price dynamics. *Journal of Economic Dynamics and Control 37*(3), 543–560.

Baumeister, C. and G. Peersman (2013a). The Role of Time-Varying Price Elasticities in Accounting for Volatility Changes in the Crude Oil Market. *Journal of Applied Econometrics 28*(7), 1087–1109.

Baumeister, C. and G. Peersman (2013b). Time-Varying Effects of Oil Supply Shocks on the US Economy. *American Economic Journal: Macroeconomics 5*(4), 1–28.

Belviso, F. and F. Milani (2006). Structural Factor-Augmented VARs (SFAVARs) and the Effects of Monetary Policy. *The B.E. Journal of Macroeconomics 6*(3).

Bernanke, B. S., J. Boivin, and P. Eliasz (2005). Measuring the Effects of Monetary Policy: A Factor-Augmented Vector Autoregressive (FAVAR) Approach. *The Quarterly Journal of Economics 120*(1), 387–422.

Bianchi, D., M. Guidolin, and F. Ravazzolo (2017). Macroeconomic factors strike back: A bayesian change-point model of time-varying risk exposures and premia in the us cross-section. *Journal of Business & Economic Statistics 35*(1), 110–129.

Bjørnland, H. C. (2000). The dynamic effects of aggregate demand, supply and oil price shocks—a comparative study. *The Manchester School 68*(5), 578–607.

Bjørnland, H. C. (2022). The effect of rising energy prices amid geopolitical developments and supply disruptions. *ECB Sintra Forum proceedings*.

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

Bjørnland, H. C., F. M. Nordvik, and M. Rohrer (2021). Supply Flexibility in the Shale Patch: Evidence from North Dakota. *Journal of Applied Econometrics 36*, 273–292.

Bjørnland, H. C. and L. A. Thorsrud (2016). Boom or Gloom? Examining the Dutch Disease in Two-speed Economies. *Economic Journal 126*(598), 2219–2256.

Bjørnland, H. C., L. A. Thorsrud, and R. Torvik (2019). Dutch Disease Dynamics Reconsidered. *European Economic Review 119*, 411–433.

Boivin, J. and M. P. Giannoni (2007). Global Forces and Monetary Policy Effectiveness. *International Dimensions of Monetary Policy*, 429–478.

Boivin, J., M. P. Giannoni, and I. Mihov (2009, 03). Sticky Prices and Monetary Policy: Evidence from Disaggregated US Data. *American Economic Review 99*(1), 350–384.

Bornstein, G., P. Krusell, and S. Rebelo (2022). A World Equilibrium Model of the Oil Market (forthcoming). *Review of Economic Studies*.

Bresnahan, T. F. and V. A. Ramey (1993). Segment Shifts and Capacity Utilization in the US Automobile Industry. *The American Economic Review 83*(2), 213–218.

Cashin, P., K. Mohaddes, M. Raissi, and M. Raissi (2014). The Differential Effects of Oil Demand and Supply Shocks on the Global Economy. *Energy Economics 44*, 113–134.

Clark, T. E. and S. J. Terry (2010). Time Variation in the Inflation Passthrough of Energy Prices. *Journal of Money, Credit and Banking 42*(7), 1419–1433.

Cogley, T. and T. J. Sargent (2001). Evolving Post-World War II U.S. Inflation Dynamics. *NBER Macroeconomics Annual 16*, 331–373.

Corden, W. M. (1984). Booming Sector and Dutch Disease Economics: Survey and Consolidation. *Oxford Economic Papers 36*(3), 359–380.

Datta, D. D., B. K. Johannsen, H. Kwon, and R. J. Vigfusson (2021). Oil, equities, and the zero lower bound. *American Economic Journal: Macroeconomics 13*(2), 214–53.

Davis, S. J. and J. Haltiwanger (2001). Sectoral Job Creation and Destruction Responses to Oil Price Changes. *Journal of Monetary Economics 48*(3), 465–512.

De Michelis, A., T. Ferreira, and M. Iacoviello (2020). Oil Prices and Consumption across Countries and US States. *International Journal of Central Banking 16*(2), 3–43.

Del Negro, M. and G. E. Primiceri (2015). Time Varying Structural Vector Autoregressions and Monetary Policy: A Corrigendum. *The Review of Economic Studies 82*(4), 1342–1345.

Ellis, C., H. Mumtaz, and P. Zabczyk (2014). What lies beneath? a time-varying FAVAR model for the UK transmission mechanism. *The Economic Journal 124*(576), 668–699.

Fetzer, T. R. (2014). Fracking Growth. *CEP Discussion Paper*.

Feyrer, J., E. T. Mansur, and B. Sacerdote (2017). Geographic Dispersion of Economic Shocks: Evidence from the Fracking Revolution. *American Economic Review 107*(4), 1313–1334.

Foroni, C., P. Guerin, and M. Marcellino (2017). Explaining the Time-Varying Effects of Oil Market Shocks on U.S. Stock Returns. *Economics Letters 155*, 84–88.

Geweke, J. (1992). Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments. *Bayesian Statistics 4*, 169–193.

Gilje, E., R. Ready, and N. Roussanov (2016). Fracking, Drilling, and Asset Pricing: Estimating the Economic Benefits of the Shale Revolution. NBER working paper no. 22914, National Bureau of Economic Research.

Guerrero-Escobar, S., G. Hernandez-del Valle, and M. Hernandez-Vega (2019). Do heterogeneous countries respond differently to oil price shocks? *Journal of Commodity Markets 16*, 100084.

Hamilton, J. D. (1983). Oil and the Macroeconomy since World War II. *Journal of Political Economy 91*(2), 228–248.

Hamilton, J. D. (2008). Oil and the Macroeconomy. *The New Palgrave Dictionary of Economics, Second Edition, 2008*.

Hamilton, J. D. (2009). Causes and Consequences of the Oil Shock of 2007-08. *Brookings Papers on Economic Activity 40*(1), 215–283.

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

Hamilton, J. D. and A. M. Herrera (2004). Comment: Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy. *Journal of Money, Credit and Banking*, 265–286.

Herrera, A. M. (2018). Oil Price Shocks, Inventories, and Macroeconomic Dynamics. *Macroeconomic Dynamics 22*(3), 620–639.

Herrera, A. M. and M. B. Karaki (2015). The effects of oil price shocks on job reallocation. *Journal of Economic Dynamics and control 61*, 95–113.

Herrera, A. M., M. B. Karaki, and S. K. Rangaraju (2017). Where do jobs go when oil prices drop? *Energy Economics 64*, 469–482.

Herrera, A. M., L. G. Lagalo, and T. Wada (2015). Asymmetries in the response of economic activity to oil price increases and decreases? *Journal of International Money and Finance 50*, 108–133.

Jiménez-Rodríguez, R. and M. Sánchez (2005). Oil price shocks and real GDP growth: empirical evidence for some OECD countries. *Applied Economics 37*(2), 201–228.

Kang, W., R. A. Ratti, and K. H. Yoon (2015, 12). Time-Varying Effect of Oil Market Shocks on the Stock Market. *Journal of Banking & Finance 61*, S150–S163. Supplement 2.

Känzig, D. R. (2021). The macroeconomic effects of oil supply news: Evidence from opec announcements. *American Economic Review 111*(4), 1092–1125.

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), 1053–1069.

Kilian, L. and D. P. Murphy (2012). Why agnostic sign restrictions are not enough: understanding the dynamics of oil market VAR models. *Journal of the European Economic Association 10*(5), 1166–1188.

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

Kilian, L. and C. Vega (2011). Do Energy Prices respond to US Macroeconomic News? A Test of the Hypothesis of Predetermined Energy Prices. *Review of Economics and Statistics 93*(2), 660–671.

Kilian, L. and R. J. Vigfusson (2011). Are the Responses of the US Economy Asymmetric in Energy Price Increases and Decreases? *Quantitative Economics 2*(3), 419–453.

Kim, S., N. Shephard, and S. Chib (1998). Stochastic volatility: likelihood inference and comparison with arch models. *The review of economic studies 65*(3), 361–393.

Koop, G. (2003). *Bayesian Econometrics*. Wiley, Chichester.

Koop, G. and D. Korobilis (2014). A new index of financial conditions. *European Economic Review 71*, 101–116.

Korobilis, D. (2013). Assessing the Transmission of Monetary Policy Using Time-Varying Parameter Dynamic Factor Models. *Oxford Bulletin of Economics and Statistics 75*(2), 157–179.

Lee, K. and S. Ni (2002). On the Dynamic Effects of oil Price Shocks: A Study Using Industry Level Data. *Journal of Monetary Economics 49*(4), 823–852.

Lippi, F. and A. Nobili (2012). Oil and the Macroeconomy: a Quantative Structural Analysis. *Journal of the European Economic Association 10*(5), 1059–1083.

NAM (2015). National Association of Manufacturers - State Manufacturing Data Reports. http://www.nam.org/Data-and-Reports/State-Manufacturing-Data/.

Peersman, G. and I. V. Robays (2012). Cross-Country Differences in the Effects of Oil Shocks. *Energy Economics 34*(5), 1532 –1547.

Primiceri, G. E. (2005). Time -Varying Structural Vector Autoregressions and Monetary Policy. *The Review of Economic Studies 72*(3), 821–852.

Stock, J. H. and M. W. Watson (1996). Evidence on Structural Instability in Macroeconomic Time Series Relations. *Journal of Business & Economic Statistics 14*(1), 11–30.

Stock, J. H. and M. W. Watson (2005). Implications of Dynamic Factor Models for VAR Analysis. *NBER Working Paper* (11467).

Stock, J. H. and M. W. Watson (2016). Dynamic Factor Models, Factor-Augmented Vector Autoregressions, and Structural Vector Autoregressions in Macroeconomics. In *Handbook of macroeconomics*, Volume 2, pp. 415–525. Elsevier.

Wang, Z. and A. Krupnick (2013). A Retrospective Review of Shale Gas Development in the United States: What Led to the Boom? *Resources for the Future Discussion Paper 13-12*.

Weber, J. G. (2012). The Effects of a Natural Gas Boom on Employment and Income in Colorado, Texas, and Wyoming. *Energy Economics 34*(5), 1580–1588.

Wu, J. C. and F. D. Xia (2016). Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound. *Journal of Money, Credit, and Banking 48*(2-3), 253–291.

—Online Appendix—

Not for publication

A Data Description

To investigate the effects of oil specific shocks on the U.S. economy, we employ a comprehensive panel consisting of both domestic and international series. A detailed breakdown of these series can be found in Table 2.

We do several transformations of the data. Non-U.S. crude oil production is determined by subtracting the U.S. crude oil production (47) from the global crude oil production (3). We construct non-oil investment series by subtracting mining exploration nonresidential fixed investment (13) from total nonresidential fixed investment (12). Further, we construct non mining employment series by subtracting mining nonfarm employment (15) from total nonfarm employment (16). To distinguish between conventional and tight U.S. oil production, we estimate conventional oil production by subtracting series (48), which represents tight oil production, from the total U.S. crude oil production denoted by series (47). The tight oil production series is obtained from the EIA data for U.S. tight oil production estimates by play, available from 2000. [22] We calculate the total U.S. tight oil production taking the sum of production in Eagle Ford (TX), Spraberry (TX Permian), Bakken (ND & MT), Wolfcamp (TX & NM Permian), Bonespring (TX & NM Permian), Niobrara-Codell (CO & WY), Mississippian (OK), Austin Chalk (LA & TX), Woodford (OK) and the "Rest of U.S. tight oil". Since the play production data is available only from 2000 and the production of most of the playes in the early 2000s are around zero or constant, we carry these levels back in time so that the transformed data equals 0 (meaning no change). Finally, we also remove mining and logging employment from the total employment in 16 U.S. states. These are Alaska, California, Colorado, Kansas, Louisiana, Montana, Mississippi, New Mexico, North Dakota, Ohio, Oklahoma, Pennsylvania, Texas, Utah, West Virginia, and Wyoming. The following series are deflated by the U.S. CPI: the oil price (2), personal income (5), S&P equity index (8), total residential fixed investments (9), total nonresidential fixed investment (12), mining exploration nonresidential fixed investment (13), as well as other dissagregate investment series (116)-(123). We also construct the real interest rate series, by deflating federal funds effective rate (10). We extend shadow federal funds rate (126) with federal funds effective rate back in time, as the former series is available only from 1990.

For properly estimation of TVP FAVAR model is it crucial to ensure stationarity,

[22] That is estimated monthly production derived from state administrative data.

comparability, and interpretability of the data. As described in the main paper Section 2.2, to ensure stationarity we transform the data into growth rates. To do so we take the log difference that allows us for a multiplicative interpretation of the changes. The data is transformed to quarter-on-quarter growth. Exceptions include the Global Economic Conditions indicator (GECON), which is stationary by construction, and interest rates series which are in differences. Since all of the data are seasonally adjusted, QoQ changes provides a clearer picture of trends and is useful for understanding immediate responses to changes in the economy. The local mean is removed from the transformed data, and the data used to extract factors are standardized. Removing the local mean and standardizing the data, ensure that all series are on a comparable scale, that is essential in FAVAR. Impulse response functions are also more interpretable, as it ensures comparability across variables.

Global Variables

Nr	Series	Start year	Source
1	Indicator of Global Economic Conditions	1974	Baumeister et al. (2020)
2	United States, National, Commodities & Energy Prices, Refiner Acquisition, Crude Oil, Domestic, Average Price, USD	1974	Macrobond
3	World, EIA, Oil & Gas, Total Oil Supply, Crude Oil & Petroleum Products, Petroleum & Other Liquids, Production, Crude Oil Including Lease Condensate, Barrels per Day	1974	Macrobond

Macro aggregates

Nr	Series	Start year	Source
4	United States, Consumer Price Index, All Urban Consumers, U.S. City Average, All Items, SA, Index	1974	Macrobond
5	United States, Income Approach, Personal Income, Total, SA, AR, USD	1974	Macrobond
6	United States, Policy Rates, Effective Rates, Federal Funds Effective Rate	1974	Macrobond
7	United States, Industrial Production, Total, Constant Prices, SA, Index	1974	Macrobond
8	United States, Equity Indices, S&P, 500, Index, Price Return, Close, USD	1974	Macrobond
9	United States, Expenditure Approach, Gross Private Domestic Investment, Fixed Investment, Residential, Total, SA, AR, USD	1974	Macrobond
10	United States, FX Indices, Federal Reserve, Nominal Advanced Foreign Economies Index	1974	Macrobond
11	United States, Expenditure Approach, Personal Consumption Expenditures, Total, Constant Prices, SA, Index	1974	Macrobond
12	United States, Expenditure Approach, Gross Private Domestic Investment, Fixed Investment, Nonresidential, Total, SA, AR, USD	1974	Macrobond
13	United States, Saving & Investment, Private Fixed Investment, Non-Residential, Structures, Mining Exploration, Shafts & Wells, SA, AR, USD	1974	Macrobond
14	United States, Industrial Production, Total, Constant Prices, SA, Index	1974	Macrobond
15	United States, Employment, Payroll, Mining & Logging, Nonfarm, Mining, Payroll, SA	1974	Macrobond
16	United States, Employment, Payroll, Nonfarm, Payroll, Total, SA	1974	Macrobond
17	United States, Foreign Trade, Crude Oil & Petroleum Products, Import, Crude Oil	1974	Macrobond

