

Ryan, Rich; Michieka, Nyakundi

## Article

# Not All Oil Price Shocks Are Alike. A Replication Study of Kilian (American Economic Review, 2009)

Journal of Comments and Replications in Economics (JCRE)

*Suggested Citation:* Ryan, Rich; Michieka, Nyakundi (2025) : Not All Oil Price Shocks Are Alike. A Replication Study of Kilian (American Economic Review, 2009), Journal of Comments and Replications in Economics (JCRE), ISSN 2749-988X, ZBW - Leibniz Information Centre for Economics, Kiel, Hamburg, Vol. 4, pp. 1-27,  
<https://doi.org/10.18718/81781.47>

This Version is available at:

<https://hdl.handle.net/10419/337522>

### Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

### Terms of use:

*Documents in EconStor may be saved and copied for your personal and scholarly purposes.*

*You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.*

*If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.*



<https://creativecommons.org/licenses/by/4.0/>

# Not All Oil Price Shocks Are Alike

## A Replication Study of Kilian (*American Economic Review*, 2009)

Rich Ryan\*      Nyakundi Michieka†

*Journal of Comments and Replications in Economics*, Volume 4, 2025-8, DOI: [10.18718/81781.47](https://doi.org/10.18718/81781.47)

*JEL: E24, E31, E32, Q41, Q43, Q48, R23*

*Keywords: Kern, Local labor market, Oil price, Real economic activity, Structural vector autoregression, Unemployment rate, Vector autoregression*

*Data Availability:* The R code and the data to reproduce the results of this replication can be downloaded at JCRE's data archive (DOI: [10.15456/j1.2024244.2317598695](https://doi.org/10.15456/j1.2024244.2317598695)).

*Please Cite As:* Ryan, R. and Michieka, N. (2025). Not All Oil Price Shocks Are Alike. A Replication Study of Kilian (2009). *Journal of Comments and Replications in Economics*, Vol.4 (2025-8). DOI: [10.18718/81781.47](https://doi.org/10.18718/81781.47)

### *Abstract*

The price of oil can rise because of a disruption to supply or an increase in demand. The nature of the price change determines the dynamic effects. As Kilian (2009a) put it: “not all oil price shocks are alike.” Using the latest available data, we extend Kilian’s analysis using the R ecosystem and provide more evidence for Kilian’s conclusions. Inference based on unknown conditional heteroskedasticity strengthens the conclusions. With the updated shocks, we assess how a local economy responds to the global oil market, an application that is relevant to policymakers concerned with the transition away from fossil fuels.

---

\*Department of Economics, California State University, Bakersfield, [richryan@csub.edu](mailto:richryan@csub.edu).

†Department of Economics, California State University, Bakersfield, [nmichieka@csub.edu](mailto:nmichieka@csub.edu).

*Declaration: We gratefully acknowledge support from the Resilient Energy Economies Initiative.*

Received September 01, 2024; Revised April 08, 2025; Accepted June 24, 2025; Published December 05, 2025.  
©Author(s) 2025. Licensed under the Creative Common License - Attribution 4.0 International (CC BY 4.0).

## 1 Introduction

In a stylized model for the oil market, the price of oil depends on supply and demand. Price can go up if demand shifts outward or supply shifts inward. Even though the nature of the shocks differs, each causes the price of oil to rise. Surprisingly, as Kilian (2009a) points out, much theoretical, empirical, and policy-oriented work focuses on the effects of oil-price fluctuations, ignoring the nature of shocks. This risks flawed policy. Fluctuations in the price of oil are an amalgam of demand and supply shocks. The dynamics generated by a demand shock differ from the dynamics generated by a supply shock. At any point in time, the composition of shocks may be far from the average, making correlations between the price of oil and US macroeconomic aggregates potentially meaningless and uninformative for policymakers (Kilian, 2008, 2009a).

In brief, here is the main issue: Kilian (2009a) establishes that the dynamic effects of an increase in the price of oil depend on the nature of the increase—whether the increase was the result of a disruption to supply or an increase in demand. The conclusions are based on a structural model of the global market for oil estimated over the period February 1973 to December 2007. Do the same conclusions hold using all available data? Are the results robust to inference that allows for conditional heteroskedasticity?

Using data on oil production and price from the US Energy Information Administration and data on global real economic activity available through FRED, we estimate the structural vector autoregressive model of the global crude-oil market proposed by Kilian (2009a). The updated results provide further evidence that Kilian's (2009a) conclusions hold. In addition, a procedure that accounts for conditional heteroskedasticity of unknown form in the VAR model, indicates that Kilian's conclusions are even stronger than initially stated (Brüggemann et al., 2016; Kilian & Lütkepohl, 2017).

The structural VAR model disentangles supply shocks from demand shocks in the global crude-oil market. In the model, supply comprises global oil production and demand comprises two components: (i) a component that reflects the global business cycle and (ii) a component that reflects oil-specific demand (buyers concerned about availability in the future will buy out of precaution, for example [Alquist & Kilian, 2010]). This allows Kilian (2009a) to show how each shock separately determines fluctuations in the real price of crude oil. Moreover, how real Gross Domestic Product responds to an unanticipated increase in the price of oil depends on the nature of the shock. The same conclusion applies to the Consumer Price Index. Kilian's results explain the macroeconomic implications of oil-price fluctuations.

In a different direction, we use the updated series of structural shocks to learn about another important policy question: how does a local labor market respond to these shocks? The question is of interest to policymakers in California. California has adopted the goal of carbon neutrality by 2045 or earlier (CARB, 2022). In Kern County, which accounts for 70 percent of California's oil production, there is concern about potential job loss and lost tax revenue as the economy transitions away from fossil fuels. Yet, given the complex linkages between oil extraction and the rest of the economy, Kern's dependence on oil is hard to assess (Golding & Kilian, 2022, describe one complex link). We document how Kern's labor market responds to identified shocks in the global market for oil. And we find that the response depends on the nature of the shock. Kern's dependence on oil suggests that economic forces will restructure the composition of regional economic activity

along the path to net-zero emissions by 2045. Past structural changes, like the transition away from coal, do not augur well for displaced workers and regions (CEA, 2024). The evidence suggests that policymakers may need to consider expanding options in Kern through, for example, place-based policy.<sup>1</sup> The exercise shows how Kilian’s (2009a) pioneering methods and conclusions apply more generally.

Questions about how oil affects the US economy are not fully answered. Further research will benefit from other contributions of our replication. We port Kilian’s original MATLAB code to R. Some researchers may prefer to work in R as the R ecosystem makes tidying data easier. (Our code embraces the `tidyverse`.) We use the `targets` package to ensure replicability. The `targets` package provides Make-like tools that automatically make any necessary changes to dependencies when any change is registered. In addition, we retrieve data from the US Energy Information Administration and FRED using code, which makes estimating the most recent structural shocks as straightforward as rerunning the code. These resources will make addressing questions easier in the ongoing debate about the effects oil-price shocks have on the US economy.<sup>2</sup>

We present a replication of Kilian’s (2009a) results in section 2. We conclude that Kilian (2009a) should be praised and thanked for writing and sharing code that makes reproducing results straightforward (Kilian, 2009b). Using the posted data and our code, we are able to exactly replicate Kilian’s (2009a) results.

The replication comprises a description of the data (section 2), a specification of the structural VAR model (section 2.1), a presentation of the estimated impulse response functions (section 2.2), a decomposition of the real price of oil into identified components (section 2.3), and a second-stage analysis of how US real GDP and the Consumer Price Index respond to the identified structural shocks. In section 3, building upon an analysis undertaken by Kilian & Lütkepohl (2017), we strengthen Kilian’s (2009a) conclusions by estimating the structural VAR model using data from February 1974 to January 2025 and constructing confidence intervals for impulse responses that are robust to arbitrary forms of conditional heteroskedasticity. To demonstrate the wider applicability of the results, in section 4, we analyze how the global oil market affects outcomes in a local labor market.

## 2 Replication

Kilian’s (2009a) pioneering model of the global crude-oil market specifies three determinants. The methodology separates a single supply shock from two demand shocks. The oil-supply shock reflects unexpected disruptions to crude-oil availability. The aggregate-demand shock reflects the consensus that shifts in global real economic activity shift current demand for crude oil (Kilian & Zhou, 2018). And the precautionary-demand shock reflects how demand may shift based on uncertainty about future unavailability of oil. Alquist & Kilian (2010) make this point, among others, in an economy that receives a random endowment of oil each period. They show that an increase

---

<sup>1</sup>What will happen in California as 2045 approaches is impossible to know with certainty. In particular, a backward-looking structural VAR model will not be a good predictor of abrupt change. Instead, we view Kern’s responses to identified structural innovations as suggestive evidence that the Kern economy depends on oil.

<sup>2</sup>An introduction to Make is provided by Jesús Fernández-Villaverde in [lecture notes](#) titled “Computational Methods for Economists,” which are available at <https://www.sas.upenn.edu/~jesusfv/teaching.html>

in the risk of future availability of oil raises current demand.

Three variables compose Kilian's (2009a) vector autoregressive model of the global crude-oil market: (i) the log difference of global crude-oil production, (ii) a measure of cyclical variation in global real economic activity, and (iii) the log of the real price of oil. Data on global crude-oil production are available from the US Energy Information Administration (EIA). Kilian's (2009a) measure of cyclical variation in global real economic activity is made available by the Federal Reserve Bank of Dallas through the FRED database. Kilian (2019) reviews the index and corrects an error that was made when the original series was constructed. The real price of oil is constructed by deflating US refiners' acquisition cost of imported crude oil (available from the EIA) by the Consumer Price Index for All Urban Consumers. Some details about the series and how we accessed them are included in an appendix in section A.1.<sup>3</sup>

## 2.1 The structural VAR model

To be precise about how the determinants of the global market for crude oil are related, consider a VAR model for  $y_t = (\Delta\text{oil production, real activity, real price of oil})'$ . The structural VAR representation is

$$B_0 y_t = \beta + B_1 y_{t-1} + \dots + B_{24} y_{t-24} + \varepsilon_t, \quad (1)$$

where the  $B_i$  are  $3 \times 3$  parameter matrices and  $\varepsilon_t \sim (0, \Sigma_\varepsilon)$ ; that is,  $\varepsilon_t$  denotes a vector of serially and mutually uncorrelated innovations. The covariance matrix  $\Sigma_\varepsilon$  is diagonal. The corresponding reduced-form model is

$$y_t = \alpha + A_1 y_{t-1} + \dots + A_{24} y_{t-24} + \epsilon_t, \quad (2)$$

where  $A_i = B_0^{-1} B_i$  and  $\epsilon_t = B_0^{-1} \varepsilon_t \sim (0, \Sigma_\epsilon)$  is a martingale-difference sequence with positive-definite covariance matrix  $\Sigma_\epsilon = (B_0^{-1}) \Sigma_\varepsilon (B_0^{-1})'$ . Going between the structural model in (1) and the reduced-form model in (2) requires an estimate of  $B_0^{-1}$ . Short-run identifying restrictions make estimating this term straightforward.

