Technical Appendix
“Reassessing the U.S. Economy’s Vulnerability to Oil Shocks”

by

Danilo Leiva-León, Giovanni Olivei, Ara Patvakanian, Jenny Tang, Egon Zakrajšek

Hamilton oil shocks: Oil prices fluctuate continuously in response to various economic forces, but not all price movements constitute meaningful economic shocks. To distinguish routine price fluctuations from potentially disruptive events, we employ the methodology developed by Hamilton (1996, 2003), which identifies oil price *increases* that are both large and unexpected relative to recent experience.

The intuition behind Hamilton’s approach is straightforward. Economic agents—households, firms, and policymakers—form expectations about oil prices based on recent observations. When prices rise but remain within the range experienced over the preceding months, agents can reasonably anticipate such movements and incorporate them into their planning. However, when prices surge beyond any level seen in recent memory, such increases are more likely to catch agents off-guard, disrupting existing plans and potentially forcing costly economic adjustments. Hamilton’s method captures this intuition by comparing current prices to a backward-looking reference price.

Formally, letting P_t denote the average real oil price in month t, the l-month look-back Hamilton oil price shock is defined as

$$\text{Shock}_t = 100 \times \max\{0, \ln P_t - \ln P_t^*\},$$

where the corresponding reference price P_t^* is given by

$$P_t^* = \max\{P_{t-1}, P_{t-2}, \ldots, P_{t-l}\}.$$

The reference price P_t^* represents the highest oil price observed over the previous l months. A Hamilton shock occurs when the current price P_t exceeds this reference price—that is, when oil prices reach levels not seen in the prior l months. The magnitude of the shock is given by the percentage increase above the reference price (that is, $100 \times (\ln P_t - \ln P_t^*)$). If the current price does not exceed the reference price, no shock is registered (that is, the shock equals zero).

Figure A-1 illustrates the construction of Hamilton shocks since 2020, using a 12-month look-back window ($l = 12$), our baseline specification. The blue line traces the actual real price of West Texas Intermediate (WTI) crude oil, while the yellow line shows the corresponding reference price—the maximum price observed over the preceding 12 months. The vertical purple bars indicate moments when Hamilton shocks occur, with the height of each bar representing the shock’s magnitude.

Several features of this approach merit emphasis. First, the method is entirely backward-looking and does not require assumptions about future price paths or structural models of the oil market. Second, by design,

