

Oil Prices and Consumption across Countries and U.S. States*

Andrea De Michelis, Thiago Ferreira, and Matteo Iacoviello
Federal Reserve Board of Governors

We study the effects of oil prices on consumption across countries and U.S. states, by exploiting the time-series and cross-sectional variation in oil dependency of these economies. We build two large data sets: one with 55 countries over the years 1975–2018, and another with all U.S. states over the period 1989–2018. We then show that oil price declines generate positive effects on consumption in oil-importing economies, while depressing consumption in oil-exporting economies. We also document that oil price increases do more harm than the good afforded by oil price decreases both in the world and in U.S. aggregates.

JEL Codes: Q43, E32, F40.

1. Introduction

The large fluctuations of oil prices in recent years have re-ignited interest regarding the macroeconomic effects of oil price shocks. The conventional wisdom of academics, policymakers, and market practitioners is that declines in oil prices boost global economic activity, as increases in consumption, especially for oil importers, outweigh the negative effects on oil producers. This wisdom has recently been reiterated by Arezki and Blanchard (2014), Bernanke (2016), and even Warren Buffett (CBS/AP 2016), amongst others. However, much of

*The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of anyone else associated with the Federal Reserve System. Erin Markievitz, Lucas Husted, Andrew Kane, Joshua Herman, Patrick Molligo, Brynne Godfrey, Charlotte Singer, and Caitlin Dutta provided outstanding research assistance. We are grateful for the comments by Carola Binder. Author e-mails: andrea.de.michelis@frb.gov; thiago.r.teixeirafferreira@frb.gov; matteo.iacoviello@frb.gov.

the research on the global effects of oil price changes has only looked at a restricted set of countries (see, e.g., Jiménez-Rodríguez and Sánchez 2005), or has focused on a set of macroeconomic variables, such as industrial production or current account, whose responses may not be easily mapped into those of gross domestic product (GDP) and consumption (see, for instance, Aastveit, Bjørnland, and Thorsrud 2015 and Kilian, Rebucci, and Spatafora 2009).

This paper studies the effects of oil prices on consumption and GDP across countries and across U.S. states. A hurdle in obtaining precise estimates of the global effects of oil price shocks is data availability. To address this shortcoming, we put together a quarterly data set containing data on GDP, consumption, and oil dependence for 55 countries. The coverage, which varies across countries, spans more than 40 years from 1975 to 2018. We go to great lengths to ensure that the data are consistent across countries and do not contain breaks. Using this unique database, we estimate the effects of oil price shocks on different countries, allowing for the effects to vary based on each country's oil dependence. We find that oil price declines have large, positive effects on the consumption and GDP of oil-importing countries, while depressing consumption and GDP of oil exporters, and that, in aggregate, the positive effects dominate. However, we also show that the boost to oil importers occurs rather gradually, while the hit to oil exporters is realized fairly quickly, thus suggesting that, in aggregate, the benefits from lower oil prices might occur only slowly. Moreover, we show that oil price decreases generate smaller boosts to the world consumption than the drag from comparable oil price increases.

We conclude by presenting complementary analysis on the heterogeneous effects of oil shocks using state-level data for the United States. To this end, we use state-level data on registrations of new cars—available quarterly from 1989 to 2018—as a proxy for state-level consumption. This data set has important advantages over state-level data on personal consumption expenditures, as the latter are still experimental, are available only annually, and cover a shorter time frame. We then show that the effects of oil price declines on consumption differ across states depending on their oil dependence, despite the fact that these states face common monetary and fiscal policy. The paper proceeds as follows. Section 2 presents the results from an aggregate vector autoregressive (VAR) model which

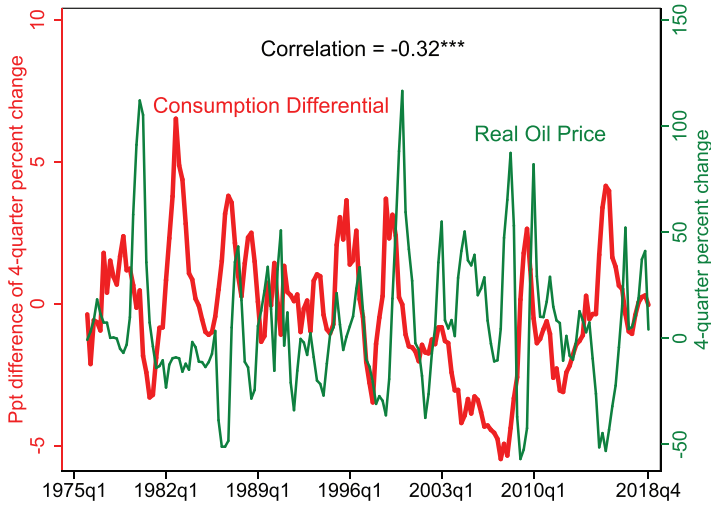
groups countries according to whether they are oil importers or oil exporters. Section 3 uses a cross-country panel VAR to dig deeper into the heterogeneous country responses. Section 4 presents the evidence from a state-level panel VAR for the United States. Section 5 relates the results from this paper to the rest of the existing literature on the macroeconomic effects of oil shocks. Section 6 concludes.

2. Oil Prices, World Consumption, and GDP

Changes in oil prices may reflect disturbances to global aggregate demand as well as disturbances that are specific to the oil market. In most of what follows, we interpret oil price changes as arising from “oil shocks,” controlling for shifts in global economic activity that simultaneously drive oil demand and oil prices. We interpret these oil shocks as reflecting disruptions in oil supply due to geopolitical or natural events, or from precautionary or speculative shifts in the demand for oil.

In order to set the stage for the empirical analysis, it is useful to briefly review the key channels by which an “exogenous” decline in oil prices should affect economic activity in both oil-producing and oil-importing countries. On the supply side, lower oil prices raise output in the non-oil sector by reducing firms’ production costs and causing investment and output to rise. This cost channel should be stronger in countries or sectors that heavily rely on oil as an input in production. Conversely, falling oil prices may depress energy-related investment across oil-producing countries, dragging down aggregate activity. On the demand side, lower oil prices transfer wealth towards oil-importing countries and away from oil exporters (e.g., Bodenstein, Erceg, and Guerrieri 2011), cause a windfall income gain for consumers, and thus shift consumption towards oil-importing countries. This wealth effect may in turn cause GDP to rise through multiplier effects, and may be larger in sectors that produce goods that are complementary to consumption of oil, such as the automobile sector.¹ Most of our analysis focuses on the consumption effects

¹See Hamilton (2008) for a discussion on several channels through which oil shocks may affect macroeconomic variables.

Figure 1. Oil Prices and Consumption Differential

Notes: Consumption differential is the aggregate consumption growth of oil importers minus the aggregate consumption growth of oil exporters (black line with y-axis on the left, or red line in online version). Real oil prices (gray line with y-axis on the right, or green line in online version) are measured in 2009 dollars. While consumption differential is measured in percentage-point difference in four-quarter percent changes, real oil prices are measured in four-quarter percent change. The correlation between consumption differential and real oil prices is statistically significant at the 99 percent confidence level (***). Exporters and importers are aggregated using weights based on each country’s GDP in constant U.S. dollars. The list of countries and their classification of importer/exporter is in table A.1 in appendix A.

of oil price changes, and exploit the differential exposure of oil producers and oil importers to these changes to highlight its empirical relevance.

Figure 1 shows our first approach to evaluating the differential effects of oil price on consumption according to the country’s status as oil importer or exporter. It plots real oil price growth against the “consumption differential,” the difference in consumption growth between importing and exporting countries. Since 1975, whenever oil price growth has fallen, the consumption differential has tended to rise, that is, consumption has grown relatively faster in importing countries. The correlation between oil price growth and the consumption differential is -0.32 and is significantly different

from zero. It is important to highlight the consumption differential because it offers a simple way to control for global demand-side effects that otherwise create a positive correlation between activity and oil prices. Indeed, separately taken, the correlations of oil prices with importers' and exporters' consumption are both positive, at 0.06 and 0.34, respectively.

2.1 Aggregate VAR: Oil Importers versus Exporters

To further explore the differential responses of oil importers and exporters, we consider a vector autoregressive model aimed at quantifying the effects of shocks to oil prices on global economic activity. The VAR uses two lags and quarterly data from 1975:Q1 to 2018:Q4 on oil prices in real terms (oil_t), global oil production ($oprod_t$), importers' and exporters' private consumption expenditure (ci_t and ce_t , respectively), and importers' and exporters' GDP ($gdpi_t$ and $gdpe_t$, respectively). Consumption and GDP are constructed by aggregating data—using constant-dollar GDP weights—for 45 oil-importing and 10 oil-exporting countries.² The countries in our sample, as well as their status as oil importer or exporter, are listed in table A.1 in appendix A.³ The structural VAR representation takes the form

$$\mathbf{A}\mathbf{x}_t = \mathbf{B}(L)\mathbf{x}_{t-1} + \boldsymbol{\varepsilon}_t, \quad (1)$$

$$\mathbf{x}_t = [gdpi_t, gdpe_t, ci_t, ce_t, oil_t, oprod_t]', \quad (2)$$

$$\mathbb{E}(\boldsymbol{\varepsilon}_t\boldsymbol{\varepsilon}_t') = \mathbf{I}, \quad (3)$$

where $\mathbf{B}(L)$ is a lag polynomial of order 2, and \mathbb{E} is the expectational operator.

We then seek to isolate the effects of oil price fluctuations stemming from oil-market-specific shocks. To do so, we use a Cholesky decomposition of the variance-covariance matrix of the reduced-form VAR residuals whereby oil price shocks affect oil prices and production on impact, and consumption and GDP with a

²The sample includes countries which together contribute to about 80 percent of world GDP.

³We take the log of all variables and then quadratically detrend them. See appendix A for details on data construction.