The recursive structure for  $B_0^{-1}$  implies that the reduced-form error  $\epsilon_t$  can be decomposed as

$$\epsilon_t \equiv \begin{bmatrix} \epsilon_t^{\Delta\text{oil production}} \\ \epsilon_t^{\text{real activity}} \\ \epsilon_t^{\text{real price of oil}} \end{bmatrix} = \begin{bmatrix} b_{11} & 0 & 0 \\ b_{21} & b_{22} & 0 \\ b_{31} & b_{32} & b_{33} \end{bmatrix} \begin{bmatrix} \epsilon_t^{\text{oil supply}} \\ \epsilon_t^{\text{aggregate demand}} \\ \epsilon_t^{\text{oil-specific demand}} \end{bmatrix}. \quad (3)$$

Within the system of three equations, the first equation indicates that the stochastic short-run supply curve is vertical. In other words, oil production does not respond to innovations in demand within the month. The maintained assumption is that decisions about oil supply are based on medium-run expectations and, because adjusting supply is costly, short-run demand innovations are postulated to have a negligible affect on medium-run expectations. Evidently, the model allows energy prices to respond to all past information. There are two demand shocks and each is identified by delay restrictions. The short-run demand curve is downward sloping and can be shifted by innovations to

---

<sup>3</sup>The appendix also reports a sensitivity analysis, in which we replace Kilian's (2009a) index with the monthly growth rate of OECD+6 industrial production, which is an index of industrial production for OECD (Organization for Economic Co-operation and Development) countries and six non-member countries (Brazil, China, India, Indonesia, the Russian Federation, and South Africa). Further discussion is provided in the appendix.

aggregate demand and innovations to oil-specific demand, such as precautionary-demand innovations (Alquist & Kilian, 2010). The delay restriction is that oil-specific demand shocks do not affect global real economic activity as measured by fluctuations in shipping rates relative to trend. These postulations are further discussed by Kilian (2008) and Kilian (2009a).

In (3),  $B_0^{-1}$  is lower triangular. A lower-triangular Cholesky decomposition can therefore be used to decompose the estimate for  $\widehat{\Sigma}_\epsilon$ , starting from the reduced-form estimates of the model in (2). Uncovering the structural shocks is possible using the estimated relationships in (3).

Inference is based on the wild bootstrap discussed by Gonçalves & Kilian (2004) with 1,000 replications. Kilian (2009a) uses 2,000 replications but the results are similar. The wild bootstrap allows for conditional heteroskedasticity of unknown form. But the asymptotic validity of the wild bootstrap is for the reduced-form impulse responses and not the asymptotic distribution of the error-covariance matrix. This feature is shown by Brüggemann et al. (2016) and discussed by Kilian & Lütkepohl (2017). In section 3, we follow Kilian & Lütkepohl (2017) and adopt the inference procedures proposed by Brüggemann et al. (2016) to investigate whether the larger confidence intervals matter in practice.

## *2.2 How global oil production, real economic activity, and the real price of oil respond to demand and supply shocks in the crude-oil market*

The structural impulse responses define the responses of elements of  $y_t$  to one-time impulses in  $\varepsilon_t$ . Figure 1 shows how global oil production, real economic activity, and the real price of oil respond to one-standard deviation structural innovations. The innovations are normalized so that each innovation will tend to raise the price of oil or, put differently and slightly imprecisely, figure 1 shows the effects of negative supply shocks and positive demand shocks. Figure 1 is comparable to figure 3 in Kilian (2009a).

Figure 1 uses data provided by Kilian (2009b). The data start in February 1973 and end in December 2007. The goal of figure 1 is to assess whether the R code written for the replication yields similar results to the MATLAB code that produces the original graphics.

Figure 1 confirms that the two sets of code produce identical results. The solid, blue lines show the point estimates of the structural impulse responses. They match those produced by the code provided by Kilian (2009a) exactly. Kilian's (2009a) original estimates are shown using red, dotted lines. These overlay the solid, blue lines, which produces a striped effect. The shaded areas depict the two-standard-error bands constructed by the wild bootstrap. The dotted, red lines show those produced by the original code. Again, there is overwhelming overlap, which gives confidence that the two sets of codes are working.

Figure 2 uses our code base and our data over the period that overlaps the original sample period. There are three differences to keep in mind when comparing figure 1 to figure 2. First, for the analysis in figure 2, all the data were retrieved by our code from the sources listed in the text. Second, the updated data on production start in February 1974. Third, the data on global real economic activity have been corrected (Kilian, 2019). Notwithstanding these differences, the

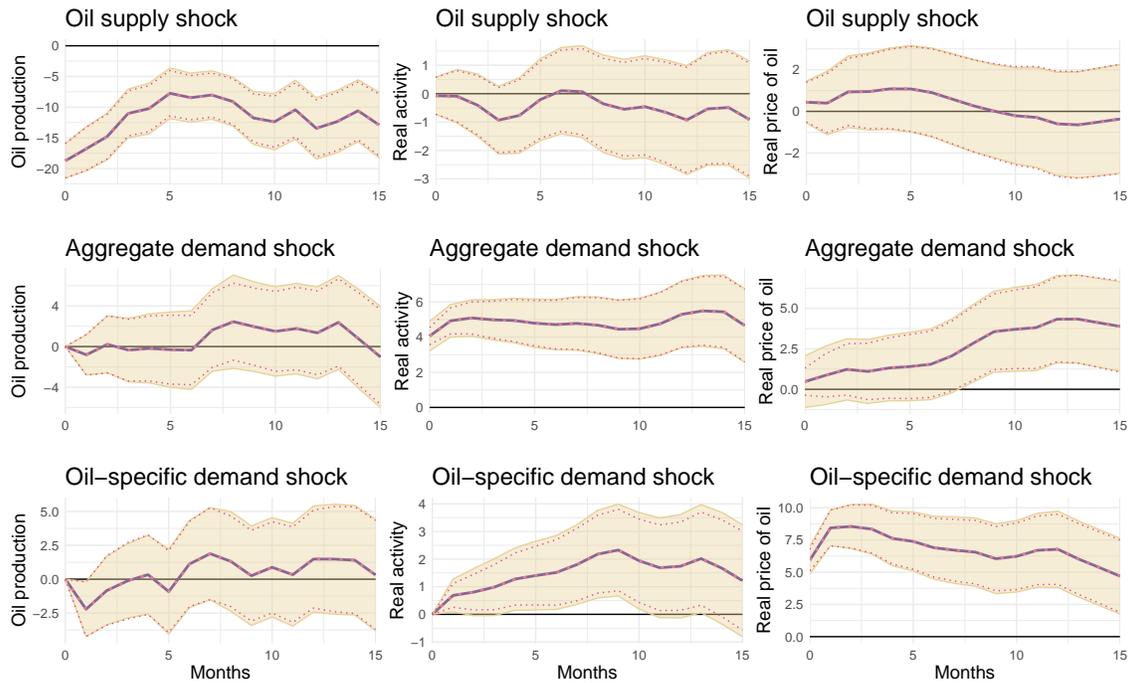


Figure 1: Responses to one-standard-deviation structural shocks (based on original data)

Note: Solid lines show the point estimates and the shaded regions show two-standard-error bands. Estimates are based on the structural VAR model in (1), using the original data (Kilian, 2009b) (February 1973–December 2007). Confidence intervals were constructed using a recursive-design wild bootstrap. Dotted lines show the point estimates and two-standard-error bands produced by Kilian’s (2009a) code. The striped effect indicates the two sets of point estimates coincide.

patterns are strikingly similar.

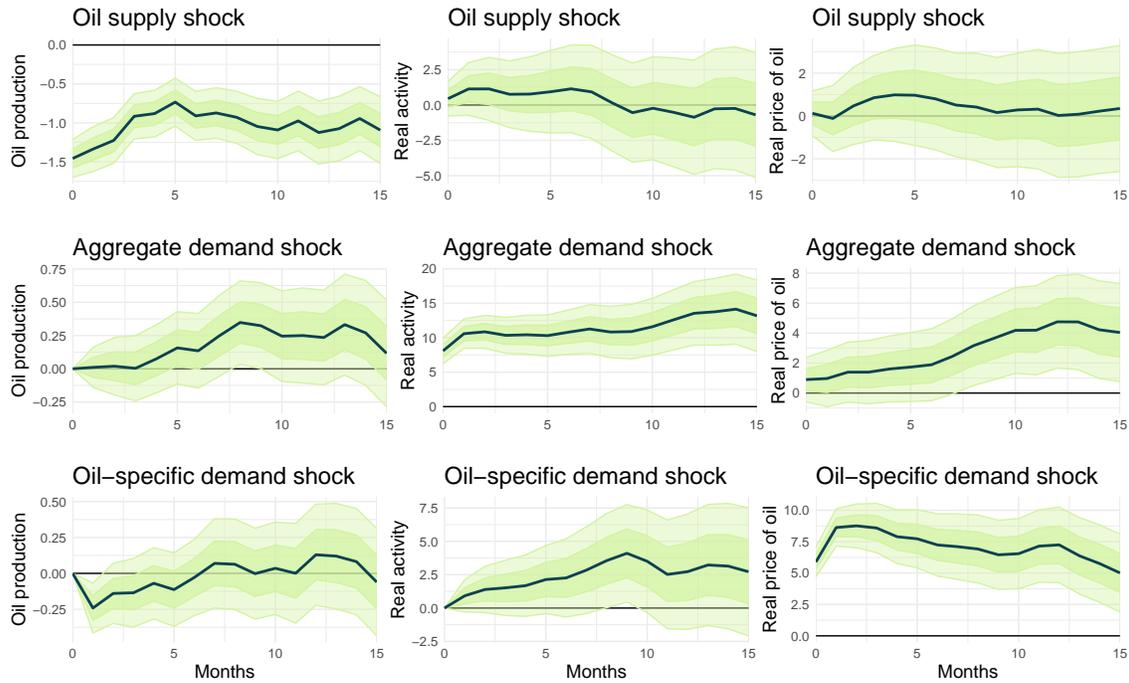


Figure 2: Responses of one-standard-deviation structural shocks (based on updated data)

Note: Solid lines show the point estimates and the shaded regions show one- and two-standard-error bands. Estimates are based on the structural VAR model in (1), using the updated data (February 1974–December 2007). Confidence intervals were constructed using a recursive-design wild bootstrap.