Figure A-1: The Hamilton filter

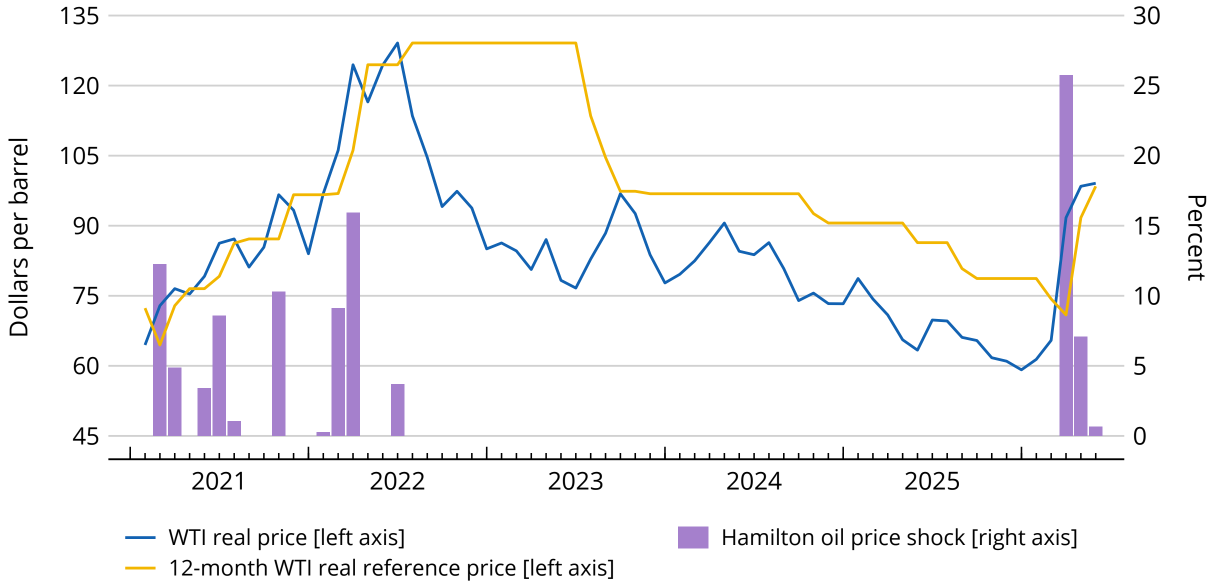

Note(s): The blue line shows the actual average monthly price of a barrel of WTI crude oil in May 2026 dollars (P_t)*, and the yellow line shows the corresponding 12-month look-back Hamilton reference price* (P_t^*)*. The vertical purple bars show the associated Hamilton oil price shocks—a shock is registered when the actual WTI price exceeds the reference price.*
Source(s): Authors' calculations using data from the CME Group and the U.S. Energy Information Administration via Haver Analytics.

Hamilton shocks identify price increases that break through recent trading ranges rather than mere returns to previous price levels. This characteristic makes the measure particularly well suited to capturing surprises that might trigger economic adjustment. Third, the choice of look-back window l involves a tradeoff. Shorter windows identify more frequent but potentially less significant price increases, while longer windows flag only the most dramatic departures from recent experience. For our baseline, we use a 12-month window, which balances these considerations and aligns with typical planning horizons for many economic decisions.

Fourth, and importantly, Hamilton shocks are asymmetric by construction—they capture only price increases, not decreases. This asymmetry reflects substantial empirical evidence that oil price increases have larger macroeconomic effects than comparable decreases, a phenomenon often attributed to adjustment costs, nominal rigidities, and sectoral reallocation frictions that bind more severely when prices rise than when they fall.

Oil shocks and state-level economic outcomes: To gauge the impact of oil shocks on state-level economic outcomes, we estimate the following panel local projection specification:

$$\nabla_h Y_{s,t+h} = \beta_1^{(h)}[\text{Shock}_t \times \text{OilShr}_{s,t-12}] + \beta_2^{(h)}\text{OilShr}_{s,t-12} + \sum_{j=1}^{12} \mathbf{X}'_{s,t-j}\boldsymbol{\gamma}_j^{(h)} + \eta_s^{(h)} + \lambda_t^{(h)} + \varepsilon_{s,t+h}^{(h)}, \quad (1)$$

where $s = 1, 2, \ldots, 50$ indexes states, $t = 1997\text{:M1}, 1997\text{:M2}, \ldots, 2024\text{:M12}$ indexes time, and the dependent variable

$$\nabla_h Y_{s,t+h} \equiv \frac{1200}{h+1} \times \ln\left(\frac{Y_{s,t+h}}{Y_{s,t-1}}\right), \quad h \geq 0,$$

denotes the annualized growth rate from month $t-1$ to month $t+h$ in the specified state-level outcome variable $Y_{s,t}$, that is, private nonfarm employment, private nonfarm employment excluding mining and logging, and house prices (Panels A through C of Figure 5 in the main text, respectively). All state-level outcome variables are seasonally adjusted. To minimize the effect of extreme observations on our estimates, we winsorize monthly growth rates of all outcome variables (that is, $\ln Y_{s,t} - \ln Y_{s,t-1}$) at the 0.5th and 99.5th percentiles before constructing the annualized growth rates for different h horizons.

The primary variables of interest are Shock_t, the Hamilton oil price shock in month t, and $\text{OilShr}_{s,t-12}$, the average share of gross state product (GSP) accounted for by the oil and gas extraction sector in state s from month $t-24$ to month $t-12$, that is, 12 months before the shock. Specifically, $\text{OilShr}_{s,t}$ is defined as the nominal gross output of the oil and gas extraction industry (NAICS 211) in state s, divided by that state's total nominal GSP. This variable captures the state's economic exposure/sensitivity to oil prices through domestic production. Because industry and aggregate GSP data are available only at annual frequency, we convert the annual shares to monthly frequency by assuming the same annual value for all 12 months of the year.

We use a 12-month lag for the oil share variable ($\text{OilShr}_{s,t-12}$) in specification (1) for two reasons. First, as noted previously, industry-level and aggregate GSP data are available only at an annual frequency, and we convert them to monthly frequency by assigning the same annual value to all 12 months of the year. The 12-month lag ensures that the oil share measure reflects data that would have been observed and reported before the shock occurs, avoiding any mechanical overlap between the measurement period of the share and the post-shock outcome period. Second, and more importantly from an econometric standpoint, the 12-month lag ensures that $\text{OilShr}_{s,t-12}$ is strictly predetermined with respect to the shock in month t and therefore cannot be influenced by contemporaneous or forward-looking responses to the shock itself. This predetermined nature of the oil share measure mitigates potential endogeneity concerns and ensures that our interaction term $[\text{Shock}_t \times \text{OilShr}_{s,t-12}]$ captures the differential impact of oil price shocks across states with different preexisting production intensities, rather than reflecting any simultaneous determination of both the shock and the production structure.

