Carbon Intensity, Productivity, and Growth*

Emanuel Moench[a,b,*], Soroosh Soofi-Siavash[c,d]

[a] *Frankfurt School of Finance & Management, Adickesallee 32-34, 60322 Frankfurt am Main, Germany*
[b] *CEPR, United Kingdom*
[c] *Bank of Lithuania, Totorių 4, 01103 Vilnius, Lithuania*
[d] *Vilnius University, Saulėtekio 9, 10222 Vilnius, Lithuania*

January 18, 2022

Abstract

Over the past several decades real GDP per capita in the US has strongly increased while CO2 emissions per capita have seen a secular decline, thus sharply reducing the carbon emission intensity of output. We identify an emission intensity shock as the innovation which explains the maximum share of variation in the ratio of CO2 emissions and real GDP at the horizon of 20 years. While the shock is associated with a persistent decline of emissions per unit of output, per capita CO2 emissions quickly revert back to their initial level. The reason is that output and its components strongly increase in response to the emission intensity shock, providing evidence of a rebound effect in aggregate US carbon emissions. The emission intensity shock triggers a substitution of fossil fuels with electricity that is mainly generated from nuclear energy. A separately identified TFP news shock is highly correlated with the emission intensity shock, and both explain similarly large fractions of GDP at longer horizons. A TFP news shock orthogonalized to the emission intensity shock only explains small fractions of GDP and its components, suggesting that innovations to the energy mix of the US economy have been a key driver of growth in recent decades. While the emission intensity shock explains the bulk of variation of emissions per unit of output, it only explains modest shares of emissions per capita. The remaining variation is driven by shocks related to the demand for oil at business cycle, and to the supply of oil at lower frequencies.

Keywords: carbon emissions, carbon intensity, news shocks, structural vector autoregressions
JEL: C32, O47, Q43, Q55

*The authors would like to thank seminar participants at the Bank of Lithuania. The views expressed here do not reflect the views of the Bank of Lithuania, or the Eurosystem.
*Corresponding author. Email: e.moench@fs.de.

1. Introduction

Climate change induced by anthropogenic emissions of carbon dioxide (CO2) and other greenhouse gases is considered the single most important threat to economic prosperity and well-being worldwide. Only a rapid reduction of carbon emissions can limit the increase of global mean temperatures to 2 degrees Celsius above pre-industrial levels (e.g. Masson-Delmotte et al., 2018). To ensure that this development does not come at the cost of sharply lower real activity, the carbon intensity of global output has to decline substantially. This will require technological change at an unrivaled pace. What impact these technological advances will have on economic growth is an open question.

While history may only be an imperfect guide for the future, our paper tries to shed light on this question using past data. Figure 1 shows that while real GDP per capita in the US has strongly increased over the last several decades, CO2 emissions per capita have seen a secular decline, reflecting a sharp reduction of the carbon emission intensity of output since the 1970s. Prior research suggests several potential explanations for this secular decline. One is that emission-intensive manufacturing goods have been increasingly imported from abroad, thus reducing carbon emissions in the US. Yet, there is at best mixed evidence that changes in the composition of goods produced by US manufacturing firms and industries can account for a reduction in their pollution (see, e.g., Levinson, 2009; Cherniwchan, 2017; Shapiro and Walker, 2018). Another potential explanation is based on environmental regulations which may have induced firms to cut their carbon emissions. However, while regulatory actions such as the Clean Air Act have resulted in a substantial reduction of criteria air pollutants, there is no regulation regarding CO2 emissions at the national level and only few such regulations at the state or municipal level. Moreover, there is little evidence of spillovers from regulations targeting other pollutants (Shapiro and Walker, 2018).[1] A third potential explanation relates to productivity as a driver of reduced carbon emissions. To the extent that productive firms use inputs including fossil fuels more efficiently, an upward trend in productivity may have resulted in a lower carbon intensity of output.

We seek to understand the economic driving forces behind and the aggregate economic effects of the secular decline of carbon intensity. To this end, we estimate a medium-sized Bayesian vector autoregression (VAR) of the US economy since 1973. We then identify an emission intensity shock as the innovation which explains the maximum share of the variation in CO2 emissions per unit of real GDP at the horizon of 20 years, and study its economic consequences. Importantly, our shock identification allows us to take an agnostic stance as to which of the above explanations drives changes in the carbon intensity.

We find that a one-standard-deviation emission intensity shock is associated with a permanent

[1] In recent years, a few environmental policies targeting CO2 emissions have been passed in different US states and counties: the Regional Greenhouse Gas Initiative (RGGI)'s cap-and-trade system of 2005 that limits CO2 emissions from the power sector in California and the Northeast states; carbon taxes introduced in Boulder, Colorado, the San Francisco Bay Area, and Montgomery County, Maryland in 2006, 2008, and 2010.

decline of about one percent of emissions per unit of output. Despite this long-lived reduction of the emission intensity, however, per capita CO2 emissions quickly revert back to their initial level. The reason is that output and its components strongly increase following an emission intensity shock. Hence, our results suggest a marked rebound effect in aggregate carbon emissions in the US, consistent with evidence of a sizable rebound effect of economy-wide energy consumption documented e.g. by Brockway, Sorrell, Semieniuk, Heun and Court (2021). Studying the energy mix of US output, we see that the emission intensity shock triggers a substitution of fossil fuels by electricity. All end-use sectors reduce their consumption of primary fossil energy sources but increase their consumption of electricity, which is mainly generated from nuclear energy. The emission intensity shock not only gives rise to a strong increase in real GDP and its main components, it is also associated with a delayed but pronounced increase in TFP and the inverse relative price of investment goods. Moreover, stock prices rise strongly and persistently in response to the shock, and inflation drops. This suggests a close link of the emission intensity shock with news about future productivity.

To analyze this link, we separately study shocks reflecting news about future TFP. We follow Kurmann and Sims (2021) and identify a TFP news shock as the innovation explaining a maximum share of TFP at the same horizon of 20 years. We find that this shock gives rise to very similar impulse responses as the emission intensity shock, with the two shocks being about 80 percent correlated. There are at least two potential explanations for this close connection. On the one hand, more productive firms could use inputs more efficiently, and thus broad-based improvements of productivity as captured by TFP may have resulted in a reduction of the carbon intensity of output. On the other hand, firms facing higher fossil fuel prices might innovate more in energy-efficient technologies. Therefore, technological advances related to the efficient use of energy could act as a key driving force of total factor productivity and thus economic growth. While our aggregate data do not allow us to fully discriminate between these two hypotheses, we separately identify a TFP news shock that is contemporaneously orthogonal to the emission intensity shock. This innovation thus reflects news about future TFP which have no direct impact on the emission intensity of output. We find that the real economic effects of this orthogonalized TFP news shock are modest at best, and only result in short-lived increases of output and its components. Moreover, while the orthogonalized shock has a larger impact on TFP in the short run than the TFP news shock, its long-term impact on TFP is negligible. Taken together, these results suggest that innovations related to the energy mix of US output have been a major source of economic growth over the past five decades.

While the emission intensity shock explains the bulk of variation of emissions per unit of output, it only explains modest shares of emissions per capita. This suggests that technological advances are insufficient to substantially reduce carbon emissions in the long-run. To understand the additional drivers of per capita carbon emissions, we follow the methodology of Angeletos, Collard and Dellas (2020) and identify two further shocks which, together with the emission intensity shock, explain essentially all of the variation of carbon emissions at business cycle and lower frequencies.

We find that the innovation driving residual business cycle variation of per capita emissions pushes the quantity and price of oil in the same direction in the short run, and as such seems to capture variations in the demand for oil. The shock which explains the residual lower frequency variation in carbon emissions instead moves the price and quantity of oil in opposite directions, and thus appears to be related to supply disturbances.

Our paper relates to several strands of literature. First, our finding that there is a negative relationship between aggregate productivity and carbon emissions per unit of output finds ample support in studies using firm-level data. For example, Bloom, Genakos, Martin and Sadun (2010) provide evidence that the management practices which are associated with higher productivity tend to use energy more efficiently and emit less CO2. Several studies rely on models in the spirit of Melitz (2003) to understand the link between firm productivity, environmental performance and trade. For example, Shapiro and Walker (2018) use plant-level data on US manufacturing and document a strong negative relation between criteria air pollution per unit of output and total factor productivity. They propose a model of heterogeneous firms with endogenous pollution abatement, in which more productive firms charge lower prices and at the same time increase investments in abatement to avoid pollution tax. Despite the strong negative correlation between pollution per unit of output and productivity documented at the plant level, they find that environmental regulation explains the bulk of variation of air pollutant emissions in the US, while changes in productivity played only a smaller role.[2]

Second, we also connect with work discussing firm-level evidence on the link between energy price changes and the direction of technological change (for an extensive literature review see Popp, Newell and Jaffe, 2010). This literature documents that on the one hand, energy-efficient technologies grew at a substantially higher rate after the oil shocks in the 1970s (Hassler, Krusell and Olovsson, 2021). On the other hand, it documents that firms exposed to other firms which undertake energy-efficient innovations (or have a history of energy-efficient innovations) tend themselves to innovate more in energy-efficient technologies in the future.[3] Taken together, this literature suggests that the technological change directed to energy efficiency has been an increasingly important factor underlying improvement in the economy-wide productivity in recent decades. Consistent with this view, we show that a shock lowering the carbon intensity of real GDP in the US is associ-

[2]Relatedly, Forslid, Okubo and Ulltveit-Moe (2018) build a model of heterogeneous firms with endogenous abatement, in which improved productivity also raises abatement investments and lowers emission intensity. The authors use firm-level data from Sweden and show that more productive firms tend to have less CO2 emissions per unit of output and invest more in abatement. For further US establishment-level data evidence on the negative correlation between firm productivity and emissions per unit of output see for example Cui, Lapan and Moschini (2016) and Holladay (2016). These authors provide evidence that when firms have access to potentially larger foreign markets, they tend to invest more in abatement given their level of productivity. In addition, Cherniwchan (2017) documents that lower trade costs following the introduction of NAFTA reduced the pollution emitted by manufacturing plants in the US during the 1990s.

[3]For evidence on the path dependence in innovations see for example Aghion, Dechezleprêtre, Hémous, Martin and Van Reenen (2016), who use firm-level data on auto industry to separate between dirty and clean patents. The authors also discussed evidence that firms facing higher tax-inclusive fuel prices tend to innovate more in clean energy.

ated with a strong and persistent increase in TFP and real output.

Third, our identification closely follows the method proposed by Kurmann and Sims (2021) who identify a news shock about future TFP as an innovation that exerts a maximum contribution to TFP in the long run, without any short-term timing constraints. This approach is consistent with the idea that new technologies spread slowly and imply predictable changes in productivity in the long run. In line with this view, several authors have provided evidence that US firms and real estate investors do not immediately use new energy efficient and cost effective technologies, such as insulation and low energy lighting, but only adopt them slowly over time.[4]

Finally, a closely related paper is by Khan, Metaxoglou, Knittel and Papineau (2019), who estimate structural VARs on US macroeconomic and emission data. They provide evidence that an investment technology shock, which leads to higher efficiency of newly produced investment goods, is a relatively important driver of per capita carbon emissions, accounting for about 25 percent of its variation. The authors find that the investment-specific technology news shock explains the largest share of the variations in per capita carbon emissions. In contrast to these authors, we start with identifying a shock driving long-run variations in the carbon emission intensity of US output, and then document its properties. Consistent with Khan, Metaxoglou, Knittel and Papineau (2019) we find that the emission intensity shock, although it explains a large fraction of the variation in carbon intensity in the long-run, captures less than one-third of the variation in per capita carbon emissions itself. We go beyond their analysis by identifying the additional drivers of per capita carbon emissions.

The remainder of this paper is organized as follows. Section 2 summarizes the econometric methodology used to identify the various shocks. Section 3 briefly reviews the data and provides details on the specification of the VAR model. In Section 4 we then present the results of our analysis and provide some robustness checks. Section 5 concludes.

2. Econometric Methodology

This section presents our econometric methodology. We first describe our identification of the emission intensity shock in a structural VAR. Our identification involves constructing a linear combination of VAR innovations that explains a maximum share of future variation of emission intensity (CO2 emissions per unit of real GDP) in the long run. The method closely follows the approach

[4]For example, tracing the spread of green building practices across US metropolitan areas Kok, McGraw and Quigley (2011) show that the fraction of energy efficient and sustainable building spaces follow the well-known S-shape diffusion curve (Griliches, 1957). This implies that it rises slowly at first, followed by rapid adoption and then slower growth once most potential owners and investors adopted the technology. A large literature has looked at the the slow diffusion of energy-efficient technologies (see the review in Popp, Newell and Jaffe, 2010). This literature is pointing to informational and principal-agent problems as one of the main explanations. For example, using data on attic insulation in the US, Hassett and Metcalf (1999) find that the realized return savings are statistically significant but less than the returns promised by engineers and product manufacturers. Principal-agent frictions arise within firms where one department or individual chooses the level of investment in energy-efficient technologies, but another one pays the energy bills. If the incentives are not aligned, the firm might decide not to make the investment.

of Uhlig (2003) which has been widely used for identification of long-run technology shocks in structural VARs (Barsky and Sims, 2011; Kurmann and Sims, 2021; Francis, Owyang, Roush and DiCecio, 2014). To the best of our knowledge, we are the first to apply this approach to the emission intensity of US output as the target variable.