An unexpected disruption to oil supply causes an immediate decline in oil production. The decline is partially reversed within the first 5 months, and “this pattern is consistent with the view that oil supply contractions in one region tend to trigger production increases elsewhere in the world” (Kilian, 2009a, 1061–1062). The fall in production causes a rise in the price of oil, but the rise is not overwhelmingly statistically significant. In figure 2, the disruption has little effect on real economic activity, which differs slightly from Kilian’s (2009a) observation that a supply disruption “causes a small temporary reduction of real economic activity that is partially statistically significant” (1,062).

Turning to the second row of figure 2, the effect of an unanticipated rise in aggregate demand causes oil production to rise about 6 months after the shock. The expansion lasts until 15 months after the shock. The unanticipated rise in aggregate demand causes a significant rise in global real economic activity. The effect is persistent. Expansions in real economic activity cause a significant rise in the real price of oil. The real price of oil rises on impact (although not statistically significantly) and the rise continues for a year.

The third row of figure 2 shows the effects of an unanticipated rise in oil-specific demand. An

unanticipated rise causes oil production to briefly fall, but over the 15-month horizon, oil-specific-demand increases do not cause oil production to change much. Real economic activity in response to the unanticipated rise in oil-specific demand expands—albeit modestly, if at all. What is most clear, though, is that the unanticipated rise in oil-specific demand causes the real price of oil to jump upward immediately and continue to rise for a few months, which is consistent with models of precautionary demand. Within this class of models, this overshooting feature is discussed by Kilian (2009a) and Alquist & Kilian (2010, figure 3b, 561). Finally, the rise in the real price of oil is highly statistically significant.

Our results are entirely consistent with the results in Kilian (2009a).

### 2.3 *The cumulative effect of oil demand and oil supply shocks on the real price of oil*

How much each structural shock explains of the historically observed price of crude oil can be assessed through a historical decomposition of the covariance-stationary VAR model. The historical decomposition decomposes  $y_t$  into structural innovations through its moving-average representation. It attributes the cumulative effect of the structural innovations in oil production, real economic activity, and oil-specific demand to the real price of oil at each point in time (Kilian & Lütkepohl, 2017).

Figure 3 shows the cumulative contribution of oil-supply, aggregate-demand, and oil-specific-demand shocks to the real price of oil in separate panels. Each panel shows 4 lines. The top panel shows the cumulative effect of oil-supply shocks. The solid, purple line is generated from the original code and original data. The dashed, beige line is generated by using the replication code together with the original data. The two series are numerically equal, which can be seen by the striped pattern, the effect of the two series overlapping. The dot-dashed, orange line is generated by using our replication code and the updated data through December 2007. The dotted, pink line is generated by using the replication code and updated data through January 2025.

The top panel of figure 3 shows that “oil supply shocks historically have made comparatively small contributions to the price of oil” (Kilian, 2009a, 1062).

The middle panel of figure 3 shows that aggregate-demand shocks cause prolonged and persistent changes in the price of oil. The bottom panel of figure 3 shows that shocks to oil-specific demand are responsible for the fluctuations in the price of oil. As Kilian (2009a) points out, this statistical understanding of crude-oil prices is inconsistent with narrative hypotheses about supply disruptions in the Middle East. There may be supply disruptions, but the disruptions affect firms that purchase oil, fearing further supply disruptions. These precautionary-demand motives cause the price of oil to fluctuate.

Figure 3 is directly comparable to Kilian’s (2009a) figure 4. The conclusions drawn from both figures agree. Therefore, going forward results will be shown for estimates produced using the full sample.

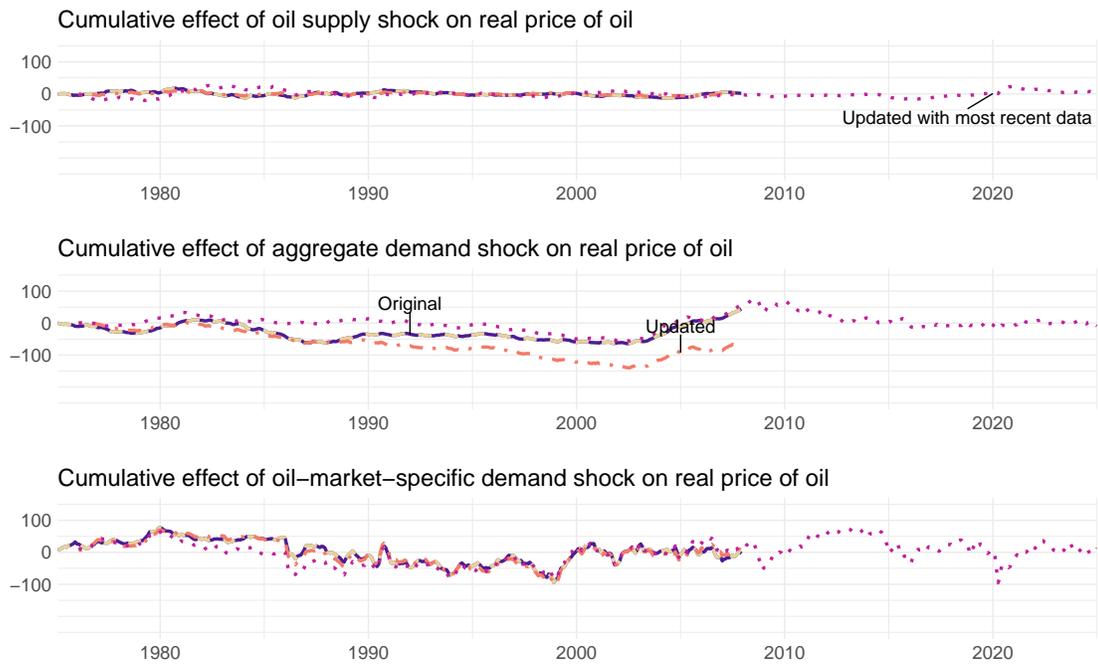


Figure 3: Historical decomposition of the real price of oil

Note: *Top*, how much oil-supply shocks explain of the historically observed fluctuations in the real price of crude oil; *middle*, how much aggregate-demand shocks explain of the historically observed fluctuations in the real price of crude oil; *bottom*, how much oil-specific-demand shocks explain of the historically observed fluctuations in the real price of crude oil. Estimates are based on the structural VAR model in (1). The solid, purple line is generated by the original code and data. The dashed, beige line, which lies on top of the purple line, is generated by the replication code and the original data. The dot-dashed, orange line is generated by the replication code using data from February 1974 to December 2007. The dotted, pink line is generated by the replication code using data through January 2025.

## 2.4 Demand and supply shocks in the global oil market affect the US economy in different ways

How does the global market for oil affect economic activity in the United States? Having identified the structural innovations using the model in 1, it is possible to relate the innovations to US statistics like the inflation rate measured by the Consumer Price Index (CPI) and the growth rate of real gross domestic product (GDP).

Because data on real GDP are available at a quarterly frequency, the structural shocks are averaged by quarter. Let  $\hat{\zeta}_{jt}$  be the  $j$ th structural shock in period  $t$ , where  $t$  now indexes quarter and  $j \in \{\text{oil supply, aggregate demand, oil-specific demand}\}$ . Put differently,  $\hat{\zeta}_{jt}$  is the quarterly average of the monthly values  $\hat{\varepsilon}_{j\tau}$ , where  $\tau$  is a month within quarter  $t$ . Kilian (2009a) formalizes the relationship between the structural innovations and US economic activity by running regressions of the form

$$z_t = \omega_j + \sum_{i=0}^{12} \phi_{ji} \hat{\zeta}_{jt-i} + u_{jt}, \quad (4)$$

where  $z_t$  is either the quarterly growth rate of real GDP or the quarterly inflation rate and  $u_{jt}$  is a potentially serially correlated error. As Kilian (2009a) notes, the term  $\phi_{jh}$  in the regression model is the impulse response at horizon  $h$  of  $z$  to structural innovation  $j$ .

The possibility for serial correlation in (4) is accounted for by constructing confidence intervals using a block-bootstrap procedure. The results presented here use a block length of 6 quarters, whereas Kilian (2009a) uses 4 quarters.

Figure 4 shows how the level of US real GDP and the level of the CPI respond to the structural shocks. The solid, gray lines show impulse responses estimated using the original data and the replication code. The associated two-standard-error bands are shaded in green. The impulse responses estimated using Kilian’s (2009a) code agree exactly, which produces a striped effect. The dotted, blue lines show the upper and lower ends of the confidence intervals using Kilian’s (2009a) code. These estimates produce slightly narrower confidence intervals, likely due to the fact that the block length used to produce the shaded region is two quarters longer. Nevertheless, the difference is minor. Figure 4 is directly comparable to Kilian’s (2009a) figure 5.

Figure 5 updates figure 4. To do this, data from February 1974 to January 2025 are first used to estimate the structural VAR model. Then the associated monthly series of structural shocks are averaged by quarter to produce the quarterly averages used in the second-stage regressions in (4).

Figure 5 shows that similar patterns emerge when the series for real economic activity is corrected and the latest data are used to construct estimates. An unanticipated oil-supply disruption lowers real GDP on impact. The effect is statistically significant when using the one-standard-error bands for up to nearly three years. The unanticipated oil-supply disruption causes the CPI to increase for nearly 9 quarters. The effect is statistically significant using the one-standard-error bands. Only after 9 quarters does the price level start to decline.

