DISCUSSION PAPER SERIES

DP19656
(v. 4)

ENERGY-SAVING TECHNOLOGY SHOCKS, EMISSIONS, AND THE MACROECONOMY

Emanuel Moench and Soroosh Soofi Siavash

MACROECONOMICS AND GROWTH
AND CLIMATE CHANGE AND THE ENVIRONMENT

CEPR

ISSN 0265-8003

ENERGY-SAVING TECHNOLOGY SHOCKS, EMISSIONS, AND THE MACROECONOMY

Emanuel Moench and Soroosh Soofi Siavash

Discussion Paper DP19656
First Published 08 November 2024
This Revision 13 June 2025

Centre for Economic Policy Research
187 boulevard Saint-Germain, 75007 Paris, France
2 Coldbath Square, London EC1R 5HL
Tel: +44 (0)20 7183 8801
www.cepr.org

ENERGY-SAVING TECHNOLOGY SHOCKS, EMISSIONS, AND THE MACROECONOMY

Abstract

We use restrictions derived from frontier models of directed technical change to identify an energy-saving technology shock in a Bayesian structural VAR of the U.S. economy. This shock is associated with a persistent reduction of the carbon intensity of output. It also leads to a delayed but strong increase of GDP which gives rise to substantial additional fossil fuel consumption and new emissions. As a result, per capita emissions fully rebound after an initial decline. These effects can largely be attributed to a substitution of fossil fuel end-use by electricity, much of which has historically been generated using fossil fuels.

JEL Classification: C32, O47, Q43, Q55

Keywords: Energy-saving technology shocks, Carbon emissions, Structural vector autoregressions

Emanuel Moench - e.moench@fs.de
Frankfurt School of Finance & Management and CEPR

Soroosh Soofi Siavash - ssoofi@lb.lt
Bank of Lithuania, Vilnius University

Acknowledgements
This paper has previously been circulated under the titles “Carbon Intensity, Productivity, and Growth” and “The dynamic effects of green and non-green technology shocks on emissions and the macroeconomy.” We thank Conny Olovsson, David Popp, Jesper Lindé, André Kurmann, Nida Çakır Melek, Niels Moeller, Ol, egs Matvejevs, and Simona Ferraro for helpful feedback. For helpful comments, we would also like to thank seminar participants at the Bank of Lithuania, the University of Hamburg, the Bundesbank, the European Central Bank, the Eurosystem Research Cluster on Climate Change, the Bank of Estonia, the Austrian National Bank, Leibniz University Han- nover, the International Monetary Fund, the Potsdam Institute for Climate Impact Research, the Leibniz Institute for Financial Research SAFE, and conference participants at CefES, IAAE, IPAG, ISEFI, and BEC. The views expressed here do not reflect the views of the Bank of Lithuania or the Eurosystem.

Energy-Saving Technology Shocks, Emissions, and the Macroeconomy*

Emanuel Moench†

Frankfurt School of Finance & Management, CEPR

Soroosh Soofi-Siavash‡

Bank of Lithuania, Vilnius University

June 5, 2025

Abstract

We use restrictions derived from frontier models of directed technical change to identify an energy-saving technology shock in a Bayesian structural VAR of the U.S. economy. This shock is associated with a persistent reduction of the carbon intensity of output. It also leads to a delayed but strong increase of GDP which gives rise to substantial additional fossil fuel consumption and new emissions. As a result, per capita emissions fully rebound after an initial decline. These effects can largely be attributed to a substitution of fossil fuel end-use by electricity, much of which has historically been generated using fossil fuels.

Keywords: energy-saving technology shocks, carbon emissions, structural vector autoregressions
JEL: C32, O47, Q43, Q55

*This paper has previously been circulated under the titles "Carbon Intensity, Productivity, and Growth" and "The dynamic effects of green and non-green technology shocks on emissions and the macroeconomy." We thank Conny Olovsson, David Popp, Jesper Lindé, André Kurmann, Nida Çakır Melek, Niels Moeller, Oļegs Matvejevs, and Simona Ferraro for helpful feedback. For helpful comments, we would also like to thank seminar participants at the Bank of Lithuania, the University of Hamburg, the Bundesbank, the European Central Bank, the Eurosystem Research Cluster on Climate Change, the Bank of Estonia, the Austrian National Bank, Leibniz University Hannover, the International Monetary Fund, the Potsdam Institute for Climate Impact Research, the Leibniz Institute for Financial Research SAFE, and conference participants at CefES, IAAE, IPAG, ISEFI, and BEC. The views expressed here do not reflect the views of the Bank of Lithuania or the Eurosystem.