We will show that the identified emission intensity shock captures the bulk of variation in CO2 emissions per unit of output, but only explains a modest share of per capita CO2 emissions. To analyze the remaining variation in carbon emissions, we identify two further shocks which explicitly capture variation at the business-cycle and medium- to longer-run frequencies. Here, we closely follow the approach of Angeletos, Collard and Dellas (2020). These authors also build on Uhlig (2003), but instead of targeting the variation at a specific horizon in the time domain, they maximize the variation of some target variable in the frequency domain.

Let X_t denote an $n \times 1$ vector of quarterly time series, containing the emission intensity and CO2 emissions, together with a number of additional variables capturing macroeconomic and financial aggregates of the US economy. The time series jointly follow a VAR representation

$$A(L)X_t = \eta_t, \tag{1}$$

where $A(L) = I - A_1 L - \cdots - A_p L^p$ is a lag polynomial matrix, and η_t is the vector of VAR innovations with mean zero and variance-covariance matrix Σ_η. From Equation (1), one obtains the reduced-form moving average representation which expresses X_t in terms of current and past values of innovations:

$$X_t = A(L)^{-1}\eta_t. \tag{2}$$

We assume that the innovations η_t summarizing the joint dynamics among the variables in X_t are linear combinations of structural shocks, denoted by the $n \times 1$ vector ν_t

$$\eta_t = H\nu_t. \tag{3}$$

The structural shocks ν_t have the variance-covariance matrix Σ_ν. Under the unit standard deviation normalization ($\Sigma_\nu = I$), one can write any matrix H as $H = Chol(\Sigma_\eta)Q$ where Q is a $n \times n$ orthonormal matrix ($Q'Q = I$), and $Chol$ denotes the Cholesky factorization. This implies the structural moving average representation

$$X_t = B(L)Q\nu_t, \text{ with } B(L) = A(L)^{-1}Chol(\Sigma_\eta), \tag{4}$$

where the impulse response function of X_t with respect to the ith shock is given by $B(L)Q_i$ with Q_i the ith column of Q. Any potential mapping from the structural shocks ν_t to the innovations η_t can thus be captured by a choice of the matrix Q.

2.1. Time-Domain Identification of Emission Intensity Shock

We can now implement our identification of an emission intensity shock as follows. Let the emission intensity shock be indexed by 1, and let the ratio of CO2 emissions over real GDP in X_t be indexed by j. Furthermore, let $B_k^{[j]}$ denote the jth row of the kth lag matrix in $B(L)$ such that $B_k^{[j]}Q_1$ is the effect of the first shock on the ratio of CO2 emissions over real GDP after k periods, where Q_1 denotes the first column of the matrix Q. We achieve identification of the first shock as the emission intensity shock by solving the following optimization problem

$$\underset{Q_1}{\operatorname{argmax}} \frac{\sum_{k=0}^{h-1}\left(B_k^{[j]}Q_1\right)^2}{\sum_{k=0}^{h-1}B_k^{[j]}B_k^{[j]\prime}} \tag{5}$$

subject to the constraint $Q_1'Q_1 = 1$ which ensures that Q_1 is an orthonormal vector. Following Uhlig (2003), the solution of this problem is characterized by the eigenvector associated with the largest eigenvalue of the $n \times n$ matrix $\frac{\sum_{k=0}^{h-1}\left(B_k^{[j]\prime}B_k^{[j]}\right)}{\sum_{k=0}^{h-1}B_k^{[j]}B_k^{[j]\prime}}$.

2.2. Frequency-Domain Identification of CO2 Emissions Shocks

The frequency-domain identification rewrites optimization problem (5) to capture the variation of variable j over a frequency band.

The contribution of the kth shock to the spectral density of X_{jt} over the frequency band $[\underline{\omega}, \overline{\omega}]$ is given by $\int_{\omega\in[\underline{\omega},\overline{\omega}]}\left(\overline{B^{[j]}\left(e^{-i\omega}\right)Q_k}B^{[j]}\left(e^{-i\omega}\right)Q_k\right)$, where $B^{[j]}\left(e^{-i\omega}\right)$ is the jth row of the lag polynomial evaluated at z represented by $z = e^{-i\omega}$ for $i = \sqrt{-1}$, Q_k is the kth column of Q, and $\bar{x}$ denotes the conjugate transpose of x. The relative contribution of shock k to the variation of variable j is then

$$\begin{aligned}\Omega_{j,k}\left(\underline{\omega},\overline{\omega}\right) &= \frac{\int_{\omega\in[\underline{\omega},\overline{\omega}]}\left(\overline{B^{[j]}\left(e^{-i\omega}\right)Q_k}B^{[j]}\left(e^{-i\omega}\right)Q_k\right)}{\int_{\omega\in[\underline{\omega},\overline{\omega}]}\left(B^{[j]}\left(e^{-i\omega}\right)\overline{B^{[j]}\left(e^{-i\omega}\right)}\right)} \\ &= Q_k'\left(\frac{\int_{\omega\in[\underline{\omega},\overline{\omega}]}\left(\overline{B^{[j]}\left(e^{-i\omega}\right)}B^{[j]}\left(e^{-i\omega}\right)\right)}{\int_{\omega\in[\underline{\omega},\overline{\omega}]}\left(B^{[j]}\left(e^{-i\omega}\right)\overline{B^{[j]}\left(e^{-i\omega}\right)}\right)}\right)Q_k \end{aligned} \tag{6}$$

or $\Omega_{j,k}\left(\underline{\omega},\overline{\omega}\right) = Q_k'S_j\left(\underline{\omega},\overline{\omega}\right)Q_k$ where $S_j\left(\underline{\omega},\overline{\omega}\right) = \frac{\int_{\omega\in[\underline{\omega},\overline{\omega}]}\left(\overline{B^{[j]}\left(e^{-i\omega}\right)}B^{[j]}\left(e^{-i\omega}\right)\right)}{\int_{\omega\in[\underline{\omega},\overline{\omega}]}\left(B^{[j]}\left(e^{-i\omega}\right)\overline{B^{[j]}\left(e^{-i\omega}\right)}\right)}$.

The identification approach from Angeletos, Collard and Dellas (2020) selects Q_k that maximizes $\Omega_{j,k}\left(\underline{\omega},\overline{\omega}\right)$ for different target variables and different frequency bands. It is then implied that Q_k is the eigenvector associated with the first eigenvalue of the $n \times n$ matrix $S_k\left(\underline{\omega},\overline{\omega}\right)$.

In our case, we seek to identify two shocks which are orthogonal to the emission intensity shock while targeting CO2 emissions. We then study the impulse responses of different variables to these

shocks. The first shock is identified by maximizing its contribution to the variation of per capita CO2 emissions over business-cycle frequencies from 6 to 32 quarters subject to the restriction that the shock is orthogonal to the emission intensity shock. The second shock is identified as the one maximizing the contribution over medium- to long-run frequencies between 32 to 80 quarters subject to two restrictions that the shock is orthogonal to the business-cycle and the emission intensity shocks.

More precisely, let CO2 emissions be indexed in X_t by j, and let the emission intensity shock and the business-cycle and the medium-run shocks to CO2 emissions be indexed in v_t by 1, 2 and 3, respectively. We first solve the following optimization problem to identify the business-cycle shock to emissions

$$\underset{Q_2}{\operatorname{argmax}}\, \Omega_{j,2}\left(2\pi/32, 2\pi/6\right), \tag{7}$$

subject to the restrictions $Q_2'Q_2 = 1$ and $Q_2'Q_1 = 0$. Identification of the medium-run shock to emissions is then achieved by solving the following problem

$$\underset{Q_3}{\operatorname{argmax}}\, \Omega_{j,3}\left(2\pi/80, 2\pi/32\right), \tag{8}$$

subject to the restrictions $Q_3'Q_3 = 1$, $Q_3'Q_2 = 0$ and $Q_3'Q_1 = 0$. Following Angeletos, Collard and Dellas (2020), the solution of the first problem is obtained as $Q_2 = ZY$ where Y is the eigenvector associated with the first eigenvalue of the matrix $Z'S_j\left(\underline{\omega}, \overline{\omega}\right)Z$ where Z is an $n \times (n-1)$ orthonormal matrix that satisfies $Z'Q_1 = 0$.[5] The solution of the second problem obtains quite similarly, finding a matrix Z that satisfies $Z'\left[Q_1, Q_2\right] = 0$.

3. Data and VAR Specification

The data we use to estimate our baseline VAR specification consist of quarterly observations on 13 variables. We measure carbon emissions as CO2 emissions per capita. We construct the emission intensity of US output as the ratio of CO2 emissions per capita over real GDP per capita, see Figure 1. The remaining 11 variables are chosen to represent different macroeconomic and financial aspects of the US economy. These are Fernald (2014)'s utilization-adjusted TFP; the inverse of the relative price of investment goods from Justiniano, Primiceri and Tambalotti (2010); the per capita levels of hours worked in the non-farm business sector, real personal consumption expenditures, real private fixed investment; the change in private inventory investment; the Federal Funds rate; PCE inflation; the real S&P 500 index, and the real price of WTI oil where the latter two are deflated by the PCE index.

Results for the impact of the identified shocks on energy consumption and prices by sources and sectors are obtained from larger VARs, each additionally including one of the following 17

[5] We compute Z in Matlab through "null(Q_1')".

variables: the per capita consumption of petroleum, natural gas, renewable energy, coal and nuclear electric power; the per capita levels of primary energy consumption of the transportation, industrial, residential, commercial and electric power sectors; the per capita levels of electricity retail sales to the transportation, industrial, residential and commercial sectors; and the PPI for coal, PPI for natural gas and PPI for refined petroleum products (the PPIs for commodities are also deflated by the PCE).

All of the emission and energy consumption data are obtained as monthly time series from the US Energy Information Administration (EIA). Moreover, we obtain monthly PPIs for commodities from the US Bureau of Labor Statistics (BLS). We first seasonally adjust all of these series using the X-12 method and then convert them to quarterly values by averaging the quarter's three monthly values. All of the macroeconomic variables were obtained from the Federal Reserve Economic Database (FRED). In our model specification, we include the inflation rate, the Federal Funds rate and the change in private inventory investment in percent and all of the remaining variables in natural logs.

The sample starts in 1973:I which is the earliest date for which the emission data are available, and ends in 2019:IV. The VARs are estimated with four lags using Bayesian methods subject to a Minnesota prior. The posterior distributions were obtained using Gibbs sampling with 50,000 draws using the first 45,000 as burn-in. In the Appendix, we show that our main findings are robust if we include two lags in the baseline VAR (selected by the AIC), estimate it by OLS and compute the confidence intervals using the parametric bootstrap instead.

4. Results

In this section we document our empirical results. First, we provide impulse response functions (IRFs) and forecast error variance decompositions (FEVDs) for the emission intensity shock in Section 4.1. In Section 4.2 we compare the emission intensity shock with a shock reflecting news about future TFP. Finally, in Section 4.3 we provide IRFs and FEVDs for two additional shocks that explain the remaining variation of emissions at business-cycle and at medium-term frequencies.

4.1. Emission Intensity Shock

The top panel of Figure 2 provides the posterior median IRF of the variables in our baseline VAR for the emission intensity shock, along with the 16-84 percent posterior coverage intervals. The top-left chart shows that the shock leads to a strongly significant one percent reduction of CO2 emissions per unit of GDP on impact. While the emission intensity of output slightly increases in the following quarters, it remains significantly compressed over the next ten years. Hence, innovations to the emission intensity of output are essentially permanent. The chart to the right shows the IRF for CO2 emissions per capita, the numerator of our emission intensity variable. While per capita emissions decline strongly on impact in response to the shock, they quickly revert back to their initial level. We thus find evidence for a rebound effect in carbon emissions: a shock which persistently reduces the emission intensity of output does not lead to lower emissions in the

long run. Such rebound effects have first been discussed by William Stanley Jevons in 1865 related to the use of coal in England during the industrial revolution. Jevons (1865) observed that coal consumption increased after the introduction of James Watt's steam engine which greatly enhanced the energy efficiency relative to earlier technologies. The more efficient machines were then widely adopted in other sectors of the economy, thus leading to an increase in the demand for coal. While a large literature discusses whether rebound effects for energy usage exist at the firm or industry level, there is only limited evidence for aggregate energy consumption at the national or global level, see Brockway, Sorrell, Semieniuk, Heun and Court (2021). Our finding of a rebound effect in carbon emissions is consistent with a rebound of fossil energy consumption.