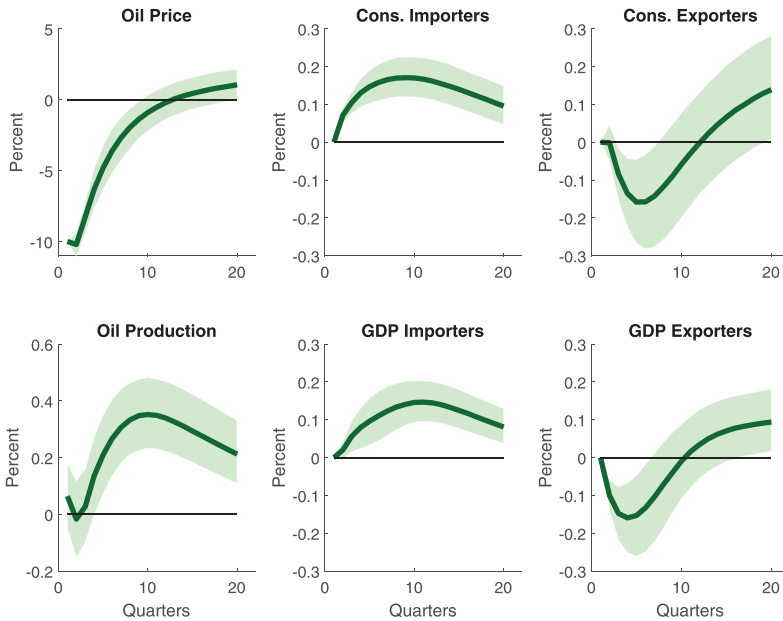
one-quarter delay. That is, the reduced-form errors can be decomposed as follows:

$$\mathbf{e}_t = \mathbf{A}^{-1} \boldsymbol{\varepsilon}_t = \begin{pmatrix} e_t^{gdpi} \\ e_t^{gdpe} \\ e_t^{ci} \\ e_t^{ce} \\ e_t^{oil} \\ e_t^{prod} \end{pmatrix} = \begin{bmatrix} a_{11} & 0 & 0 & 0 & \mathbf{0} & 0 \\ a_{21} & a_{22} & 0 & 0 & \mathbf{0} & 0 \\ a_{31} & a_{32} & a_{33} & 0 & \mathbf{0} & 0 \\ a_{41} & a_{42} & a_{43} & a_{44} & \mathbf{0} & 0 \\ a_{51} & a_{52} & a_{53} & a_{54} & \mathbf{a}_{55} & 0 \\ a_{61} & a_{62} & a_{63} & a_{64} & \mathbf{a}_{65} & a_{66} \end{bmatrix} \begin{pmatrix} \varepsilon_t^1 \\ \varepsilon_t^2 \\ \varepsilon_t^3 \\ \varepsilon_t^4 \\ \varepsilon_t^{oil} \\ \varepsilon_t^6 \end{pmatrix}, \quad (4)$$

where oil price shocks ε_t^{oil} affect the economy according to the fifth column of the \mathbf{A}^{-1} matrix. This assumption is in line with a host of empirical VAR studies that place asset prices after the macroeconomic variables in a Cholesky ordering of the residuals of a VAR, on grounds that asset prices are highly sensitive to contemporaneous economic news or shocks (see, for instance, Bernanke, Boivin, and Elias 2005 and Stock and Watson 2005).

Figure 2 shows the impulse response functions (IRFs) to an identified oil price shock that is normalized to imply a 10 percent decline in real oil prices upon the shock impact. This fall in oil prices boosts importing countries' GDP, with a peak effect reached 10 quarters after the shock. Importers' consumption also rises slowly, with a magnitude slightly larger than GDP. The larger increase in consumption relative to GDP suggests that wealth effects may be an important channel by which oil price declines are transmitted to activity in oil-importing countries. By contrast, exporters' consumption drops rapidly and exporters' GDP declines even more so, bottoming out four and three quarters after the oil price shock, respectively. The larger decline in GDP relative to consumption in oil-exporting countries is consistent with a rapid deterioration in investment in the energy sector. Finally, oil production moves little on impact and

Figure 2. Impulse Response Functions to Oil Price Shock in Aggregate VAR



Notes: The figure shows the IRFs to oil price shocks in the aggregate VAR (section 2.1). Solid lines represent median IRFs. Shaded areas represent 16th and 84th percentiles constructed with 5,000 bootstrap replications. Variables are expressed as percent deviation from quadratic trend.


rises persistently in the subsequent quarters, peaking at 0.3 percent above baseline after about two years. Such response supports the interpretation of the oil price shock as being caused by news about higher future supply.

All told, the evidence from the aggregate VAR is consistent with the role of wealth effects on consumption for oil importers, and suggestive of the importance of direct effects on energy-related investment for oil exporters.

2.2 Aggregate VAR under Different Identification Strategies

Our benchmark VAR (equations (1)–(4)) orders oil prices before oil production in the Cholesky decomposition of the variance-covariance matrix of the residuals, with this ordering often being referred to

Figure 3. Comparison of Oil Shocks Identified under Different Strategies



Notes: “Oil Price Shocks (Aggregate VAR)” represent the oil shocks identified by the aggregate VAR from section 2.1 of this paper. “KOD Shocks” represent the “oil-specific demand” shocks identified by Kilian (2009) averaged to the quarterly frequency. The correlation between the two series of shocks is statistically significant at the 99 percent confidence level (***)

as “exogenous oil price assumption” (e.g., Stock and Watson 2016). This assumption regards unexpected changes in oil prices, once world factors are controlled for, as reflecting international developments specific to oil markets, such as exogenous changes in oil supply conditions. A different approach is to order oil production before oil prices, as in Kilian (2009). With this different ordering, it is possible to distinguish oil supply shocks—the residuals of the oil production equation in the VAR—from oil-specific demand shocks—the residuals of the oil price equation in the VAR. In this second ordering, oil-specific demand shocks are a key driver of oil prices and are interpreted as shifts in demand due to concerns about future availability of oil supply.

Figure 3 shows that the two ordering assumptions identify very similar shocks. In particular, the correlation between our identified oil price shocks and Kilian’s oil-specific demand shocks is 0.41, and

is statistically different from zero. An important reason for the similarity between our estimated shocks and those in Kilian (2009) is that there is little contemporaneous correlation in the data between oil prices and oil production. Therefore, once global conditions are controlled for, it does not seem to matter whether one orders prices before or after production.⁴

Our VAR ordering does not impose any restriction on the contemporaneous co-movement of oil production and oil prices in response to an oil price shock. If one interprets the identified shock as stemming from oil supply, the ratio between the IRF at impact of oil production and the IRF at impact of oil prices measures the implied short-run elasticity of oil demand to oil prices. Given the IRFs from our VAR shown in figure 2, production rises by 0.06 percent when prices drop by 10 percent, thus implying a short-run oil demand elasticity of 0.006, which is in the low end of the estimates of Caldara, Cavallo, and Iacoviello (2019) (henceforth abbreviated to CCI).

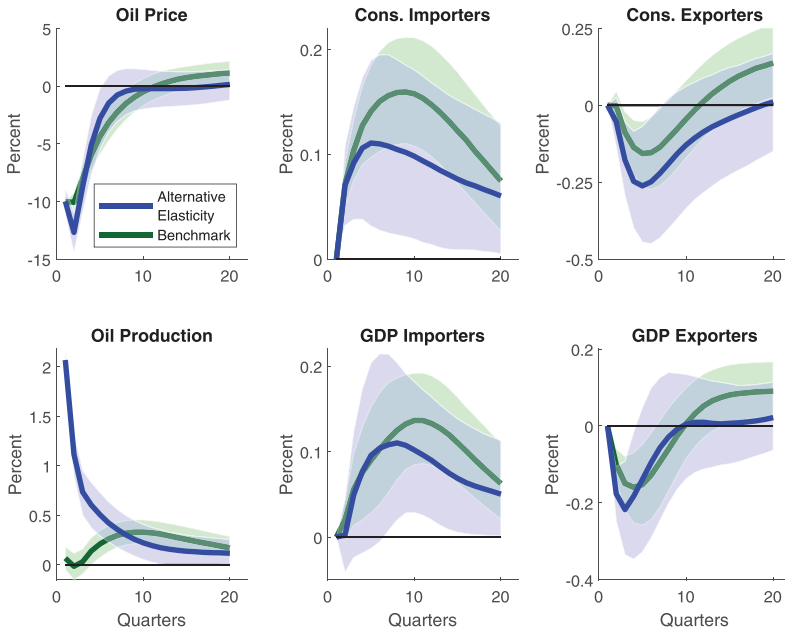
We then examine the robustness of our benchmark results by adopting an alternative identification strategy. We retain the assumption that oil shocks affect GDP and consumption with a one-period delay, but we decompose the variance-covariance matrix of the VAR residuals assuming joint oil demand and oil supply elasticities (the 2×2 block in the bottom right of matrix \mathbf{A}^{-1}) in line with existing studies, such as CCI. In particular, we assume an impact oil supply elasticity of 0.1, which in turn implies a VAR-consistent oil demand elasticity of -0.21 .⁵ As shown in figure 4, when this newly identified oil “supply” shock is scaled so that its effect on oil prices is the same as that of our benchmark specification, the response of oil production is now larger. Nevertheless, the responses of GDP and consumption of oil importers and exporters are similar to those estimated using the benchmark aggregate VAR.

Since the identification strategies from Kilian (2009) and CCI yield results similar to ours, we interpret our results as relatively robust and we maintain a similar order of variables in the VARs

⁴For a detailed discussion, see the appendix in Caldara, Cavallo, and Iacoviello (2019).

⁵Notice that the Kilian (2009) ordering of variables implies a short-run oil supply elasticity of zero, which is not consistent with studies such as CCI.

Figure 4. Impulse Response Functions to Oil Shock under Alternative Oil Supply Elasticity



Notes: “Benchmark” IRFs are calculated using the aggregate VAR of section 2.1 and are shown in green (see online version for colors). “Alternative Elasticity” IRFs, shown in blue, are calculated under the assumption of a world oil demand elasticity of -0.21 and world supply elasticity of 0.1 , as explained in section 2.2. Solid lines represent median IRFs. Shaded areas represent 16th and 84th percentiles constructed using 5,000 bootstrap replications. Variables are expressed as percent deviation from quadratic trend.

of the next sections. Importantly, we focus on identifying oil price movements that are driven by shocks that are specific to the oil market, and thus we do not try to separate out fluctuations driven by demand or supply shocks.

3. Oil Dependency across Countries: Panel VAR

In the previous section, we assumed that the effects of an oil price shock only depend on whether a country is an oil importer or exporter, disregarding any cross-sectional and time-series variations

in countries' oil dependency. In this section, we account for such variation and show that GDP and consumption of countries with higher oil dependency increase more after exogenous oil price declines. We also show that oil price increases do more harm than the good afforded by oil price decreases.