†Email: e.moench@fs.de

‡Email: ssoofi@lb.lt

1 Introduction

Climate change induced by anthropogenic emissions of carbon dioxide (CO2) and other greenhouse gases is considered as one of the most important threats to economic prosperity and well-being worldwide. Only a rapid reduction of carbon emissions can limit the increase of global mean temperatures to two 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 output has to decline substantially. This will require fast technological progress aimed at reducing the use of fossil fuels in economic activities. Understanding the implications of this technological transition for economic growth and the path of emissions is of first-order importance.

In this paper, we study the effects of energy-saving technological innovations on emissions and the macroeconomy. We first derive identifying restrictions from frontier models of directed technical change by Hassler et al. (2021, 2022) and Casey (2023). We then use these restrictions to identify an energy-saving technology shock in a structural vector autoregressive (VAR) model of the U.S. economy. Specifically, we recover two innovations that jointly explain the bulk of low frequency variation in total factor productivity (TFP) and energy intensity. We then restrict these innovations to satisfy sign restrictions implied by the model. The energy-saving technology (EST) shock raises TFP and is required to lower the energy intensity of output as well as the income share of energy in the short to medium run. We contrast this shock with an orthogonal technology shock which is required to raise TFP and the energy income share. We label this a non-energy-saving technology (NEST) shock. Methodologically, our identification builds on Angeletos, Collard and Dellas (2020) who extend the max-share approach of Uhlig (2003) to the frequency domain.

We find that the EST shock is associated with a persistent reduction of the fossil fuel intensity of U.S. output, defined as fossil fuel consumption relative to GDP. Importantly, however, fossil fuel consumption only declines temporarily and then quickly reverts. This rebound is driven by real output which responds little on impact, but then increases strongly and remains elevated, leading to additional energy consumption. Hence, we document a *Jevons paradox* in aggregate U.S. data: a technology shock which persistently lowers the energy intensity of output is associated with a temporary *increase* of fossil fuel consumption.

The response to a NEST shock is very different. It is associated with an immediate hump-shaped increase of fossil fuel consumption which dissipates after around five years. This is driven by temporarily higher real output as well as TFP. Accordingly, we find that the NEST shock explains larger fractions of TFP and output at short and intermediate horizons, while the EST shock is a much more important driver of TFP and output at longer horizons. Importantly, while Hassler et al. (2021) and others document that fossil fuel price increases have historically led to increased energy-saving innovation, we show that our fossil energy-saving technology shock is not contaminated by energy price shocks. Orthogonalizing with respect to a shock that explains the bulk of business cycle variation in real energy prices delivers the same results.

Moreover, while our baseline analysis relies on the intensity of fossil energy and its income share as target variables, our results are essentially unchanged when we use a broader measure of energy which additionally includes renewable and nuclear energy and excludes electrical system energy losses.

We corroborate the identified EST shock as capturing innovations which lead to a more efficient use of energy in two ways. First, we show that the shock series correlates strongly with observable measures of energy-related innovations based on triadic patent data as well as government research, development and demonstration (RD&D) in energy-related technologies. Second, we estimate an alternative structural VAR which directly includes indicators of input-saving technologies. These are backed out from a constant elasticity of substitution (CES) function in a capital/labor composite and fossil energy following Hassler et al. (2021). We replace TFP and fossil energy intensity with the two input-saving technology series and achieve identification by maximizing their long-run variation and applying sign restrictions. This alternative identification yields impulse responses and variance contributions for the EST and NEST shocks which closely mimic those obtained in our baseline analysis.