Why are carbon emissions not persistently reduced when the emission intensity of output is lowered permanently? The reason is that output increases persistently in response to the emission intensity shock, as shown in the third chart in the first row of Figure 2. While the initial response of GDP per capita is small, it rises sharply over the first three years after the shock hits, peaks at a one percent higher level and then declines somewhat over the next several years. Ten years after the shock, output is still about 0.8 percent higher, highlighting the long-term impact of emission intensity shocks on the real economy. The strong and persistent increase of output is driven by several of its major components. Consumption increases sharply over the first three years after the shock and levels out at a little less than one percent after ten years. Fixed investment shows a more hump-shaped response. It initially rises sharply, peaks at two percent above its initial level after three years and then declines, reaching a one percent higher level after ten years. Interestingly, the strong increase of output on impact is partially driven by a negative response of inventory investment in the first few quarters after the shock hits. Hence, firms reduce their inventories at the same time as they increase their fixed investment, suggesting that much of this investment is replacing existing capital stock.

Hours worked increases persistently for a number of years, with a peak response about 0.4 percent higher after three years and then slowly reverting back to its initial level. Utilization-adjusted TFP initially shows only a small response but then starts to increase strongly and flattens out about 0.4 percent above its initial level. The shock thus implies a negative conditional correlation between emission intensity and TFP growth. The relative price of investment increases strongly and persistently over the ten year period after the shock. This suggests that technology in the production of investment goods grows at a higher rate than that for consumption goods, consistent with investment-specific technology being the main driver behind productivity in the long run (Fisher, 2006; Ben Zeev and Khan, 2015). Finally, PCE inflation declines on impact and reverts back to its initial level after about five years. This suggests that firms pass on some of the cost reductions through lower energy prices to their customers. Taken together, these results are in line with e.g. Shapiro and Walker (2018) who find a strong link between productivity growth and the emission intensity of output at the firm-level. We document such a negative correlation exists at the national level in the medium to long run.

Let us now look at the economic importance of the emission intensity shock by studying the

FEVDs in the bottom panel of Figure 2. The shock initially explains about 30 percent of the emissions per unit of output, but that fraction increases to almost 60 percent ten years after impact. Interestingly, only about 20 percent of total emissions per capita are explained by the emission intensity shock across forecast horizons. We will analyze the remaining variation in Section 4.3. Strikingly, while the shock explains little of the variation of real GDP per capita on impact, the fraction rises sharply across forecast horizons to more than 60 percent after ten years. Hence, a large share of the variation in real GDP is driven by the emission intensity shock. This is mirrored in the FEVD for the main components of GDP, equally showing that about three quarters of the variation of consumption and about half of the variation of fixed investment are explained by the emission intensity shock at the ten year horizon. The persistence and magnitude of its impact on key measures of real economic activity is in line with that found in response to news shocks about future TFP (see, e.g. Beaudry and Portier, 2006; Barsky and Sims, 2011; Kurmann and Otrok, 2013; Kurmann and Sims, 2021). Consistent with these studies we find that, while the shock initially explains only small fractions of the variability of TFP and the relative price of investment goods, these fractions rise sharply with the forecast horizon and reach about 50 percent after ten years. Hence, the emission intensity shock is closely related to future productivity growth. We will analyze the link between shocks to emission intensity and TFP in the next section.

Before doing so, we would like to better understand the lack of persistence in the negative response of carbon emissions after shocks which lower the emission intensity of output. To this end, it is instructive to look at the energy mix that the US economy consumes. Figure 3 displays the aggregate energy consumption in the US by sources and sectors for the year 2020. This information is obtained from the EIA from which we also obtain our CO2 emissions and commodities price and quantity data. The EIA measures energy consumption across sources in terms of British Thermal Units (BTU).[6] In 2020, total energy produced in the US amounted to more than 90 quadrillion BTU. Petroleum and natural gas each accounted for a little over one third of the total energy production. The remaining one third was roughly equally split between renewable energy (12%), coal (10%), and nuclear energy (9%).

The primary consumption of energy in the four main sectors was 69.7 quadrillion BTU in 2020, as shown in the right column of Figure 3. The industrial and transportation sectors each consumed a bit more than one third of the total, with the residential (17%) and commercial (12%) sectors accounting for the remainder. Looking at the sources of energy consumption for each sector, we see that one third and 41 percent of the industrial energy consumption was in petroleum and natural gas, respectively. The remaining 25 percent came from electricity (12%), coal (4%) and renewable sources (9%). The transportation sector, in turn, relies on a quite different energy mix with 90 percent of its energy consumption in petroleum. In contrast, the residential and commercial sectors primarily use natural gas and electricity as inputs.

[6]One BTU is the quantity of heat needed to increase the temperature of one pound of liquid water from 39 to 40 degrees Fahrenheit.

The difference of 23.2 quadrillion BTU between total consumption of primary energy inputs and total energy consumption by the four major sectors of the economy is accounted for by energy losses of the electrical system. This is a staggering two thirds of all energy transformed from primary sources into electricity. The electricity sector relies on natural gas (33%), nuclear energy (23%), coal (23%) and renewable energy (19%) as primary inputs.

Of course, the energy mix of the US economy has changed over time. While a full account of this transition is beyond the scope of this paper, we can parse the effect of an emission intensity shock on energy consumption in total as well as in the four main sectors of the economy using our modeling framework. Specifically, we obtain the impulse responses of these additional variables by re-estimating our baseline VAR, adding one at a time the different sources of energy and the end-use sectors'consumption of primary energy sources as well as electricity. Figure 4 provides the impulse responses. Looking first at the different primary energy sources, we see that the emission intensity shock initially leads to a drop in the consumption of petroleum and natural gas. While oil consumption takes about two years to revert back to its initial level, natural gas remains subdued for at least three years.[7] Interestingly, coal consumption increases somewhat albeit not statistically significantly in response to the emission intensity shock, while renewable energy is initially unaffected and then drops significantly for a number of years. The only primary source of energy that increases significantly in response to an emission intensity shock is nuclear energy. Nuclear energy consumption jumps by about 1.5 percent on impact and remains persistently above its initial level thereafter. Combining these primary energy inputs, electricity consumption initially does not respond to the emission intensity shock, but then rises persistently after a few quarters and remains significantly above its initial level.

Looking at the energy consumption by end-use sectors on the right-hand side of Figure 4, we see that the industrial sector initially requires one percent less energy in response to the emission intensity shock. Energy consumption then reverts back to its initial level after three years before falling again by around half a percentage point. The electricity consumption of the industrial sector shows a hump-shaped response, increasing for a few years but not in the longer run. The transportation sector, which predominantly relies on petroleum as energy source, also initially reduces its energy consumption in response to an emission intensity shock but then quickly reverts back to its initial level and even slightly exceeds it for a few years. While the electricity consumption of the transportation sector permanently increases after an emission intensity shock, electricity only accounts for a tiny fraction of the energy used by that sector.

The residential and commercial sectors both significantly reduce their primary energy consumption in response to an emission intensity shock. While this reduction is strongest on impact and recovers somewhat in the first few quarters after the shock hits, it is very persistent and remains

[7] We show the impulse response of petroleum as a primary source of energy. This comprises crude oil and natural gas plant liquids. In 2020, crude oil accounted for more than three quarters of petroleum production in the US. For the sake of simplicity and to be consistent with the literature on the macroeconomic effects of oil shocks, we use the terms oil and petroleum interchangeably.

below its initial level even ten years after the shock. Importantly, however, both of these sectors strongly and persistently increase their use of electricity which accounts for about half of their total energy consumption.

To summarize, these impulse responses show that shocks which persistently reduce the emission intensity of output are associated with a permanent reduction of fossil energy sources in all sectors except for transportation. At the same time, all sectors permanently increase their consumption of electricity. The additionally required electricity, in turn, is predominantly provided by the nuclear sector. Hence, innovations associated with an increased electrification of the US economy and higher nuclear energy production are associated with a persistent and significantly lower carbon emission intensity of US output over the past five decades. In the next section, we study the link between the identified emission intensity shock and news shock about future TFP which has been shown to be a key long-run driver of macroeconomic variables.

4.2. Emission Intensity and Total Factor Productivity

We have seen above that the identified emission intensity shock has initially small but then strongly rising persistent effects on real output and consumption but also on total factor productivity and the relative price of investment goods. The protracted response of these variables is reminiscent of the responses found for identified news shocks about future productivity (see, e.g. Kurmann and Sims, 2021). Here, we analyze to what extent TFP news shocks and emission intensity shocks are related to one another.

We follow Kurmann and Sims (2021) and identify a TFP news shock as the innovation which maximizes the share of forecast error variance of utilization-adjusted TFP at the horizon of 20 years. As discussed by these authors, this identification avoids imposing short-run zero restrictions which are sensitive to the cyclical properties of utilization-adjusted TFP which itself varies across different data vintages and detrending approaches. Hence, the identification of the TFP news shock is essentially the same as the identification of the emission intensity shock, except that the target variable is TFP rather than CO2 emissions per unit of output.

Figure 5 provides the median posterior impulse responses for a one standard deviation TFP news shock along with their 16-84 percent posterior coverage intervals in red dashed lines for the main variables included in our baseline VAR. We superimpose the emission intensity shock for comparison. Focus first on the response of emissions per unit of output in the top-left panel. The TFP news shock leads to a somewhat smaller and on impact insignificant response of emissions per unit of output compared to the emission intensity shock. In the longer run, however, both shocks reduce the emission intensity of output significantly. The IRF of per capita CO2 emissions is also qualitatively similar across the two shocks with the TFP news shock leading to a more pronounced hump-shaped response. The IRFs for the other variables included in our model, in particular for GDP per capita, its components consumption and investment, TFP and the relative price of investment goods, are all highly similar across the two shocks. This suggests a close link between news about future TFP and the emission intensity of output.

The close similarity between the two shocks also extends to the IRFs of energy consumption by source and end-use sector, shown in Figure 6. Both shocks lead to essentially identical on-impact reductions of oil and natural gas consumption which revert back to their initial levels after a few quarters. Both shocks also lead to higher consumption of coal and especially nuclear energy where only the latter is statistically significant across most forecast horizons. The increase in coal and nuclear energy consumption is in line with a pronounced hump-shaped increase in electricity production of which all end-use sectors increase their consumption.

To further compare the two shocks, we show the FEVDs for the TFP news and the emission intensity shock for CO2 per unit of output and for CO2 emissions per capita in the upper panel of Figure 7. By construction, the emission intensity shock explains large shares of CO2 emissions per unit of output across forecast horizons, ranging from around 30 percent on impact to about 70 percent ten years out. The TFP news shock explains somewhat lower, but still sizable, variance shares, especially at longer horizons. Moving to the FEVDs for emissions per capita on the bottom right, we see that both shocks explain similarly small shares of CO2 emissions per capita which do not exceed 30 percent. Hence, innovations other than those to the emission intensity of output or productivity more broadly account for the bulk of variation in CO2 emissions per capita. We study the additional drivers of CO2 emissions in Section 4.3 below.

4.2.1. Purging TFP News from Innovations to Emission Intensity

The previous results document a strong similarity between the emission intensity shock and the TFP news shock. This is confirmed in the lower panel of Figure 7 which provides a scatter plot of the two innovations. The correlation is close to 80 percent, suggesting that although both shocks have been identified using entirely different target variables, they capture much of the same information. As discussed above, there are two competing explanations for this finding. On the one hand, productivity growth enables firms to produce more output for a given level of fossil energy inputs and thus carbon emissions, therefore lowering emission intensity (see e.g. Shapiro and Walker, 2018). One the other hand, historical increases in fossil fuel prices might have redirected technological change away from energy-intensive toward more energy-efficient innovations, so that economy-wide technological change has raised energy efficiency in recent decades (see e.g. Aghion, Dechezleprêtre, Hémous, Martin and Van Reenen, 2016 and Hassler, Krusell and Olovsson, 2021).

With the aggregate data we have available, it is difficult to disentangle the two competing explanations. To shed some light on the economic importance of improvements in emission intensity for aggregate productivity, we conduct the following simple exercise. Specifically, we identify a TFP news shock the same way as before but under the additional restriction that it be orthogonal to the emission intensity shock. This orthogonalized shock thus captures innovations to overall technology that do not lower the emission intensity of output. Figure 8 shows the IRFs to the orthogonalized TFP news shock and superimposes the unrestricted TFP news shock for comparison. Looking first at the top-left chart in the figure, we see that the orthogonalized shock leads to higher

CO2 emissions per unit of output two to three years after the shock. This is in sharp contrast to the unrestricted TFP shock which lowers the emission intensity of output. In line with this, CO2 emissions increase immediately after the orthogonalized shock hits and not with a delay as is the case for the unrestricted shock.

More strikingly, real GDP per capita increases much less strongly and reverts back to its initial level more quickly in response to the orthogonalized shock. This muted response is matched by its consumption and fixed investment components as well as by hours worked. Hence, innovations to TFP which do not lower the emission intensity of output have much smaller effects on real activity than innovations which reduce CO2 emissions per unit of output.

Interestingly, the two shocks also have quite different effects on utilization-adjusted TFP and the relative price of investment. While the TFP news shock raises TFP primarily at longer horizons, the orthogonalized shock has a strong positive effect on impact which then subsides over time. Moreover, the relative price of investment, the real S&P500 index, and PCE inflation all show a much more muted response to the orthogonalized TFP news shock. These results suggest that technological advances unrelated to the emission intensity did not substantially contribute to longer-run productivity growth. Rather, they support the notion that innovations related to energy use and emission intensity have been a key driving force of longer-run productivity and of real economic output in the US over the past five decades.