3.1 Oil Dependency across Countries

Figure 5 documents important differences in oil dependence, both across countries and within country over time. Such differences are obviously lost when we use the simple distinction between oil importers and oil exporters. For a country i at quarter t , we measure oil dependence as the ratio of net oil imports to total oil consumption expressed in percent:

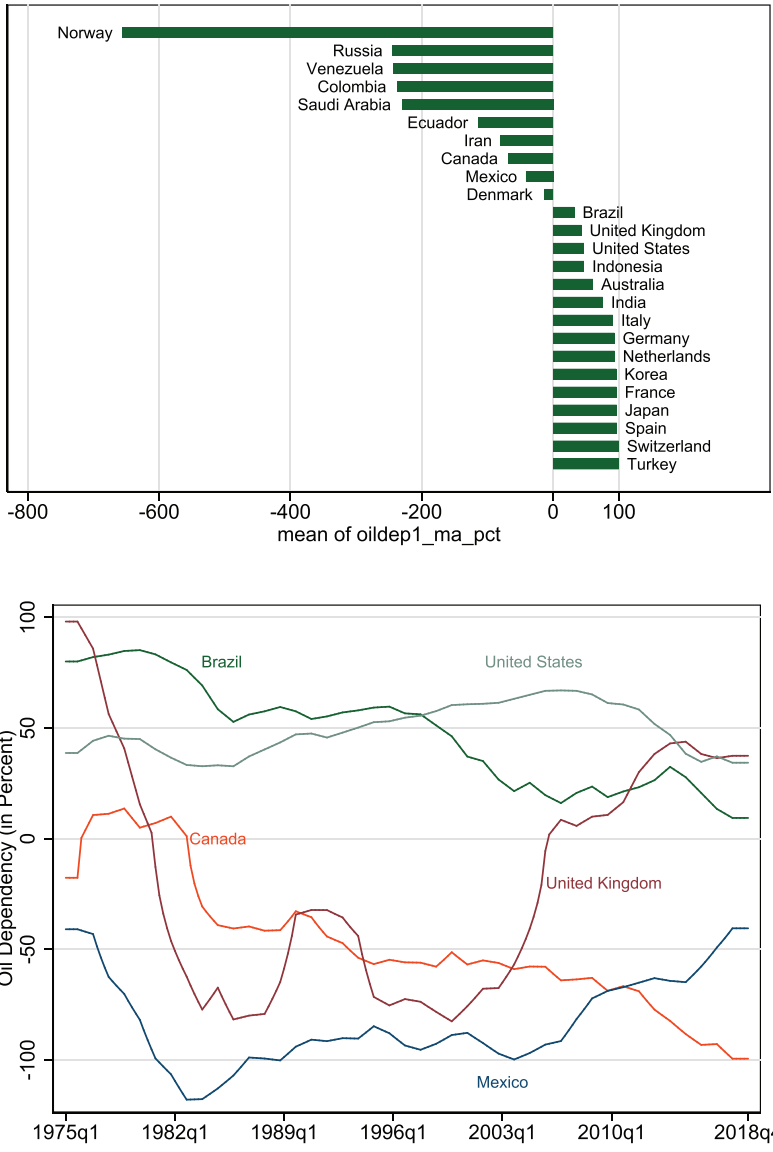
$$d_{i,t} \equiv 100 \times \frac{ocons_{i,t} - oprod_{i,t}}{ocons_{i,t}}, \quad (5)$$

where $ocons_{i,t}$ is the amount of oil consumed and $oprod_{i,t}$ is the amount of oil produced.

The top panel of figure 5 shows considerable differences in oil dependence across countries in, say, 2013 (similar differences show up when looking at other points in time). Among oil importers, Japan and the large majority of countries in the euro area had an oil dependence of nearly 100 percent, while the United States had an oil dependence of about 50 percent. Among oil exporters, there was even greater heterogeneity according to this metric, ranging from Canada's (net) oil exports of about 70 percent of its oil consumption to Norway's 650 percent.


The bottom panel of figure 5 presents evidence on the variation of oil dependency over time for a select group of countries. For instance, Canada was a net importer in the late 1970s and early 1980s, but now exports a large share of its oil production. Additionally, oil dependency in the United States has fallen sharply since 2010, following the massive growth in shale oil production. By contrast, the United Kingdom was a large exporter throughout the 1980s and the 1990s, but is now a net importer. Table A.1 in appendix A provides more details on oil dependency across countries.

Figure 5. Oil Dependency across Countries and over Time



Notes: The top panel shows oil dependency in 2013 for a subset of the largest countries in the sample (largest exporters and countries with share of world GDP in 2013 larger than 1 percent), where oil dependency is measured by oil imports as a percentage of oil consumption (equation (5)). The bottom panel plots oil dependency over time for a select group of economies.

Figure 6. Oil Dependency and Long-Run Elasticity of Consumption to Oil Prices



Notes: On the horizontal axis, oil dependency is calculated as the sample average, where oil dependency is measured by oil imports as a percentage of oil consumption (equation (5)). The vertical axis features long-run elasticities of consumption to changes in oil prices. These elasticities are estimated either by country-by-country regressions (6) with fitted β_i 's represented by circles, or by panel regression (7) with fitted values represented by the solid line (red in online version).

Our strategy is to exploit this heterogeneity to more precisely estimate how economies respond to fluctuation in oil prices. Indeed, figure 6 illustrates why a country's oil dependence may go a long way in shaping the link between changes in oil prices and changes in consumption expenditures.

First, we run country-specific ordinary least squares regressions of consumption ($c_{i,t}$) on its lag, lagged real oil prices (oil_{t-1}), and world GDP (Y_t):

$$c_{i,t} = \rho_i c_{i,t-1} + \beta_i (1 - \rho_i) oil_{t-1} + \eta_i Y_t + u_{i,t}, \quad (6)$$

where the country-specific coefficient β_i can be interpreted as the long-run elasticity of consumption to changes in oil prices.⁶ If all countries were to respond similarly to changes in oil prices, one should not find a relationship between elasticities β_i 's and countries' oil dependence. Yet, as shown in figure 6, the elasticity of consumption to oil prices (β_i 's in circles) is negatively correlated with oil dependency across countries.

Then, we propose a parsimonious approach to capture the heterogeneous responses of consumption to oil prices seen above. We estimate a panel regression in which the long-run elasticity of consumption to changes in oil prices, $\tilde{\beta}_{i,t}$, depends on the country's oil dependency, $d_{i,t}$, after a flexible transformation $g(d_{i,t})$:

$$c_{i,t} = \rho c_{i,t-1} + \tilde{\beta}_{i,t} (1 - \rho) oil_{t-1} + \eta Y_t + u_{i,t}, \quad (7)$$

$$\tilde{\beta}_{i,t} = b_0 + b_1 g(d_{i,t}; b_2, b_3). \quad (8)$$

Specifically, we assume that function $g(d_{i,t})$ follows a strictly increasing Gompertz transformation:

$$g(d_{i,t}; b_2, b_3) = \exp \{ - \exp [- b_2 (d_{i,t} - b_3)] \}, \quad b_2 > 0, \quad b_3 \in \mathbb{R}. \quad (9)$$

By using this transformation, we can capture a wide range of relationships between an economy's oil dependency, $d_{i,t}$, and its long-run elasticity of consumption to oil prices, $\tilde{\beta}_{i,t}$: from a linear one to a nonlinear one in which lower oil dependency yields less than proportional increases in the elasticity to oil prices.⁷ Indeed, the solid line (red in online version) of figure 6 shows that as countries become less oil dependent, the elasticity of consumption to oil prices predicted by regression (7) increases by a magnitude less than proportional to what would be predicted by a simple linear relationship.

⁶We take logs of all variables in this regression and detrend them with a cubic trend, with the exact detrending mattering little for the results. Also, the actual regression has two lags on consumption.

⁷We estimate the model (7)–(9) by nonlinear least squares. Consistent with figure 6, we also impose that $b_0 > 0$ and $b_1 < 0$.

3.2 Cross-Country Panel VAR

Motivated by the cross-country heterogeneity documented in the previous section, we estimate a hybrid panel VAR that enriches the benchmark VAR of section 2 by controlling for how the degree of oil dependency affects the macroeconomic response to an oil price shock. The data used are the same as in section 2: quarterly data for the 1975–2018 period, covering 55 countries, and totaling more than 6,700 observations. We specify the hybrid panel VAR as follows:

$$\begin{aligned}
 y_{i,t} &= \gamma_{yy}y_{i,t-1} + \gamma_{yY}Y_{t-1} + \gamma_{yc}c_{i,t-1} + \gamma_{yC}C_{t-1} + \gamma_{yz}z_{t-1} \\
 &\quad + (\gamma_{yo} + \gamma_{ydg}(d_{i,t-1}))oil_{t-1} \\
 &\quad + (\delta_{yo} + \delta_{ydg}(d_{i,t-1}))noil_{t-1} + \varepsilon_{i,t}^y, \tag{10}
 \end{aligned}$$

$$\begin{aligned}
 c_{i,t} &= \alpha_{cy}y_{i,t} + \alpha_{cY}Y_t + \gamma_{cy}y_{i,t-1} + \gamma_{cY}Y_{t-1} + \gamma_{cc}c_{i,t-1} + \gamma_{cC}C_{t-1} \\
 &\quad + \gamma_{cz}z_{t-1} + (\gamma_{co} + \gamma_{cdg}(d_{i,t-1}))oil_{t-1} \\
 &\quad + (\delta_{co} + \delta_{cdg}(d_{i,t-1}))noil_{t-1} + \varepsilon_{i,t}^c, \tag{11}
 \end{aligned}$$

$$\begin{aligned}
 z_t &= \alpha_{zY}Y_t + \alpha_{zC}C_t + \gamma_{zY}Y_{t-1} + \gamma_{zC}C_{t-1} + \gamma_{zz}z_{t-1} \\
 &\quad + \gamma_{zo}oil_{t-1} + \varepsilon_t^z, \tag{12}
 \end{aligned}$$

$$\begin{aligned}
 oil_t &= \alpha_{oY}Y_t + \alpha_{oC}C_t + \alpha_{oz}z_t + \gamma_{oY}Y_{t-1} + \gamma_{oC}C_{t-1} + \gamma_{oz}z_{t-1} \\
 &\quad + \gamma_{oo}oil_{t-1} + \varepsilon_t^{oil}, \tag{13}
 \end{aligned}$$

$$noil_t \equiv \max(0, oil_t - \max(oil_{t-1}, oil_{t-2}, oil_{t-3}, oil_{t-4})), \tag{14}$$

where z_t is the Conference Board's composite index of leading indicators (CLI) of business cycles conditions,⁸ oil_t is the real price of oil, $c_{i,t}$ is country-specific consumption, $y_{i,t}$ is country-specific GDP, and $C_t \equiv \sum_{i=1}^N \omega_{i,t} \cdot c_{i,t}$ and $Y_t \equiv \sum_{i=1}^N \omega_{i,t} \cdot y_{i,t}$ are the corresponding world aggregates, with $\omega_{i,t}$ denoting country weights.⁹ All variables are measured in percent deviation from a quadratic trend.

⁸See Camacho and Perez-Quiros (2002) for the CLI's performance in anticipating business cycles conditions. The VAR results are similar if we use instead the Conference Board's composite index of coincident indicators or global oil production.

⁹The estimated VAR has an additional lag for every lagged variable and intercepts. We omit these variables in equations (10)–(13) to save on notation.