18	United States, Foreign Trade, Crude Oil & Petroleum Products, Export, Total Petroleum	1974	Macrobond

Disaggregate Industrial Production

Nr	Series	Start year	Source
19	United States, Industrial Production, Industry Group, Manufacturing, Total (SIC), Constant Prices, SA, Index	1974	Macrobond
20	United States, Industrial Production, Industry Group, Manufacturing, Wood Products, Constant Prices, SA, Index	1974	Macrobond
21	United States, Industrial Production, Industry Group, Manufacturing, Non-metallic Mineral Products, Constant Prices, SA, Index	1974	Macrobond
22	United States, Industrial Production, Industry Group, Manufacturing, Primary Metal, Constant Prices, SA, Index	1974	Macrobond
23	United States, Industrial Production, Industry Group, Manufacturing, Fabricated Metal Products, Constant Prices, SA, Index	1974	Macrobond
24	United States, Industrial Production, Industry Group, Manufacturing, Machinery, Constant Prices, SA, Index	1974	Macrobond
25	United States, Industrial Production, Industry Group, Manufacturing, Computer & Electronic Products, Constant Prices, SA, Index	1974	Macrobond
26	United States, Industrial Production, Industry Group, Manufacturing, Electrical Equipment, Appliances & Components, Constant Prices, SA, Index	1974	Macrobond
27	United States, Industrial Production, Industry Group, Manufacturing, Motor Vehicles & Parts, Constant Prices, SA, Index	1974	Macrobond
28	United States, Industrial Production, Industry Group, Manufacturing, Aerospace & Miscellaneous Transportation Equipment, Constant Prices, SA, Index	1974	Macrobond
29	United States, Industrial Production, Industry Group, Manufacturing, Furniture & Related Products, Constant Prices, SA, Index	1974	Macrobond
30	United States, Industrial Production, Industry Group, Manufacturing, Miscellaneous, Constant Prices, SA, Index	1974	Macrobond
31	United States, Industrial Production, Industry Group, Manufacturing, Food, Beverage & Tobacco, Constant Prices, SA, Index	1974	Macrobond
32	United States, Industrial Production, Industry Group, Manufacturing, Textiles & Products, Constant Prices, SA, Index	1974	Macrobond
33	United States, Industrial Production, Industry Group, Manufacturing, Apparel & Leather Goods, Constant Prices, SA, Index	1974	Macrobond
34	United States, Industrial Production, Industry Group, Manufacturing, Paper, Constant Prices, SA, Index	1974	Macrobond
35	United States, Industrial Production, Industry Group, Manufacturing, Printing & Related Support Activities, Constant Prices, SA, Index	1974	Macrobond
36	United States, Industrial Production, Industry Group, Manufacturing, Petroleum & Coal Products, Constant Prices, SA, Index	1974	Macrobond
37	United States, Industrial Production, Industry Group, Manufacturing, Chemicals, Constant Prices, SA, Index	1974	Macrobond
38	United States, Industrial Production, Industry Group, Manufacturing, Plastics & Rubber Products, Constant Prices, SA, Index	1974	Macrobond
39	United States, Industrial Production, Industry Group, Mining, Constant Prices, SA, Index	1974	Macrobond
40	United States, Industrial Production, Industry Group, Electric & Gas Utilities, Constant Prices, SA, Index	1974	Macrobond
41	United States, Industrial Production, Market Group Summary, Products & Non-Industrial Supplies, Final, Consumer Goods, Total, Constant Prices, SA, Index	1974	Macrobond

42	United States, Industrial Production, Market Group Summary, Products & Non-Industrial Supplies, Final, Equipment, Business, Total, Constant Prices, SA, Index	1974	Macrobond
43	United States, Industrial Production, Market Group Summary, Products & Non-Industrial Supplies, Final, Equipment, Defense & Space Equipment, Constant Prices, SA, Index	1974	Macrobond
44	United States, Industrial Production, Market Group Summary, Products & Non-Industrial Supplies, Nonindustrial Supplies, Construction, Constant Prices, SA, Index	1974	Macrobond
45	United States, Industrial Production, Market Group Summary, Materials, Total, Constant Prices, SA, Index	1974	Macrobond
46	United States, Industrial Production, Market Group Summary, Products & Non-Industrial Supplies, Nonindustrial Supplies, Other Business, Total, Constant Prices, SA, Index	1974	Macrobond

U.S. Oil Production Variables

Nr	Series	Start year	Source
47	United States, EIA, Oil & Gas, Total Oil Supply, Crude Oil & Petroleum Products, Petroleum & Other Liquids, Production, Crude Oil Including Lease Condensate, Barrels per Day	1974	Macrobond
48	Tight oil production estimates by play, Million Barrels per Day	2000	EIA

Nonfarm Employment - States

Nr	Series	Start year	Source
49	United States, Employment, By State, Total Nonfarm, Alabama, SA	1990	Macrobond
50	United States, Employment, By State, Total Nonfarm, Alaska, SA	1990	Macrobond
51	United States, Employment, By State, Total Nonfarm, Arizona, SA	1990	Macrobond
52	United States, Employment, By State, Total Nonfarm, Arkansas, SA	1990	Macrobond
53	United States, Employment, By State, Total Nonfarm, California, SA	1990	Macrobond
54	United States, Employment, By State, Total Nonfarm, Colorado, SA	1990	Macrobond
55	United States, Employment, By State, Total Nonfarm, Connecticut, SA	1990	Macrobond
56	United States, Employment, By State, Total Nonfarm, Delaware, SA	1990	Macrobond
57	United States, Employment, By State, Total Nonfarm, District of Columbia, SA	1990	Macrobond
58	United States, Employment, By State, Total Nonfarm, Florida, SA	1990	Macrobond
59	United States, Employment, By State, Total Nonfarm, Georgia, SA	1990	Macrobond
60	United States, Employment, By State, Total Nonfarm, Hawaii, SA	1990	Macrobond
61	United States, Employment, By State, Total Nonfarm, Idaho, SA	1990	Macrobond
62	United States, Employment, By State, Total Nonfarm, Illinois, SA	1990	Macrobond
63	United States, Employment, By State, Total Nonfarm, Indiana, SA	1990	Macrobond
64	United States, Employment, By State, Total Nonfarm, Iowa, SA	1990	Macrobond
65	United States, Employment, By State, Total Nonfarm, Kansas, SA	1990	Macrobond
66	United States, Employment, By State, Total Nonfarm, Kentucky, SA	1990	Macrobond
67	United States, Employment, By State, Total Nonfarm, Louisiana, SA	1990	Macrobond
68	United States, Employment, By State, Total Nonfarm, Maine, SA	1990	Macrobond
69	United States, Employment, By State, Total Nonfarm, Maryland, SA	1990	Macrobond
70	United States, Employment, By State, Total Nonfarm, Massachusetts, SA	1990	Macrobond
71	United States, Employment, By State, Total Nonfarm, Michigan, SA	1990	Macrobond
72	United States, Employment, By State, Total Nonfarm, Minnesota, SA	1990	Macrobond
73	United States, Employment, By State, Total Nonfarm, Mississippi, SA	1990	Macrobond
74	United States, Employment, By State, Total Nonfarm, Missouri, SA	1990	Macrobond

75	United States, Employment, By State, Total Nonfarm, Montana, SA	1990	Macrobond
76	United States, Employment, By State, Total Nonfarm, Nebraska, SA	1990	Macrobond
77	United States, Employment, By State, Total Nonfarm, Nevada, SA	1990	Macrobond
78	United States, Employment, By State, Total Nonfarm, New Hampshire, SA	1990	Macrobond
79	United States, Employment, By State, Total Nonfarm, New Jersey, SA	1990	Macrobond
80	United States, Employment, By State, Total Nonfarm, New Mexico, SA	1990	Macrobond
81	United States, Employment, By State, Total Nonfarm, New York, SA	1990	Macrobond
82	United States, Employment, By State, Total Nonfarm, North Carolina, SA	1990	Macrobond
83	United States, Employment, By State, Total Nonfarm, North Dakota, SA	1990	Macrobond
84	United States, Employment, By State, Total Nonfarm, Ohio, SA	1990	Macrobond
85	United States, Employment, By State, Total Nonfarm, Oklahoma, SA	1990	Macrobond
86	United States, Employment, By State, Total Nonfarm, Oregon, SA	1990	Macrobond
87	United States, Employment, By State, Total Nonfarm, Pennsylvania, SA	1990	Macrobond
88	United States, Employment, By State, Total Nonfarm, Rhode Island, SA	1990	Macrobond
89	United States, Employment, By State, Total Nonfarm, South Carolina, SA	1990	Macrobond
90	United States, Employment, By State, Total Nonfarm, South Dakota, SA	1990	Macrobond
91	United States, Employment, By State, Total Nonfarm, Tennessee, SA	1990	Macrobond
92	United States, Employment, By State, Total Nonfarm, Texas, SA	1990	Macrobond
93	United States, Employment, By State, Total Nonfarm, Utah, SA	1990	Macrobond
94	United States, Employment, By State, Total Nonfarm, Vermont, SA	1990	Macrobond
95	United States, Employment, By State, Total Nonfarm, Virginia, SA	1990	Macrobond
96	United States, Employment, By State, Total Nonfarm, Washington, SA	1990	Macrobond
97	United States, Employment, By State, Total Nonfarm, West Virginia, SA	1990	Macrobond
98	United States, Employment, By State, Total Nonfarm, Wisconsin, SA	1990	Macrobond
99	United States, Employment, By State, Total Nonfarm, Wyoming, SA	1990	Macrobond
100	United States, Alaska, BLS, CES State & Metro Area, All Employees, Mining & Logging, Total, SA	1990	Macrobond
101	United States, California, BLS, CES State & Metro Area, All Employees, Mining & Logging, Total, SA	1990	Macrobond
102	United States, Colorado, BLS, CES State & Metro Area, All Employees, Mining & Logging, Total, SA	1990	Macrobond
103	United States, Kansas, BLS, CES State & Metro Area, All Employees, Mining & Logging, Total, SA	1990	Macrobond
104	United States, Louisiana, BLS, CES State & Metro Area, All Employees, Mining & Logging, Total, SA	1990	Macrobond
105	United States, Mississippi, BLS, CES State & Metro Area, All Employees, Mining & Logging, Total, SA	1990	Macrobond
106	United States, Montana, BLS, CES State & Metro Area, All Employees, Mining & Logging, Total, SA	1990	Macrobond
107	United States, New Mexico, BLS, CES State & Metro Area, All Employees, Mining & Logging, Total, SA	1990	Macrobond
108	United States, North Dakota, BLS, CES State & Metro Area, All Employees, Mining & Logging, Total, SA	1990	Macrobond
109	United States, Ohio, BLS, CES State & Metro Area, All Employees, Mining & Logging, Total, SA	1990	Macrobond
110	United States, Oklahoma, BLS, CES State & Metro Area, All Employees, Mining & Logging, Total, SA	1990	Macrobond
111	United States, Pennsylvania, BLS, CES State & Metro Area, All Employees, Mining & Logging, Total, SA	1990	Macrobond
112	United States, Texas, BLS, CES State & Metro Area, All Employees, Mining & Logging, Total, SA	1990	Macrobond
113	United States, Utah, BLS, CES State & Metro Area, All Employees, Mining & Logging, Total, SA	1990	Macrobond
114	United States, West Virginia, BLS, CES State & Metro Area, All Employees, Mining & Logging, Total, SA	1990	Macrobond
115	United States, Wyoming, BLS, CES State & Metro Area, All Employees, Mining & Logging, Total, SA	1990	Macrobond

	Additional Variables		
Nr	**Series**	**Start year**	**Source**
	Series used in Section F.1		
116	United States, Saving & Investment, Private Fixed Investment, Non-Residential, Structures, Manufacturing, SA, AR, USD	1974	Macrobond
117	United States, Saving & Investment, Private Fixed Investment, Non-Residential, Structures, Commercial & Health Care, SA, AR, USD	1974	Macrobond
118	United States, Saving & Investment, Private Fixed Investment, Non-Residential, Structures, Power & Communication, SA, AR, USD	1974	Macrobond
119	United States, Saving & Investment, Private Fixed Investment, Non-Residential, Structures, Mining Exploration, Shafts & Wells, SA, AR, USD	1974	Macrobond
120	United States, Saving & Investment, Private Fixed Investment, Non-Residential, Structures, Other Structures, SA, AR, USD	1974	Macrobond
121	United States, Saving & Investment, Private Fixed Investment, Non-Residential, Equipment, Industrial Equipment, Total, SA, AR, USD	1974	Macrobond
122	United States, Saving & Investment, Private Fixed Investment, Non-Residential, Equipment, Transportation Equipment, Total, SA, AR, USD	1974	Macrobond
123	United States, Saving & Investment, Private Fixed Investment, Non-Residential, Equipment, Other Equipment, Total, SA, AR, USD	1974	Macrobond
	Series used to test alternative hypothesis (see Section 4.1)		
124	World, World Bank, Maize, Average Price, End of Period, USD	1974	Macrobond
125	United States, Producer Price Index, Commodity, Fuels & Related Products & Power, Coal	1974	Macrobond
126	United States, Policy Rates, Target Rates, Shadow Federal Funds Rate (Wu-Xia)	1990	Macrobond
127	United States, Federal Reserve Bank of Philadelphia, Survey of Professional Forecasters, Median, GDP Inflation Rate, Short-Term, 1 Year, Estimate, AR	1974	Macrobond
	Series used for robustness (see Section 4.2)		
128	World, Crude Oil, WTI, Global Spot, Close, USD	1974	Macrobond
129	World Industrial Production Index	1974	*Baumeister and Hamilton (2019)*
	Series used for VAR (see Section 1)		
130	United Kingdom, Industrial Production, Total, Constant Prices, SA, Index	1974	Macrobond

Table 2. Data description

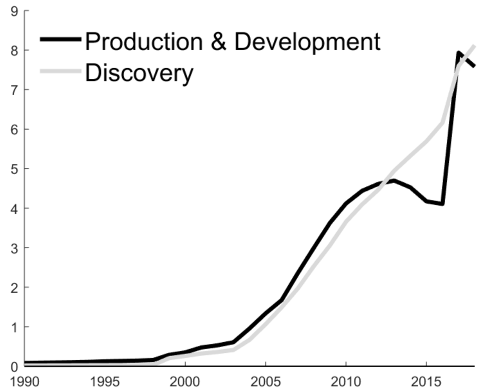

(a) Present Value of Discoveries and Production+Development for Unconventional Crude Oil, (% of U.S. GDP)

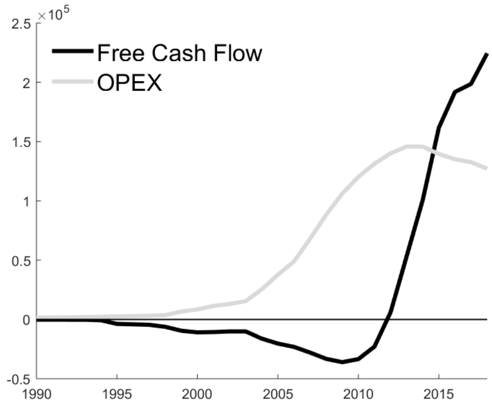

(b) Present Value of OPEX and Free Cash Flow for production of Unconventional Crude Oil, Millions USD

Figure 10. Present value of discovery of oil

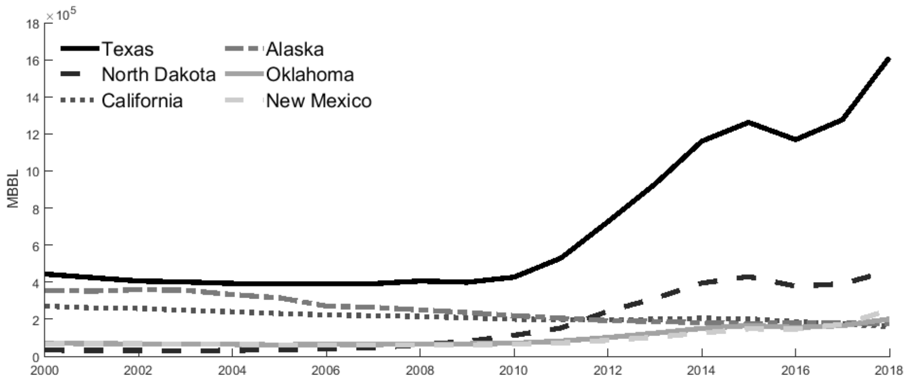

(a) Annual Crude Oil Production in States Exceeding 100,000 MBBL annually on average from 2010 to 2018