The lower-right chart of Figure 7 shows a scatter plot of the TFP news shock and the orthogonalized TFP news shock. The correlation is slightly above 60 percent which is considerably lower than the correlation between the TFP news shock and the emission intensity shock. This underscores that a substantial share of innovations to TFP are related to energy usage and thus have an effect on emission intensity.

4.3. Business Cycle and Medium Term Variation in CO2 Emissions

Thus far, we have shown that the emission intensity shock explains a large share of the variation in output and its components and is highly correlated with news about future TFP. At the same time, while both shocks explain large shares of the longer-run variation of carbon emissions per unit of output, they only explain moderate shares of per capita CO2 emissions across forecast horizons.

We now study the remaining drivers of carbon emissions. To this end, we identify two additional shocks with the aim of capturing the residual variation of emissions not explained by the emission intensity shock. Specifically, we follow Angeletos, Collard and Dellas (2020) and identify two shocks explaining per capita CO2 emissions at different frequencies. The first shock is designed to explain a maximum share of variation of CO2 emissions at business cycle frequencies, i.e. between 6 and 32 quarters. We impose orthogonality with respect to the emission intensity shock to ensure that we only capture residual variation not explained by that shock. The second shock captures medium-term variation in emissions, i.e. from 32 to 80 quarters, and is made orthogonal to both

the emission intensity and the business-cycle emissions shock.[8]

Figure 9 provides the variance shares for CO2 emissions per capita for the three shocks, individually and combined. The emission intensity shock explains less than 30 percent of the variation across forecast horizons, as discussed earlier. The business-cycle shock captures larger fractions of the FEV of CO2 emissions, dropping from 60 percent on impact to about 30 percent at the five year horizon. In contrast, the medium-term shock explains none of the near-term variation but up to 40 percent of the variation at intermediate and longer horizons. Combined, the three shocks capture essentially all of the variation of per capita CO2 emissions across forecast horizons.

To study the economic underpinnings of the two additional shocks we now analyze their associated impulse response functions. Figure 10 reports them for a one-standard-deviation business-cycle shock. For comparison, we superimpose the corresponding IRFs for a one-standard-deviation emission intensity shock. Similar to the emission intensity shock, the business-cycle shock is associated with a sharp drop of about one percent in emissions on impact, but then reverts back to its initial level somewhat more slowly. In contrast to the emission intensity shock, however, the real output response is negative on impact and essentially zero after a few quarters. This is mirrored by the IRFs for consumption and investment. The real price of oil drops significantly following the business-cycle shock, consistent with the notion that it captures recessionary forces which exert a short-term drag on demand. The FEVDs associated with the business-cycle shock are provided in the bottom panel of the figure. They show that while its impact on CO2 emissions and the real price of oil is sizable, the shock only explains small shares of variation of output and its components. This suggests that it captures variations in energy demand that are largely independent of broader demand variations.

Figure 11 provides the IRFs and FEVDs for the medium-run shock. The shock leads to a hump-shaped negative response of emission intensity, with the peak decline of 0.5 percent after about ten quarters. This decline is associated with a more pronounced decline of CO2 emissions of about one percent, and a concurrent reduction of GDP per capita of a bit less than 0.5 percent. All three variables revert back to their initial levels after around five to seven years. The hump-shaped decline of output is mirrored by its two main components, with fixed investment declining more strongly than consumption. Consistent with its well-documented procyclical behavior, hours worked also shows a similar hump-shaped negative response. Utilization-adjusted TFP responds negatively to the medium-run shock, persistently declining by about 0.2 percent. Moreover, real stock prices decline by about two percent on impact and take about three years to recover. Consumer prices, however, rise sharply in response to the medium-run shock. PCE inflation increases by about 0.2 percent on impact and then subsides. The real price of crude oil rises by about three percent on impact and takes several years to return to its initial level. These impulse responses are thus consistent with an interpretation of the medium-run shock as a supply shock, possibly driven by

[8]In unreported results, we find that the results are robust to changing the ordering of the two frequency-based shocks.

exogenous commodity-price increases.

The FEVDs, provided in the lower panel of Figure 11, show that the medium-run shock explains sizable shares of the variation in CO2 emissions per capita for horizons of several years out. At the same time, it captures only modest shares of real GDP, consumption and investment, suggesting that even large reductions of CO2 emissions triggered by fossil energy price increases have had only a small effect on output over the past five decades.

To verify our interpretation of the business-cycle and the medium-term shocks as innovations capturing the demand and supply of fossil energy sources, we compute the IRFs and FEVDs of the consumption and prices of petroleum, natural gas and coal to the two shocks. The results are provided in Figure 12. Focusing first on the business-cycle shock in the two columns on the left, we see that consumption of all three primary fossil energy sources declines significantly by one to two percent on impact in response to the business-cycle shock. At the same time, the price of oil and natural gas drops significantly by two and five percent, respectively, and takes more than five years to recover. The price of coal also drops persistently, but with a delay of several quarters. Hence, with prices and quantities moving in the same direction, the business-cycle shock appears to capture demand-type forces in the market for primary fossil energy inputs. The FEVDs in the bottom panel of the figure show that between 20 and 40 percent of the variation in the price and demand for these fossil fuels are captured by the business cycle shock at different forecast horizons.

The right-hand columns of Figure 12 provide the corresponding results for the medium-term shock. The shock leads to a strongly significant immediate increase in the price of petroleum of about three percent, and is associated with a sizable hump-shaped decline of petroleum consumption, taking more than five years to return to its initial level. The price of natural gas also increases somewhat on impact and subsequently declines, but there is only a mild and statistically insignificant response of natural gas consumption. The price of coal increases significantly for a few years, possibly driven by some degree of substitutability of petroleum and coal. But as for natural gas the consumption of coal is not significantly affected by these price increases. The FEVDs in the bottom panel of the figure show that the medium-run shock explains sizable fractions of petroleum consumption several years out, reaching 50 percent at the three year horizon. The shares of variance explained are considerably smaller for natural gas and coal, however. In sum, these results lead us to interpret the medium-run shock as an oil supply shock.

To summarize, this section has shown that in addition to the emission intensity shock, two further disturbances are needed to capture all of the variation of per capita CO2 emissions: a business-cycle shock which appears to capture innovations to the demand for fossil energy sources; and a medium-run shock which is associated with a persistent, opposite-signed response of the quantity and price of petroleum used by the US economy.

5. Conclusion

In this paper, we have documented that a shock which permanently reduces the carbon emission intensity of US output has large and persistent positive effects on TFP and GDP. The emission intensity shock is strongly correlated with news about future TFP. TFP news uncorrelated with the emission intensity shock, in turn, have only modest and transitory effects on real output. The emission intensity shock is associated with lower consumption of fossil fuels, especially petroleum and natural gas, but an increased consumption of electricity which is mainly produced by the nuclear sector. These results suggest that innovations related to the energy mix of the US economy have been a key driving force of longer-run productivity and real economic output over the past five decades.

While these innovations permanently reduce the carbon emission intensity of US output, they do not lead to a permanent drop of per capita carbon emissions. The reason is that the strong increase of output triggered by the innovations to emission intensity also leads the consumption of fossil fuels to revert back to their initial levels. Hence, we document a sizable rebound effect of carbon emissions in aggregate US data, in line with a rebound effect for overall energy consumption.

As a result, the emission intensity shock only explains modest shares of per capita CO2 emissions across forecast horizons. The remaining variation of emissions is captured by two additional shocks. A business cycle shock which reflects innovations in the demand for fossil fuels, and a medium-term shock which captures supply-side innovations in particular in the market for petroleum.

While our results provide a coherent account of the dynamics of carbon emissions and output in the US over the past few decades, a general equilibrium model would be needed to infer robust policy conclusions. That said, our findings might inform the debate about how the transition to a net-zero carbon economy can be achieved. Most importantly, if history is a guide, technological advances lowering the carbon intensity of output alone are unlikely to considerably and persistently reduce carbon emissions. Policies or societal trends affecting the supply of and demand for fossil fuels will also have to play a key role. At the same time, technological innovations affecting the efficient use of energy might have substantial positive spillover effects for productivity and growth of the broader economy.

References

Aghion, P., Dechezleprêtre, A., Hémous, D., Martin, R., Van Reenen, J., 2016. Carbon Taxes, Path Dependency, and Directed Technical Change: Evidence from the Auto Industry. Journal of Political Economy 124, 1–51.

Angeletos, G.M., Collard, F., Dellas, H., 2020. Business-Cycle Anatomy. American Economic Review 110, 3030–70.

Barsky, R.B., Sims, E.R., 2011. News Shocks and Business Cycles. Journal of Monetary Economics 58, 273 – 289.

Beaudry, P., Portier, F., 2006. Stock Prices, News, and Economic Fluctuations. American Economic Review 96, 1293–1307.

Ben Zeev, N., Khan, H., 2015. Investment-Specific News Shocks and US Business Cycles. Journal of Money, Credit and Banking 47, 1443–1464.

Bloom, N., Genakos, C., Martin, R., Sadun, R., 2010. Modern Management: Good for the Environment or Just Hot Air? The economic journal 120, 551–572.

Brockway, P.E., Sorrell, S., Semieniuk, G., Heun, M.K., Court, V., 2021. Energy Efficiency and Economy-Wide Rebound Effects: A Review of the Evidence and Its Implications. Renewable and Sustainable Energy Reviews , 110781.

Cherniwchan, J., 2017. Trade Liberalization and the Environment: Evidence from NAFTA and US Manufacturing. Journal of International Economics 105, 130–149.

Cui, J., Lapan, H., Moschini, G., 2016. Productivity, Export, and Environmental Performance: Air Pollutants in the United States. American Journal of Agricultural Economics 98, 447–467.

Fernald, J., 2014. A Quarterly, Utilization-Adjusted Series on Total Factor Productivity, Federal Reserve Bank of San Francisco Working Paper 2012-19.

Fisher, J.D., 2006. The Dynamic Effects of Neutral and Investment-Specific Technology Shocks. Journal of political Economy 114, 413–451.

Forslid, R., Okubo, T., Ulltveit-Moe, K.H., 2018. Why Are Firms That Export Cleaner? International Trade, Abatement and Environmental Emissions. Journal of Environmental Economics and Management 91, 166–183.

Francis, N., Owyang, M.T., Roush, J.E., DiCecio, R., 2014. A Flexible Finite-Horizon Alternative to Long-Run Restrictions with an Application to Technology Shocks. The Review of Economics and Statistics 96, 638–647.

Griliches, Z., 1957. Hybrid Corn: An Exploration in the Economics of Technological Change. Econometrica 25, 501–522.

Hassett, K.A., Metcalf, G.E., 1999. Measuring the Energy Savings from Home Improvement Investments: Evidence from Monthly Billing Data. Review of economics and statistics 81, 516–528.

Hassler, J., Krusell, P., Olovsson, C., 2021. Directed Technical Change as a Response to Natural Resource Scarcity. Journal of Political Economy 129, 3039–3072.

Holladay, J.S., 2016. Exporters and the Environment. Canadian Journal of Economics 49, 147–172.

Jevons, W.S., 1865. The Coal Question; An Inquiry Concerning the Progress of the Nation and the Probable Exhaustion of Our Coal-Mines. London: Macmillan and Co.

Justiniano, A., Primiceri, G.E., Tambalotti, A., 2010. Investment Shocks and Business Cycles. Journal of Monetary Economics 57, 132–145.

Khan, H., Metaxoglou, K., Knittel, C.R., Papineau, M., 2019. Carbon Emissions and Business Cycles. Journal of Macroeconomics 60, 1–19.

Kok, N., McGraw, M., Quigley, J.M., 2011. The Diffusion of Energy Efficiency in Building. American Economic Review 101, 77–82.

Kurmann, A., Otrok, C., 2013. News Shocks and the Slope of the Term Structure of Interest Rates. American Economic Review 103, 2612–32.

Kurmann, A., Sims, E., 2021. Revisions in Utilization-Adjusted TFP and Robust Identification of News Shocks. Review of Economics and Statistics 103, 216–235.

Levinson, A., 2009. Technology, International Trade, and Pollution from US Manufacturing. American Economic Review 99, 2177–92.

Masson-Delmotte, V., Zhai, P., Pörtner, H.O., Roberts, D., Skea, J., Shukla, P.R., Pirani, A., Moufouma-Okia, W., Péan, C., Pidcock, R., et al., 2018. Global Warming of 1.5 c. An IPCC Special Report on the impacts of global warming of 1.

Melitz, M.J., 2003. The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. Econometrica 71, 1695–1725.

Popp, D., Newell, R.G., Jaffe, A.B., 2010. Energy, the Environment, and Technological Change. Handbook of the Economics of Innovation 2, 873–937.

Shapiro, J.S., Walker, R., 2018. Why is Pollution from US Manufacturing Declining? The Roles of Environmental Regulation, Productivity, and Trade. American Economic Review 108, 3814–54.