Turning to the second row of figure 5, an unanticipated expansion in aggregate demand causes real GDP to increase on impact, although the effect is marginally statistically significant. After 5

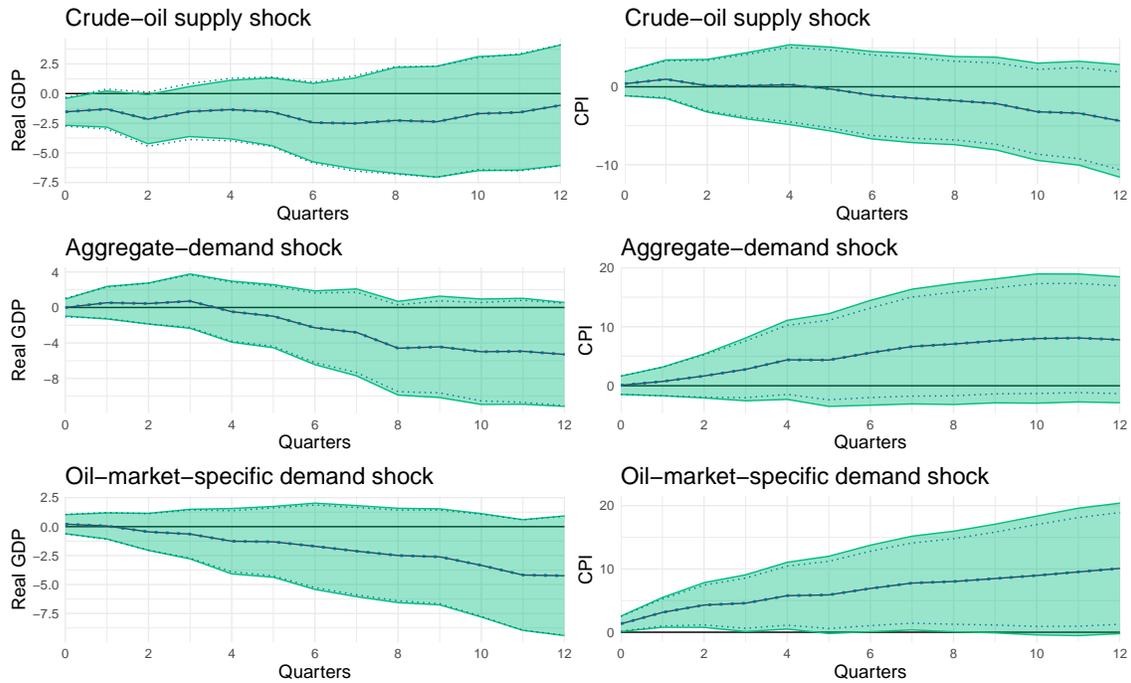


Figure 4: Responses of US real GDP and the CPI to identified structural shocks (based on original data)

Note: Clockwise from top right, cumulative response of the price level to an unanticipated disruption in crude-oil supply; cumulative response of the price level to an unanticipated increase in aggregate demand; cumulative response of the price level to an unanticipated increase in oil-specific demand; cumulative response of real GDP to an unanticipated increase in oil-specific demand; cumulative response of real GDP to an unanticipated increase in aggregate demand; cumulative response of real GDP to an unanticipated disruption in crude-oil supply. Solid lines are based on the point estimates from the model in (4). Shaded areas show two-standard-error bands constructed using a block bootstrap procedure. Dotted lines show the point estimates and two-standard error bands produced by Kilian's (2009a) code. All statistics in the figure are based on the original data (Kilian, 2009b).

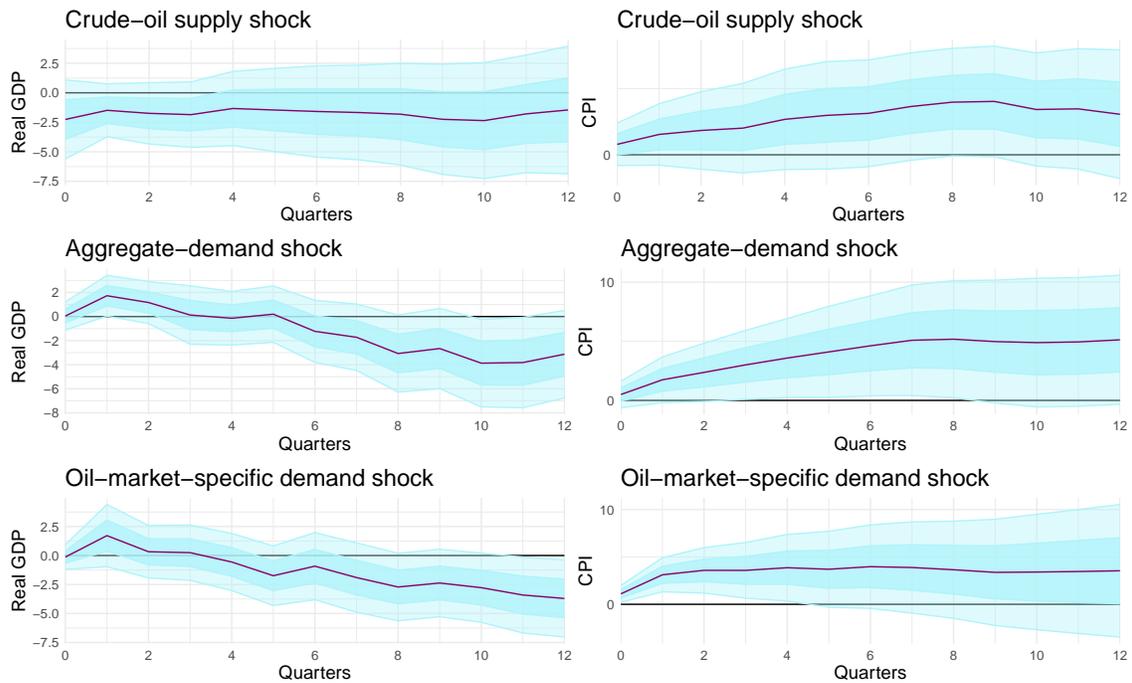


Figure 5: Responses of US real GDP and the CPI to identified structural shocks

Note: Clockwise from top right, cumulative response of the price level to an unanticipated disruption in crude-oil supply; cumulative response of the price level to an unanticipated increase in aggregate demand; cumulative response of the price level to an unanticipated increase in oil-specific demand; cumulative response of real GDP to an unanticipated increase in oil-specific demand; cumulative response of real GDP to an unanticipated increase in aggregate demand; cumulative response of real GDP to an unanticipated disruption in crude-oil supply. Solid lines are based on the point estimates from the model in (4). Shaded areas show one- and two-standard-error bands constructed using a block bootstrap procedure. The estimated structural shocks used to produce the responses are based on data from February 1974 to January 2025.

quarters, real GDP falls below its initial level. The same unanticipated expansion causes the price level to rise for the first 7 quarters after the shock. The response is statistically significant using the one-standard-error bands. This pattern is consistent with the increase in global aggregate demand causing an initial increase in US growth, but the growth is eventually hampered by higher prices of global commodities, including the price of oil, causing real GDP to fall. [Kilian \(2009a\)](#) provides further discussion.

The third row of figure 5 shows that an unanticipated increase in oil-specific demand causes real GDP and the price level to rise. The increase in the price level is highly statistically significant. Eventually the higher prices cause real GDP to fall. Again, [Kilian \(2009a\)](#) provides further discussion.

In summary, we are able to replicate [Kilian's \(2009a\)](#) substantive results. Patterns uncovered by the historical decomposition and the second-stage impulse response functions hold using data updated through January 2025.

### *3 Updated responses of oil production, real economic activity, and the real price of oil to demand and supply shocks*

Section 2.3 provides evidence that conclusions about the determinants of the real price of crude oil apply through January 2025; namely, disruptions to oil supply contribute little to fluctuations in price, whereas oil-specific demand shocks are responsible for the ups and downs and aggregate-demand shocks cause “long swings” in price ([Kilian, 2009a](#), 1062). Section 2.4 provides evidence that the US economy responds similarly to changes in the oil market over the later period.

In this section we use updated data to estimate the impulse response functions. We follow [Kilian & Lütkepohl \(2017\)](#) and construct confidence intervals using the moving-block-bootstrap procedure introduced by [Brüggemann et al. \(2016\)](#). [Kilian & Lütkepohl \(2017, pp 415–416\)](#) use this procedure to assess [Kilian's \(2009a\)](#) conclusions. (We are doing nothing new—we are using the same inference procedures, only inference is being carried out over a different sample period.) [Brüggemann et al.'s \(2016\)](#) recursive-design residual-block-bootstrap procedure is robust to conditional heteroskedasticity. In addition, the procedure is designed to replicate the asymptotic distribution of both the slope parameters and the variance–covariance matrix. In effect, the wild bootstrap “produces confidence intervals that are too narrow and have coverage accuracy that is too low asymptotically” ([Kilian & Lütkepohl, 2017, 346–347](#)). An excellent discussion of the issue is provided by [Kilian & Lütkepohl \(2017\)](#). To implement the moving-block bootstrap procedure, a block length of 36 months is chosen.

Figure 6 shows the responses of production, real economic activity, and the real price of oil to one-standard-deviation structural innovations. When comparing figure 6 to figure 2, there are three differences to keep in mind. First, the structural VAR model in (1) is estimated using data retrieved by our code from the sources listed in the text. Second, the data cover the period February 1974 to January 2025. Third, the confidence intervals are constructed using the procedure introduced by [Brüggemann et al. \(2016\)](#). Nevertheless, results are remarkably similar. An unexpected disruption to oil supply causes an immediate decline in oil production and the decline persists. The unexpected

fall in oil production initially has no effect on bulk-dry-cargo shipping rates relative to trend, but this measure of real economic activity picks up after 4 months. The rise could reflect higher fuel costs or expected higher fuel costs (relative to the US consumer price index, which is used to deflate the series). For example, the next panel shows that the unexpected disruption to supply causes an even more pronounced rise in the price of oil. The rise in the price of oil is statistically significant at all horizons based on the one-standard-error bands.

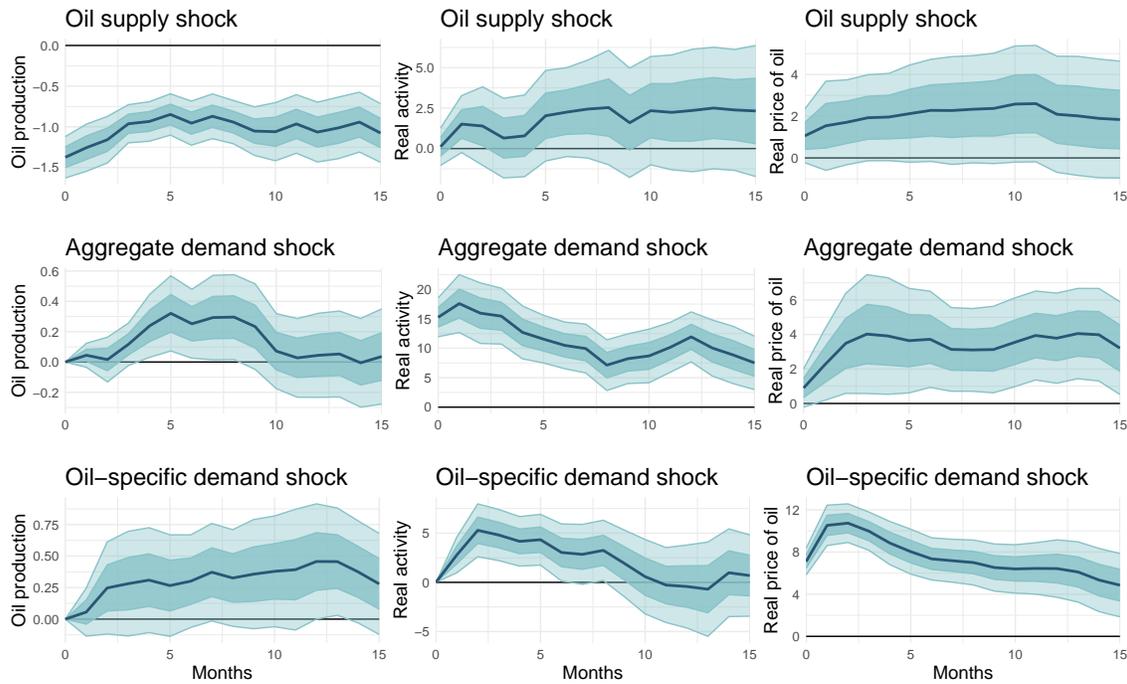


Figure 6: Responses to one-standard-deviation structural shocks

Note: Solid lines show the point estimates and the shaded regions show one- and two-standard-error bands. Estimates are based on the structural VAR model in (1), using the updated data (February 1974–January 2025). Confidence intervals were constructed using a residual block bootstrap procedure (Brüggemann et al., 2016).