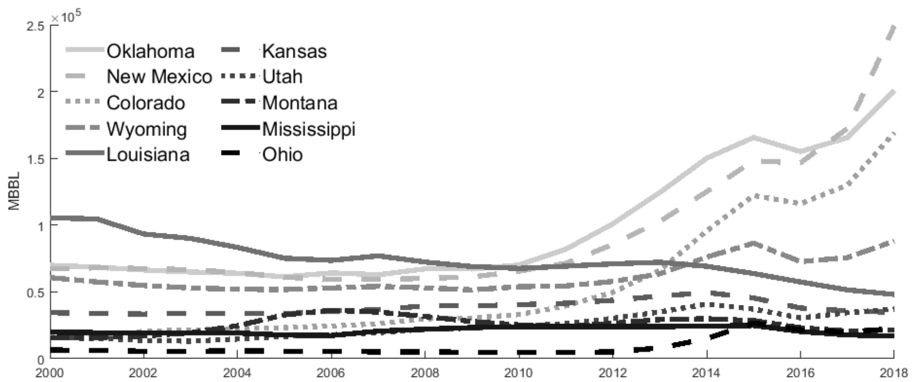

(b) Annual crude oil production in the medium sized oil-producing states with averages exceeding 10,000 MBBL (10 million barrels) from 2010 to 2018, compared to two major producers (from plot a): Oklahoma and New Mexico

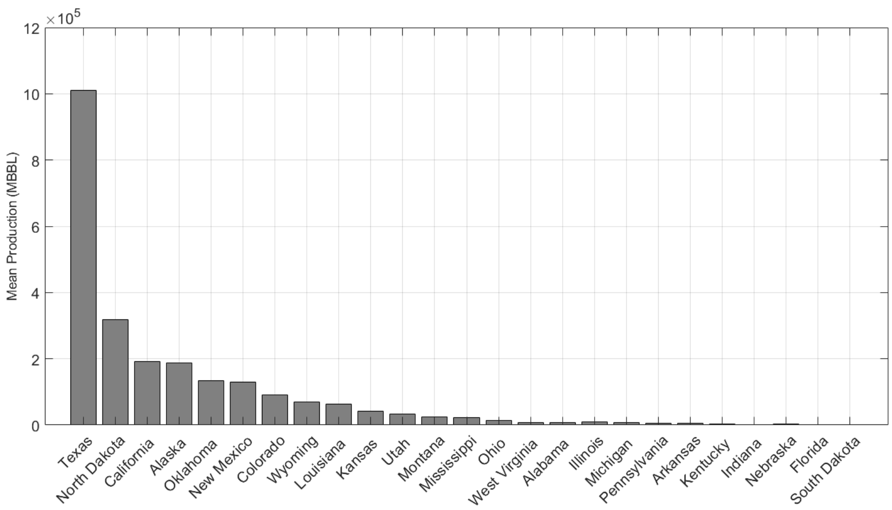

(c) Average crude oil production (MBBL) across states during 2010-2018.

Figure 11. Crude oil production by state. Note 1 MBBL is equal to 1000 barrels

Series	Factor 1	Factor 2
Manufacturing	0.98	0.04
IP index	0.95	0.21
Business supplies	0.91	-0.09
Construction supplies	0.90	-0.08
Materiels	0.89	0.30
Fabricated metal products	0.86	0.09
Electronical equipment appliances	0.83	0.01
Furniture	0.81	-0.16
Consumer goods	0.80	-0.17
Plastic and rubber products	0.79	-0.17
Employment total	0.78	0.10
Nonmetallic mineral products	0.78	-0.09
Business equipment	0.76	0.46
Chemical	0.74	0.11
Textile	0.71	-0.32
Miscellaneous	0.69	-0.16
Residential investment	0.68	-0.16
Paper	0.67	-0.19
Wood products	0.67	-0.31
Machinery	0.66	0.27
Motor vehicles	0.64	-0.13
Primery metal	0.63	-0.06
Printing related	0.62	-0.22
Computer electronic products	0.61	0.04
Apparel and leather goods	0.56	-0.17
Nonresidential non-oil investment	0.56	0.35
Consumption	0.49	-0.09
Personal income	0.41	0.32
FED	0.40	0.02

Table 3. Data series with correlation above 0.4 with first factor. The first column shows correlation between current time series and the first factor, while the second column shows correlation between current time series and the second factor.

Series	Factor 1	Factor 2
Mining	0,25	0,80
Aerospace	0,16	0,69
Employment mining	0,09	0,59
Investment Oil	0,07	0,54
Business equipment	0,76	0,46
Conventional oil	-0,01	0,42
Petroleum export	-0,03	0,42
Tight oil	-0,06	0,40

Table 4. Data series with correlation above 0.4 with second factor. The first column shows correlation between current time series and the first factor, while the second column shows correlation between current time series and the second factor.

B Estimation of a FAVAR model - Two step estimator approach

In Section 2 of the main paper we described the benchmark model. Here we provide a more detailed overview of how the model is estimated. We start by repeating the main system equations. Recall the two main two equations in our model are the factor Equation (10) and the VAR Equation (11):

$$X_t = \Lambda F_t + e_t, \tag{10}$$

$$F_t = c_t + b_{1t} F_{t-1} + ... + b_{pt} F_{t-p} + A_t^{-1} \Sigma_t \varepsilon_t. \tag{11}$$

where the common factors F_t contain both the unobservables latent factors,f_t and the observables factors y_t: $F_t = \begin{pmatrix} y_t \\ f_t \end{pmatrix}$.

The time-varying parameters and covariances of the model follow random walk processes:

$$B_t = B_{t-1} + \eta_t^B \tag{12}$$

$$\alpha_t = \alpha_{t-1} + \eta_t^\alpha \tag{13}$$

$$log\sigma_t = log\sigma_{t-1} + \eta_t^\sigma \tag{14}$$

where B_t is the vector of all R.H.S. coefficients in (11), α_t is the vector of non-zero and no-none elements of the matrix A_t, and σ_t is the vector containing the diagonal elements of Σ_t .

The innovations in the model are assumed to be normally distributed with the following assumptions on the variance covariance matrix:

$$Var\left(\begin{bmatrix} e_t \\ \varepsilon_t \\ \eta_t^B \\ \eta_t^\alpha \\ \eta_t^\sigma \end{bmatrix}\right) = \begin{bmatrix} R & 0 & 0 & 0 & 0 \\ 0 & I_m & 0 & 0 & 0 \\ 0 & 0 & Q & 0 & 0 \\ 0 & 0 & 0 & S & 0 \\ 0 & 0 & 0 & 0 & W \end{bmatrix} \tag{15}$$

We also define V as

$$V = \begin{bmatrix} I_m & 0 & 0 & 0 \\ 0 & Q & 0 & 0 \\ 0 & 0 & S & 0 \\ 0 & 0 & 0 & W \end{bmatrix} \tag{16}$$

The system is then estimated in two steps. In the first step we estimate the unobservable factors f_t, while in the second step we estimate model parameters, conditional on the factors. Below we describe each step in greater detail.

B.1 Step1: Latent factor estimation

We start by extracting k principal components from X_t and obtain estimates of the latent factors, f_t. In doing so, we do not impose a constraint whereby the observable factors y_t are the common component. So if the variables in y_t are common components, they should be captured by the principal components. To remove y_t from the space covered by the principal components, we follow the approach advocated by Boivin and Giannoni (2007), and impose the constraint that observable variables are two of the factors in the first-step estimation. We denote the initial estimate of f_t by f_t^0, and iterate through the following steps:

1. RegressX_t on f_t^0 and the observed factors y_t and obtain $\hat{\lambda}_y^0$
2. Compute $\hat{X}_t^0 = X_t - \hat{\lambda}_y^0 y_t$
3. Estimate f_t^1 as the k principal components of $\hat{X}_t^0$
4. Repeat the procedure multiple times

This procedure guarantees that the estimated latent factors will recover dimensions of the common dynamics not already captured by the observable variables, y_t. Given the factors, F_t, we can estimate parameters in (10) and (11) independently of each other.

B.2 Step 2: The Gibbs sampling approach - Estimation of model parameters

Once the factors have been extracted we simulate the distribution of the parameters of interest, given the data and the priors specified in D. Gibbs sampling is carried out in four steps. In the first block we draw the parameters from the factor equation, while in the remaining three blocks we follow Primiceri (2005) and draw parameters from the VAR part of model.

Block 1: Draw factor loading states $\Lambda|X^T, F^T, R$ and hyperparameters $R|\Lambda, X^T, F^T$

Since the covariance matrix of the error terms in (10) is diagonal, we can estimate all the parameters equation-by-equation. The parameters are sampled using standard arguments for linear regression models (see Koop (2003)).

$$\overline{\lambda_i} = (\overline{Var(\lambda_i)})^{-1}((\underline{Var(\lambda_i)})^{-1}\underline{\lambda_i} + R_{i,i}^{-1}F'X_i)$$
$$\overline{Var(\lambda_i)} = \left((\underline{Var(\lambda_i)})^{-1} + R_{i,i}^{-1}X_i'X_i\right)^{-1}$$

for $i = 1, ..., n$.

Once Λ has been drawn, we can compute the residuals sum of squares from Equation 10, SSE_i^{Λ} and draw the conditional posterior for R from inverse Gamma distribution:

$$R_{i,i}|\cdots \sim IG\left(\frac{\nu_1}{2}, \frac{\delta_1^{(i)}}{2}\right) \text{ for } i = 1, \ldots, n. \tag{17}$$

where $\nu_1 = \nu_0 + T$ and $\delta_1^{(i)} = \delta_0 + SSE_i^{\Lambda}$

Block 2: Drawing coefficient states $B^T|y^T, A^T, \Sigma^T, V$ and hyperparameters $Q|B^T$

As shown in Primiceri (2005), conditional on A^T, Σ^T and V, the observation Equation 11 is linear and has Gaussian innovations with known variance.

Stacking in a vector B_t all R.H.S. coefficients from 11, and defining

$$G_t = \begin{bmatrix} g_{t,1} & 0 & \ldots & 0 \\ 0 & g_{t,2} & \ldots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \ldots & 0 & g_{t,m} \end{bmatrix} \tag{18}$$

where $g_{t,i} = [\mathbf{1}, F'_{t-1}, \ldots, F'_{t-p}]$, for $i = 1, \ldots m$

Equation 11 can be written as

$$F_t = G_t B_t + e_t \sim \mathcal{N}(0, \Psi) \tag{19}$$

where Ψ is a block diagonal matrix given by:

$$\Psi = \begin{bmatrix} \Omega_1 & 0 & \ldots & 0 \\ 0 & \Omega_2 & \ldots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \ldots & 0 & \Omega_m \end{bmatrix}$$

Given the state space form 19 and 12, we can use Carter and Kohn's Gibbs sampling approach, see Section B.3, to sample $B's$ from:

$$B_T|F^T, A^T, \Sigma^T, V \sim \mathcal{N}(B_{T|T}, P^B_{T|T})$$
$$B_t|B_{t+1}, F^t, A^T, \Sigma^T, V \sim \mathcal{N}(B_{t|t+1}, P^B_{t|t+1}), \text{ for } t = T-1, T-2, \ldots, 1$$

Once B^T has been drawn, the innovations in Equation 12 are observables. We can then compute the residuals sum of squares, SSE^B and Q can be easy drawn from inverse Wishart posterior distribution:

$$Q|\dots \sim IW(\Gamma^Q, Var(Q))$$

where $\Gamma^Q = SSE^B + \underline{\Gamma}^Q$

Block 3: Drawing covariance states $A^T|y^T, B^T, \Sigma^T, V$ and hyperparameters $S|A^T$

We use B^T to compute $\hat{F}_t = F_t - G_t' B_t$

Equation 19 can then be written as:

$$A_t \hat{F}_t = \Sigma_t \varepsilon_t \tag{20}$$

where, taking B_T as given, $\hat{F}_t$ is observable. Since A_t is a lower triangular matrix with ones on the main diagonal, the above equation can be rewritten as:

$$\hat{F}_t = \Phi_t \alpha_t + \Sigma_t \varepsilon_t \tag{21}$$

where α_t is defined Equation 13 and Φ_t is the following $m \times m(m-1)/2$ matrix:

$$\Phi_t = \begin{bmatrix} 0 & \dots & 0 \\ -\hat{F}_{1,t} & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & -\hat{F}_{[1,\dots,m-1],t} \end{bmatrix}$$

where, abusing notation, $\hat{F}_{[1,\dots,i],t}$ denotes the row vector $\hat{F}_{1t}, \hat{F}_{2t}, \dots, \hat{F}_{it}$.

As it is noted in Primiceri (2005), the model given by 21 and 13 has a Gaussian but nonlinear state space representation. However, due to assumption of block diagonality of S, this problem can be solved by applying Kalman filter and the backward recursion equation by equation. We draw each block of α_t, where i-th block corresponds to the i-th equation in 21. As in block 3, we can use Carter and Kohn's Gibbs sampling approach, see Section B.3, to sample $\alpha_{i,t}$ recursively:

$$\alpha_{iT}|F^T, B^T, \Sigma^T, V \sim \mathcal{N}(\alpha_{i,T|T}, P_{T|T}^{\alpha_i})$$

$$\alpha_{it}|\alpha_{i,t+1}, F^t, B^T, \Sigma^T, V \sim \mathcal{N}(\alpha_{i,t|t+1}, P_{t|t+1}^{\alpha_i}), \text{ for } t = T-1, T-2, \dots, 1$$

Once A^T has been drawn, the innovations in Equation 13 are observables. We can then compute the residuals sum of squares, SSE^A and S_i can be easy drawn from inverse Wishart posterior distribution:

$$S_i|\cdots \sim IW(\Gamma^{S_i}, Var(S_i))$$

where $\Gamma^{S_i} = SSE^{\alpha_i} + \underline{\Gamma}^{S_i}$, for $i = 2, \ldots, m$

Block 4: Drawing covariance states $\Sigma^T|y^T, B^T, A^T, V$ and hyperparameters $Q|\Sigma^T$

Define $F_t^* = A_t\hat{F}_t$, such that we can rewrite 20 as following:

$$F^* = \Sigma_t\varepsilon_t \tag{22}$$

Taking B^T and A^T as given, F_t^* is observable. Together with the transition Equation 14, the observations in 22 constitutes a nonlinear state space system. This nonlinearity can be converted into linear system, by squaring and taking logarithms of every element of 22, yielding:

$$F_t^{**} = 2h_t + \omega_t \tag{23}$$

$$h_t = h_t - 1 + \nu_t \tag{24}$$

where $F^{**}it = log[(F^*it^2) + \bar{c}]$; $\bar{c}$ represents the offset constant (assigned a value of 0.0001) to prevent taking the log of zero; $\omega_{it} = log(\varepsilon_{it}^2)$; and $h_{it} = log\sigma_{it}$. While the system in 23 and 24 is now linear, it adopts a non-Gaussian state space form due to the innovations in the observation equation being distributed as log $\chi^2(1)$. To convert the system into a Gaussian form, we employ a mixture of normals approximation for the log $\chi^2(1)$ distribution. As suggested by Kim et al. (1998), we opt for a mixture of seven normal densities with component probabilities q_j, means $m_j - 1.2704$, and variance v_j^2, where $j = 1, \ldots, 7$. The constants q_j, m_j, v_j^2 are determined to align with several moments of the log $\chi^2(1)$ distribution. It's worth noting that the variance-covariance matrix of the ω's is diagonal since the variance-covariance matrix of the ε's is the identity matrix. Thus, we can represent each s_{it} in the selection matrix $s^T = [s_1, \ldots, S_T]'$ as:

$$Pr(s_{it} = j|F_{it}^{**}, h_{it}) \propto q_j f_N(F_{it}^{**}|2h_{it} + m_j - 1.2704, v_j^2), j = 1, \ldots, 7, i = 1, \ldots, m$$

This is then used to select which member of the mixture of the normal approximation should be applied to each element of ω, subsequently adjusting the mean of F_{it}^{**} at every point in time. Conditional on B^T, A^T, V, and s^T, the system in 23 and 24 can be approximated as having a linear and Gaussian state space form. Consistent with earlier steps in the sampler, this method enables the recursive derivation of h_t for $t = 1, \ldots, T$ through the Carter and Kohn algorithm.

$$h_T|F^T, A^T, V, s^T \sim \mathcal{N}(h_{T|T}, P^h_{T|T})$$

$$h_t|h_{t+1}, F^t, A^T, V, s^T \sim \mathcal{N}(h_{t|t+1}, P^h_{t|t+1}), \text{ for } t = T-1, T-2, \ldots, 1$$

Similarly, conditional on h^T, we can compute the residuals sum of squares, SSE^h, from Equation 14 and draw the posterior of W from the Inverse-Wishart distribution:

$$W|\cdots \sim IW(\Gamma^W, Var(W))$$

where $\Gamma^W = SSE^h + \underline{\Gamma}^W$, for $i = 2, \ldots, m$

B.3 The Carter and Kohn algorithm

Consider a measurement equation:

$$y_t = H_t\beta_t + \varepsilon_t \sim \mathcal{N}(0, R_t) \tag{25}$$

and a transition equation

$$\beta_t = F\beta_{t-1} + u_t \sim \mathcal{N}(0, Q) \tag{26}$$

The errors from the measurement equation and from the transition equation are assumed to be uncorrelated across time and with each other.