Uhlig, H., 2003. What Moves Real GNP?, Unpublished working paper, Euro Area Business Cycle Network.

Figure 1: US GDP and CO2 Emissions per Capita Since the 1970s

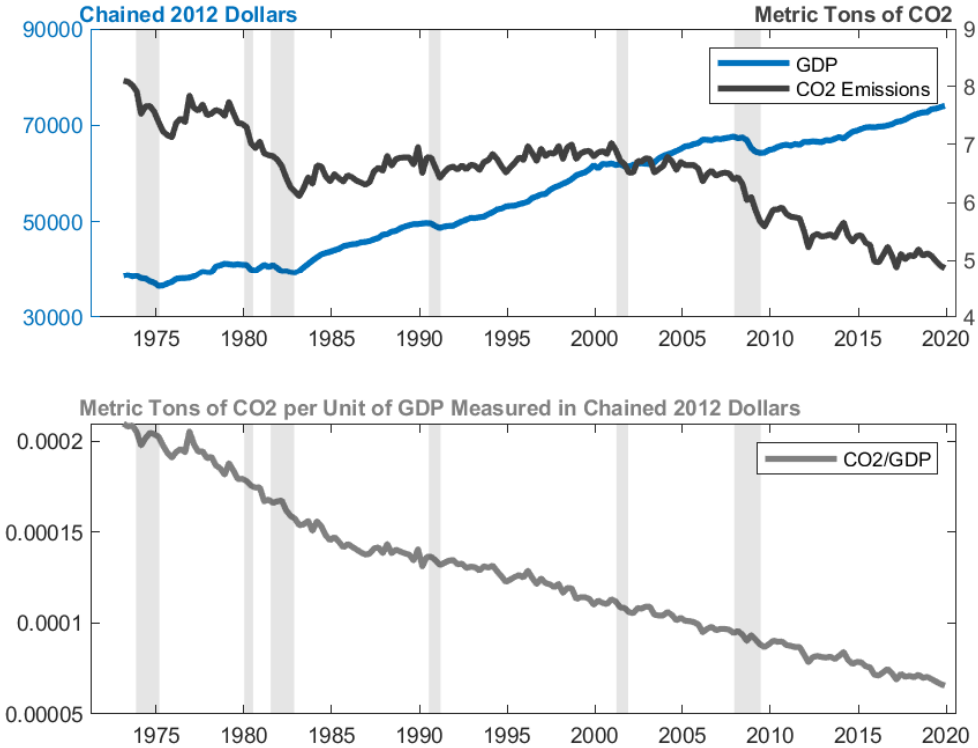

Note: This figure depicts the real GDP per capita (blue solid), the CO2 emissions per capita (black solid), and the CO2 emissions per unit of real GDP in the US (gray solid). The shaded vertical bars denote the NBER-dated recessions.

Figure 2: IRFs and FEVDs for Emission Intensity Shock

(a) IRFs

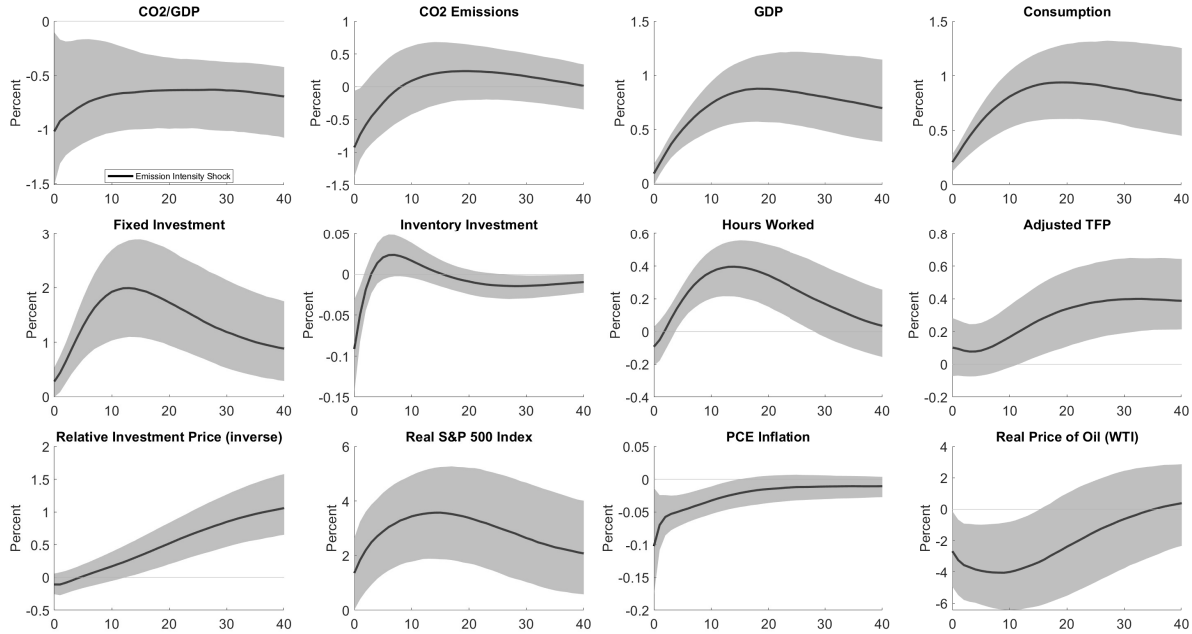

(b) FEVDs

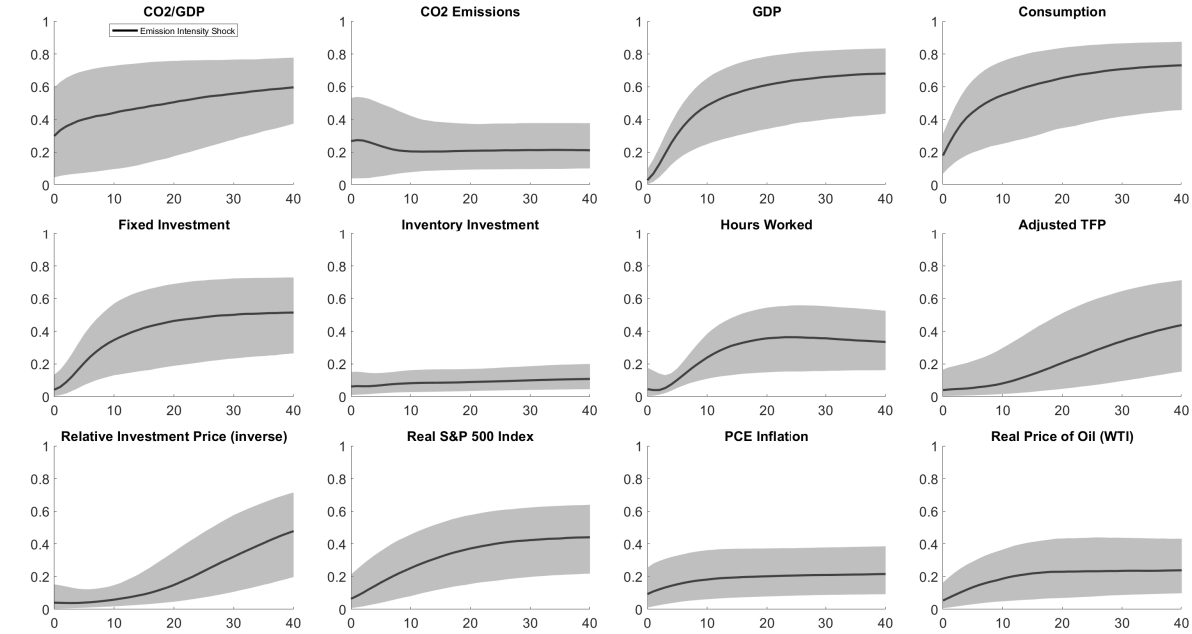

Note: The top panel shows the IRFs for the emission intensity shock from the structural VAR. The emission intensity shock is reported as a one-standard-deviation impulse. The bottom panel displays the corresponding FEVDs. The shaded bands correspond to the 16 to 84 percent posterior coverage intervals.

Figure 3: US Energy Consumption by Sources and Sectors, 2020

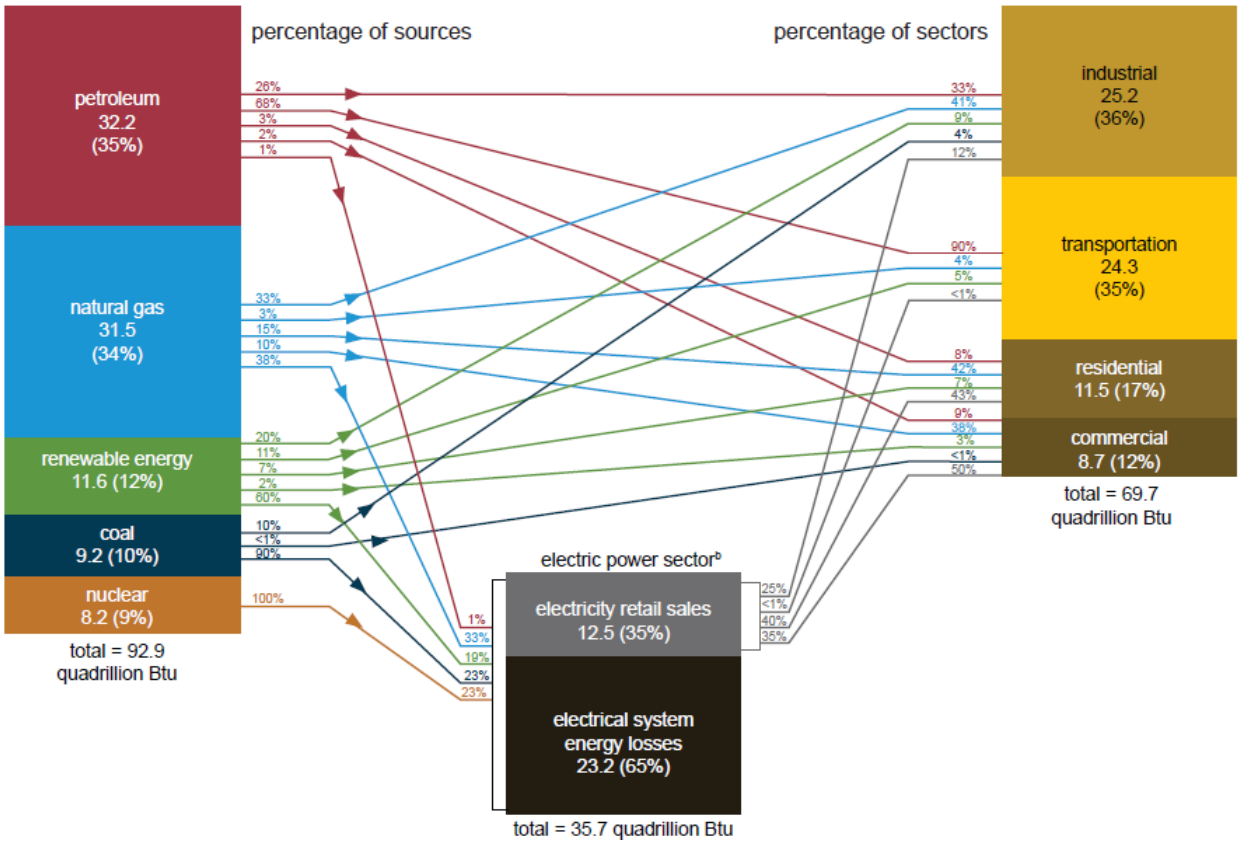

Source: U.S. Energy Information Administration (EIA), Monthly Energy Review (April 2021), Tables 1.3 and 2.1-2.6. A detailed description of this figure can be found here: https://www.eia.gov/energyexplained/us-energy-facts/.