Turning to the second row of figure 6, over the entire sample, oil production seems to respond more immediately to an unanticipated increase in aggregate demand, although a months-long delay remains. The unanticipated rise in aggregate demand causes a large, significant jump in global real economic activity, but the effect is less persistent compared to the effect estimated over only the earlier period. And, compared to only the earlier period, there is a more pronounced rise in the real price of oil in response to an unanticipated increase in aggregate demand. Nevertheless, the patterns are overwhelmingly consistent with Kilian’s (2009a) substantive conclusions.

The bottom row of figure 6 shows that an unanticipated rise in oil-specific demand causes an immediate rise in the real price of oil. Again there is evidence of a precautionary-demand motive. Like over the earlier period, the unanticipated rise in oil-specific demand is also associated with an

increase in real economic activity and a rise in oil production. Perhaps, unlike in the earlier period, oil production responds more to surprises in oil-specific demand, although there is a delay of two months and the effect is significant using the one-standard-error bands only.

In summary, the substantive conclusions of [Kilian \(2009a\)](#) hold when all available data are used and confidence sets are constructed using a method robust to potential conditional heteroskedasticity.

#### *4 Application to the clean-energy transition*

Much previous interest has centered around the role that oil-supply and -demand shocks have on macroeconomic outcomes (for example, [Kilian, 2008, 2009a](#), looked at macroeconomic outcomes in response to oil-demand and oil-supply shocks). Another compelling reason to study these types of shocks is understanding how local economies depend on oil.

Certain local economies may be especially hard hit as the US transitions away from fossil fuels to sources of clean energy. Previous structural transitions—such as the transition away from coal and the transition to work in the United States being replaced by machines and jobs elsewhere—are causes for concern. Workers displaced from jobs that are no longer performed are especially hard-hit. And concern extends to the communities in which these workers live ([Hanson, 2022](#)).<sup>4</sup>

In California, policymakers have adopted the goal of carbon neutrality by 2045 or earlier ([CARB, 2022](#)). Within California, Kern County accounts for 70 percent of all oil produced. As California transitions away from oil, major concerns for Kern include lost jobs and lost tax revenue. Many workers there are primarily engaged in the exploration, development, and production of oil. In addition, the assessed value of oil-field properties fluctuates with factors primarily outside the control of Kern policymakers ([Espinoza, 2024](#)). In particular, fluctuations in oil futures prices determine valuations used to assess property taxes. Despite efforts by Kern policymakers to smooth the ups and downs, shocks from the global oil market may propagate to other sectors.

Linkages between labor-market outcomes in oil-producing counties and the oil market have been investigated previously. [Michieka & Gearhart \(2019, 2021\)](#) investigate relationships between nominal regional oil prices and (i) employment and (ii) nominal wages in oil-producing counties. The dynamic-panel models they use are based on the idea that along with the number of time periods getting large, the number of cross-sectional groups is also relatively large. A common “long-run relationship” between regional oil prices and other variables is posited to exist across counties.

---

<sup>4</sup>The forces that restructure the composition of economic activity and the consequences of these forces are not fully understood. Job loss in the coal industry is discussed by [Weber \(2020\)](#). Between 2011 and 2016, the number of people working in mines declined by 45 percent. Each lost coal-mining job lowered county-level income by \$100,000 annually. Retirement of coal-fired power plants in the United States is discussed by [Davis et al. \(2022\)](#) and the negative effects on workers of shutting down coal-fired power plants are discussed by [Hanson \(2022\)](#); [Colmer et al. \(2024\)](#). Work being automated is discussed by [Acemoglu & Autor \(2011\)](#); [Acemoglu & Restrepo \(2020, 2022\)](#). World trade and China’s participation in world trade is discussed by [Autor et al. \(2021, 2016\)](#). The change in composition associated with the hollowing out of middle-pay jobs and when this reallocation takes place over the business cycle is discussed by [Autor & Dorn \(2013\)](#); [Foote & Ryan \(2014\)](#); [Howes \(2022\)](#). [Herrendorf et al. \(2014\)](#) provide a framework for structural change within the tradition of economic growth (see also [Ngai & Pissarides, 2007](#)). Within this framework, [Rogerson \(2015\)](#) provides a way to think about evaluating environmental regulation.

Importantly, they focused on the causal short- and long-run effects of oil prices (consistent with the econometric specification here that assumes oil prices are exogenous with respect to outcomes in Kern). In contrast, we are interested in the impulse response functions, which trace out the responses of labor-market outcomes in Kern to one-time structural innovations in the global oil market. We are also interested in the labor-market dynamics associated with a disruption to supply and a shift in demand, including whether the shift represents a structural innovation to global real economic activity or precautionary demand.<sup>5</sup>

It is nearly impossible to definitively state how much of Kern’s economy depends on oil. Yet, one way to understand the linkages between Kern and the global market for oil is by assessing how Kern’s labor market responds to structural innovations identified by the statistical model in (1), using the second-stage methodology proposed by Kilian (2009a). Specifically, we regress seasonally adjusted labor-market statistics on lags of the estimated structural shocks:

$$k_t = \kappa + \sum_{i=0}^{12} \psi_{ji} \hat{\varepsilon}_{jt-i} + e_{jt}, \quad (5)$$

where  $k_t$  is either the unemployment rate or the monthly growth rate of employment in Kern County, California,  $\hat{\varepsilon}_{jt}$  refers to the estimated residual in (1) for the  $j$ th structural shock in month  $t$ , and  $e_{jt}$  is a potentially serially correlated error. The coefficient  $\psi_{jh}$  is the impulse response at horizon  $h$  of series  $k$  to structural innovation  $j$ . Inference is based on a block-bootstrap procedure, where the block length is chosen to be half a year.

Data on the level of employment and the unemployment rate in Kern County, California, are available from the Local Area Unemployment Statistics program of the Bureau of Labor Statistics starting in January 1990. The data extend through January 2025. Both series are retrieved from FRED and seasonally adjusted using the US Census Bureau’s X-13ARIMA-SETS procedure.

Figure 7 shows the estimated coefficients. The solid, beige lines show the impulse-response coefficients. The shaded areas show the one- and two-standard-error bands. Because employment is entered into the statistical model as the change in log employment, we sum the estimated coefficients from (5) to estimate the cumulative response of employment.

The top, left panel shows that, upon impact, an unanticipated disruption in oil production causes employment to rise. This evidence is consistent with the idea that “oil supply contractions in one region tend to trigger production increases elsewhere in the world” (Kilian, 2009a, 1060–1062). The evidence is also consistent with the notion that an increase in the price of oil represents a windfall boost to income from the perspective of some in Kern’s economy. The windfall may propagate throughout the local economy. Employment rises and, turning to the top, right panel, expanded opportunities increase participation in the labor market, as the unemployment rate rises. The effect on the unemployment rate, however, is not statistically significant based on the one- and two-standard-error bands.

---

<sup>5</sup>Michieka & Gearhart (2015) posit a cointegrating relationship between the nominal price of oil and employment in Kern County and estimate a vector error correction model. The rationale for imposing the cointegrating relationship is statistical rather than economic, which is one reason to examine the sensitivity to this modeling assumption. Michieka & Gearhart (2015) do not investigate impulse responses, which is our interest.

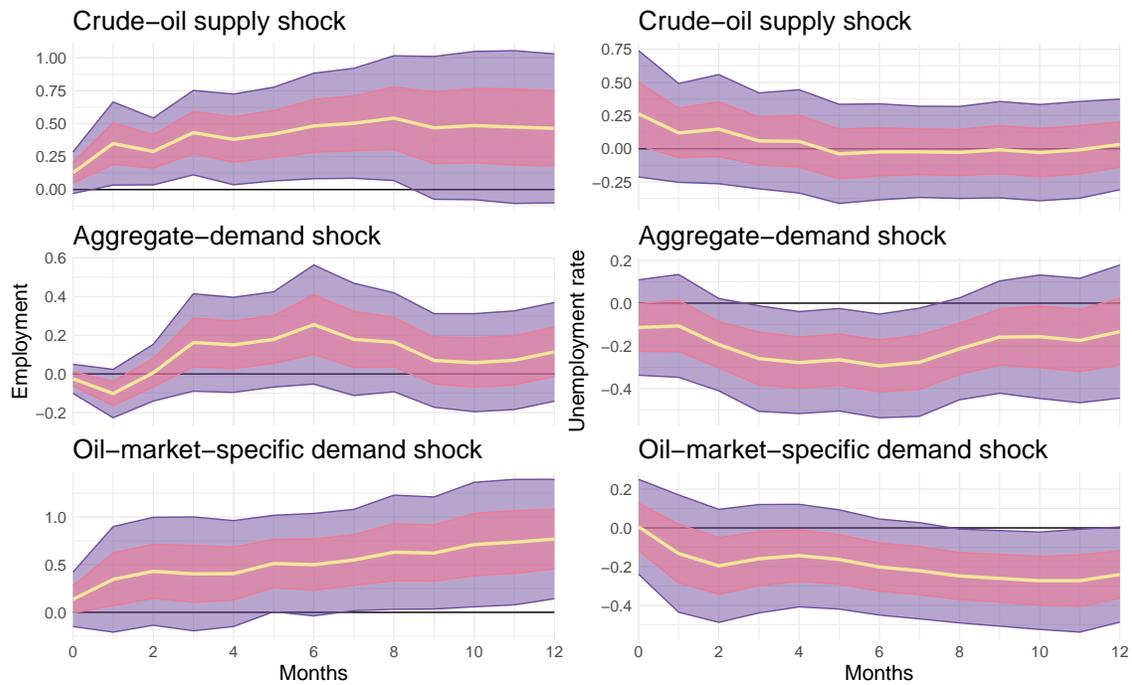


Figure 7: Responses of employment and the unemployment rate in Kern County, California, to identified structural shocks

*Note:* Clockwise from top right, response of the unemployment rate to an unanticipated disruption in oil supply; response of the unemployment rate to an unanticipated increase in aggregate demand; response of the unemployment rate to an unanticipated increase in oil-specific demand; cumulative response of employment to an unanticipated increase in oil-specific demand; cumulative response of employment to an unanticipated increase in aggregate demand; cumulative response of employment to an unanticipated disruption in oil supply. Solid lines are based on the point estimates from the model in (5). Shaded areas show one- and two-standard-error bands constructed using a block bootstrap procedure. The estimated structural shocks used to produce the responses are based on data from February 1974 to January 2025. The regressions that use Kern's unemployment rate begin in January 1990 and extend through January 2025. The regressions that use the change in employment begin one month later.