Given that the state space model given in 25 and 26 is linear and Gaussian, the distribution of β_t given $\tilde{y}^T$ and that of β_t given β_{t+1} and $\tilde{y}_t$ for $t = T-1, \ldots, 1$ are also Gaussian:

$$\beta_T|\tilde{y_T} \sim \mathcal{N}(\beta_{T|T}, P_{T|T}), \text{ for } t = T \tag{27}$$

$$\beta_t|\tilde{y}_t, \beta_{t+1} \sim \mathcal{N}(\beta_{t|t,\beta_{t+1}}, P_{t|t,\beta_{t+1}}), \text{ for } t = T-1, T-2, \ldots, 1 \tag{28}$$

where

$$\beta_{T|T} = E(\beta_T|\tilde{y_T})$$

$$P_{T|T} = Cov(\beta_T|\tilde{y_T})$$

$$\beta_{t|t,\beta_{t+1}} = E(\beta_t|\tilde{y}_t, \beta_{t|t+1})$$

$$P_{t|t,\beta_{t+1}} = Cov(\beta_t|\tilde{y}_t, \beta_{t|t+1})$$

Given $\beta_{0|0}$ and $P_{0|0}$, the Gaussian Kalman filter delivers:

$$\beta_{t|t-1} = F\beta_{t-1|t-1} \tag{29}$$

$$P_{t|t-1} = FP_{t-1|t-1}F' + Q \tag{30}$$

$$K_t = P_{t|t-1}H_t'(H_tP_{t|t-1}H_t' + R_t)^{-1} \tag{31}$$

$$\beta_{t|t} = \beta_{t|t-1} + K_t(y_t - H_t\beta_{t|t-1}) \tag{32}$$

$$P_{t|t} = P_{t|t-1} + K_tH_tP_{t|t-1} \tag{33}$$

The last elements of the recursion, that are $\beta_{T|T}$ and $P_{T|T}$, which are the mean and the variance of the normal distribution (see Equation 27) can be used to make a draw for β_T. The draw β_T and the output from the kalman filter are now used for the first step of the backward recursion, which provides $\beta_{T|T-1}$ and $P_{T|T-1}$, used to draw β_{T-1} according to 28, where $\beta_{t|t,\beta_{t+1}}$ and $P_{t|t,\beta_{t+1}}$ are generated from the following updating equations:

$$\beta_{t|t,\beta_{t+1}} = \beta_{t|t} + P_{t|t}F'P_{t+1|t}^{-1}(\beta_{t+1} - F\beta_{t|t}) \tag{34}$$

$$P_{t|t,\beta_{t+1}} = P_{t|t} + P_{t|t}F'P_{t+1|t}^{-1}FP_{t|t} \tag{35}$$

The backward recursion continues until time zero.

C Number of lags

As we have described in Section 2.2, the VAR part of our benchmark model incorporates four lags ($p = 4$). Although a lag length shorter than one year is often deemed too restrictive to capture the dynamics in the oil–macro relationship (see e.g. Hamilton and Herrera (2004)), we explore two alternative models specification with two ($p = 2$) and six ($p = 6$) number of lags. We conduct a quasi-real-time forecasting experiment, similar to that of Bjørnland and Thorsrud (2016). For the period 1991:Q1 –2018:Q4, we estimate the model with varying lags. We then assess out-of-sample forecasting from 1992:Q2–2018:Q4 using root mean forecasting errors (RMSE). As our goal is to compare nested structural models, we maintain consistent model estimations across different forecast datasets. For a detailed discussion on this forecasting experiment for model selection, we refer to Bjørnland and Thorsrud (2016). The results of the forecasting experiment are presented in Table 5.

The table reveals that the benchmark model surpasses the alternative models, boasting a lower RMSE for approximately 2/3 of the variables across all forecasting horizons for the alternative with six lags, and outperforming the model with two lags for almost all of the variables. Moreover, augmenting the number of lags increases computational complexity and, as highlighted in Hamilton and Herrera (2004), employing an excessively restrictive lag structure in modeling the oil market can yield misleading outcomes. Accordingly, we report the results for $p = 4$.

Horizon		**Model**				
		IP index	**Business equipment**	**Mining**	**Non-oil investment**	**All variables**
1	Benchmark ($p = 4$)	0,38	0,60	0,95	0,70	32
	Alternative ($p = 2$)	0,49	0,67	0,98	0,70	2
	Alternative ($p = 6$)	0,41	0,62	0,97	0,70	13
3	Benchmark ($p = 4$)	0,57	0,72	1,01	0,72	30
	Alternative ($p = 2$)	0,67	0,79	1,02	0,74	1
	Alternative ($p = 6$)	0,61	0,74	1,02	0,70	16

Table 5. Quasi-out-of-sample forecasting results. The benchmark is the main model used in the article, where the number of lags is set to four. The alternative models are the models estimated with two and six lags respectively. For each model, variable, and horizon the reported number are the RMSE values. The last column reports how many times current model is ranked as the best model according to RMSE values when the performance is evaluated across all variables.

D Prior specification

Following the methodology of Primiceri (2005) and Bianchi et al. (2017), we aim to diminish the impact of prior beliefs on posterior estimates. For this purpose, we utilize the initial 16 years of our dataset (1974:Q1-1990:Q4) as a training sample to derive an informative prior. The mean and variance of $\underline{B}$ and $\underline{\alpha}$, see Equations 12 and 13 respectively, are set to the OLS point estimates and four times the variance of these estimates from the initial subsample. For $\log \underline{\sigma}$, see Equation 14, the mean is determined by the logarithm of the OLS point estimates of the standard errors from the corresponding time-invariant VAR, while its variance-covariance matrix is assumed to be an identity matrix. In a similar vein, the mean and variance of the factor loadings from (10), $\underline{\Lambda}$, see Equation 10, are set to the OLS point estimates and four times the variance of these estimates derived from the training sample.

$$\underline{B} \sim \mathcal{N}(\hat{B}_{OLS}, 4Var(\hat{B}_{OLS}))$$
$$\underline{\alpha} \sim \mathcal{N}(\hat{\alpha}_{OLS}, 4Var(\hat{\alpha}_{OLS}))$$
$$log\underline{\sigma} \sim \mathcal{N}(log\hat{\sigma}_{OLS}, I_n)$$
$$\underline{\Lambda} \sim \mathcal{N}(\hat{\Lambda}_{OLS}, 4Var(\hat{\Lambda}_{OLS}))$$

We use prior from Inverted Gamma distribution for variance-covariance matrix R .

$$R \sim IG(\frac{\nu_1}{2}, \frac{\delta_1}{2})$$

where $\nu_1 = \nu_0 + T$ and $\delta_1 = \delta_0 + (X - \hat{\Lambda}^{post}F)$. The priors for the remaining hyperparameters are all from the Inverse-Wishart distribution:

$$Q \sim IW(\underline{\Gamma}^Q, Var(Q))$$
$$W \sim IW(\underline{\Gamma}^W, Var(W))$$
$$S_i \sim IW(\underline{\Gamma}_i^S, Var(S_i))$$

where

$$\underline{\Gamma}^Q = k_Q^2(1 + dim_B)Var(\hat{B}_{OLS})$$
$$\underline{\Gamma}^W = k_W^2(1 + dim_W)I_p$$
$$\underline{\Gamma}_i^S = k_S^2(1 + dim_{S_i})Var(\hat{A}_{i,OLS})$$

while $Var(Q) = 1 + dim_B$, $Var(W) = 1 + dim_W$, and $Var(S_i) = 1 + dim_{S_i}$.

Following Korobilis (2013) the degrees of freedom are set to $dim_B = m \times m \times p$, $dim_W = m$ and $dim_{S_i} = 1, .., m-1$, and are larger than the dimension of the corresponding matrices, required to achieve a proper Inverse-Wishart distribution.

The benchmark results in this paper are obtained using the following values: $k_Q = 0.1$, $k_S = 0.1$, $k_W = 0.1$ and $\nu_0 = 10$, $\delta_0 = 10$.

In Section H.2 in the Online Appendix we justify this choice and demonstrate the robustness of our conclusions to alternative prior specifications. We focus on alternative specifications of k_Q, k_S, and k_W, since the choice for other priors seems to be of minor importance (see e.g. Primiceri (2005)).

E Convergence of the Markov Chain Monte Carlo algorithm

We perform 90,000 iterations of the Gibbs sampler. The first 40,000 draws are discarded and only every tenth of the remaining iterations is used for inference. The produced results are not sensitive to the number of discarded draws or the number of passes used for inferences. Following Primiceri (2005) and Baumeister and Peersman (2013b), we ascertain that our Markov chain has converged based on the inefficiency factors (IFs) for the posterior estimates of the parameters, that is the inverse of the relative numerical efficiency (RNE) measure proposed by Geweke (1992). Here the estimates are performed by employing a four percent tapered window used in computation of the RNE. As was noticed by Primiceri (2005), values of the IFs below or around 20 are regarded as satisfactory. As can be seen from the summary of the distribution of the inefficiency factors for different set of parameters, reported in Table 6, the sample seems to have converged. That is, all mean IF values are below 5 and 90 percent of the IFs are below at most 21, indicating modest autocorrelation for all elements.

	Median	Mean	Min	Max	10-th Percentile	90-th Percentile
B_t	1,92	1,80	0,56	8,39	1,21	2,75
Λ	0,96	0,94	0,49	1,93	0,68	1,30
Σ_t	5,49	4,04	1,05	21,09	2,27	10,34
A_t	9,92	3,46	0,97	92,83	1,37	21,14
V	3,99	3,83	1,30	40,57	2,64	5,48
R	0,98	0,93	0,53	2,00	0,71	1,37

Table 6. Summery of the distribution of the IFs for the benchmark model. Table includes different set of parameters, where V is the set of hyperparameters $\{Q, S, W\}$

F Additional results

F.1 Impulse responses for investment series

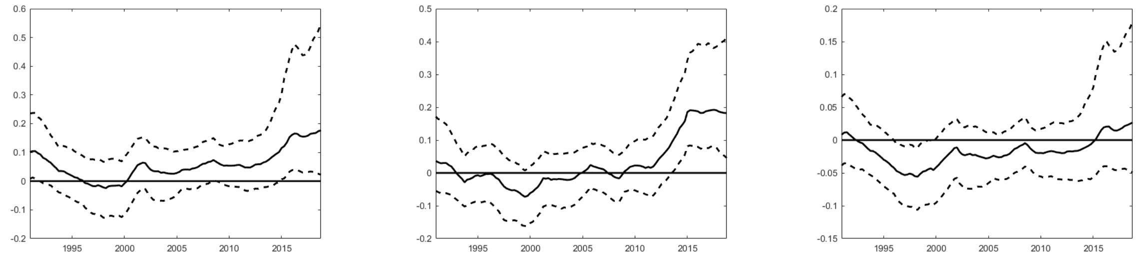

(a) Manufacturing **(b)** Commercial and Health Care **(c)** Power and Communication

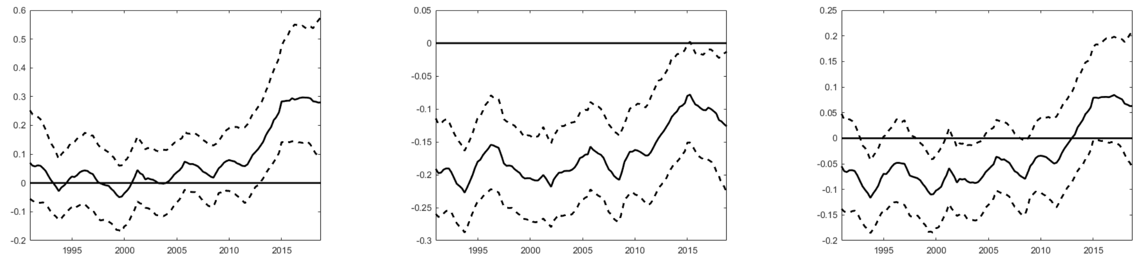

(d) Industrial Equipment **(e)** Transportation Equipment **(f)** Other Equipment

Figure 12. The effect of an oil-specific shock: Impulse responses for investment in structures (upper row): Manufacturing, commercial and health care, and power and communication; and in equipment (lower row): Industrial, transportation and other equipment. The initial shock is normalized to increase oil prices by 10%. All responses are reported in levels (of the standardized data). We report impulse responses across the sample two quarters after the shock. The dashed lines represent 68% posterior probability bands. The solid line is the median estimate.