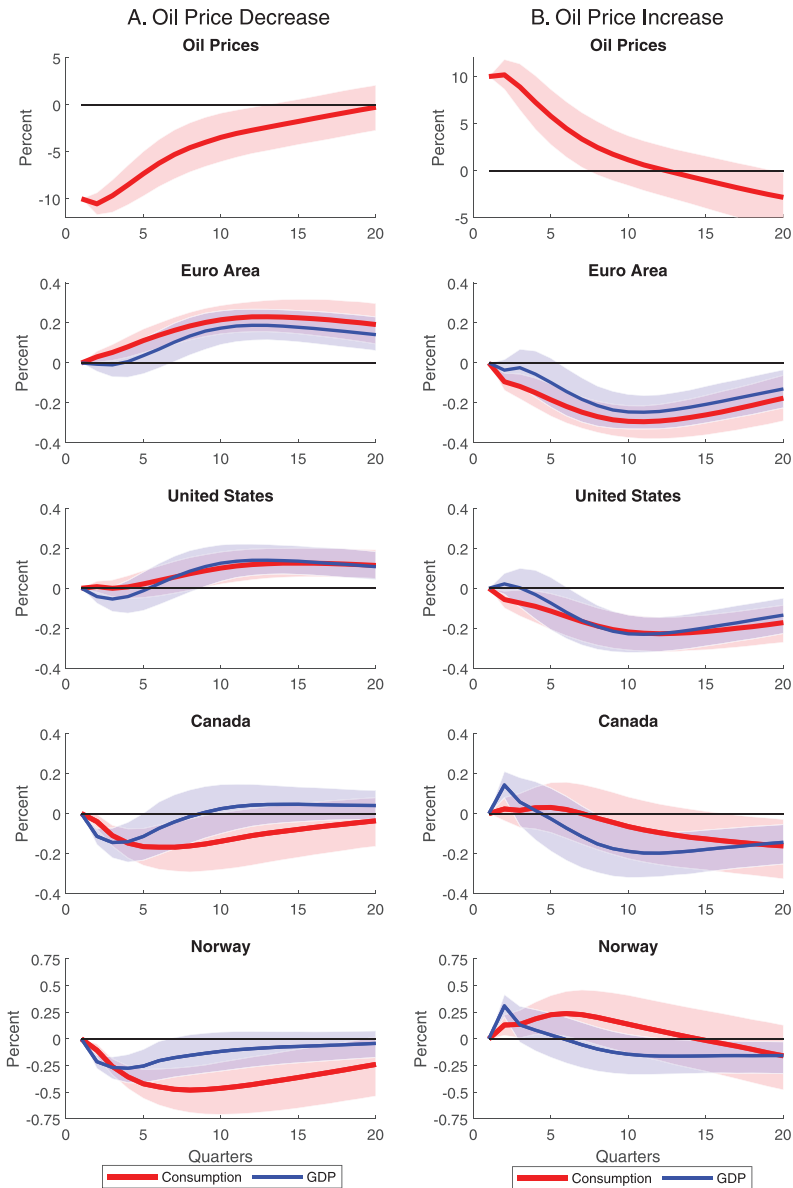
Note that this specification assumes that shocks to oil prices affect consumption, GDP, and commodity prices with a one-period delay.¹⁰

The specification above departs from a standard panel VAR (Holtz-Eakin, Newey, and Rosen 1988) in three ways. First, we allow countries' responses of consumption and GDP to oil price shocks to be a function of their time-varying oil dependence $d_{i,t}$. The specific functional form is the transformation $g(d_{i,t}; b_2, b_3)$ estimated in the panel model (7)–(9). Second, we allow individual countries' consumption and GDP to depend not only on their own lags but also on lags of the world aggregates, thus modeling dynamic interdependencies across countries. The world aggregates are averages of the country-specific variables, as done for instance in the global VAR approach proposed by Pesaran (2006) and discussed in Canova and Ciccarelli (2013) and Chudik and Pesaran (2016). Third, we augment the panel VAR with the net real oil price increase variable ($noil_t$), following Kilian and Vigfusson (2011) and similar to Hamilton (1996), thus allowing for increases and decreases in oil prices to yield asymmetric effects on economic activity, following the lead of Davis and Haltiwanger (2001) and Hamilton (2003).

Figure 7 plots the IRFs of consumption and GDP to oil price shocks using the 2018 values of countries' oil dependencies. The choice of countries showcases the heterogeneous responses arising from these differing oil dependencies. When oil prices decline (figure 7A), consumption and GDP of oil-importer countries, such as those in the euro area, increase. However, these increases take time to materialize, reaching peaks after roughly 12 quarters. As in section 2, consumption responds before GDP, suggesting that wealth effects are larger than supply-side effects. Conversely, consumption and GDP of oil-exporter countries, such as Canada and Norway, quickly decrease, bottoming out six and two quarters after the shock, respectively. These results are consistent with a rapid tumble in investment, presumably in the energy sector, which then compounds with a protracted fall in consumption, possibly due to

¹⁰To see this, notice that only the lags of oil prices (oil_{t-1}) and net oil price increases ($noil_{t-1}$) are present in equations (10) and (11). However, differently from a standard VAR, the cross-equations restrictions embedded in the panel VAR ($C_t \equiv \sum_{i=1}^N \omega_{i,t} \cdot c_{i,t}$ and $Y_t \equiv \sum_{i=1}^N \omega_{i,t} \cdot y_{i,t}$) produce errors that are generally correlated across equations. See Canova and Ciccarelli (2013) for discussion.

Figure 7. Impulse Response Functions to Oil Price Shocks to Cross-Country Panel VAR



Notes: The figure shows the IRFs to oil price shocks in the cross-country panel VAR (section 3.2). Solid lines represent median IRFs. Shaded areas represent 16th and 84th percentiles constructed with 5,000 bootstrap replications. Variables are expressed as percent deviation from quadratic trend.

Table 1. Responses of Consumption and GDP in Cross-Country Panel VAR

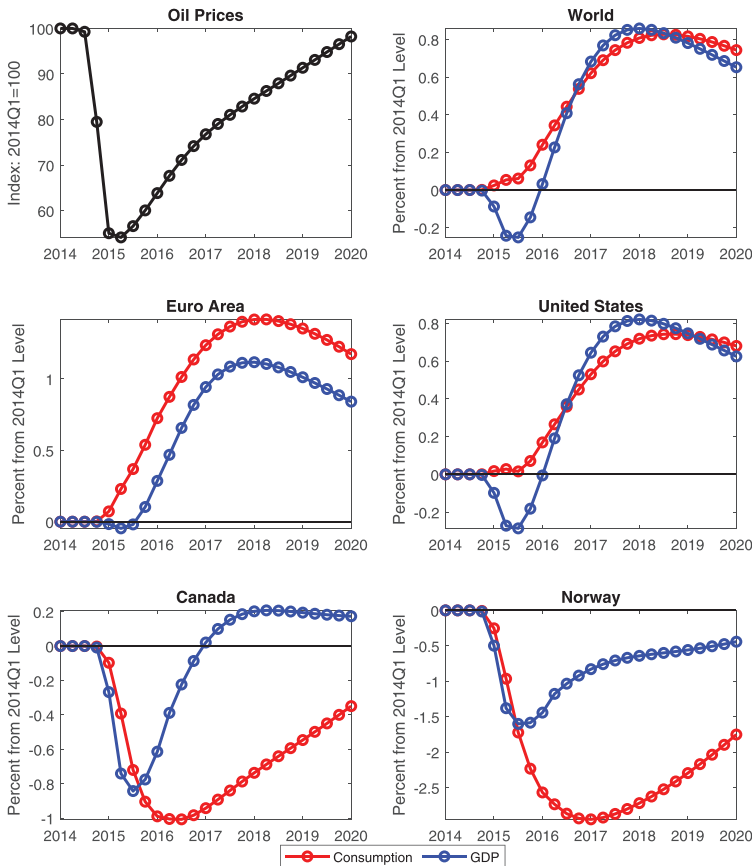
	Euro Area	United States	Canada	Norway	World
10% Oil Price Decline					
Consumption	0.16	0.06	-0.13	-0.40	0.07
GDP	0.10	0.06	-0.04	-0.17	0.06
10% Oil Price Increase					
Consumption	-0.23	-0.16	-0.02	0.16	-0.17
GDP	-0.16	-0.14	-0.09	-0.03	-0.14

Notes: The consumption and GDP responses are measured as the average of their impulse response functions during the first three years (from quarter 1 through quarter 12) after a shock moving oil prices by 10 percent (in quarter 0). These IRFs are calculated using the cross-country panel VAR from section 3.2.

negative wealth effects. The United States, neither a large net oil importer nor a large net oil exporter in 2018, experiences more mixed effects: its GDP leans towards a contraction in the periods immediately after the shock, and only rises above trend many quarters later, pulled by a modest rise in consumption.

Figure 7 also shows that oil price increases (figure 7B) do more harm than the good afforded by oil price decreases (figure 7A). In oil-importing economies, such as those in the euro area, the negative responses of GDP and consumption to oil price increases are larger in magnitude than the positive responses to oil price decreases, as in Hamilton (2003). In large oil-exporting economies, such as Norway, GDP and consumption experience short-lived expansions after an oil price increase, with these variables returning to their original level after six quarters. This result contrasts with the protracted drag to economic activity in the aftermath of oil price declines. One possible reason for such a small boost to oil-exporter economies from oil price increases is that these economies could suffer large drags from the downturn experienced by their oil-importer trading partners (which account for a much larger share of the world GDP). Table 1 summarizes these results, showing that the world GDP experiences only small boosts after oil price declines and much larger drags after oil price increases.

Figure 8. Contribution of the 2014–15 Oil Price Slump to GDP and Consumption



Notes: The figure shows the historical contribution of oil price shocks to consumption and GDP of the world aggregate and selected countries. Oil price shocks are those estimated for the periods 2014:Q3 through 2015:Q1 using the cross-country panel VAR (section 3.2). Variables are expressed as percent deviation from the 2014:Q1 level, with the exception of oil prices, which is measured as an index with level 100 at 2014:Q1.

In figure 8, we use the estimates of the cross-country panel VAR to quantify how the “exogenous” portion of the 2014–15 oil price slump affected consumption and GDP across countries. First and foremost, the VAR attributes a large chunk of the oil price

decline throughout that period to surprise innovations to the oil price equation (ε_t^{oil} in equation (13)) in periods 2014:Q3–2015:Q1. Other shocks, as well as delayed effects from other shocks taking place before 2014:Q3, play a much more limited role in driving oil prices in 2014–15. Accordingly, the response of oil prices—shown in the top-left panel in terms of the 2014:Q1 level—closely mirrors the actual path of oil prices.

The remaining panels of figure 8 show the response of consumption and GDP across countries. The top-right panel shows the response for the world. Consumption rises gradually, peaking at about 0.8 percent above baseline in mid-2018. The response of GDP is initially slightly negative, likely reflecting some drag from declines in investment in the oil sector, but also builds to about 0.8 percent by the end of 2017. The middle panel compares the euro area and the United States. While the U.S. response mirrors that of the world as a whole, the response of the euro area to the oil price decline is larger than that of the United States, reflecting the euro zone’s larger oil dependence. In the bottom panel, GDP and consumption fall markedly in Canada and Norway, consistent with the fact that both economies are large oil exporters.