In contrast, an unanticipated one-standard-deviation increase in aggregate demand causes the unemployment rate to fall on impact. The effect is statistically significant after 2 months based on the two-standard-error bands. The effect on employment, however, is less pronounced. On impact the effect cannot be distinguished from zero. And while employment expands after two months, the effect is statistically significant based on one-standard-error bands only. These conclusions can be seen in the middle panel of figure 7.

The bottom panel shows that an unanticipated increase in oil-specific demand expands employment opportunities in Kern. On impact, employment rises. And continues to rise until by 7 months the effect is statistically different from zero based on the two-standard-error bands. Employment is higher a year after the initial shock. The right panel shows that the unemployment rate falls on impact. The effect is statistically different from zero after around 9 months. These effects are consistent with the notion that oil production in Kern may respond to oil-specific demand and local revenue from oil production will rise. Linkages between oil extraction and the regional economy propagate to expanded opportunities throughout Kern's labor market.

The effects suggest that Kern depends on oil in a meaningful way. As the US economy transitions away from dependence on fossil fuels toward green energy, Kern's economy may be drastically reshaped. What had previously been a major component of regional economic activity may no longer provide opportunities. Policymakers may need to consider expanding options in Kern through, for example, place-based policy. Further research on this topic will be aided by having the option to work within the R ecosystem.

Figure 7 makes clear, though, that Kilian's (2009a) point about identifying the origin of the shock matters. Even though the price of oil may be determined independently of economic activity in Kern County, the dynamic effect of a negative supply disruption differs from the dynamic effect of a positive demand innovation. A negative supply disruption expands labor-market opportunities more rapidly than the expanded opportunities associated with an oil-specific-demand shock. The effects of a positive demand innovation are visible in the data after 1 year.

## 5 Conclusion

In this paper we successfully replicated the substantive conclusions made by Kilian (2009a); namely, not all oil-price shocks are alike. Understanding the origins of such shocks is necessary for policymakers. As Kilian (2009a) points out, a positive structural innovation to the global business cycle will cause national income in the United States to increase—as the United States directly benefits from an increase in global economic activity—but the expansion of US income will be hampered by the rising cost of oil that inevitably follows. A negative innovation to crude-oil production will cause national income in the US to fall and cause the CPI to rise. Both innovations initially cause the price of crude oil to rise, but the dynamic effects differ meaningfully. Policymakers will need to understand the origin of the shock in order to respond appropriately.

Kilian (2009a) should be commended for his excellent code. The original replication package reproduces the results even when we port the code to R (Kilian, 2009b). In fact, the original replication package offers an excellent starting point for further analysis.

We furthered the analysis in several ways:

- We ported the code to R, which will benefit researchers interested in working within the R ecosystem. For many, tidying data in R is much easier than tidying data in MATLAB.
- Our replication and paper are written using the `targets` package. We demonstrate how to use this Make-like tool to ensure reproducibility. The `targets` package can also be used to exploit parallel computing, which could be used to speed the construction of bootstrapped confidence intervals.
- Data from the US Energy Information Administration and the FRED database are retrieved by code. This makes updating the structural shocks straightforward.
- We provided an application of Kilian's (2009a) procedure for identifying structural shocks. These shocks are used learn about how a local labor market responds to the global oil market. Understanding linkages like these may well be more important as society transitions toward sustainable energy sources.

### *Bibliography*

- Acemoglu, D. & Autor, D. (2011).** “Skills, Tasks and Technologies: Implications for Employment and Earnings”. In **Card, D. & Ashenfelter, O.**, editors, “Handbook of Labor Economics”, volume 4, Part B, pages 1043–1171. Elsevier. DOI: [10.1016/S0169-7218\(11\)02410-5](https://doi.org/10.1016/S0169-7218(11)02410-5).
- Acemoglu, D. & Restrepo, P. (2020).** “Robots and Jobs: Evidence from US Labor Markets”. *Journal of Political Economy*, 128(6): 2188–2244. DOI: [10.1086/705716](https://doi.org/10.1086/705716).
- Acemoglu, D. & Restrepo, P. (2022).** “Tasks, Automation, and the Rise in U.S. Wage Inequality”. *Econometrica*, 90(5): 1973–2016. ISSN 0012-9682. DOI: [10.3982/ecta19815](https://doi.org/10.3982/ecta19815).
- Alquist, R. & Kilian, L. (2010).** “What do we learn from the price of crude oil futures?” *Journal of Applied Econometrics*, 25(4): 539–573. ISSN 1099-1255. DOI: [10.1002/jae.1159](https://doi.org/10.1002/jae.1159).
- Autor, D., Dorn, D. & Hanson, G. (2021).** “On the Persistence of the China Shock”. *Brookings Papers on Economic Activity*, Fall: 381–447. DOI: [10.3386/w29401](https://doi.org/10.3386/w29401).
- Autor, D. H. & Dorn, D. (2013).** “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market”. *American Economic Review*, 103(5): 1553–1597. ISSN 0002-8282. DOI: [10.1257/aer.103.5.1553](https://doi.org/10.1257/aer.103.5.1553).
- Autor, D. H., Dorn, D. & Hanson, G. H. (2016).** “The China Shock: Learning from Labor-Market Adjustment to Large Changes in Trade”. *Annual Review of Economics*, 8(1): 205–240. ISSN 1941-1391. DOI: [10.1146/annurev-economics-080315-015041](https://doi.org/10.1146/annurev-economics-080315-015041).
- Barsky, R. B. & Kilian, L. (2001).** “Do We Really Know That Oil Caused the Great Stagflation? A Monetary Alternative”. *NBER Macroeconomics Annual*, 16: 137–183. ISSN 1537-2642. DOI: [10.1086/654439](https://doi.org/10.1086/654439).

- Baumeister, C. & Hamilton, J. D. (2019).** “Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks”. *American Economic Review*, 109(5): 1873–1910. ISSN 0002-8282. DOI: [10.1257/aer.20151569](https://doi.org/10.1257/aer.20151569).
- Baumeister, C. & Kilian, L. (2014).** “What Central Bankers Need to Know About Forecasting Oil Prices”. *International Economic Review*, 55(3): 869–889. ISSN 0020-6598. DOI: [10.1111/iere.12074](https://doi.org/10.1111/iere.12074).
- Brüggemann, R., Jentsch, C. & Trenkler, C. (2016).** “Inference in VARs with conditional heteroskedasticity of unknown form”. *Journal of Econometrics*, 191(1): 69–85. ISSN 0304-4076. DOI: [10.1016/j.jeconom.2015.10.004](https://doi.org/10.1016/j.jeconom.2015.10.004).
- California Air Resources Board (2022).** “2022 Scoping Plan for Achieving Carbon Neutrality”. <https://ww2.arb.ca.gov/our-work/programs/ab-32-climate-change-scoping-plan/2022-scoping-plan-documents> (hereafter cited in text as CARB).
- Colmer, J., Krause, E., Lyubich, E. & Voorheis, J. (2024).** “Transitional Costs and the Decline of Coal: Worker-Level Evidence”. Working Paper. Environmental Inequality Lab. URL: <https://www.census.gov/library/working-papers/2024/adrm/CES-WP-24-53.html>.
- Council of Economic Advisers (2024).** “Economic Report of the President transmitted to congress March 2024 together with the Annual Report of the Council of Economic Advisers”. <https://www.whitehouse.gov/cea/written-materials/2024/03/21/the-2024-economic-report-of-the-president/> (hereafter cited in text as CEA).
- Davis, R. J., Holladay, J. S. & Sims, C. (2022).** “Coal-Fired Power Plant Retirements in the United States”. *Environmental and Energy Policy and the Economy*, 3: 4–36. ISSN 2689-7865. DOI: [10.1086/717217](https://doi.org/10.1086/717217).
- Espinoza, A. X. (2024).** “Annual Comprehensive Financial Report”. Prepared by the Office of Aimee X. Espinoza, Auditor-Controller-County Clerk. <https://www.auditor.co.kern.ca.us/cafr/24CAFR.pdf>.
- Foote, C. L. & Ryan, R. W. (2014).** “Labor-Market Polarization over the Business Cycle”. *NBER Macroeconomics Annual*, 29(1): 371–413. URL: <https://www.nber.org/papers/w21030>.
- Golding, G. & Kilian, L. (2022).** “Don’t Look to Oil Companies to Lower High Retail Gasoline Prices”. Federal Reserve Bank of Dallas. <https://www.dallasfed.org/research/economics/2022/0510>.
- Gonçalves, S. & Kilian, L. (2004).** “Bootstrapping autoregressions with conditional heteroskedasticity of unknown form”. *Journal of Econometrics*, 123(1): 89 – 120. doi:<http://dx.doi.org/10.1016/j.jeconom.2003.10.030>. ISSN 0304-4076. URL <http://www.sciencedirect.com/science/article/pii/S0304407603002847>.
- Hanson, G. (2022).** “Local Labor Market Impacts of the Energy Transition: Prospects and Policies”. In **Kearney, M. S. & Ganz, A.**, editors, “Economic Policy in a More Uncertain World”, pages 156–199. Washington, DC: The Aspen Institute.
- Herrendorf, B., Rogerson, R. & Valentinyi, Á. (2014).** Growth and Structural Transformation, pages 855–941. Elsevier. DOI: [10.1016/b978-0-444-53540-5.00006-9](https://doi.org/10.1016/b978-0-444-53540-5.00006-9).