Figure 4: IRFs of Energy Consumption by Sources and Sectors

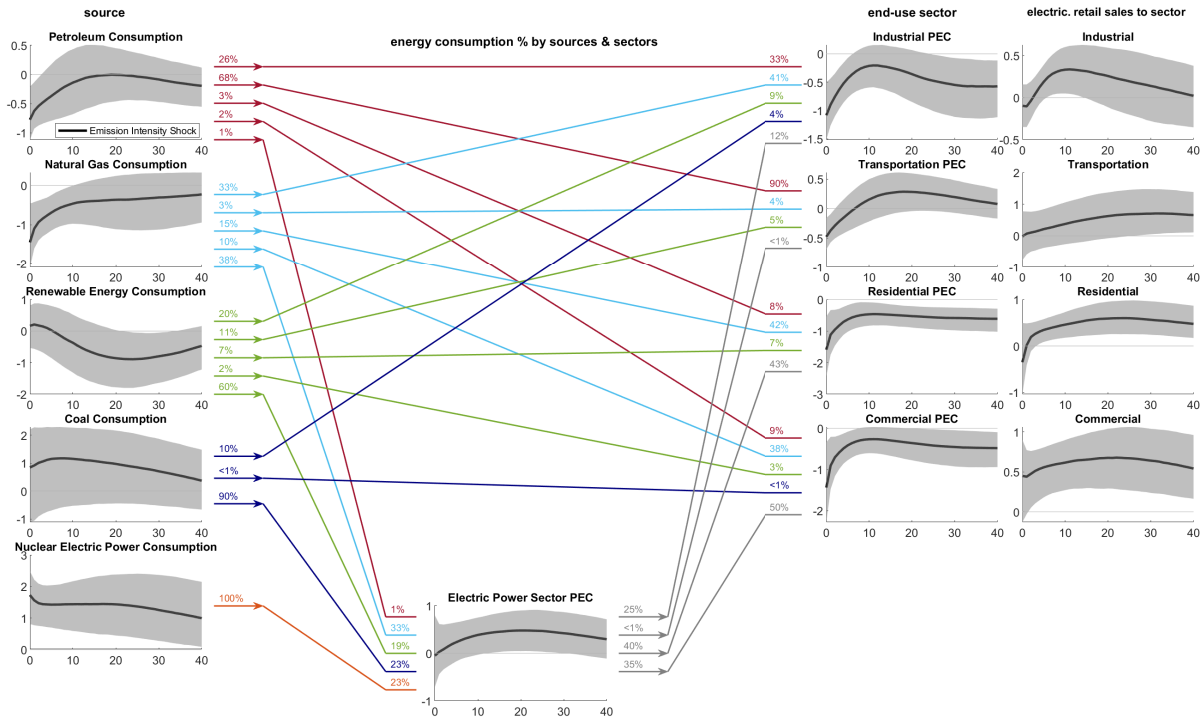

Note: This figure shows the IRFs for the emission intensity shock with the 16 to 84 percent posterior coverage intervals from the structural VAR. The emission intensity shock is reported as a one-standard-deviation impulse

Figure 5: IRFs for Emission Intensity and TFP News Shocks

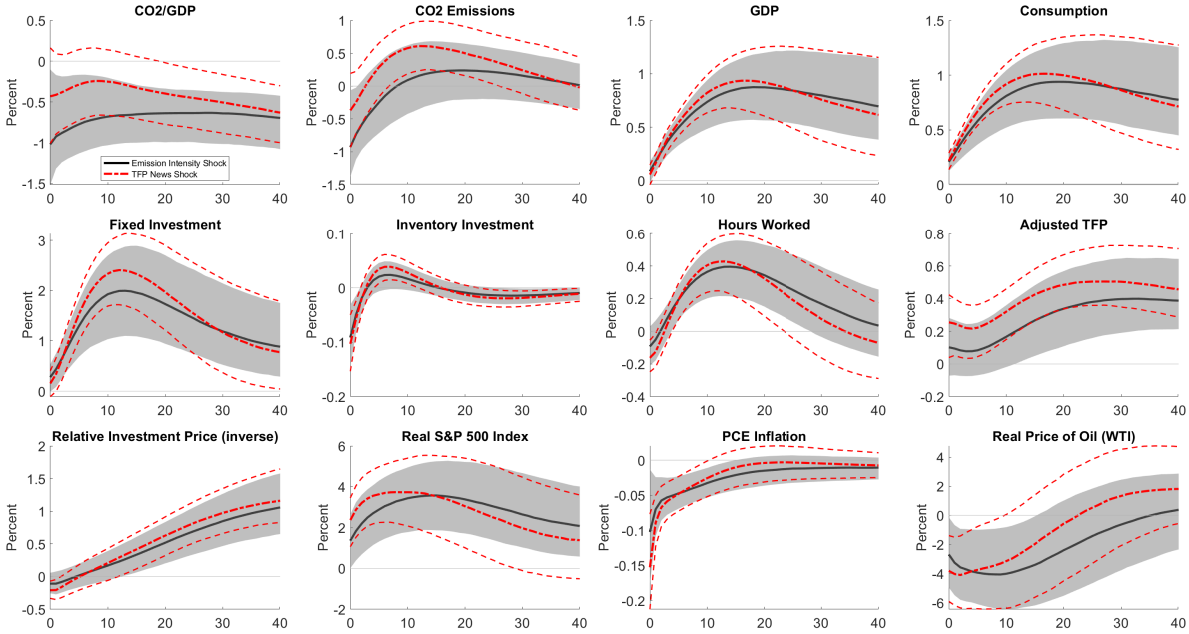

Note: This figure shows the IRFs for the emission intensity shock (black solid), and the TFP news shock (red dashed) from the structural VAR. Each of the shocks is reported as a one-standard-deviation impulse. The shaded bands and the dashed lines correspond to the 16 to 84 percent posterior coverage intervals.

Figure 6: IRFs of Energy Consumption by Sources and Sectors for Emission Intensity and TFP News Shocks

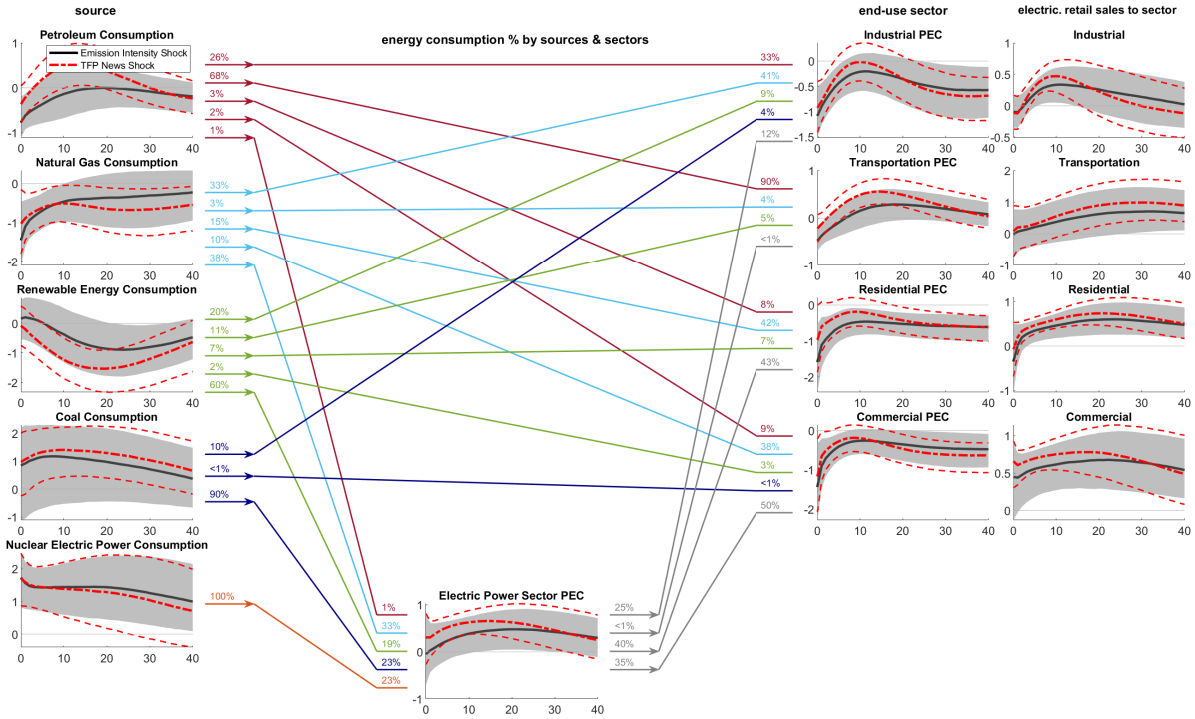

Note: This figure shows the IRFs for the emission intensity shock (black solid), and the TFP news shock (red dashed) from the structural VAR. Each of the shocks is reported as a one-standard-deviation impulse. The shaded bands and the dashed lines correspond to the 16 to 84 percent posterior coverage intervals.

Figure 7: Comparing Emission Intensity and TFP News Shocks

(a) FEVDs

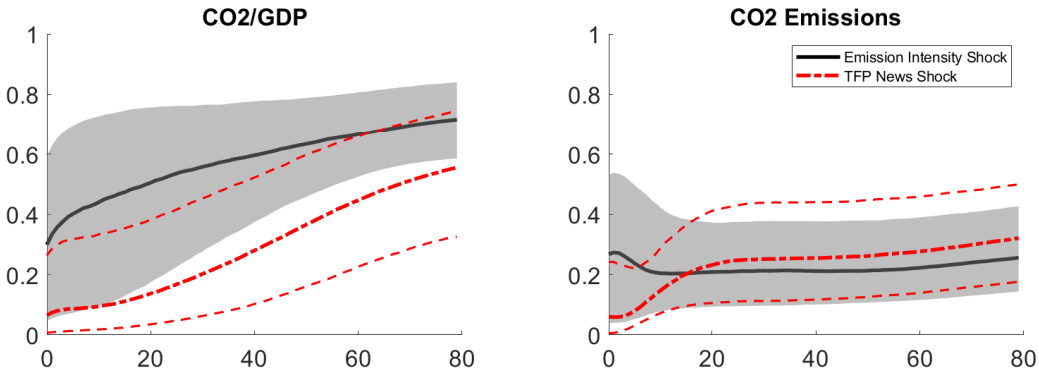

(b) Correlations among estimated shocks

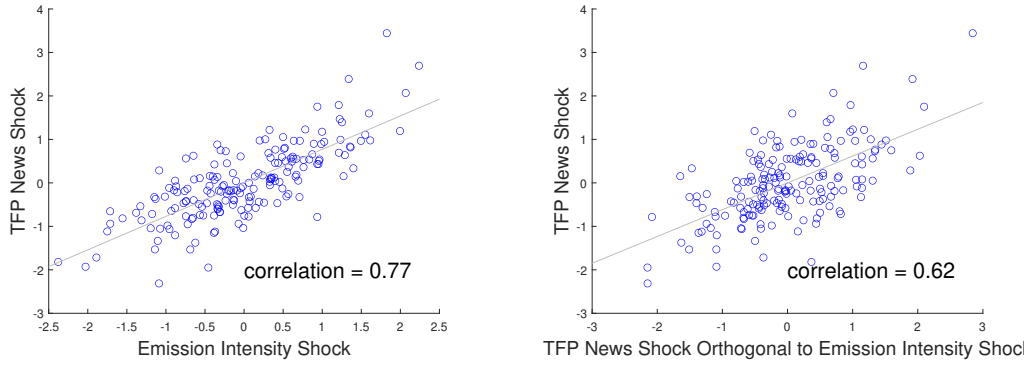

Note: The top panel of this figure shows the FEVDs the emission intensity shock (black solid), and the TFP news shock (red dashed) from the structural VAR. The shaded bands and the dashed lines correspond to the 16 to 84 percent posterior coverage intervals. The bottom panel provides the scatterplot of the two identified shocks against one another, as well as that for the identified TFP news shock against the TFP news shock that is made orthogonal to the emission intensity shock.