In appendix B, we conduct two robustness exercises. First, we compare the oil price shocks ε_t^{oil} estimated by our cross-country panel VAR (equations (10)–(14)) with the oil-specific demand shocks identified by Kilian (2009). Figure B.1 shows that these shocks are similar, exhibiting a correlation of 0.49, which is higher than the one found by the analogous comparison done in section 2.2. Second, we estimate our panel VAR starting in 1986, given a possible structural break in oil dynamics around this period (e.g., Baumeister and Peersman 2013). The impulse responses for the restricted sample, shown in figure B.2, are similar to those from our baseline model.

4. Oil Dependency across U.S. States: State-Level Panel VAR

In this section, we analyze how economic activity reacts to oil price shocks across U.S. states by exploiting the cross-section and time-series variation in states’ oil dependence. We show that consumption (proxied by car registrations) of states with higher oil dependency

increases more after exogenous oil price declines. We also show that changes in oil prices generate asymmetric effects: oil price decreases generate smaller boosts to aggregate U.S. consumption than the drag from oil price increases.

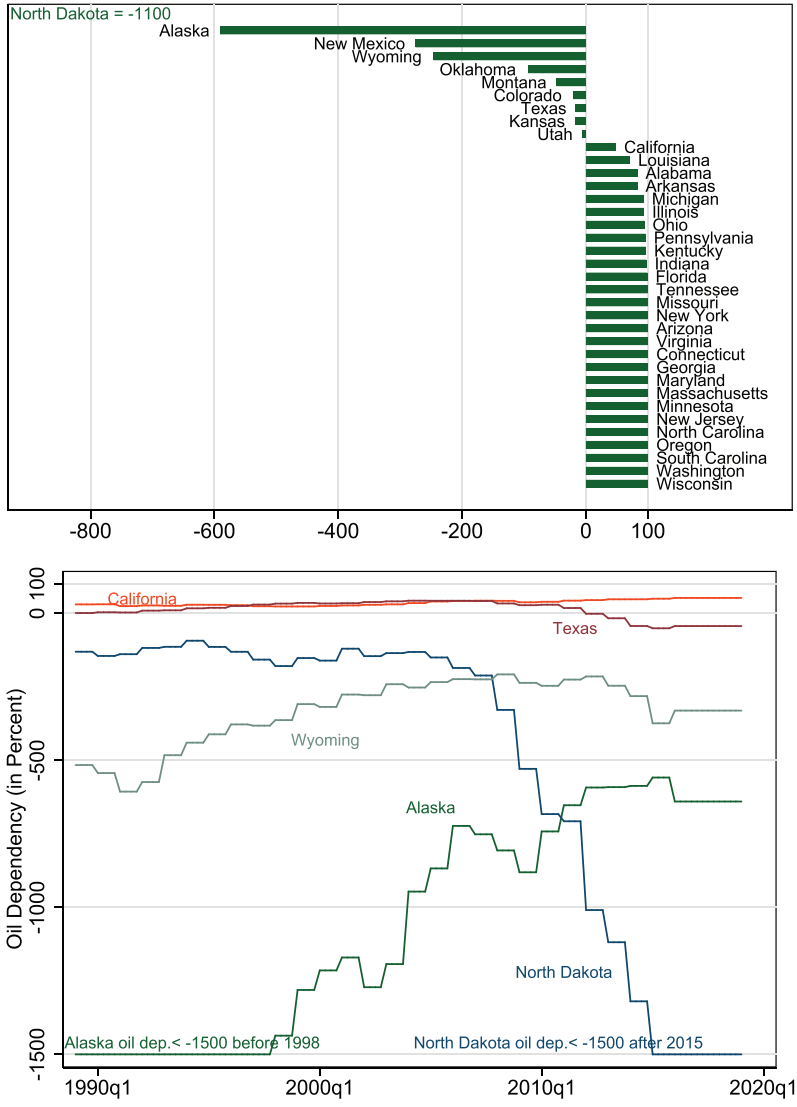
4.1 Oil Dependency across the United States

Just like countries in the world, U.S. states are very heterogeneous in their oil dependence.¹¹ Figure 9 provides information about cross-sectional and time-series variation in U.S. states' oil dependency, using the measure defined in equation (5). The top panel of figure 9 shows that in 2013 states in the Northeast had an oil dependence of nearly 100 percent, while Texas was roughly oil independent. By contrast, states such as Alaska and North Dakota produced substantially more oil than consumed, with very large negative oil dependencies. The bottom panel of figure 9 shows that U.S. states also exhibit considerable time-series variation in oil dependence, with North Dakota exporting increasing amounts of oil over time and Alaska exporting decreasing amounts. Table A.2 in appendix A provides more details on oil dependency across U.S. states.

In order to estimate the response of consumption to oil price shocks at the state level, we use data on retail new car registrations as a proxy for consumption. Available at a quarterly frequency from 1989:Q1 to 2018:Q4, this data set constructed by Polk's National Vehicle Population Profile has important advantages over state-level data on personal consumption expenditures, as the latter are still experimental, are available only annually, and cover a shorter time frame. Appendix A describes the state-level car registrations data and shows that retail registrations at the national level are highly correlated with NIPA (national income and product accounts) expenditures on motor vehicles and parts.

¹¹In fact, using the cross-sectional standard deviation of oil dependency as a measure of heterogeneity, the U.S. states are actually more heterogeneous than the world's countries in terms of oil dependency. In 2013, the cross-sectional standard deviation of oil dependency across U.S. states was 200. The corresponding number for the world's countries was 130.

Figure 9. Oil Dependency across U.S. States and over Time



Notes: The top panel shows oil dependency in 2013 for a subset of the largest states in the United States (oil producers and states with share of car purchases in 2013 larger than 1 percent), where oil dependency is measured by oil imports as a percentage of oil consumption (equation (5)). The bottom panel plots oil dependency over time for a select group of U.S. states.

4.2 State-Level Panel VAR

We set up a panel VAR that mirrors the cross-country analysis discussed in the previous section. To do so, we use unemployment rates to measure overall economic activity at the state level. Thus, for each state i , let $u_{i,t}$ denote the unemployment rate in quarter t and $car_{i,t}$ the car registrations, with $U_t \equiv \sum_{i=1}^N \omega_{i,t} u_{i,t}$ referring to the national unemployment, $Car_t \equiv \sum_{i=1}^N \omega_{i,t} car_{i,t}$ to the national car registrations, and $\omega_{i,t}$ to state-specific weights. The rest of the notation is the same as in section 3: oil_t is the real price of oil, $noil_t$ is the real net oil price increase (equation (14)), and z_t is the CLI. All variables are measured in percent deviation from a quadratic trend, with the exception of unemployment rates, which are measured as percentage-point deviation from a quadratic trend. With a specification analogous to the one of section 3.2, the state-level panel VAR allows for heterogeneous regional responses to an oil shock as a function of oil dependence:

$$\begin{aligned} u_{i,t} = & \gamma_{uu}u_{i,t-1} + \gamma_{uU}U_{t-1} + \gamma_{uc}car_{i,t-1} + \gamma_{uC}Car_{t-1} \\ & + \gamma_{uz}z_{t-1} + (\gamma_{uo} + \gamma_{udg}(d_{i,t-1}))oil_{t-1} \\ & + (\delta_{uo} + \delta_{udg}(d_{i,t-1}))noil_{t-1} + \varepsilon_{i,t}^u, \end{aligned} \quad (15)$$

$$\begin{aligned} car_{i,t} = & \alpha_{cu}u_{i,t} + \alpha_{cU}U_t + \gamma_{cu}u_{i,t-1} + \gamma_{cU}U_{t-1} + \gamma_{cc}car_{i,t-1} \\ & + \gamma_{cC}Car_{t-1} + \gamma_{cz}z_{t-1} + (\gamma_{co} + \gamma_{cdg}(d_{i,t-1}))oil_{t-1} \\ & + (\delta_{co} + \delta_{cdg}(d_{i,t-1}))noil_{t-1} + \varepsilon_{i,t}^c, \end{aligned} \quad (16)$$

$$\begin{aligned} z_t = & \alpha_{zU}U_t + \alpha_{zC}Car_t + \gamma_{zU}U_{t-1} + \gamma_{zC}Car_{t-1} + \gamma_{zz}z_{t-1} \\ & + \gamma_{zo}oil_{t-1} + \varepsilon_t^z, \end{aligned} \quad (17)$$

$$\begin{aligned} oil_t = & \alpha_{oU}U_t + \alpha_{oC}Car_t + \alpha_{oz}z_t + \gamma_{oU}U_{t-1} + \gamma_{oC}Car_{t-1} \\ & + \gamma_{oz}z_{t-1} + \gamma_{oo}oil_{t-1} + \varepsilon_t^{oil}, \end{aligned} \quad (18)$$

where the functional form of $g(d_{i,t})$ is described in equation (9) with coefficients estimated by a regression analogous to (7).¹²

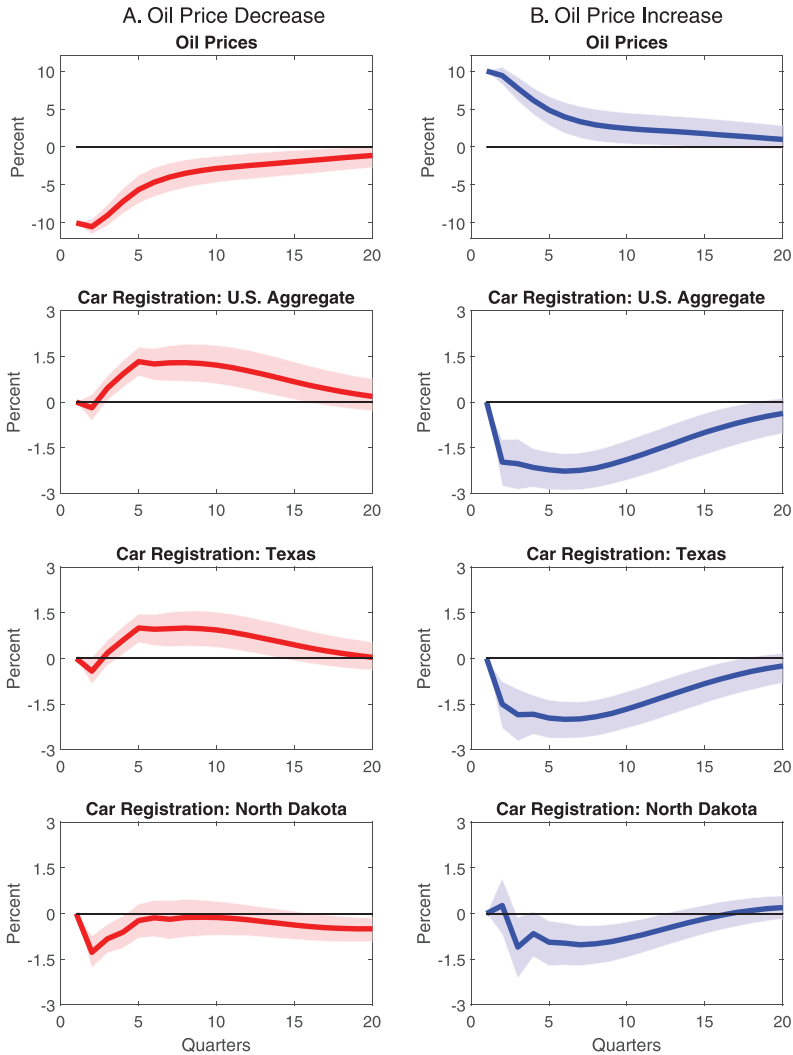
¹²The estimated VAR has an additional lag for every lagged variable and intercepts. We omit these variables in equations (15)–(18) to save on notation.