- Howes, C. (2022).** “Why does structural change accelerate in recessions? The credit reallocation channel”. *Journal of Financial Economics*, 144(3): 933–952. ISSN 0304-405X. DOI: [10.1016/j.jfineco.2021.07.009](https://doi.org/10.1016/j.jfineco.2021.07.009).
- Kilian, L. (2008).** “The Economic Effects of Energy Price Shocks”. *Journal of Economic Literature*, 46(4): 871–909. ISSN 00220515. URL <http://www.jstor.org/stable/27647084>.
- Kilian, L. (2009a).** “Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market”. *American Economic Review*, 99(3): 1053–1069. DOI: [10.1257/aer.99.3.1053](https://doi.org/10.1257/aer.99.3.1053).
- Kilian, L. (2009b).** “Replication data for: Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market”. Nashville, TN: American Economic Association [publisher], 2009. Ann Arbor, MI: Inter-university Consortium for Political and Social Research [distributor], 2019-10-12. <https://doi.org/10.3886/E113299V1>.
- Kilian, L. (2019).** “Measuring global real economic activity: Do recent critiques hold up to scrutiny?” *Economics Letters*, 178: 106–110. DOI: [10.1016/j.econlet.2019.03.001](https://doi.org/10.1016/j.econlet.2019.03.001).
- Kilian, L. (2022).** “Facts and fiction in oil market modeling”. *Energy Economics*, 110: 105973. ISSN 0140-9883. DOI: [10.1016/j.eneco.2022.105973](https://doi.org/10.1016/j.eneco.2022.105973).
- Kilian, L. & Lütkepohl, H. (2017).** *Structural Vector Autoregressive Analysis*. New York, NY: Cambridge University Press. DOI: [10.1017/9781108164818](https://doi.org/10.1017/9781108164818).
- Kilian, L. & Murphy, D. P. (2014).** “The Role of Inventories and Speculative Trading in the Global Market for Crude Oil”. *Journal of Applied Econometrics*, 29(3): 454–478. ISSN 1099-1255. DOI: [10.1002/jae.2322](https://doi.org/10.1002/jae.2322).
- Kilian, L. & Zhou, X. (2018).** “Modeling fluctuations in the global demand for commodities”. *Journal of International Money and Finance*, 88: 54–78. DOI: [10.1016/j.jimonfin.2018.07.001](https://doi.org/10.1016/j.jimonfin.2018.07.001).
- Michieka, N. M. & Gearhart, R. (2015).** “Oil price fluctuations and employment in Kern County: A Vector Error Correction approach”. *Energy Policy*, 87: 584–590. DOI: [10.1016/j.enpol.2015.09.043](https://doi.org/10.1016/j.enpol.2015.09.043).
- Michieka, N. M. & Gearhart, R. S. (2019).** “Oil price dynamics and sectoral employment in the U.S.” *Economic Analysis and Policy*, 62: 140–149. ISSN 0313-5926. DOI: [10.1016/j.eap.2019.02.001](https://doi.org/10.1016/j.eap.2019.02.001).
- Michieka, N. M. & Gearhart, R. S. (2021).** “Oil price changes and wages: a nonlinear and asymmetric approach”. *Economic Change and Restructuring*, 55(1): 1–71. ISSN 1574-0277. DOI: [10.1007/s10644-020-09314-4](https://doi.org/10.1007/s10644-020-09314-4).
- Ngai, L. R. & Pissarides, C. A. (2007).** “Structural Change in a Multisector Model of Growth”. *American Economic Review*, 97(1): 429–443. ISSN 0002-8282. DOI: [10.1257/aer.97.1.429](https://doi.org/10.1257/aer.97.1.429).
- Rogerson, R. (2015).** “A Macroeconomic Perspective on Evaluating Environmental Regulations”. *Review of Environmental Economics and Policy*, 9(2): 219–238. ISSN 1750-6824. DOI: [10.1093/reep/rev005](https://doi.org/10.1093/reep/rev005).

**Weber, J. G. (2020).** “How Should We Think about Environmental Policy and Jobs? An Analogy with Trade Policy and an Illustration from U.S. Coal Mining”. *Review of Environmental Economics and Policy*, 14(1): 44–66. ISSN 1750-6824. DOI: [10.1093/reep/rez016](https://doi.org/10.1093/reep/rez016).

**Zhou, X. (2020).** “Refining the workhorse oil market model”. *Journal of Applied Econometrics*, 35(1): 130–140. ISSN 1099-1255. DOI: [10.1002/jae.2743](https://doi.org/10.1002/jae.2743).

## A Appendix

### A.1 Data

Kilian’s (2009a) structural VAR model includes three variables: (1) the log difference of global crude-oil production, (2) a measure of cyclical variation in global real economic activity, and (3) the log of the real price of oil. Data on production and price are available from the [US Energy Information Administration](#) (EIA). The index of global real economic activity is made available by the Federal Reserve Bank of Dallas through the FRED database. Details about accessing the data are discussed below.

Data on production are available from the EIA under [Geography](#).<sup>6</sup> Data on price are also available from the EIA under the series [Refiner Acquisition Cost of Crude Oil, Imported](#).<sup>7</sup> The EIA website includes links for retrieving the data through a web application programming interface or API. Through the API, it is possible to send structured HTTP requests that return data with JSON format or JavaScript Object Notation format. The code that gathers our data uses the EIA’s API to retrieve the latest data for production and price.

The nominal price series is converted to constant dollars using the Consumer Price Index for All Urban Consumers.<sup>8</sup>

The time series used as an indicator of the global business cycle is the Index of Global Real Economic Activity. The series is available through FRED.<sup>9</sup> It “is derived from a panel of dollar-denominated global bulk dry cargo shipping rates and may be viewed as a proxy for the volume of shipping in global industrial commodity markets.” Kilian (2009a) initially used this measure of the global business cycle. Kilian (2019) reviews the index and Kilian & Zhou (2018, 55) discuss the merits of the series, concluding “that the Kilian (2009a) index and indices based on real commodity prices have clear advantages in modeling industrial commodity markets compared with alternative indicators of global real economic activity, including proxies for global industrial production.”

Figure 8 depicts the three series used in the structural VAR model. All series are monthly. The top panel depicts the change in global crude-oil production. The solid, blue line shows the data

---

<sup>6</sup>Using the [link](#) in the text, the data can be found under “Petroleum and Other Liquids” and “Crude oil including lease condensate.”

<sup>7</sup>Further details are provided in Section 9 of the EIA’s *Monthly Energy Review*.

<sup>8</sup>U.S. Bureau of Labor Statistics, Consumer Price Index for All Urban Consumers: All Items in U.S. City Average [CPIAUCSL], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/CPIAUCSL>.

<sup>9</sup>Federal Reserve Bank of Dallas, Index of Global Real Economic Activity [IGREA], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/IGREA>.

used by Kilian (2009a), and the dashed, red line shows data retrieved using the code for this paper (see Kilian, 2009b, for the original data and code). The original series begin in February 1973 and go through December 2007. The updated series used for the replication begin in February 1974 and go through January 2025. Kilian (2009a) extends the series through 1973 using a procedure described by Barsky & Kilian (2001), a procedure that is not undertaken here. The two series largely coincide. The striped effect indicates that the series overlap.

The middle panel in figure 8 shows the time series of global real economic activity. Again, the solid, blue line depicts the original data and the dashed, red line depicts the updated data. A few years before 2005, the two series diverge. This discrepancy has to do with a coding error in the original analysis described by Kilian (2019).

The bottom panel in figure 8 shows the time series for the real price of crude oil. The solid, blue line shows the series used by Kilian (2009a). The dashed, red line shows data retrieved by our code. The discrepancy between the series is accounted for by demeaning over different periods. The dotted, green line shows our updated data used in the replication demeaned over the period that overlaps with the original data. The striped effect indicates the two series agree.

As Kilian (2009a) demonstrates, the three time series in figure 8 are key determinants of the global market for crude oil. These are the three series used to estimate the structural VAR in the paper.

## *A.2 A sensitivity analysis that uses an index of global industrial production*

Kilian (2009a) uses a measure of global shipping rates as an indicator of global real economic activity, which reflects global demand for all industrial commodities (Kilian, 2019, provides an update). We focus on Kilian's (2009a) index in the text for two reasons. First, it was used in Kilian's (2009a) original specification of the oil market and replicating that analysis was our goal. Second, Kilian & Zhou (2018) make a persuasive case that Kilian's (2009a) index is the appropriate index to use in this application: Kilian's (2009a) index has "the advantage of being [a leading indicator] with respect to changes in latent global real output and of accounting for the role of expectations about the business cycle" (Kilian & Zhou, 2018, 75).

An alternative indicator of global real economic activity is the OECD+6 index of industrial production. This is an index of industrial production for OECD (Organization for Economic Cooperation and Development) countries and, starting in 2006, six non-member countries (Brazil, China, India, Indonesia, the Russian Federation, and South Africa). It serves as a proxy for global industrial production. Kilian & Zhou (2018) point out that it is a coincident indicator (as opposed to a leading indicator) and it fails to account for Chinese demand prior to 2006. The OECD+6 industrial production index was first explored by Baumeister & Kilian (2014) and has been used by Baumeister & Hamilton (2019). For further discussion see Kilian (2022).

We follow Zhou (2020), who conducts a similar sensitivity analysis for a refinement of Kilian & Murphy's (2014) model of the global oil market, and replace Kilian's (2009a) index with the first difference of the log of the OECD+6 index. (Zhou [2020] uses the log-linearly detrended OECD+6 index.)