F.2 Results with states employment, 1990-2018

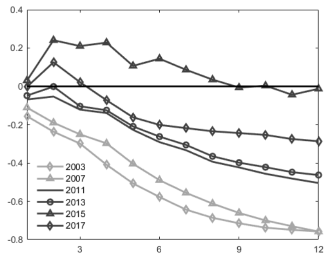

(a) Non-oil Investment (median)

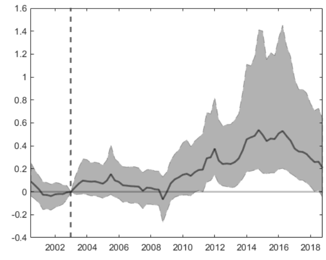

(b) Non-oil Investment (diff two quarters)

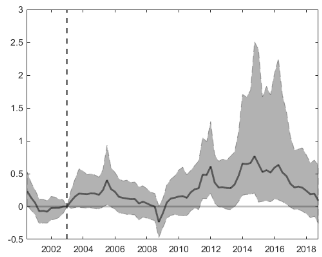

(c) Investment(diff four quarters)

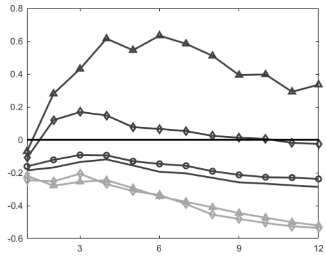

(d) Income (median)

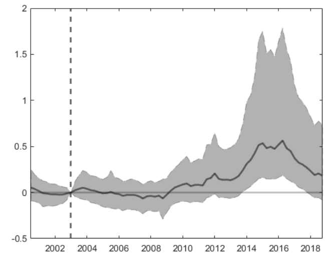

(e) Income (diff two quarters)

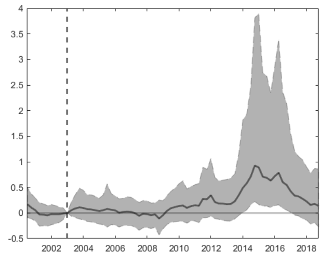

(f) Income (diff four quarters)

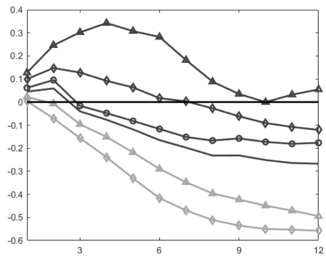

(g) IP Index (median)

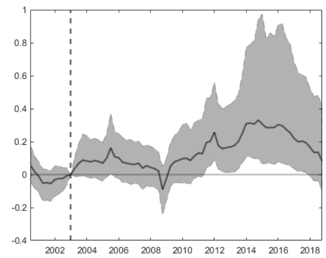

(h) IP Index (diff two quarters)

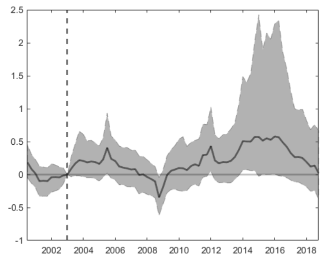

(i) IP Index (diff four quarters)

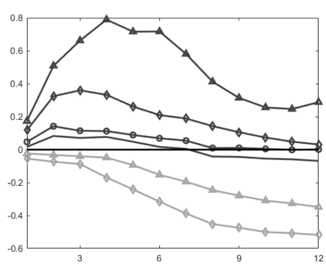

(j) Business Equipment (median)

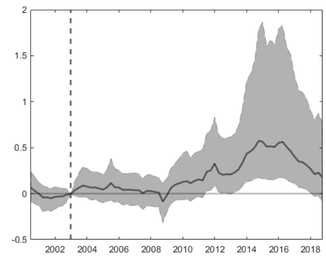

(k) Business Equipment (diff two quarters)

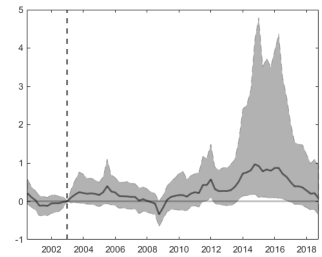

(l) Business Equipment (diff four quarters)

Figure 13. The effect of an oil-specific shock in a model with state employment: Impulse responses for non-residential investment, income, industrial production and business equipment. The initial shock is normalized to increase oil prices by 10%. All responses are reported in levels (of the standardized data). Left column: posterior median of impulse responses at six different periods in time. Middle and right columns: the difference between the responses in period 2000:Q3-2018:Q4 and the responses in 2002:Q1 after two and four quarters respectively. The solid line is the difference between the median estimates. The shaded area represents 68% posterior probability bands for the difference in impulse responses.

G Alternative hypothesis

G.1 Impulse responses, level and difference: Flow supply shock

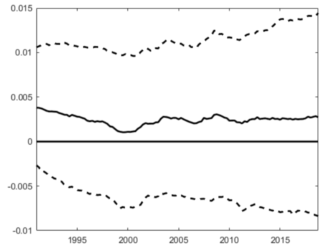

(a) Oil Price

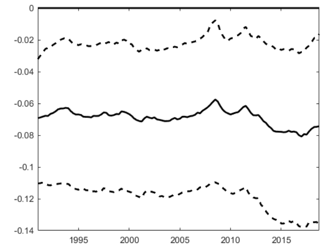

(b) Income

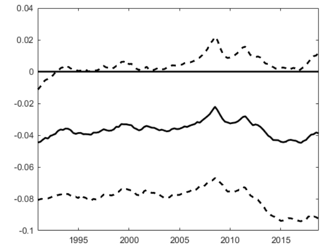

(c) Non-oil Investment

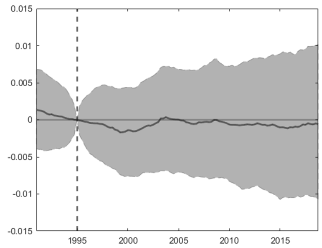

(d) Oil Price

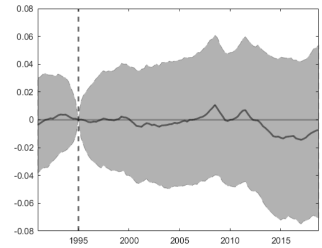

(e) Income

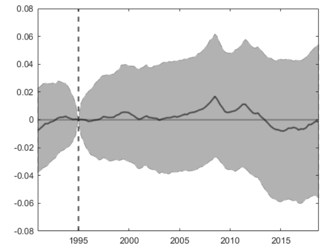

(f) Non-oil Investment

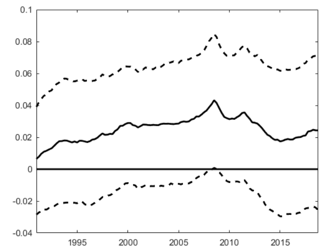

(g) Total Employment

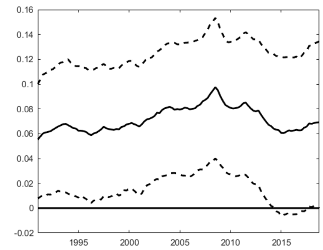

(h) IP Index

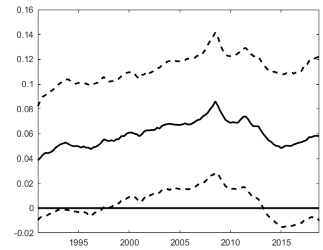

(i) Manufacturing

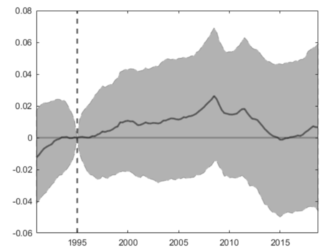

(j) Total Employment

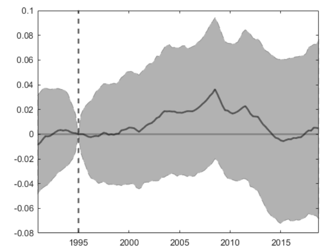

(k) IP Index

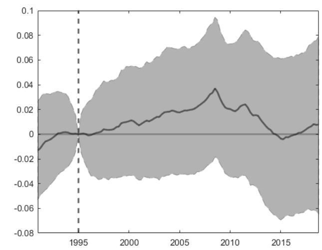

(l) Manufacturing

Figure 14. The effect of a flow supply shock: Impulse responses for the real price of oil and selected indicators of the U.S. economy. The initial shock is normalized to decrease non-U.S. oil production by 1%. First and third rows: Impulse responses across the sample two quarters after the shock. The solid line is the median estimate. The dashed lines represent 68% posterior probability bands. Second and fourth rows: The difference between the responses in period 1991:Q1-2018:Q4 and the responses in 1995:Q1 after two quarters respectively. The solid line is the difference between the median estimates. The shaded area represents 68% posterior probability bands for the difference in impulse responses.

G.2 Impulse responses, level and difference: Flow demand shock

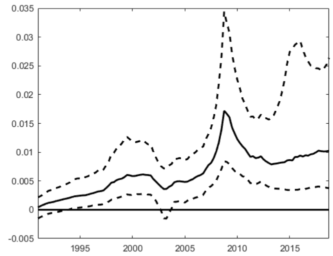

(a) Oil Price

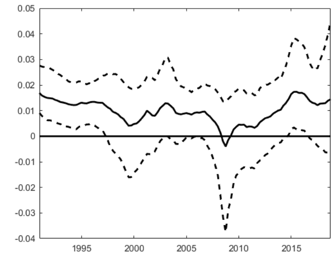

(b) Income

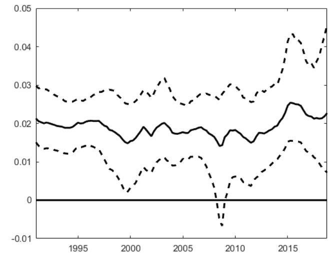

(c) Non-oil Investment

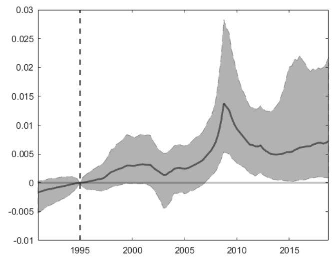

(d) Oil Price

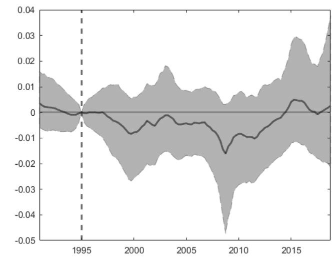

(e) Income

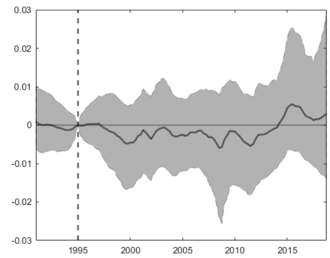

(f) Non-oil Investment

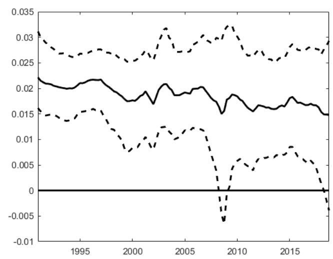

(g) Total Employment

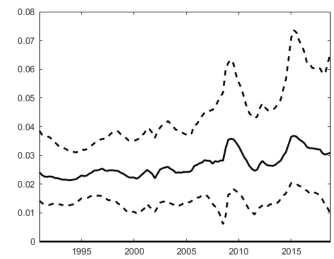

(h) IP Index

(i) Manufacturing

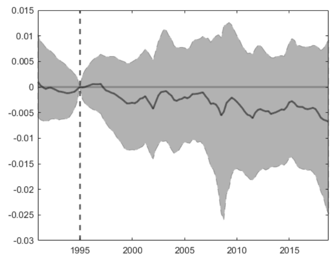

(j) Total Employment

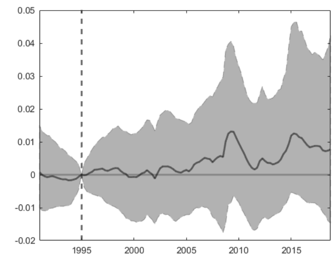

(k) IP Index

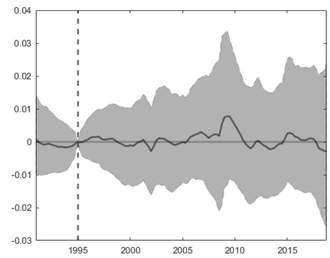

(l) Manufacturing

Figure 15. The effect of flow demand shock: Impulse responses for the real price of oil and selected indicators of the U.S. economy. The initial shock is normalized to increase global activity by 1%. First and third rows: Impulse responses across the sample two quarters after the shock. The solid line is the median estimate. The dashed lines represent 68% posterior probability bands. Second and fourth rows: The difference between the responses in period 1991:Q1-2018:Q4 and the responses in 1995:Q1 after two quarters respectively. The solid line is the difference between the median estimates. The shaded area represents 68% posterior probability bands for the difference in impulse responses.

G.3 Impulse responses to an oil-specific shock for GECON and real oil prices

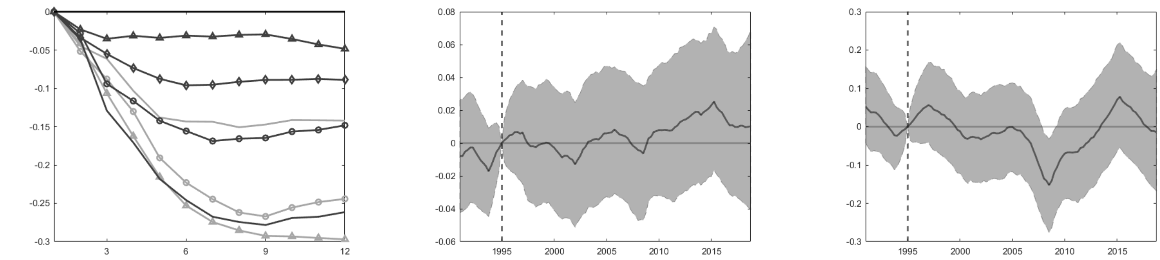

(a) GECON (median) **(b)** GECON (after two quarter) **(c)** GECON (after four quarter)

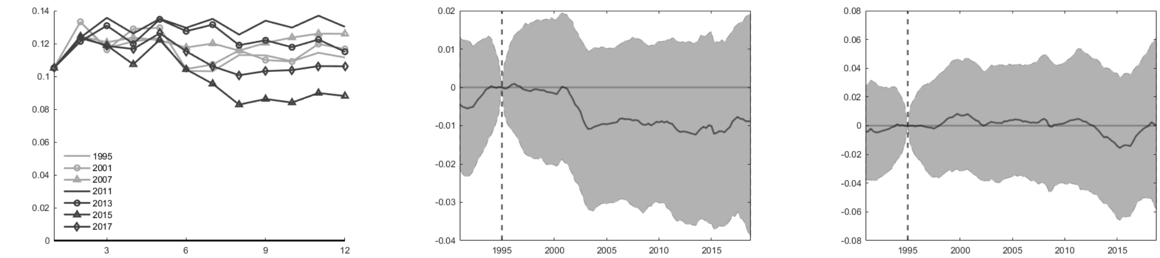

(d) Oil Price (median) **(e)** Oil Price (after two quarter) **(f)** Oil Price (after four quarter)

Figure 16. The effect of an oil specific shock: Impulse responses for GECON and oil price. The initial shock is normalized to increase oil prices by 10%. All responses are reported in levels (of the standardized data). Left column: posterior median of impulse responses at seven different periods in time. Middle and right columns: the difference between the responses in period 1991:Q1-2018:Q4 and the responses in 1995:Q1 after two and four quarters respectively. The solid line is the difference between the median estimates. The shaded area represents 68% posterior probability bands for the difference in impulse responses.

G.4 Including shadow rate, expectations and real interest rate

We analyse if the period of low interest rates has changed the relationship between oil prices and macroeconomic variables. In particular, as pointed out by Datta et al. (2021), as the central bank can not respond to inflationary pressures by changing interest rates at the zero lower bound (ZLB), changes in inflation can affect the real rate of interest differently. Thus, the impact of an oil specific shock on variables such as output, consumption and equity prices, may also be different at the ZLB. To account for this, we add a shadow rate defined by Wu and Xia (2016) to our dataset, in addition to inflation expectations and the real interest rate, and re-estimate the model. We show that the responses for the interest rates and inflation expectations are as expected and show no evidence of time-varying changes (c.f., Figure 17), and that results are robust to the inclusion of these variables (c.f., Figure 18).

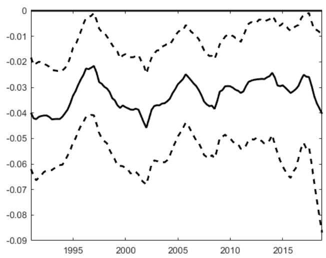

(a) Shadow Rate (after two quarter)

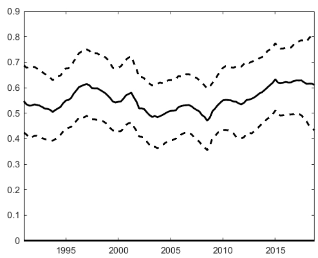

(b) Professionals Forecasters Expectation - GDP Inflation (after two quarter)

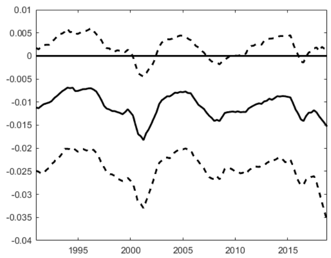

(c) Real Interest Rate (after two quarter)

Figure 17. The effect of an oil specific shock in a model with monetary variables: Impulse responses for shadow interest rate, inflation expectations and real interest rate. The initial shock is normalized to increase oil prices by 10%. All responses are reported in levels (of the standardized data). We report impulse responses across the sample two quarters after the shock. The dashed lines represent 68% posterior probability bands. The solid line is the median estimate.