Figure 10A shows the predicted effect on car registrations of a 10 percent oil price decline triggered by an ε_t^{oil} shock. For the United States in aggregate, such shock is associated with a substantial increase in car registrations, peaking at 1.3 percent a year after the shock, sustaining this increase for another year, and then experiencing fading effects. The large response of car registrations is consistent with the channels that emphasize the large oil price elasticity of demand for goods that are complementary with oil, and lines up well with the findings in the literature. For example, Edelstein and Kilian (2009) find that vehicle expenditures are highly sensitive to movements in the price of gasoline. Coglianesi et al. (2017) find an estimate of -0.37 for the price elasticity of gasoline demand, although their numbers do not appear directly comparable to ours, as they look at the cumulative effects of gasoline price changes.

We then turn to the state-specific responses of car registrations to oil price drops (figure 10A). Oil-importing states face effects similar to those seen for the United States in the aggregate, and thus are omitted from the figure. Oil-independent Texas has a response similar to the U.S. aggregate, but with a magnitude slightly smaller. By contrast, major oil-exporting states such as North Dakota suffer a substantial decline in car registrations, with a trough of -1.3 percent in the quarter right after the shock. North Dakota's car registrations soon rise, but only to a level not statistically different from zero. This rise may reflect the fact that, because cars run on gas, lower oil prices make car-buying more attractive, even in oil-intensive states. Moreover, the subsequent return of car registrations to North Dakota's original level may also reflect the increase in overall U.S. aggregate demand, with the close trade and financial linkages within the U.S. states dampening the effects of heterogeneity in oil production.

Figure 10B then shows that the benefits from oil price decreases are smaller than the losses from oil price increases also across U.S. states. We see that shocks pushing up oil prices are followed by a national decrease in car registrations. This decrease is larger than the increase in registrations following a comparable oil price drop (figure 10A). The results for Texas are qualitatively similar to those for the national aggregate, although with smaller magnitudes. In the case of oil-rich North Dakota, car registrations also fall after an oil price increase, possibly due to the U.S. aggregate demand effect. Figure B.4 in the appendix with supplementary results (appendix B)

Figure 10. Impulse Response Functions to Oil Price Shocks in State-Level Panel VAR



Notes: The figure shows the IRFs to oil price shocks in the U.S. state-level panel VAR (section 4). Solid lines represent median IRFs. Shaded areas represent 16th and 84th percentiles constructed using 5,000 bootstrap replications. Variables are expressed as percent deviation from quadratic trend.

shows that results for unemployment rates are similar to those for car registrations, with the main difference being that unemployment rates react a bit more slowly to oil price shocks than car registrations.

All told, the state-level evidence supports the findings of the country-level VAR. It highlights once again how varying oil dependencies may imply differences in consumption responses across regions, even if these regions face common monetary and fiscal policies. In fact, as is well known in the literature on fiscal and monetary unions, a common policy response may either amplify or dampen the heterogeneity of regional responses, especially when shocks present authorities with stabilization tradeoffs. Consider for instance an oil shock that lowers inflation and activity in North Dakota but lowers inflation and boosts activity in Connecticut. If the policy goal is to stabilize activity, such shock calls for expansionary policies in North Dakota and for contractionary policies in Connecticut. Thus, heterogeneous policy responses would imply responses of activity in North Dakota and Connecticut being more similar than under the case of a common policy response.

5. Literature Review

There is a vast literature that analyzes how fluctuations in oil prices affect macroeconomic outcomes. In particular, our paper belongs to a growing set of studies that investigate the responses to oil price shocks across different economies. These studies have explored variation in geographical location, institutional development, and net-exporting position of energy across different countries and states of the United States (e.g., Aastveit, Bjørnland, and Thorsrud 2015; Bjørnland and Zhulanova 2018; Guerrero-Escobar, Hernandez-del-Valle, and Hernandez-Vega 2019; and Peersman and Van Robays 2012). In addition, Baumeister, Peersman, and Van Robays (2010) (henceforth abbreviated as BPVR) is particularly related to our work because of its cross-country perspective. BPVR estimate country-specific VARs for eight different countries and document that net-energy importing economies suffer declines in economic activity following an oil supply shock, with these effects being not significant or positive for net-energy exporters. BPVR also show qualitative evidence that improvements in the net-energy position

of countries are associated with more muted responses to oil supply shocks.

Our paper makes four contributions. First, we explore a much larger set of economies than previous papers by using two panel data sets: one with 55 countries and another with all U.S. states. Second, we exploit the variation in oil dependency both across countries and over time—beyond the status of oil importer/exporter—thus better quantifying how changes in oil dependency affect an economy's response to oil price shocks. Third, we measure the effect of oil shocks on aggregate consumption, a variable often overlooked in the literature. We show that oil price declines generate large and positive effects on consumption and economic activity (GDP and unemployment rate) in oil-importing economies, while depressing consumption and activity in oil-exporting economies. Fourth, we allow asymmetric responses to oil price shocks, with macroeconomic effects from oil price increases being potentially different from oil price decreases. We then show that oil price increases do more harm than the good afforded by oil price decreases across both the world and the United States.

6. Conclusions

In this paper, we provide empirical evidence that lower oil prices boost consumption in oil-importing countries and depress it in oil-exporting ones. This effect varies in proportion to the degree of oil dependency of an economy. For instance, a 10 percent decline in oil prices boosts euro-area consumption and GDP by about 0.2 percent, while leading to declines of similar magnitudes in consumption and GDP in oil-exporting countries such as Canada. In between, there are countries like the United States where the effects are more mixed. While in the short run U.S. GDP might temporarily decline, presumably reflecting weaker investment in the oil industry, a gradual rise in U.S. consumption pushes its GDP toward a modest expansion. For the world aggregate, oil price declines boost economic activity, although by smaller magnitudes than similar increases in oil prices. We complement our cross-country results by showing analogous findings for the U.S. states and the U.S. aggregate.

Our results have important implications for the design of monetary policy in the presence of oil price fluctuations. A large and growing literature documents how optimal monetary policy may entail different responses to oil price increases depending on the source of the shock, say an increase in foreign demand or a glut in foreign supply. For instance, Bodenstein, Guerrieri, and Kilian (2012) develop and estimate a dynamic stochastic general equilibrium model with oil in which multiple shocks drive oil price fluctuations. In their model, the peak response of output and consumption to oil shocks happens in the same period of the shock. They find that no two shocks induce the same policy response, even controlling for the same path of oil prices implied by the shocks. They further show that, in the wake of fluctuations in oil prices, a monetary policy rule that places a higher weight on stabilizing wage inflation fosters the stabilization of core inflation. Our evidence reinforces their argument, by showing that the same shock may imply different macroeconomic responses depending on the oil dependency of a particular country. It also illustrates an important challenge for extant models of the oil market and the macroeconomy, which may want to incorporate features such as habits and adjustment costs in order to capture the delays in which oil shocks affect economic activity.

Our evidence shows that changes in oil dependence may imply nontrivial effects of changes in oil prices, that the effects of oil shocks on consumption may take time to materialize, and that the response of GDP to an oil shock is not equal to the mere contribution of consumption, presumably because investment in the energy sector may be sensitive to changes in oil prices too. We view such evidence, gathered from a large set of countries over a long time period, as providing a useful empirical benchmark to gauge the plausibility of the policy recommendations of optimizing monetary models of the business cycle that incorporate an oil sector.

Appendix A. The Data

A.1 Country-Level Data

The sample includes 55 countries for which it was possible to assemble sufficiently long and reliable data on quarterly real consumption and quarterly real GDP from the country's official statistics. Of the

richest countries in the world (measured by total GDP in U.S. dollars), we do not include China because of the lack of consumption data at a quarterly frequency. By the same reason, we exclude some of the largest oil producers, such as United Arab Emirates, Iraq, Kuwait, Nigeria, Qatar, Angola, Algeria, and Kazakhstan.

In the aggregate VAR analysis, we construct GDP and consumption using data for real consumption and real GDP using weights based on each country's GDP in constant U.S. dollars from the World Bank World Development Indicators (if the data are missing, the weight is set at zero, and the weights of other countries are changed accordingly). We list in table A.1 all countries in our sample, their oil dependency, their sample weights for the year 2013, and their data coverage.

Oil dependency is constructed using annual data on oil production and consumption from the BP Statistical Review of World Energy (with data dating back to 1965). An advantage of these data is that, barring small changes in year-to-year inventories, they satisfy the identity that oil production must be equal oil consumption over very long horizons. For countries that are small oil producers, the BP Statistical Review does not report data (shown as missing in table A.1). In these cases, we search for oil production and consumption in the CIA World Factbook and the International Energy Agency (IEA). If we find these oil statistics, we set oil dependency to a constant implied by the data for all years. If we are not able to find the data in either the BP Statistical Review, the CIA World Factbook, or the IEA, we set oil dependency to 1 in all years. In all these cases of small oil producers, the implied oil dependency is above 0.65.