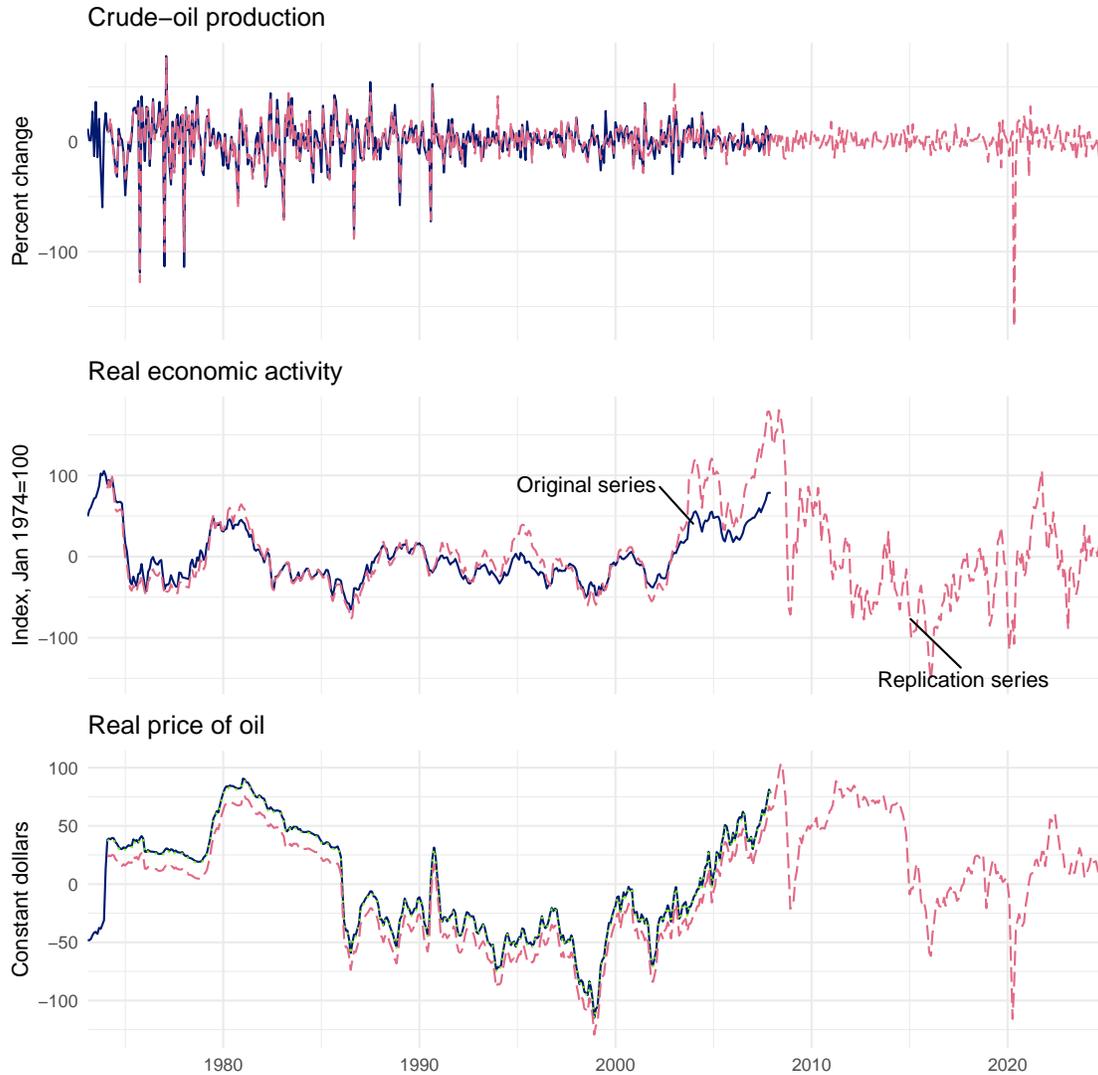


Figure 8: Determinants of the global market for crude oil

Note: *Top*, change in global crude-oil production; *middle*, global real economic activity measured using the index proposed by Kilian (2009a) and revised by Kilian (2019); *bottom*, real price of oil measured by the acquisition cost of refiners and deflated using the Consumer Price Index for All Urban Consumers. The solid, blue lines show the data originally used by Kilian (2009a). The dashed, red lines show the updated data. The dotted, green line in the bottom panel shows the updated series for the real price of oil demeaned over the period February 1974 to December 2007, which overlaps the original sample period after 1974. The striped effect indicates the two series coincide. All data are monthly. The updated data begin in February 1974 and go through January 2025. The original data begin one year earlier.

Sources: US Energy Information Administration, Federal Reserve Bank of Dallas, and US Bureau of Labor Statistics.

For our analysis, we use the index published on Christiane Baumeister’s website.<sup>10</sup>

### A.2.1 Impulse responses based on the structural VAR

The structural VAR representation of this model is the same as in (1) with  $y_t$  replaced by  $\tilde{y}_t = (\Delta \text{oil production}, \Delta \text{OECD+6 industrial production}, \text{real price of oil})'$ . The model is estimated using the same timing assumptions as the ones specified in (3). Figure 9 shows the structural impulse response functions, which are the responses of elements of  $\tilde{y}_t$  to one-time impulses in structural innovations. Because industrial production is measured as the difference of logs, all responses for this variable are summed to approximate the cumulative percent change. Again, the innovations are normalized so that each innovation will tend to raise the price of oil.

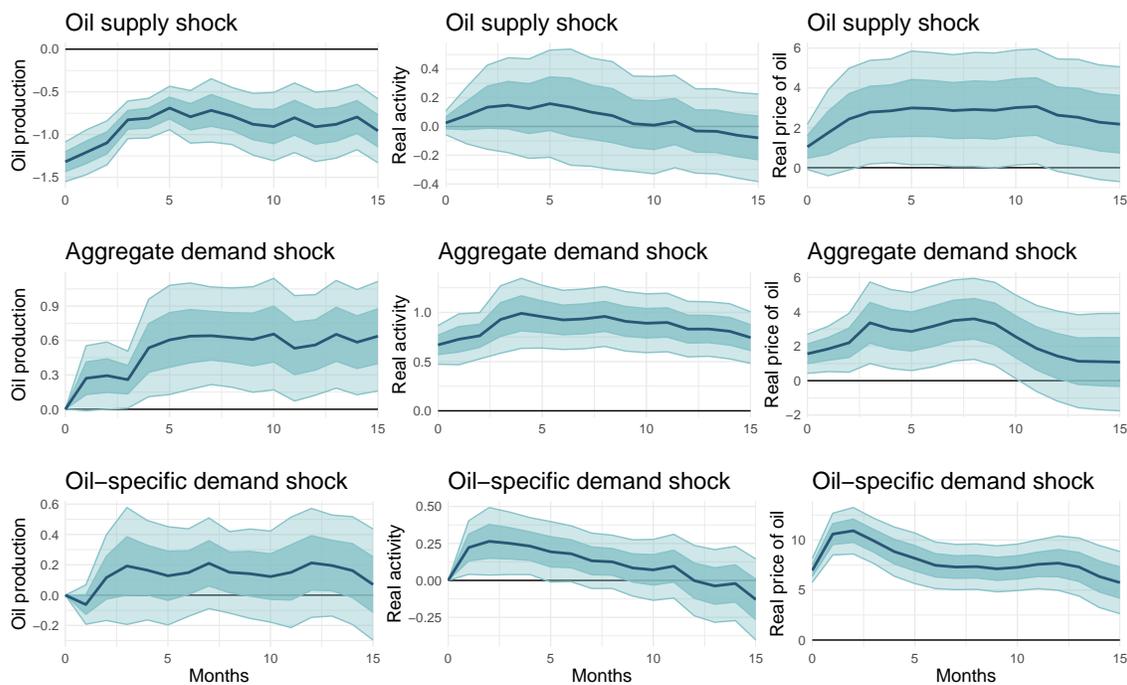


Figure 9: Responses of one-standard-deviation structural shocks (based on an index of industrial production)

Note: Solid lines show the point estimates and the shaded regions show one- and two-standard-error bands. Estimates are based on the structural VAR model in (1), using the updated data (February 1974–January 2025). Confidence intervals were constructed using a residual block bootstrap procedure (Brüggemann et al., 2016).

Figure 9, which is based on OECD+6 industrial production, is comparable to figure 6. In brief, the broad patterns are remarkably similar. Our results are in line with the results reported by Zhou (2020).

<sup>10</sup>Which is accessible at <https://sites.google.com/site/cjsbaumeister/datasets>.

### A.2.2 Historical decomposition

Given that the impulse response functions are very similar, it is not surprising that the historical decomposition is similar as well. Figure 10 shows the cumulative contribution of oil-supply, aggregate-demand, and oil-specific-demand shocks to the real price of oil in separate panels. The contributions based on the industrial-production index are shown using solid, blue lines. The contributions shown with dotted, purple lines are based on Kilian’s (2009a) index. (These are copied from figure 3.)

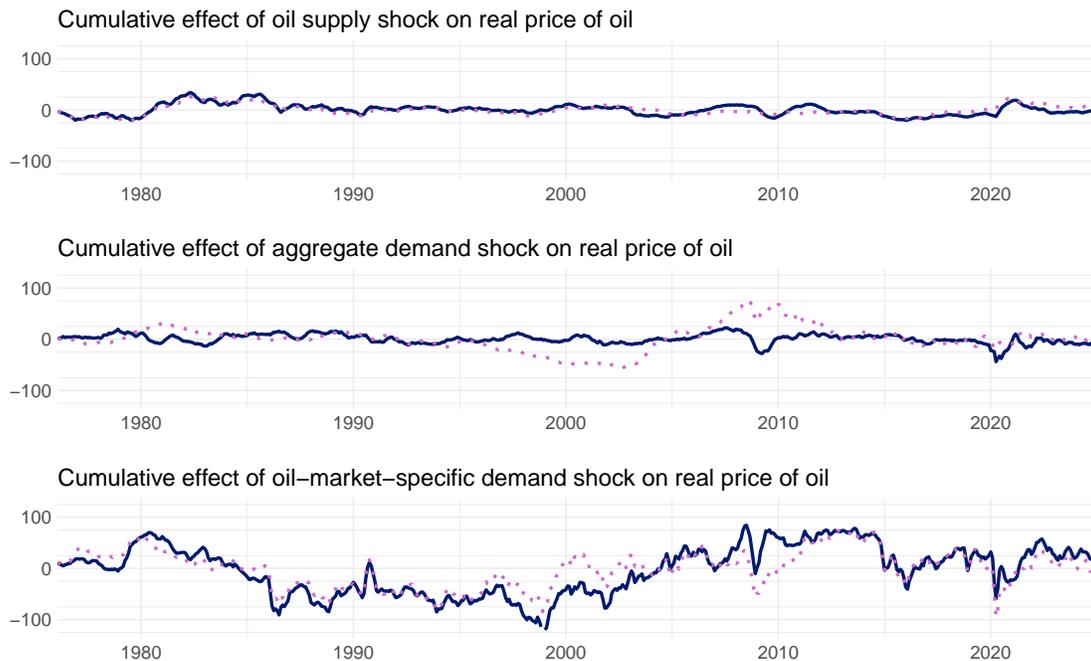


Figure 10: Historical decomposition of the real price of oil

Note: *Top*, how much oil-supply shocks explain of the historically observed fluctuations in the real price of crude oil; *middle*, how much aggregate-demand shocks explain of the historically observed fluctuations in the real price of crude oil; *bottom*, how much oil-specific-demand shocks explain of the historically observed fluctuations in the real price of crude oil. Estimates are based on the structural VAR model in (1), using data from February 1974 to January 2025. The solid, blue lines are based on the growth rate of industrial production instead of using the index of global real economic activity described in the main text (Kilian, 2009a, 2019). The dotted, purple lines correspond to statistics from figure 3, which are based on Kilian’s (2009a) index of global real economic activity.

The cumulative contributions—whether based on Kilian’s (2009a) index or the index of industrial production—are broadly similar. Some differences are expected, however. As Zhou (2020) explains, the timing of the two indicators of global real economic activity differ. Kilian’s (2009a) index is “a leading indicator with respect to global real output;” whereas, the industrial-production index is a coincident indicator of global real output and therefore “[misidentifies] the timing of shifts in aggregate demand in commodity markets” (Kilian & Zhou, 2018, 58). Regardless, insights

made by [Kilian \(2009a\)](#) continue to hold: oil-supply shocks have made small contributions to the real price of oil and oil-specific-demand shocks have contributed the ups and downs to the real price of oil.

### *A.2.3 Responses in Kern County, California, to the global oil market*

We also investigated the robustness of the second-stage regressions that investigate the responses of labor-market outcomes in Kern County, California. The patterns based on shocks estimated using industrial production are similar to those shown in figure 7. We therefore skip reporting the results here and, instead, report the figure that corresponds to figure 7 in our replication materials.