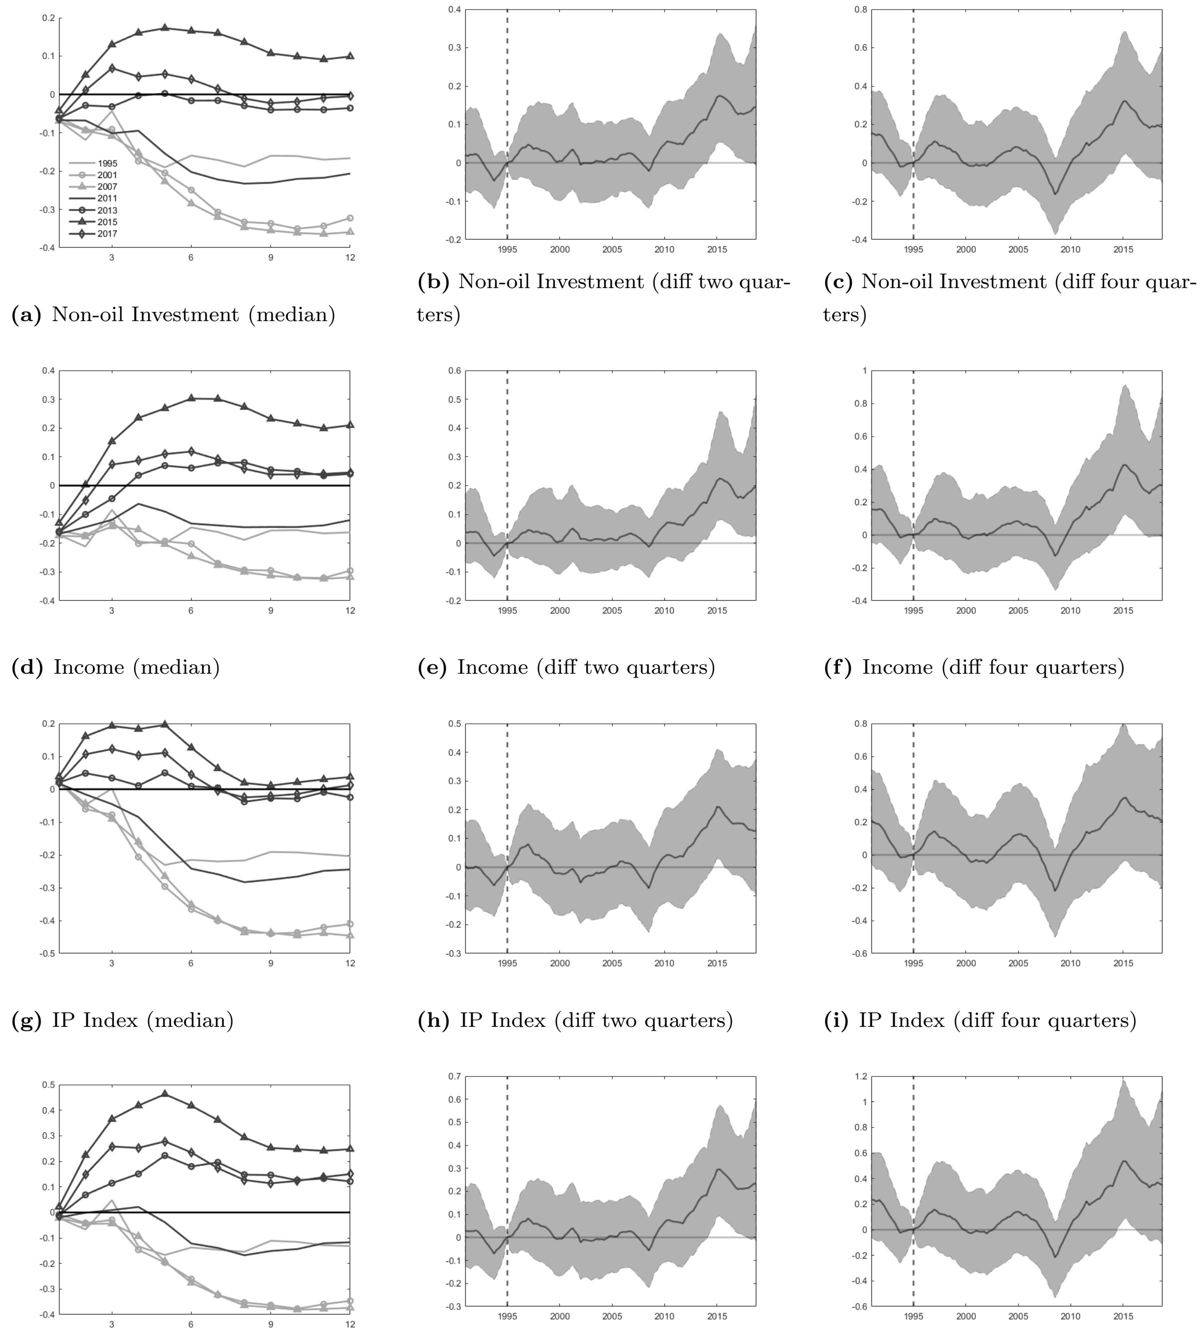

(a) Non-oil Investment (median)

(b) Non-oil Investment (diff two quarters)

(c) Non-oil Investment (diff four quarters)

(d) Income (median)

(e) Income (diff two quarters)

(f) Income (diff four quarters)

(g) IP Index (median)

(h) IP Index (diff two quarters)

(i) IP Index (diff four quarters)

(j) Business Equipment (median)

(k) Business Equipment (diff two quarters)

(l) Business Equipment(diff four quarters)

Figure 18. The effect of an oil specific shock in a model with monetary variables: Impulse responses for non-residential investment, income, industrial production and business equipment. The initial shock is normalized to increase oil prices by 10%. All responses are reported in levels (of the normalized data). Left column: posterior median of impulse responses at seven different periods in time. Middle and right columns: the difference between the responses in period 1991:Q1-2018:Q4 and the responses in 1995:Q1 after two and four quarters respectively. The solid line is the difference between the median estimates. The shaded area represents 68% posterior probability bands for the difference in impulse responses.

G.5 Other commodity prices

We examine if our results could be due to the fact that there are other commodity prices, say corn and coal prices, correlated with oil prices, that also show evidence of time varying behavior. Hence, we may have found significant effects on employment in states with high coal and corn production. In order to address this issue, we include prices of coal and corn in our dataset and reestimate our model. As can be seen from Figure 19, the effects of an oil price shock on either coal or corn prices are stable over time. Furthermore, the results for the other variables remain robust, see Figure 20.

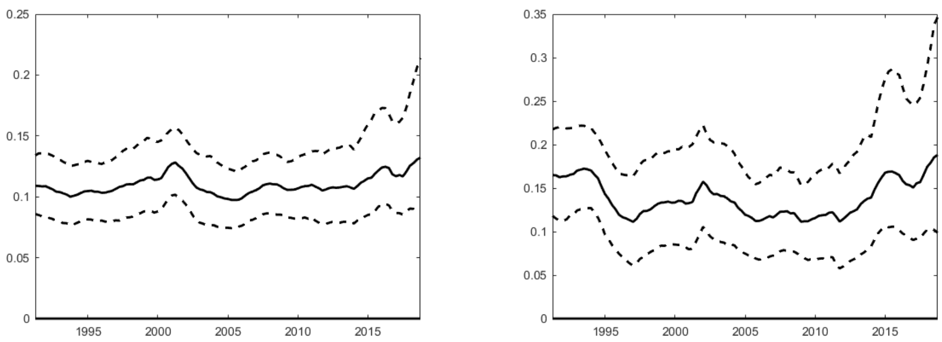

(a) Corn Prices (after two quarter)

(b) Coal Prices (after two quarter)

Figure 19. The effect of an oil specific shock in a model with other commodity prices: Impulse responses for corn and coal prices. The initial shock is normalized to increase oil prices by 10%. All responses are reported in levels (of the standardized data). We report impulse responses across the sample two quarters after the shock. The dashed lines represent 68% posterior probability bands. The solid line is the median estimate.

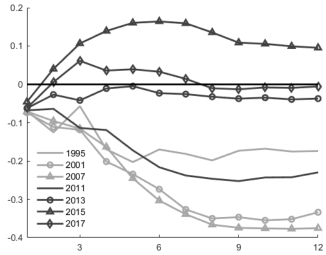

(a) Non-oil Investment (median)

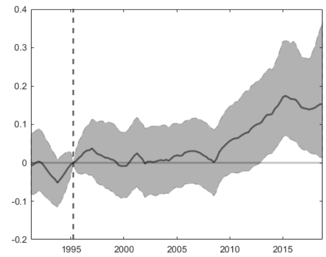

(b) Non-oil Investment (diff two quarters)

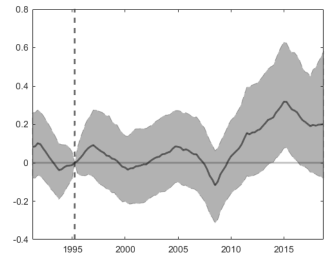

(c) Non-oil Investment (diff four quarters)

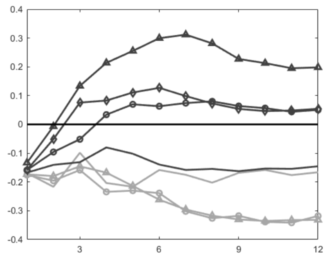

(d) Income (median)

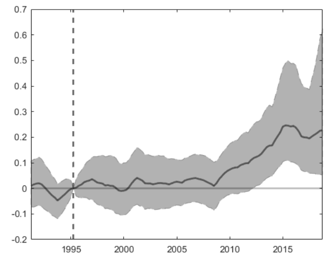

(e) Income (diff two quarters)

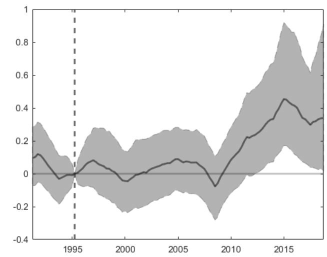

(f) Income (diff four quarters)

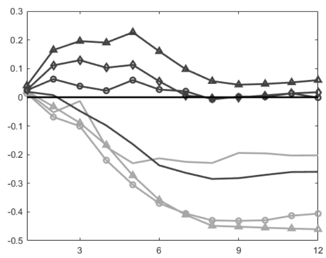

(g) IP Index (median)

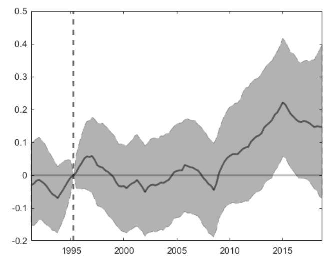

(h) IP Index (diff two quarters)

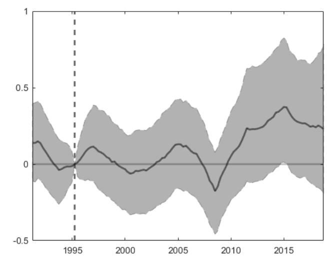

(i) IP Index (diff four quarters)

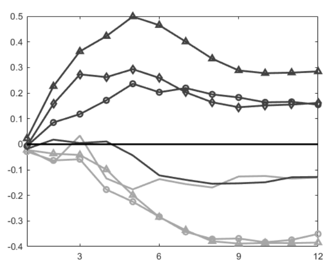

(j) Business Equipment (median)

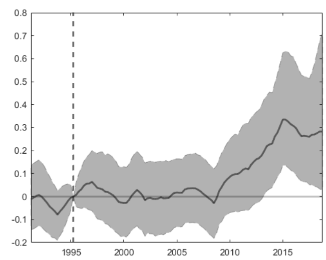

(k) Business Equipment (diff two quarters)

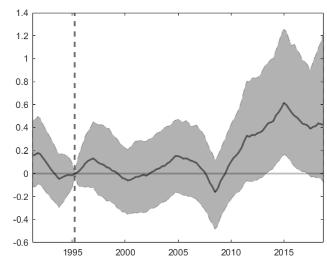

(l) Business Equipment (diff four quarters)

Figure 20. The effect of an oil specific shock in a model with several commodity prices: Impulse responses for non-residential investment, income, industrial production and business equipment. The initial shock is normalized to increase oil prices by 10%. All responses are reported in levels (of the normalized data). Left column: posterior median of impulse responses at seven different periods in time. Middle and right columns: the difference between the responses in period 1991:Q1-2018:Q4 and the responses in 1995:Q1 after 2 and 4 quarters respectively. The solid line is the difference between the median estimates. The shaded area represents 68% posterior probability bands for the difference in impulse responses.

H Robustness

H.1 Model Selection

Our findings, presented in Section 3, show a clear evidence of time-varying effects. However, pinpointing the primary source of this variation—whether from coefficients or the variance-covariance matrix of innovations — is challenging. For instance, it could be argued that our main findings are influenced more by heteroskedaxtic shocks than by genuine structural changes in the economy.

To address these issues and illustrate how different time-varying components influence our main results, we have followed Bianchi et al. (2017) and examined three model variations. In each model we allow only one time-varying component at a time.

1. TVB Model: Only the coefficients, B_t, are time-varying, with a constant variance-covariance matrix. I.e. time variation is coming only from Equation 6.
2. TVA Model: Only elements of A_t vary, indicating time-dependent effects of innovations on model variables. I.e. time variation is coming only from Equation 7.
3. SV Model: Only standard deviations are time-varying, with fixed coefficients B_t and elements of A_t, i.e time variation is coming only from Equation 8

Figure 21 compares the benchmark model coefficients $b_i^{benchmark}$ with TVB coefficients b_i^{TVB}, taken from Equation 5 where the left hand side factor is the real oil price.[23]

Interestingly, when the variance-covariance matrix remains static, coefficients show more fluctuation, especially around the 2007/2009 financial crisis. This suggests that the crisis-induced volatility might overshadow genuine structural changes in the oil market, when we do not account for heteroskedastity in standard deviations.

Comparing the impulse responses of the benchmark and TVB models (Figure 22 first and second rows), we observe more noise and fluctuation in the latter, although the core findings remain consistent.[24]

[23] To save the space we only show coefficients from the oil equation in the VAR part of the model. However, the behavior of coefficient in other equations are very similar, and would not change the discussion following below.

[24] As discussed in Primiceri (2005) and Bianchi et al. (2017), the random walk assumption for the evolution of the coefficients has some undesirable implications. For this reason we have also re-estimated the model assuming that parameters follow an AR(1) process instead of random walk, where we followed Primiceri (2005) and set autoregrassive coefficients to 0.95. Although, this has contributed to a more disciplined behavior of the B_t, still the model captures many temporary parameter shifts. The results from this model are available upon request.

The SV model, focusing solely on time-varying standard deviations, shows minimal time variation in responses (see Figure 22 third row). When we assessed the forecasting performance of these models, the benchmark consistently outperformed the SV model. Based on these results we conclude that our main findings are not driven by heteroskedastic shocks.

The bottom row in Figure 22 shows impulse responses from the TVA model, which allows only for time-varying simultaneous relationships. As it is noted in Primiceri (2005), allowing the matrix A_t to vary over time is crucial for a time varying structural VAR. A constant A_t would imply that an innovation to for instance oil price has a time invariant effect on the estimated factors. This is clearly undesirable as simultaneous interactions among variables are fundamental in our quarterly model. From the figure, similar to the case of the TVB model, the absence of stochastic volatility yields noisier impulse responses. Concurrently, we note indications of more positive responses throughout the estimated sample period, advocating for the presence of time-varying simultaneous relationships.

Our analysis underscores the presence of time variation in our model. Structural changes, rather than just temporary shifts, play a significant role in our findings. While both time-varying coefficients and simultaneous relationships are crucial, it's also essential to consider changes in volatilities to differentiate between temporary and permanent shifts. Ignoring these nuances can lead to a noisier model and less definitive results.