A.2 State-Level Data

The key variable included in our state-level panel vector autoregression is “new car registrations, retail” broken down by state of registration. This measure is constructed using data from Polk's National Vehicle Population Profile, which is a census of all currently registered passenger cars and light-duty trucks in the United States. Polk's New Registration Data provides detailed indicators for new vehicle registrations (including where the vehicle was registered). The “retail” file represents new vehicles purchased by individuals,

Table A.1. Cross-Country Oil Dependency in 2013 and Data Availability

Country	GDP Share	Oil Dependency	Oil Consumption	Oil Production	Exporter	Start Date	End Date
United States	27.03	47	18,961	10,071	0	1975:Q1	2018:Q4
Japan	10.05	97	4,516		0	1975:Q1	2018:Q4
Germany	6.10	93	2,408		0	1975:Q1	2018:Q4
France	4.65	97	1,664		0	1975:Q1	2018:Q4
United Kingdom	4.40	43	1,518	864	0	1975:Q1	2018:Q4
Brazil	4.11	32	3,124	2,110	0	1990:Q1	2018:Q4
Italy	3.48	91	1,260	114	0	1975:Q1	2018:Q4
India	3.37	76	3,727	906	0	1996:Q2	2018:Q4
Canada	2.96	-68	2,383	4,000	1	1975:Q1	2018:Q4
Russia	2.89	-245	3,135	10,809	1	1995:Q1	2018:Q4
Spain	2.31	97	1,195		0	1975:Q1	2018:Q4
Australia	2.13	61	1,034	407	0	1975:Q1	2018:Q4
Korea	2.04	96	2,455		0	1975:Q1	2018:Q4
Mexico	1.96	-41	2,034	2,875	1	1980:Q1	2018:Q4
Turkey	1.66	100	757		0	1987:Q1	2018:Q4
Indonesia	1.53	47	1,663	882	0	1983:Q1	2018:Q4
Netherlands	1.45	94	898		0	1975:Q1	2018:Q4
Saudi Arabia	1.07	-230	3,451	11,393	1	2005:Q1	2018:Q4
Switzerland	1.04	100	249		0	1975:Q1	2018:Q4
Poland	0.88	93	538		0	1995:Q1	2018:Q4
Sweden	0.86	96	309		0	1975:Q1	2018:Q4
Belgium	0.84	98	636		0	1975:Q1	2018:Q4

(continued)

Table A.1. (Continued)

Country	GDP Share	Oil Dependency	Oil Consumption	Oil Production	Exporter	Start Date	End Date
Taiwan	0.83	100	1,010		0	1975:Q1	2018:Q4
Iran	0.79	-80	2,011	3,617	1	2004:Q2	2018:Q2
Argentina	0.78	5	683	647	0	1993:Q1	2018:Q4
Norway	0.77	-656	243	1,838	1	1975:Q1	2018:Q4
Venezuela	0.75	-243	782	2,680	1	1997:Q1	2015:Q4
Austria	0.69	87	264		0	1975:Q1	2018:Q4
South Africa	0.69	73	572		0	1975:Q1	2018:Q4
Thailand	0.65	65	1,299	452	0	1993:Q1	2018:Q4
Colombia	0.57	-237	298	1,004	1	2000:Q1	2018:Q4
Denmark	0.56	-13	158	178	1	1975:Q1	2018:Q4
Malaysia	0.51	22	803	626	0	1991:Q1	2018:Q4
Singapore	0.47	98	1,225		0	1975:Q1	2018:Q4
Israel	0.44	98	247		0	1995:Q1	2018:Q4
Chile	0.43	95	362		0	1996:Q1	2018:Q4
Hong Kong	0.43	100	311		0	1990:Q1	2018:Q4
Finland	0.42	95	191		0	1975:Q1	2018:Q4
Greece	0.42	98	295		0	1975:Q1	2018:Q4
Philippines	0.40	94	326		0	1981:Q1	2018:Q4
Ireland	0.40	99	138		0	1975:Q1	2018:Q4
Portugal	0.38	97	239		0	1975:Q1	2018:Q4
Czech Republic	0.36	94	184		0	1996:Q1	2018:Q4
Peru	0.30	26	227	167	0	1980:Q1	2018:Q4
New Zealand	0.27	66	151		0	1975:Q1	2018:Q4
Hungary	0.23	82	129		0	1995:Q1	2018:Q4

(continued)

Table A.1. (Continued)

Country	GDP Share	Oil Dependency	Oil Consumption	Oil Production	Exporter	Start Date	End Date
Slovakia	0.16	89	75		0	1995:Q1	2018:Q4
Ecuador	0.14	-113	247	527	1	1990:Q1	2018:Q4
Luxembourg	0.10	100	57		0	1975:Q1	2018:Q4
Slovenia	0.08	99	50		0	1995:Q1	2018:Q4
Latvia	0.05	97	33		0	1995:Q1	2018:Q4
Estonia	0.04	65	31		0	1995:Q1	2018:Q4
El Salvador	0.03	100	25		0	1990:Q1	2018:Q4
Botswana	0.03	100	17		0	1994:Q1	2018:Q4
Iceland	0.03	100	15		0	1997:Q1	2018:Q4

Notes: Oil dependency is calculated as oil consumed minus oil produced as a percentage of oil consumed (source: BP Statistical Review of World Energy, downloaded in June 2019). Oil production is production of crude oil, tight oil, oil sands, and natural gas liquids (NGLs) (thousand barrels daily). Oil consumption is inland demand plus international aviation and marine bunkers and refinery fuel and loss (thousand barrels daily). In the aggregate VAR, the “exporters” are countries with negative oil dependency in 2013.

Table A.2. U.S. State-Level Data in 2013

State	Car Share	Oil Dependency	Oil Consumption	Oil Production
California	11.45	48	1,713	891
Texas	8.55	-16	3,564	4,153
Florida	7.14	99	824	10
New York	5.77	100	658	2
Pennsylvania	4.45	96	618	24
Illinois	4.26	93	627	43
Michigan	4.21	93	449	32
Ohio	4.17	94	589	36
New Jersey	3.55	100	516	0
Georgia	2.94	100	471	0
North Carolina	2.86	100	437	0
Virginia	2.58	100	425	0
Massachusetts	2.25	100	298	0
Maryland	2.17	100	253	0
Arizona	1.88	100	267	0
Indiana	1.86	97	398	11
Missouri	1.86	100	325	1
Tennessee	1.84	100	352	1
Wisconsin	1.81	100	279	0
Washington	1.71	100	369	0
Louisiana	1.62	71	1,102	322
Colorado	1.48	-21	246	297
Alabama	1.48	83	268	47
Minnesota	1.47	100	318	0
South Carolina	1.36	100	262	0
Connecticut	1.27	100	167	0
Oklahoma	1.14	-93	267	516
Kentucky	1.10	97	303	10
Oregon	1.09	100	172	0
Arkansas	0.94	83	172	30
Iowa	0.92	100	241	0
Kansas	0.79	-16	180	210
Mississippi	0.79	50	219	109
Nevada	0.77	99	119	1
West Virginia	0.65	69	106	32
Utah	0.63	-6	148	157
New Hampshire	0.60	100	80	0
New Mexico	0.60	-275	123	460
Nebraska	0.55	90	123	13
Maine	0.44	100	99	0
Hawaii	0.40	100	116	0
Delaware	0.36	100	51	0

(continued)

Table A.2. (Continued)

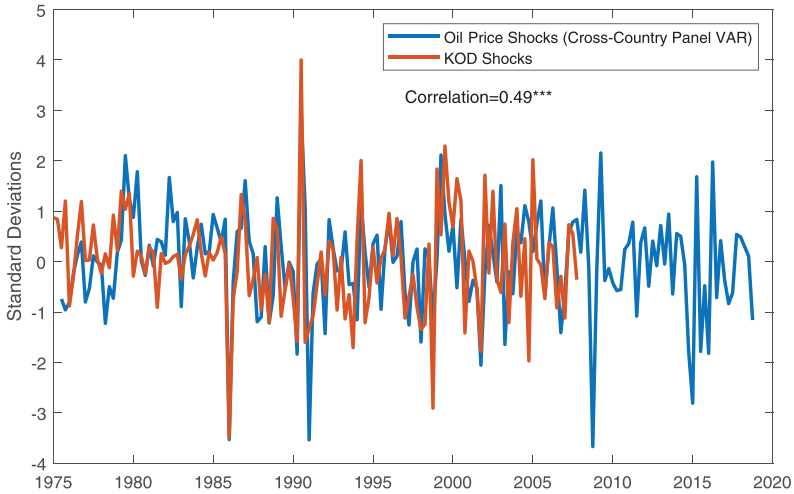
State	Car Share	Oil Dependency	Oil Consumption	Oil Production
Idaho	0.36	100	84	0
Rhode Island	0.33	100	44	0
Vermont	0.28	100	44	0
Montana	0.28	-47	89	131
South Dakota	0.24	86	61	8
North Dakota	0.22	-1,119	115	1,398
Alaska	0.20	-591	122	841
Wyoming	0.19	-246	82	284
District of Columbia	0.14	100	10	0

Notes: Oil dependency is calculated as oil consumed minus oil produced as a percentage of oil consumed (source: U.S. Energy Information Administration: State Profiles and Energy Estimates). Oil production is production of crude oil by state (*sumcrudestate*). Production of crude oil by state is scaled up by the ratio of total U.S. oil production from the BP Statistical Review (which includes tight oil, oil sands, and other NGLs and including offshore production) to *sumcrudestate* (which is about 1.5), so as to facilitate comparison with oil-dependency measures by country. Oil consumption includes all petroleum products consumed, and lines up with oil consumption from the BP Statistical Review. Production and consumption are in thousands of barrels daily.

thus excluding registrations of rental, fleet, government, and other commercial use vehicles. From 1989 to 2002, Polk's data is computed by marketing regions, with information also available for states. We then splice this data set with the one for 2002 onward, which is computed directly at the state level. As shown by figure B.3, retail registrations at the national level are very highly correlated with durable goods, namely motor vehicles and parts from the NIPA.