The coefficient of interest in specification (1) is $\beta_1^{(h)}$ on the interaction term $[\text{Shock}_t \times \text{OilShr}_{s,t-1}]$, which captures the differential effect of an oil price shock on state s depending on its oil production intensity. A positive estimate of $\beta_1^{(h)}$ indicates that states with higher oil and gas extraction shares experience relatively stronger growth in the outcome variable following an oil shock. The interaction term specification allows us to estimate how the marginal effect of an oil shock varies continuously with states' oil production intensity, rather than imposing arbitrary cutoffs to define "oil-producing" versus "non-oil-producing" states.

The vector $\mathbf{X}_{s,t-j}$ $(j = 1, 2, \ldots, 12)$ includes 12 lags of relevant state-level variables to control for economic conditions prior to the shock. When the outcome variable $Y_{s,t}$ is private nonfarm employment, $\mathbf{X}_{s,t-j}$ comprises the log-level of private nonfarm employment and the log-level of house prices; when $Y_{s,t}$ is private nonfarm employment excluding mining and logging, $\mathbf{X}_{s,t-j}$ comprises the log-level of private nonfarm employment excluding mining and logging and the log-level of house prices; and when $Y_{s,t}$ is house prices, $\mathbf{X}_{s,t-j}$ comprises the log-level of house prices and the log-level of private nonfarm employment.

Importantly, specification (1) also includes state fixed effects ($\eta_s^{(h)}$) and time fixed effects ($\lambda_t^{(h)}$). The state fixed effects control for time-invariant state characteristics—such as geographic, demographic, and institutional factors—that may influence both economic outcomes and oil production intensity. The time fixed effects absorb common monthly factors that affect all states simultaneously, such as monetary policy, federal fiscal policy, or global economic conditions. This two-way fixed effects structure ensures that we identify $\beta_1^{(h)}$ from cross-sectional variation in oil production intensity interacted with the time series of oil shocks, conditional on state-specific trends and common time effects.

The estimated coefficient $\hat{\beta}_1^{(h)}$ can be interpreted as follows: For a 1 percentage point increase in a state's oil and gas extraction share of GSP, a 1 percent Hamilton oil shock induces a $\hat{\beta}_1^{(h)}$ percentage point differential in the annualized growth rate of the outcome variable at horizon h, relative to a state with zero oil production. Note that our specification estimates only the *relative* effect of oil shocks across states with different oil production intensities, not the absolute effect of oil shocks on any individual state. This is because the time fixed effects $\lambda_t^{(h)}$ absorb all common variation across states in month t, including the average effect of the oil shock itself. Consequently, the coefficient $\beta_1^{(h)}$ captures how states with higher oil production shares respond *differentially* compared with states with lower shares, conditional on the common national response.

To illustrate this interpretation concretely, Figure 5 in the main text presents estimated effects for Massachusetts and Texas relative to a hypothetical state with the average oil production share, rather than relative to a no-shock counterfactual. This relative identification strategy is well suited to our research question—understanding how domestic oil production creates regional offsets that buffer aggregate employment—but it does not allow us to quantify the total effect of oil shocks on individual states' economies.

We estimate specification (1) by ordinary least squares (OLS) separately for each horizon $h \in \{3, 6, 12, 24\}$ months. Standard errors are clustered two-way by state and time to account for potential serial correlation within states over time and cross-sectional correlation across states within each month. This clustering approach produces robust inference in the presence of arbitrary within-state autocorrelation and arbitrary cross-state correlation patterns (see Cameron, Gelbach, and Miller, 2011).

References

Cameron, A. Colin, Jonas B. Gelbach, and Douglas L. Miller. 2011. "Robust inference with multi-way clustering." *Journal of Business and Economic Statistics* 29 (2): 238–249.

Hamilton, James D. 1996. "This is what happened to the oil price-macroeconomy relationship." *Journal of Monetary Economics* 38 (2): 215–220.

———. 2003. “What is an oil shock?” *Journal of Econometrics* 113 (2): 363–398.