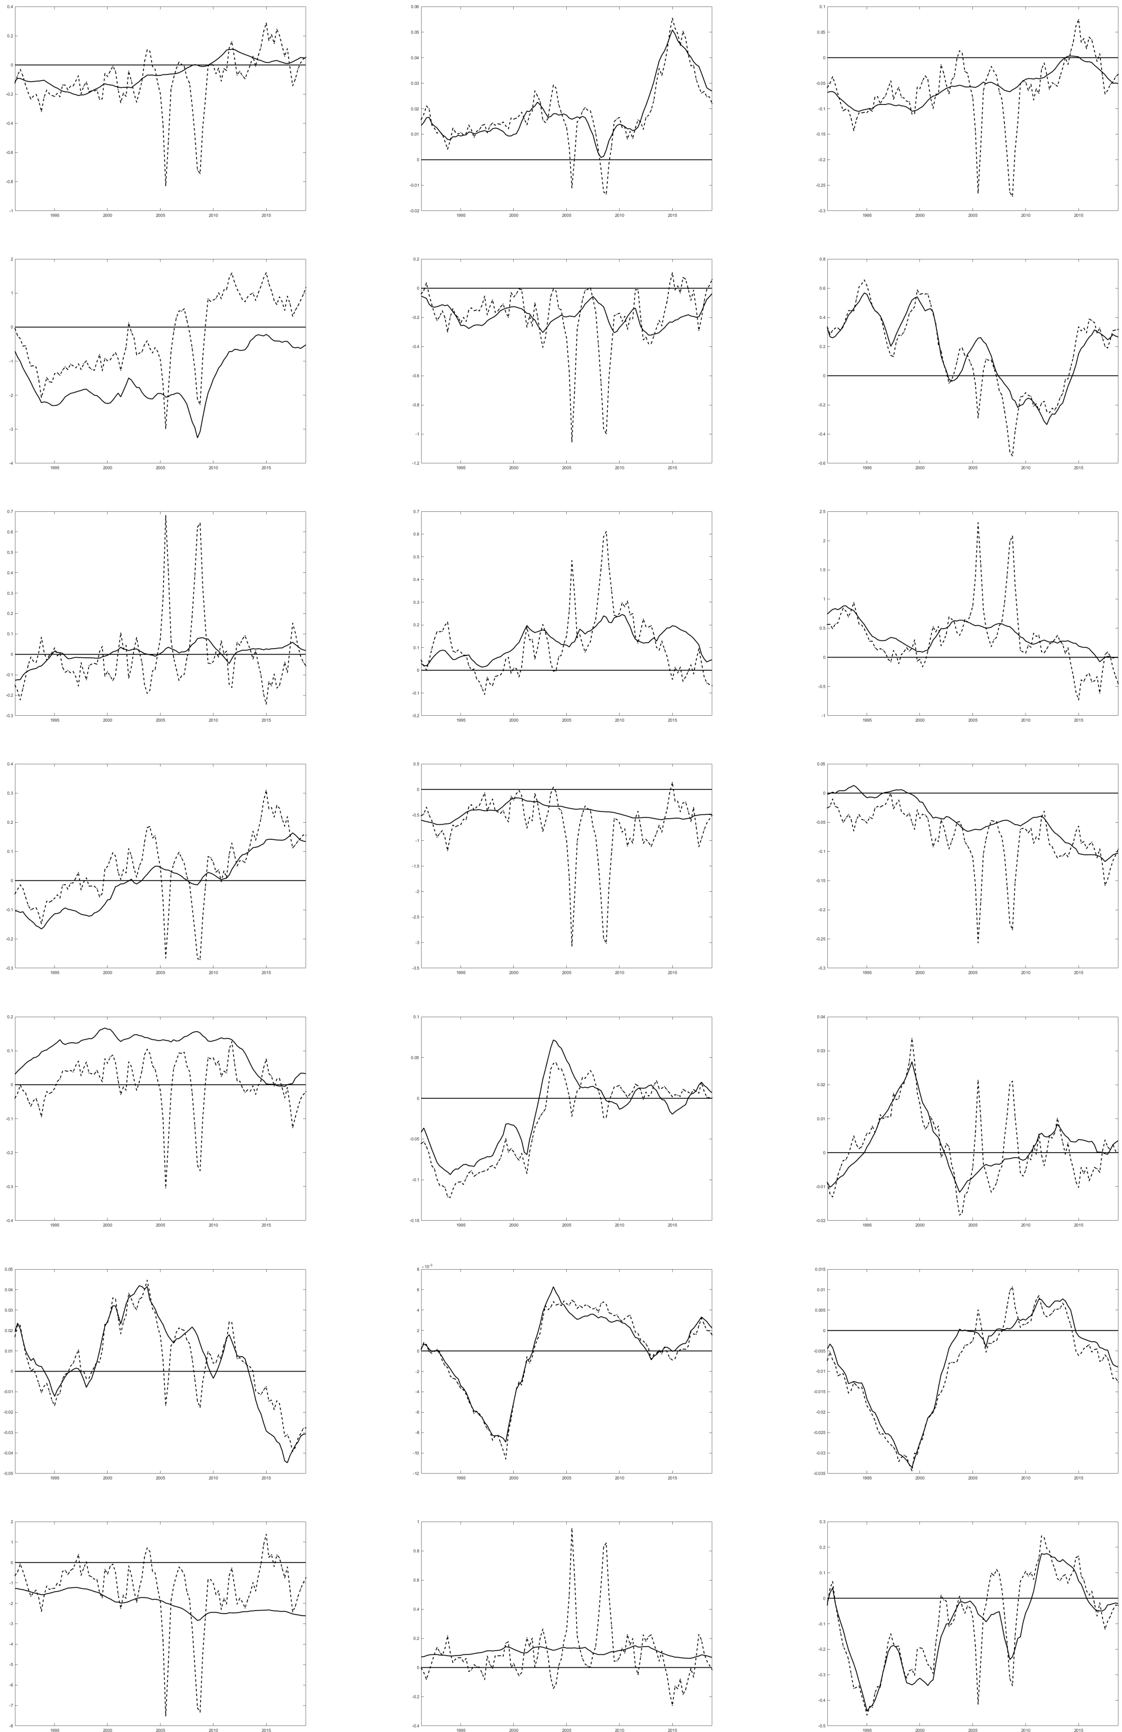

Figure 21. Posterior mean of time-varying coefficients, B_t, from Equation 5. Solid line: Benchmark model; dashed line: TVB model. To save space, the 21 coefficients, i.e., number of lags ($p = 4$) times number of variables in the VAR part of the model ($m = 5$) plus the constant, are taken from the oil equation (the equation where oil price is the left-hand side variable). Coefficients from other equations can be obtained upon request.

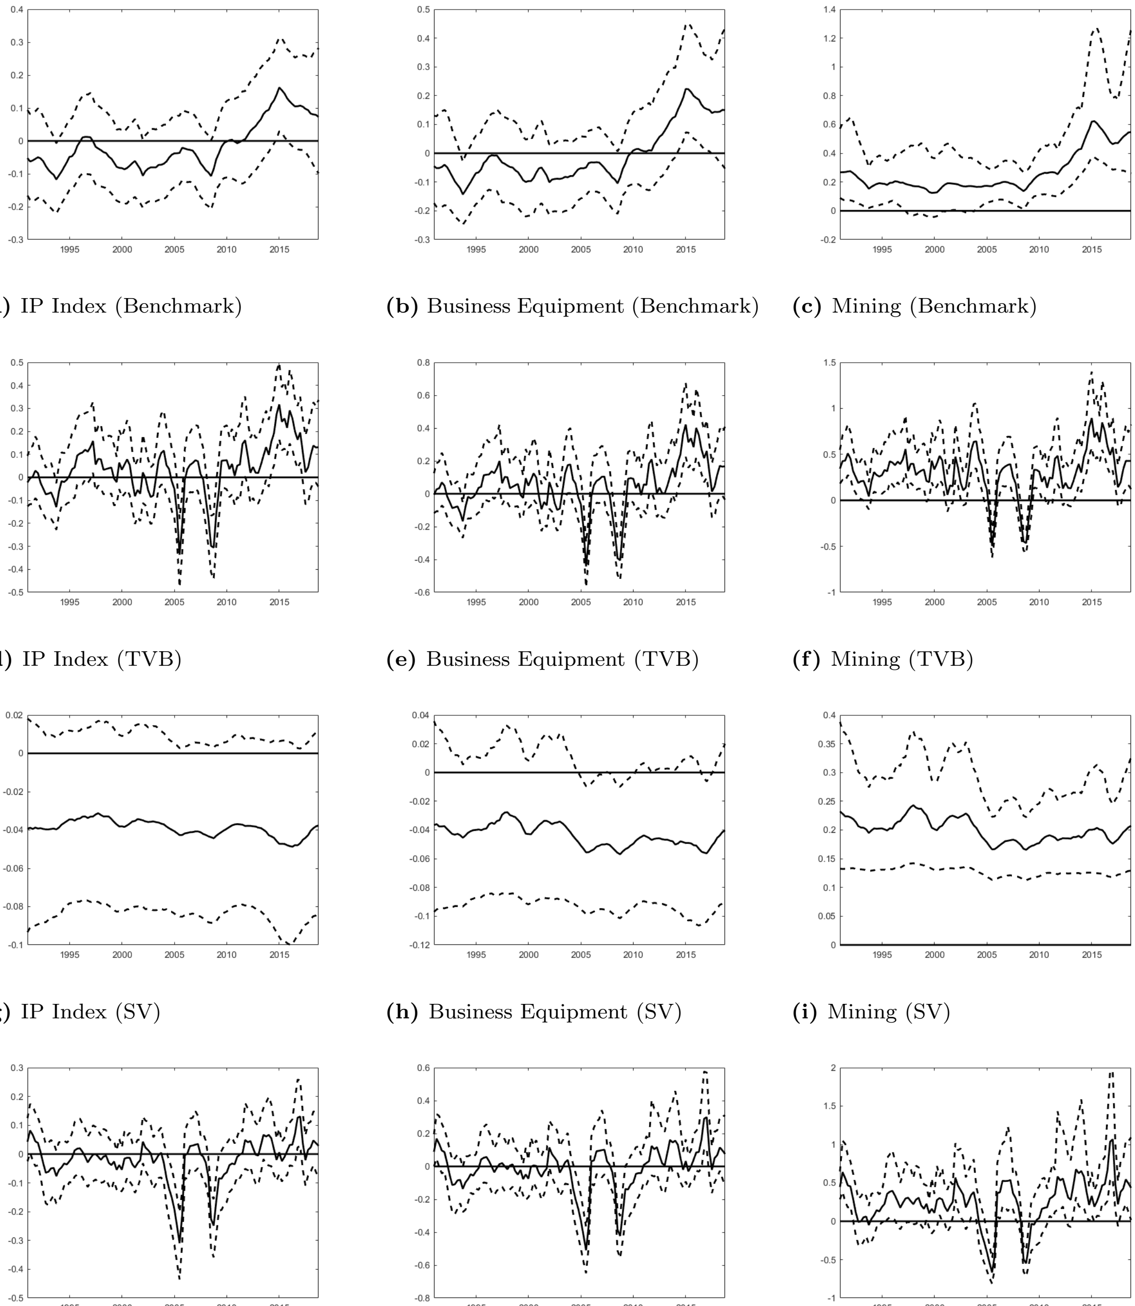

(a) IP Index (Benchmark) **(b)** Business Equipment (Benchmark) **(c)** Mining (Benchmark)

(d) IP Index (TVB) **(e)** Business Equipment (TVB) **(f)** Mining (TVB)

(g) IP Index (SV) **(h)** Business Equipment (SV) **(i)** Mining (SV)

(j) IP Index (TVA) **(k)** Business Equipment (TVA) **(l)** Mining (TVA)

Figure 22. The effect of an oil-specific shock: Impulse responses from the benchmark model (a-c); time-varying coefficients model (d-f); stochastic volatility model (g-i); time-varying "innovation" model (j-l); for selected indicators of the U.S. economy. The initial shock is normalized to increase oil prices by 10%. All responses are reported in levels (of the standardized data). We report impulse responses across the sample two quarters after the shock. The dashed lines represent 68% posterior probability bands. The solid line is the median estimate.

H.2 Priors Selection

As highlighted by Primiceri (2005), the parameters k_Q, k_S, and k_W do not directly parameterize time variation. Instead, they shape our prior beliefs regarding the extent of time variation in parameters. Specifically, k_Q informs our beliefs about time variation in the coefficients of Equation 5. Meanwhile, k_W and k_S define our beliefs about time variation in the equation's stochastic volatility. It's crucial to recognize the trade-off: a high k_Q combined with low k_W and k_S values will primarily capture model fit through the time-varying coefficients, B_t. Conversely, elevating k_S and k_W while diminishing k_Q will nearly eliminate variation in B_t. As these extreme cases are explored in previous section, where we evaluate models with constant coefficients and stochastic volatility, in this section we focus on cases where k_Q, k_S, and k_W are set to lower values, as commonly seen in literature.[25]

To rigorously evaluate our model against alternative specifications, we adopt a quasi-real-time forecasting experiment, similar to that in Section C. In particular, for the period 1991:Q1 –2018:Q4, we estimate the model with varying priors. We then assess out-of-sample forecasting from 1992:Q2–2018:Q4 using root mean forecasting errors (RMSE). As our goal is to compare nested structural models, we maintain consistent model estimations across different forecast datasets. For a detailed discussion on this forecasting experiment for model selection, we refer to Bjørnland and Thorsrud (2016).

Horizon		**Model**				
		IP Index	**Business equipment**	**Mining**	**Non-oil investment**	**All variables**
1	Benchmark	0,38	0,60	0,95	0,70	20
	k_Q	0,56	0,70	1,04	0,71	0
	k_S	0,39	0,60	0,96	0,70	1
	k_W	0,43	0,65	0,94	0,71	26
3	Benchmark	0,57	0,72	1,01	0,72	18
	k_Q	0,75	0,87	1,07	0,78	0
	k_S	0,59	0,73	1,01	0,72	1
	k_W	0,56	0,70	1,00	0,69	28

Table 7. Quasi-out-of-sample forecasting results. The benchmark is the main model used in the article. k_Q, k_S, and k_W denotes which of the parameters have been changed in current version of the model (see text the main text for additional explanations.) For each model, variable, and horizon the reported number are the RMSE values. The last column reports how many times current model is ranked as the best model according to RMSE values when the performance is evaluated across all variables.

Table 7 presents the forecasting experiment results, and Figure 23 illustrates the impulse responses across different models. From Figure 23, we see that models with reduced

[25] For instance, Korobilis (2013), Primiceri (2005), Cogley and Sargent (2001), and Stock and Watson (1996) set the k_Q prior to 0.01; Primiceri (2005) also sets k_W to 0.01.

time variation in coefficients $(B's)$ yield smoother impulse responses. Yet, we can discern variations in how industrial and aerospace production react to oil price increases. As Table 7 indicates, the benchmark model surpasses this specification for nearly all variables. A model with a lower k_S value outperforms the benchmark in certain cases. However, the impulse responses between them are nearly indistinguishable, suggesting results are less sensitive to this parameter choice. The table also shows enhanced forecasting performance with a decreased value of k_W. Yet, the impulse responses become more erratic, especially noticeable during financial crises. These observations align with findings in Section H.1 for models with constant volatilities and time-varying coefficients. Such models seem to inadequately capture the dynamics during crises, whereas time-varying coefficients adeptly account for the heightened volatilities, leading to more accurate predictions based on RMSE values. Given the significant shifts in the volatility of oil prices and other international business cycle shocks over the past decade, as documented by Baumeister and Peersman (2013a) and Baumeister and Peersman (2013b), we posit that setting k_W to 0.1 represents an economically sound choice.

In conclusion, our analysis emphasizes the need for a balanced approach in setting priors for time-varying parameters. At the same time as we show that the time variation identified in this paper persists across most of alternative model specifications.