Figure 8: IRFs and FEVDs for Emission Intensity Shock

(a) IRFs

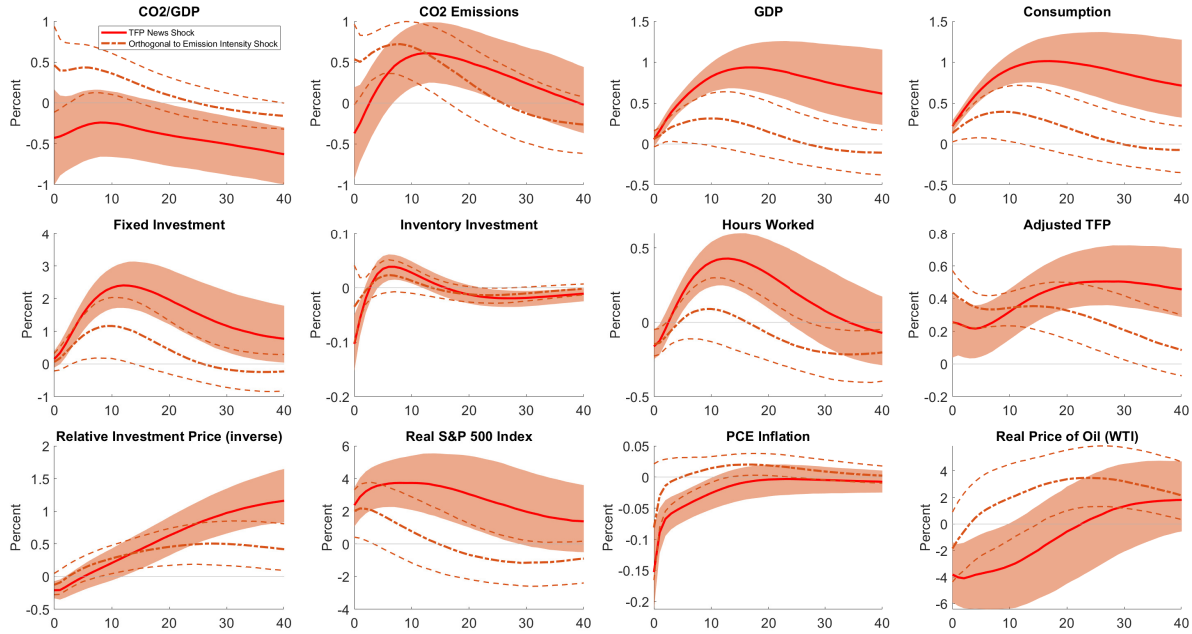

(b) FEVDs

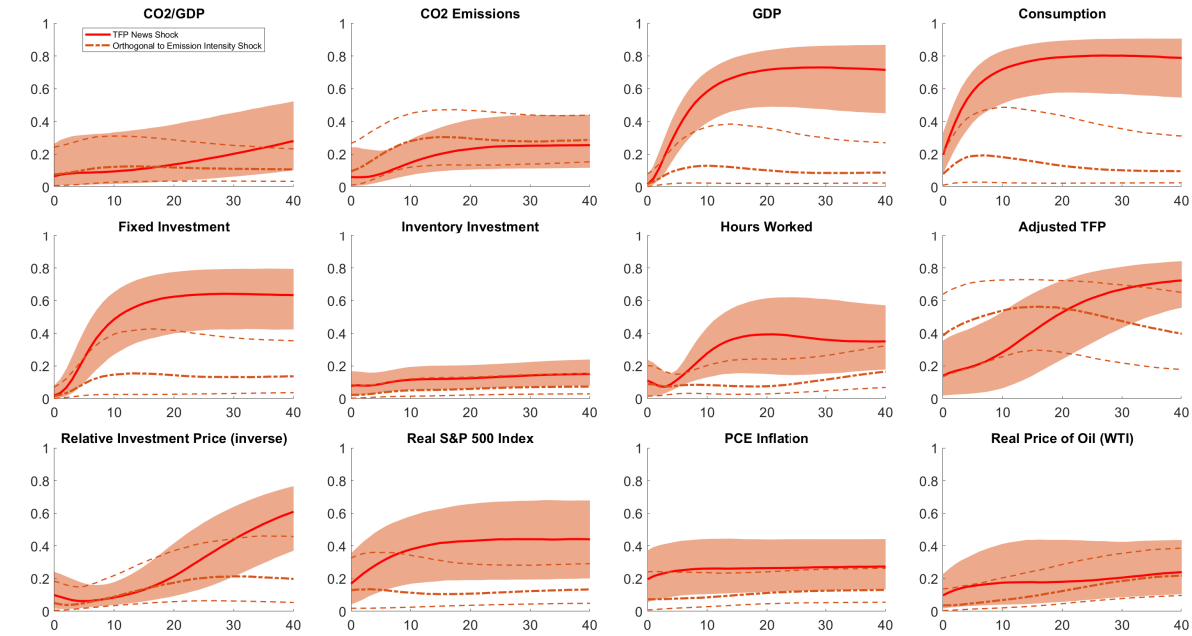

Note: The top panel shows the IRFs for the TFP news shock (red solid), and the TFP news shock that is made orthogonal to the emission intensity shock (orange dotted-dashed) from the structural VAR. Each of the shocks is reported as a one-standard-deviation impulse. The bottom panel displays the corresponding FEVDs. The shaded bands and the dashed lines correspond to the 16 to 84 percent posterior coverage intervals.

Figure 9: FEVDs for Three Shocks

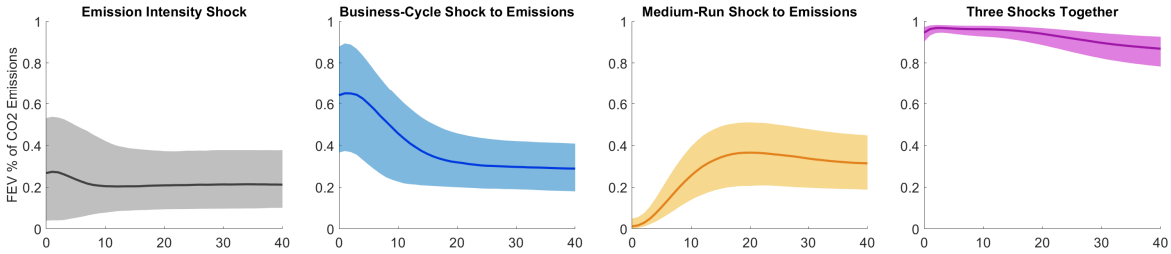

Note: This figure shows the FEVDs for the emission intensity shock (black solid), the business-cycle shock to emissions (blue solid), the medium-run shock to emissions (yellow solid), as well as the three shocks combined (purple solid). The shaded bands correspond to the 16 to 84 percent posterior coverage intervals.

Figure 10: IRFs and FEVDs for Business-Cycle Shock to Emissions

(a) IRFs

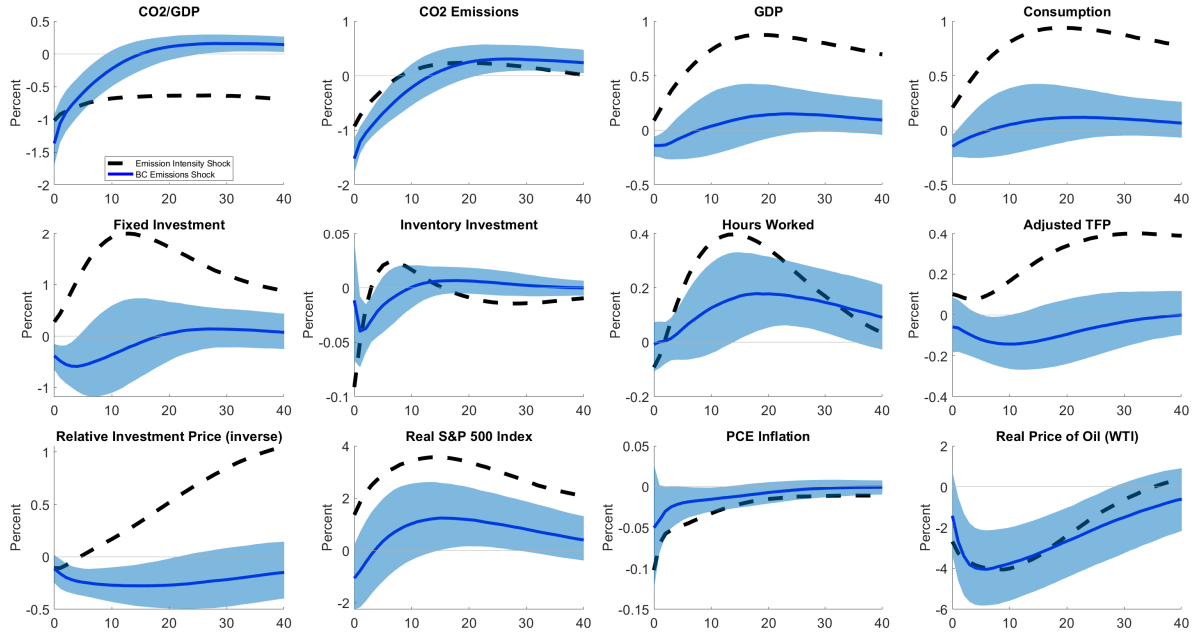

(b) FEVDs

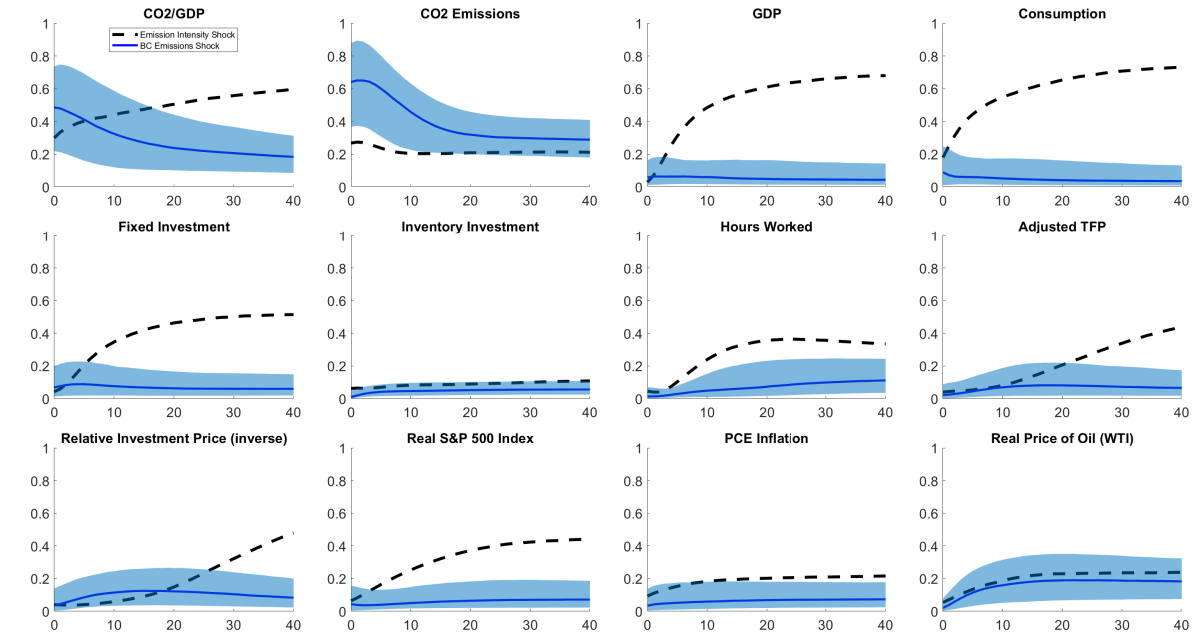

Note: The top panel shows the IRFs for the emission intensity shock (black dashed), and the business-cycle shock to emissions (blue solid) from the structural VAR. Each of the shocks is reported as a one-standard-deviation impulse. The bottom panel displays the corresponding FEVDs. The shaded bands correspond to the 16 to 84 percent posterior coverage intervals.

Figure 11: IRFs and FEVDs for Medium-Run Shock to Emissions

(a) IRFs

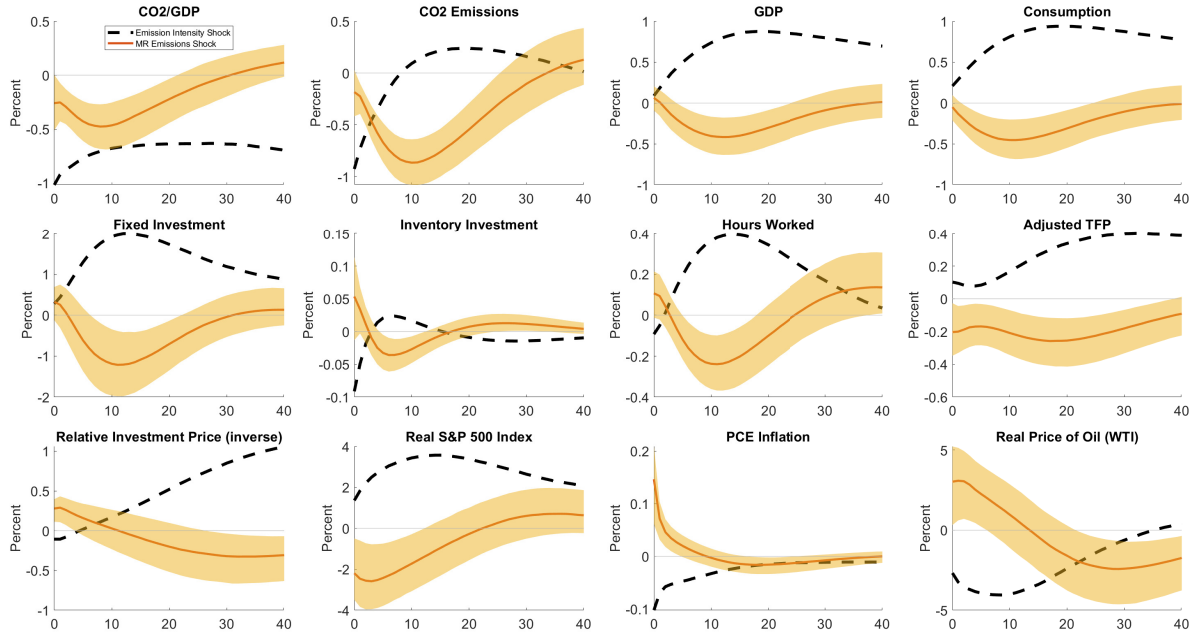

(b) FEVDs

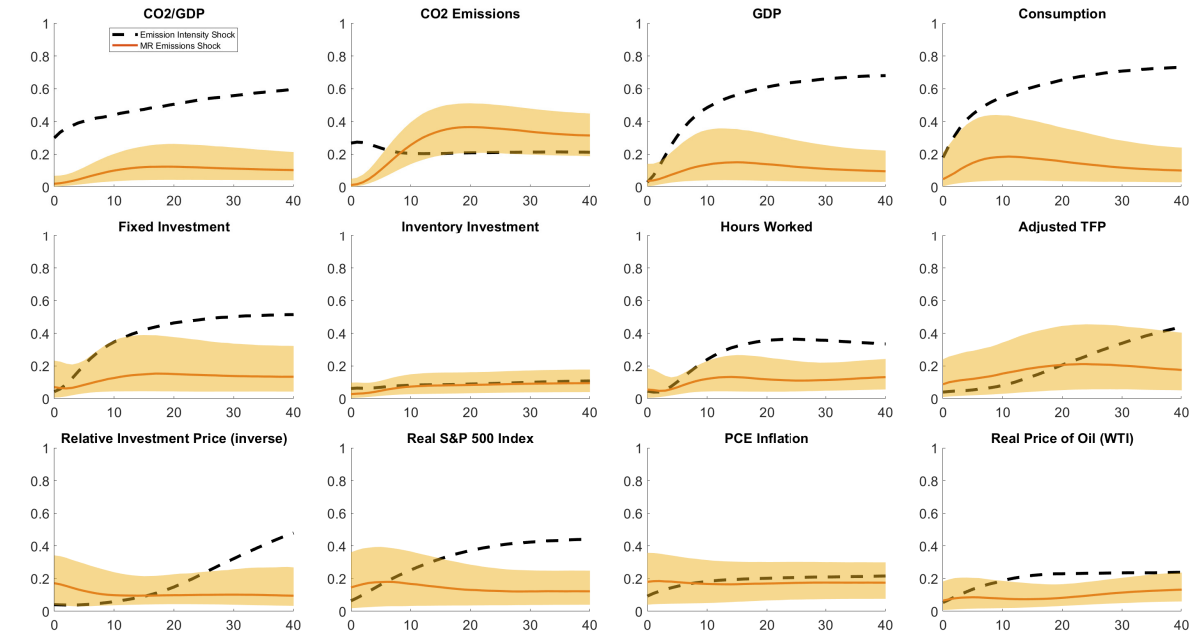

Note: The top panel shows the IRFs for the emission intensity shock (black dashed), and the medium-run shock to emissions (yellow solid) from the structural VAR. Each of the shocks is reported as a one-standard-deviation impulse. The bottom panel displays the corresponding FEVDs. The shaded bands correspond to the 16 to 84 percent posterior coverage intervals.