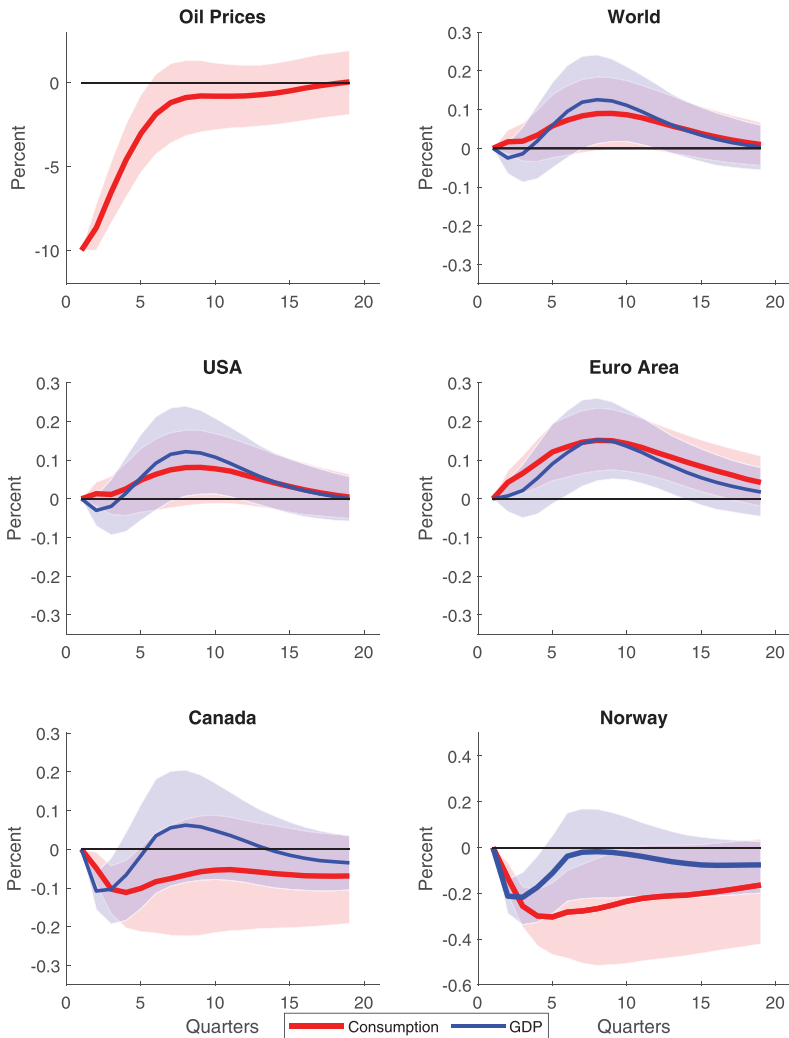
Appendix B. Supplementary Results

Figure B.1. Comparison of Oil Shocks Identified under Different Strategies



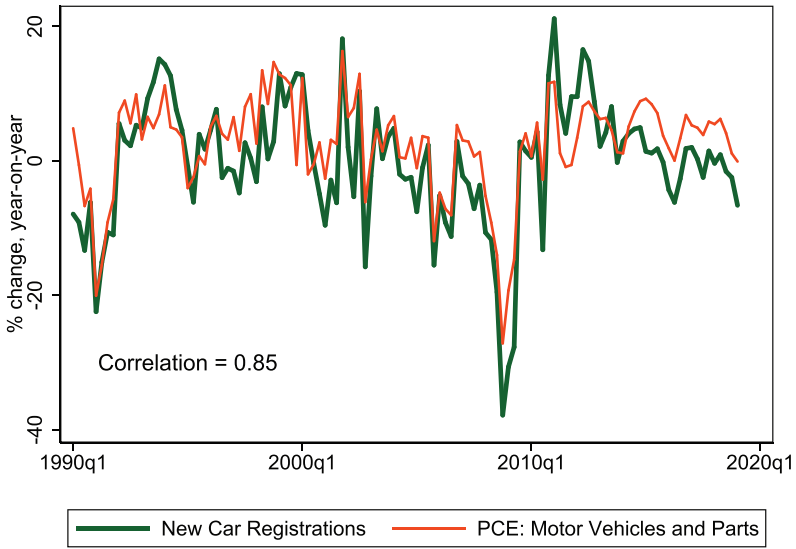
Notes: “Oil Price Shocks (Cross-Country Panel VAR)” represent the oil shocks identified by the panel VAR from section 3.2 of this paper. “KOD Shocks” represent the “oil-specific demand” shocks identified by Kilian (2009) averaged to the quarterly frequency. The correlation between the two series of shocks is statistically significant at the 99 percent confidence level (***)

Figure B.2. Impulse Response Functions to Oil Price Shocks in Cross-Country Panel VAR, Data Post-1986



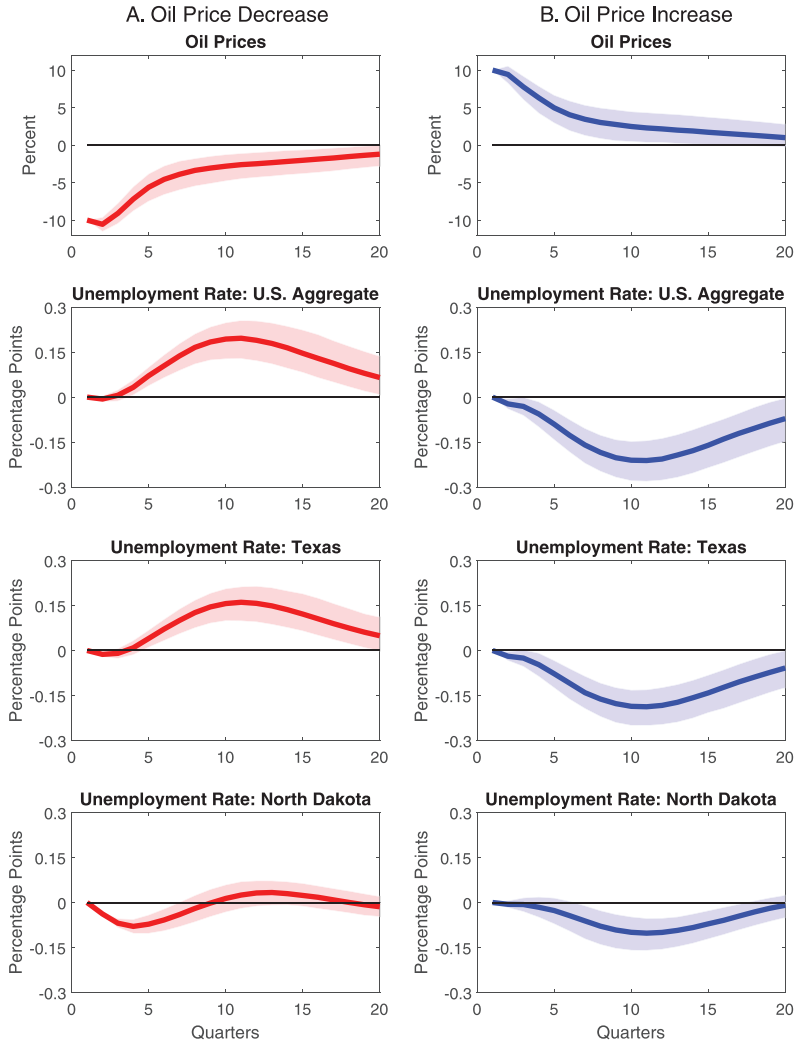
Notes: The figure shows the IRFs to oil price shocks in the cross-country panel VAR (section 3.2) using data from 1986 onward. Solid lines represent median IRFs. Shaded areas represent 16th and 84th percentiles constructed with 5,000 bootstrap replications. Variables are expressed as percent deviation from quadratic trend.

Figure B.3. New Cars' Registrations and Motor Vehicles Expenditures from NIPA



Note: This figure compares NIPA expenditures on motor vehicles (in real terms) with the aggregate data on new cars' registrations from Polk's National Vehicle Population Profile.

Figure B.4. Impulse Response Functions to Oil Price Shocks in State-Level Panel VAR



Notes: This figure shows the IRFs to oil price shocks in the U.S. state-level panel VAR (section 4). Solid lines represent median IRFs. Shaded areas represent 16th and 84th percentiles constructed using 5,000 bootstrap replications. Variables are expressed as percent deviation from quadratic trend.

References

- 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): 1013–28.
- Arezki, R., and O. Blanchard. 2014. "Seven Questions about the Recent Oil Price Slump." IMF Blog, December 22.
- Baumeister, C., and G. Peersman. 2013. "The Role of Time-Varying Price Elasticities in Accounting for Volatility Changes in the Crude Oil Market." *Journal of Applied Econometrics* 28 (7): 1087–1109.
- Baumeister, C., G. Peersman, and I. Van Robays. 2010. "The Economic Consequences of Oil Shocks: Differences across Countries and Time." In *Inflation in an Era of Relative Price Shock*, 91–128. Reserve Bank of Australia.
- Bernanke, B. S. 2016. "The Relationship between Stocks and Oil Prices." Brookings Institution.
- Bernanke, B. S., J. Boivin, and P. Eliasch. 2005. "Measuring the Effects of Monetary Policy: A Factor-Augmented Vector Autoregressive (FAVAR) Approach." *Quarterly Journal of Economics* 120 (1): 387–422.
- Bjørnland, H. C., and J. Zhulanova. 2018. "The Shale Oil Boom and the US Economy: Spillovers and Time-Varying Effects." Working Paper No. 8/2018, Centre for Applied Macro- and Petroleum Economics (CAMP), BI Norwegian Business School.
- Bodenstein, M., C. J. Erceg, and L. Guerrieri. 2011. "Oil Shocks and External Adjustment." *Journal of International Economics* 83 (2): 168–84.
- Bodenstein, M., L. Guerrieri, and L. Kilian. 2012. "Monetary Policy Responses to Oil Price Fluctuations." *IMF Economic Review* 60 (4): 470–504.
- Caldara, D., M. Cavallo, and M. Iacoviello. 2019. "Oil Price Elasticities and Oil Price Fluctuations." *Journal of Monetary Economics* 103 (May): 1–20.
- Camacho, M., and G. Perez-Quiros. 2002. "This Is What the Leading Indicators Lead." *Journal of Applied Econometrics* 17 (1): 61–80.
- Canova, F., and M. Ciccarelli. 2013. "Panel Vector Autoregressive Models: A Survey." ECB Working Paper No. 1507.

- CBS/AP. 2016. "Buffett Says U.S. Economy Weaker Than He Expected but Growing." CBS MoneyWatch, February 29.
- Chudik, A., and M. H. Pesaran. 2016. "Theory and Practice of GVAR Modelling." *Journal of Economic Surveys* 30 (1): 165–97.
- Coglianesi, J., L. W. Davis, L. Kilian, and J. H. Stock. 2017. "Anticipation, Tax Avoidance, and the Price Elasticity of Gasoline Demand." *Journal of Applied Econometrics* 32 (1): 1–15.
- 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.
- Edelstein, P., and L. Kilian. 2009. "How Sensitive Are Consumer Expenditures to Retail Energy Prices?" *Journal of Monetary Economics* 56 (6): 766–79.
- 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 (December).
- Hamilton, J. D. 1996. "This Is What Happened to the Oil Price-Macroeconomy Relationship." *Journal of Monetary Economics* 38 (2): 215–20.
- . 2003. "What Is an Oil Shock?" *Journal of Econometrics* 113 (2): 363–98.
- . 2008. "Oil and the Macroeconomy." In *The New Palgrave Dictionary of Economics*, ed. S. Durlauf and L. Blume. Palgrave MacMillan.
- Holtz-Eakin, D., W. Newey, and H. S. Rosen. 1988. "Estimating Vector Autoregressions with Panel Data." *Econometrica* 56 (6): 1371–95.
- Jiménez-Rodríguez, R., and M. Sánchez. 2005. "Oil Price Shocks and Real GDP Growth: Empirical Evidence for Some OECD Countries." *Applied Economics* 37 (2): 201–28.
- Kilian, L. 2009. "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market." *American Economic Review* 99 (3): 1053–69.
- Kilian, L., A. Rebucci, and N. Spatafora. 2009. "Oil Shocks and External Balances." *Journal of International Economics* 77 (2): 181–94.

- Kilian, L., and R. J. Vigfusson. 2011. "Are the Responses of the U.S. Economy Asymmetric in Energy Price Increases and Decreases?" *Quantitative Economics* 2 (3): 419–53.
- Peersman, G., and I. Van Robays. 2012. "Cross-country Differences in the Effects of Oil Shocks." *Energy Economics* 34 (5): 1532–47.
- Pesaran, M. H. 2006. "Estimation and Inference in Large Heterogeneous Panels with a Multifactor Error Structure." *Econometrica* 74 (4): 967–1012.
- Stock, J. H., and M. W. Watson. 2005. "Implications of Dynamic Factor Models for VAR Analysis." NBER Working Paper No. 11467.
- . 2016. "Dynamic Factor Models, Factor-Augmented Vector Autoregressions, and Structural Vector Autoregressions in Macroeconomics." In *Handbook of Macroeconomics*, Vol. 2A, ed. J. B. Taylor and H. Uhlig, 415–525 (chapter 8). Elsevier.