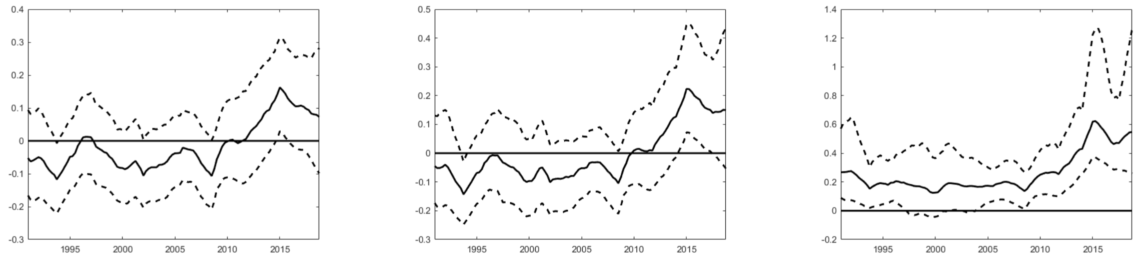

(a) IP Index (Benchmark) **(b)** Business Equipment (Benchmark) **(c)** Mining (Benchmark)

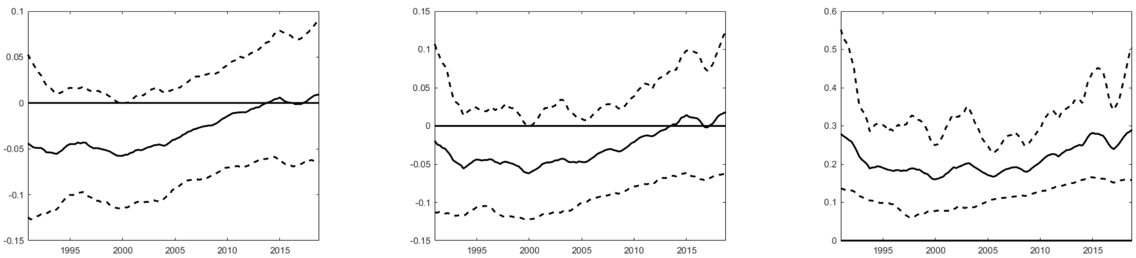

(d) IP Index ($k_Q = 0.01$) **(e)** Business Equipment ($k_Q = 0.01$) **(f)** Mining ($k_Q = 0.01$)

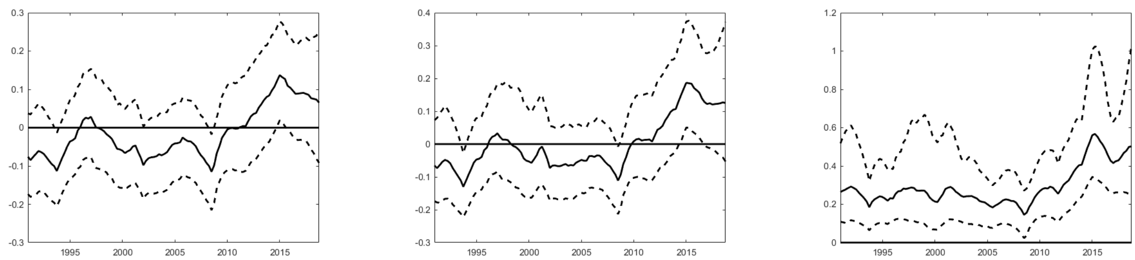

(g) IP Index ($k_S = 0.01$) **(h)** Business Equipment ($k_S = 0.01$) **(i)** Mining ($k_S = 0.01$)

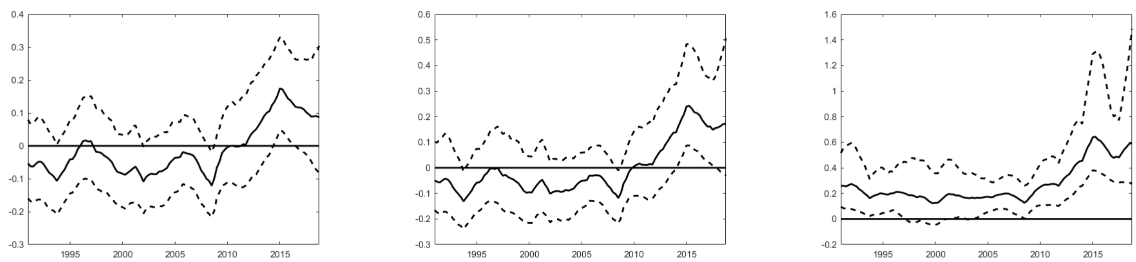

(j) IP Index ($k_W = 0.01$) **(k)** Business Equipment ($k_W = 0.01$) **(l)** Mining ($k_W = 0.01$)

Figure 23. The effect of an oil-specific shock: Impulse responses from the benchmark model and alternative model specifications, where k_Q, k_S, and k_W denotes which of the parameters have been changed in current version of the model, for selected indicators of the U.S. economy. The initial shock is normalized to increase oil prices by 10%. All responses are reported in levels (of the standardized data). We report impulse responses across the sample two quarters after the shock. The dashed lines represent 68% posterior probability bands. The solid line is the median estimate.

H.3 Monthly TVP-FAVAR

We re-estimate the baseline model using monthly data. This required us to reduce our dataset since variables like investment and personal income aren't available at a monthly frequency. We further limited our dataset to only encompass industrial production divided by market groups, excluding NAICS industry groups. This results in a panel of 24 series.

To align with the baseline model, we retained the same sample period as the training sample and incorporated 12 lags in the VAR segment. Further, we set the number of factors extracted from the domestic datablock to two, consistent with the baseline model. Similar to the quarterly model, we ensured stationarity by transforming data into growth rates, taking log differences, and adjusting for local mean. The data used for factor extraction was standardized. The estimation procedure follows that of the quarterly model, as detailed in Section 2.4.

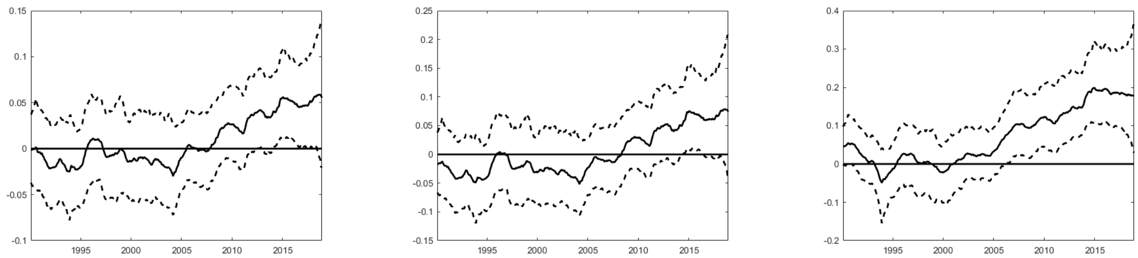

(a) IP Index **(b)** Business Equipment **(c)** Mining

Figure 24. The effect of an oil-specific shock using monthly data: Impulse responses for IP index, business equipment, and mining production. The initial shock is normalized to increase oil prices by 10%. All responses are reported in levels (of the normalized data). We report impulse responses across the sample two quarters after the shock. The dashed lines represent 68% posterior probability bands. The solid line is the median estimate.

H.4 Robustness to the choice of global activity variable

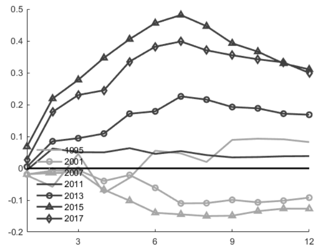

(a) Non-oil Investment (median)

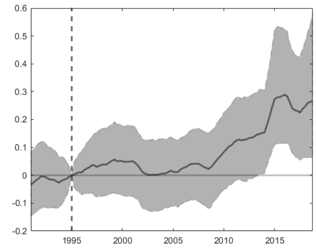

(b) Non-oil Investment (diff two quarters)

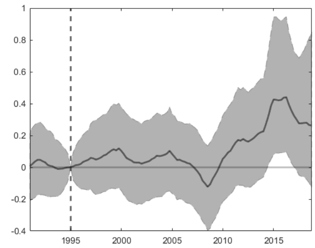

(c) Non-oil Investment (diff four quarters)

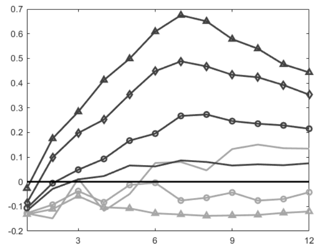

(d) Income (median)

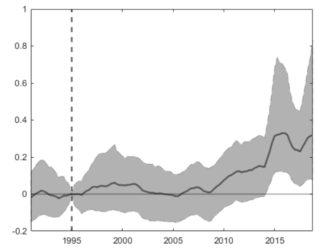

(e) Income (diff two quarters)

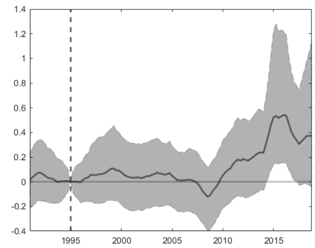

(f) Income (diff four quarters)

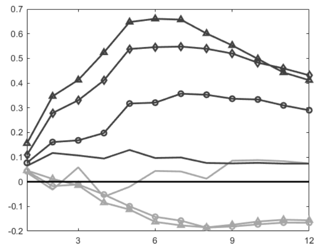

(g) IP Index (median)

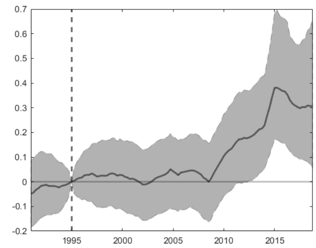

(h) IP Index (diff two quarters)

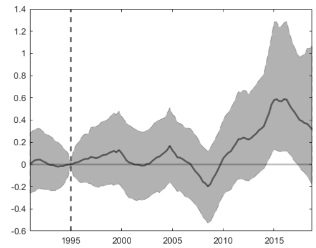

(i) IP Index (diff four quarters)

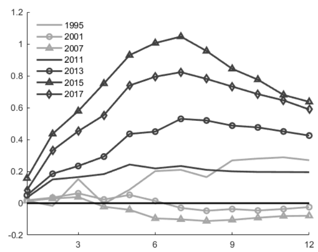

(j) Business Equipment (median)

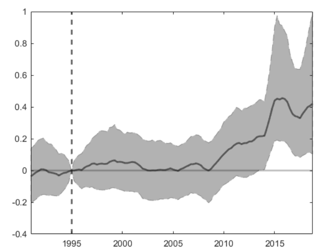

(k) Business Equipment (diff two quarters)

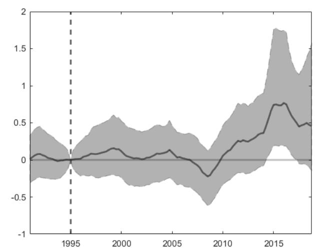

(l) Business Equipment (diff four quarters)

Figure 25. The effect of an oil-specific shock using an estimate of industrial production for the OECD plus other major countries as global activity, see Baumeister and Hamilton (2019): Impulse responses for non-residential investment, income, industrial production and business equipment. The initial shock is normalized to increase oil prices by 10%. All responses are reported in levels (of the normalized data). Left column: posterior median of impulse responses at seven different periods in time. Middle and right columns: the difference between the responses in period 1991:Q1-2018:Q4 and the responses in 1995:Q1 after two and four quarters respectively. The solid line is the difference between the median estimates. The shaded area represents 68% posterior probability bands for the difference in impulse responses.

H.5 Robustness to the choice of oil price variable

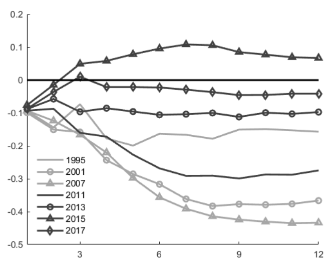

(a) Non-oil Investment (median)

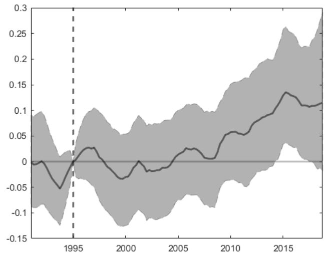

(b) Non-oil Investment (diff two quarters)

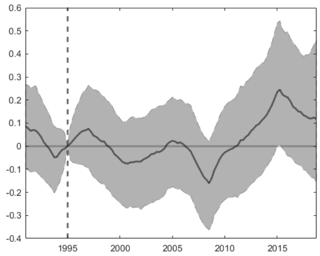

(c) Non-oil Investment (diff four quarters)

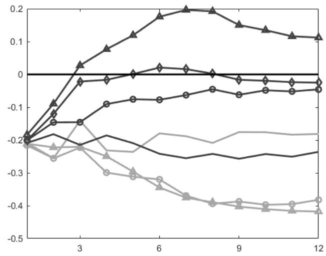

(d) Income (median)

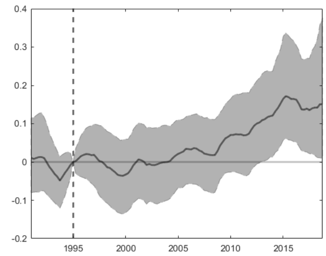

(e) Income (diff two quarters)

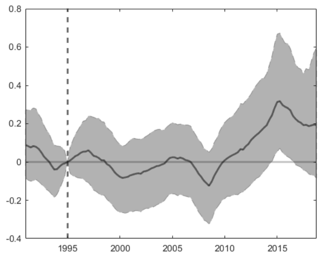

(f) Income (diff four quarters)

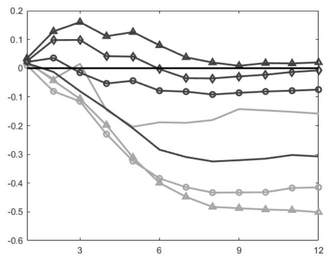

(g) IP Index (median)

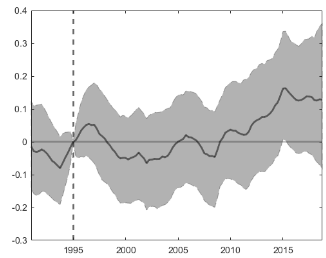

(h) IP Index (diff two quarters)

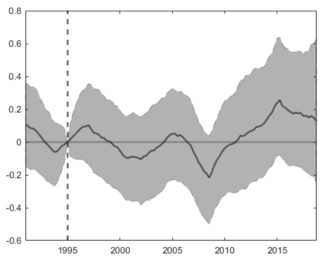

(i) IP Index (diff four quarters)

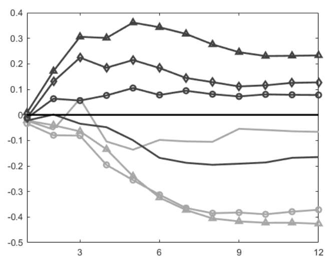

(j) Business Equipment (median)

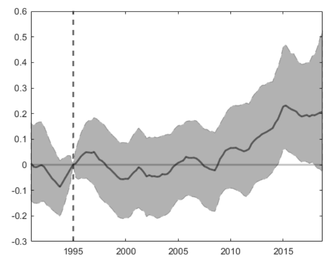

(k) Business Equipment (diff two quarters)

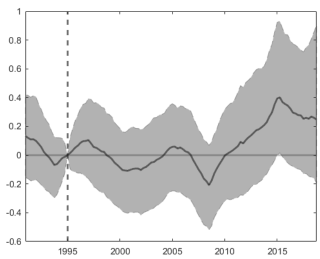

(l) Business Equipment(diff four quarters)

Figure 26. The effect of an oil-specific shock using WTI as the oil price variable: Impulse responses for non-residential investment, income, industrial production and business equipment. The initial shock is normalized to increase oil prices by 10%. All responses are reported in levels (of the normalized data). Left column: posterior median of impulse responses at seven different periods in time. Middle and right columns: the difference between the responses in period 1991:Q1-2018:Q4 and the responses in 1995:Q1 after two and four quarters respectively. The solid line is the difference between the median estimates. The shaded area represents 68% posterior probability bands for the difference in impulse responses.

Centre for Applied Macro - and Petroleum economics (CAMP)
will bring together economists working on applied macroeconomic issues, with special emphasis on petroleum economics.

BI Norwegian Business School
Centre for Applied Macro - Petroleum economics (CAMP)
N-0442 Oslo

http://www.bi.no/camp

CAMP Working Paper Series
ISSN: 1892-2198