Figure 12: IRFs of Price and Consumption of Fossil Energy

(a) IRFs

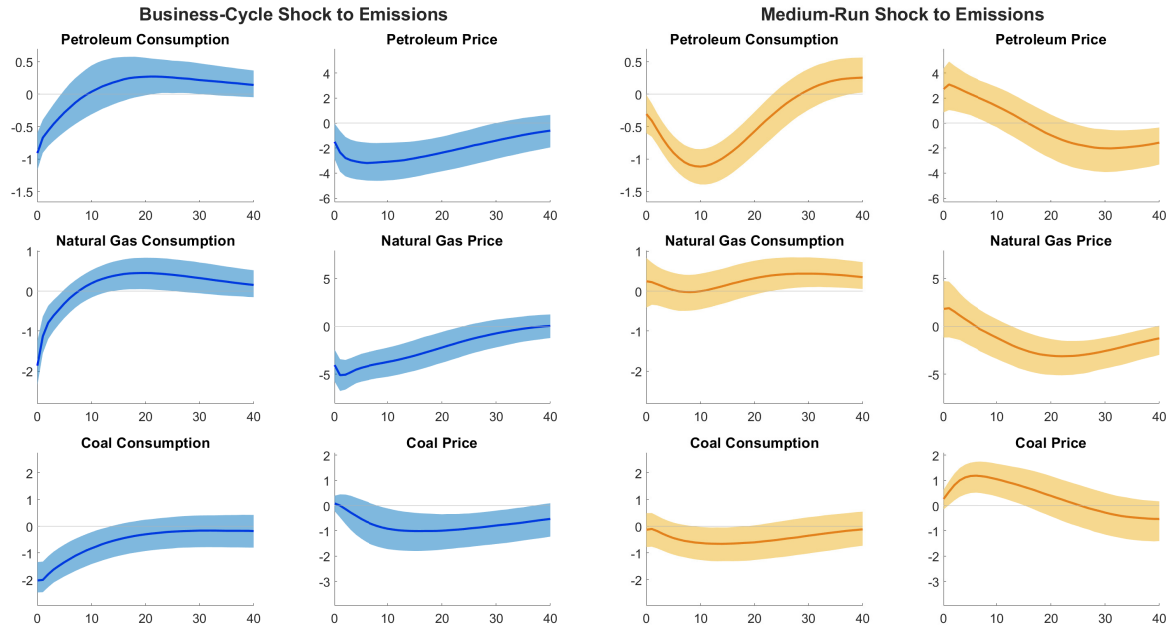

(b) FEVDs

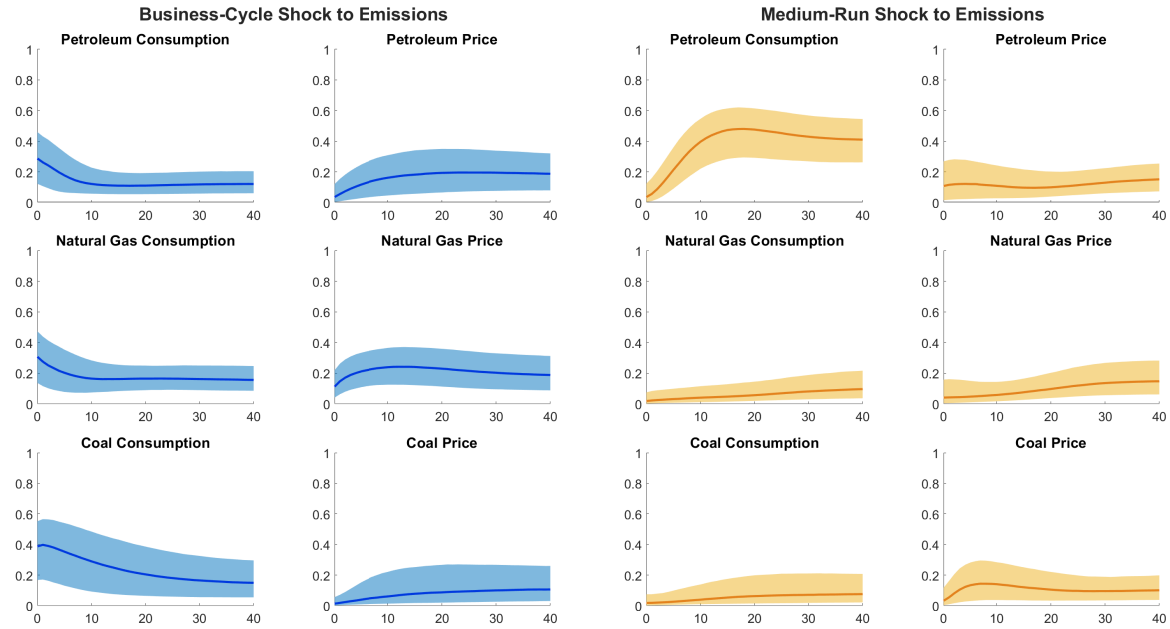

Note: The top panel shows the IRFs for the business-cycle shock to emissions (blue solid), and the medium-run shock to emissions (yellow solid) from the structural VAR. Each of the shocks is reported as a one-standard-deviation impulse. The bottom panel displays the corresponding FEVDs. The shaded bands correspond to the 16 to 84 percent posterior coverage intervals.

Appendix A. Robustness: Results from OLS Estimation

In the main text, because of the large number of the variables used in our VAR specifications, we have used Bayesian methods imposing a Minnesota prior to improve precision. Here, we document that our main findings are robust if we estimate the baseline VAR via OLS and compute the confidence intervals using the parametric bootstrap instead. Instead of four lags used in our baseline specification above, we now use two lags, which is the number selected by the AIC. For a comparison of the results for the emission intensity shock and the TFP news shock obtained via Bayesian vs. OLS estimation, see Figure A.1 on the next page.

Figure A.1: Robustness: Bayesian vs. OLS Estimation

(a) IRFs for Emission Intensity Shock

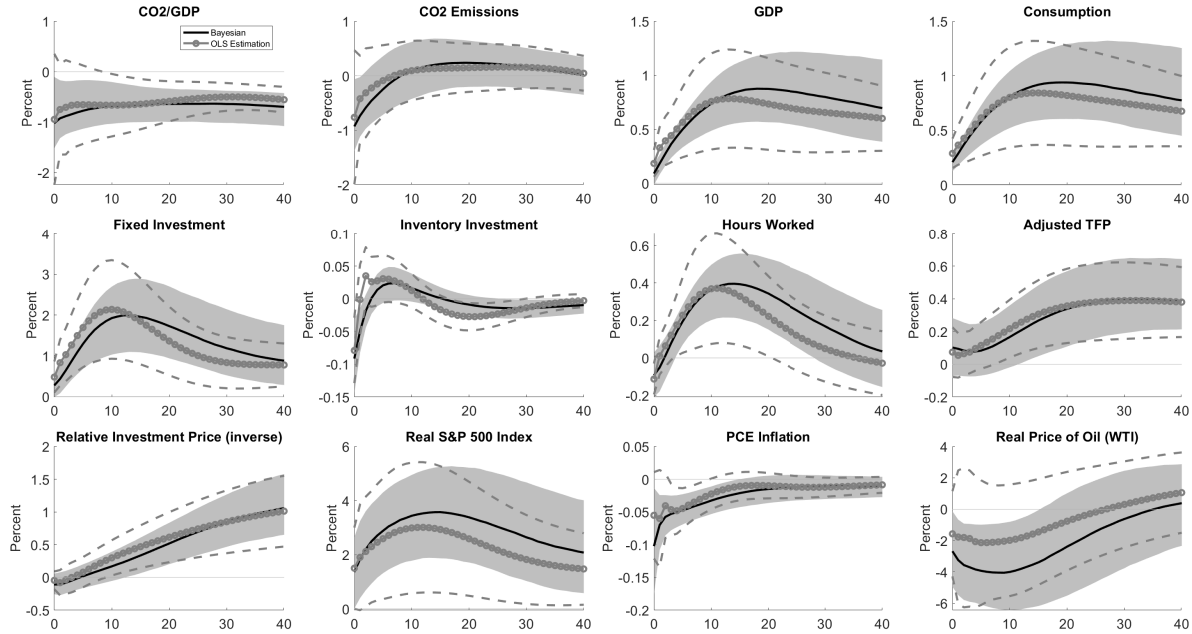

(b) IRFs for TFP News Shock

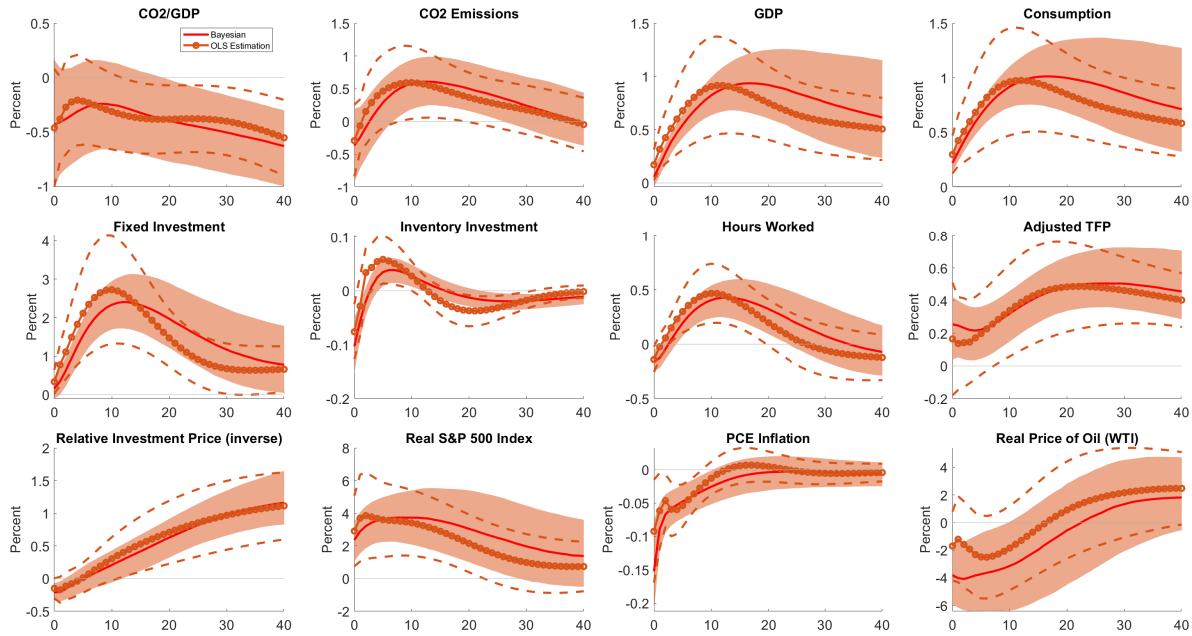

Note: The top panel shows the IRFs for the emission intensity shock from the structural VAR which we estimated both via Bayesian method subject to a Minnesota prior and OLS. The bottom panel displays the corresponding results for the TFP news shock. Each of the shocks is reported as a one-standard-deviation impulse. The shaded bands correspond to the 16 to 84 percent posterior coverage intervals. The dashed lines correspond to the one-standard-error bands which we computed by the parametric bootstrap.