We shed further light on our finding of a rebound effect in fossil fuel consumption following an EST shock by asking two questions. First, what role do changes in the energy mix of the U.S. economy play in explaining this rebound? Second, what did the rebound imply for carbon emissions? To answer these questions, we decompose the carbon emission intensity for the aggregate U.S. economy into the emission intensity of fossil fuel consumption, fossil fuels embodied in electricity relative to those directly consumed by end-use sectors, and the fossil fuel end-use intensity of output. We show that while the U.S. economy has reduced its reliance on direct fossil fuels over the past several decades, it has increasingly used them in electricity production. Moreover, while the emission intensity of fossil fuel use was fairly stable over most of our 1973-2019 sample, it has substantially declined in recent years, consistent with the shale gas boom. Not surprisingly, the responses of emission intensity and its components to the EST shock closely match those of the fossil fuel intensity. The decomposition further shows that the EST shock is followed by a persistent increase of fossil fuel consumption in electricity production. Our results thus highlight that energy-saving technological innovations have primarily led to a substitution away from direct fossil fuel end-use to an indirect use in the generation of electricity.

We complement our baseline findings based on aggregate U.S. data with a sectoral analysis. Specifically, we decompose emissions per capita in each of the four end-use sectors "industrial", "residential", "commercial", and "transportation" into three components: the emissions associated with total fossil fuel use in the respective sector, the fossil fuel embodied in electricity used by the sector, and the sector's direct fossil fuel consumption per capita. We then estimate auxiliary VAR models, individually adding these components. We find that in all sectors except for Transportation, per capita emissions initially decline in response to an EST shock, but then follow a hump-shaped increase. This rebound of per capita emissions mimics the response of per capita fossil fuel consumption. Moreover, the use of fossil fuels embodied

in electricity consumption persistently increases in all end-use sectors, again with the exception of Transportation. In that sector, electricity accounted for a very small share of energy consumed over our sample period. Hence, while the EST shock is associated with a substitution away from fossil fuels towards electricity, historically a sizable fraction of this additional electricity has been produced using fossil fuels. Consequently, this compositional change as a result of technological innovations did not contribute to the reduction of emissions in our sample.

Importantly, while we find that energy-saving technological advances did not lead to a reduction of emissions in the medium run, we also show that they only explain less than a third of their medium-run variation. Hence, other factors than technology must have caused the observed decline of per capita emissions in the U.S. economy in recent decades. Prior research suggests several potential alternative drivers of this decline. First, U.S. consumers may have shifted their consumption towards less carbon-intensive goods and services (Edenhofer et al. 2014). Second, high-income countries such as the U.S. have increasingly outsourced emission-intensive production and reduced emissions at home (Copeland et al. 2022). Finally, although there have been limited regulations with regards to CO2 emissions at the national level in our sample period, U.S. firms have been increasingly required to install effective pollution abatement technologies by federal and state agencies (Henderson 1996; Chay and Greenstone 2005). Since different types of pollution are correlated, these regulations might have indirectly contributed to a reduction in carbon emissions (Shapiro and Walker 2018).

We subject our findings to a host of robustness checks. First, we show that orthogonalizing the two technology shocks with respect to a business cycle frequency fossil energy-price shock does not change our results. Second, we document that they are also essentially unchanged when we rely on a broad measure of energy consumption that additionally includes renewable and nuclear energy and excludes electrical system energy losses. Third, we illustrate that our results are robust to considering departures from the baseline model in terms of the values for the hyperparameters used in the Bayesian estimation of the VAR, the horizon used for imposing sign restrictions, and considering different subsamples. Fourth, we also show robustness with respect to recovering the long-run innovations in the time rather than the frequency domain. Specifically, we follow Francis, Owyang, Roush and DiCecio (2014) and Kurmann and Sims (2021) and alternatively recover two innovations as the major driver of fossil energy intensity and TFP at a long but finite horizon. The resulting macroeconomic dynamics are essentially identical to those implied by our baseline frequency domain identification.

Our paper is related to several strands of the literature. The first uses microeconomic data to document a strong negative relationship between productivity and emissions, albeit without disentangling between different forms of input-saving technological changes (Bloom, Genakos, Martin and Sadun 2010; Cui, Lapan and Moschini 2016; Holladay 2016; Shapiro and Walker 2018; Forslid, Okubo and Ulltveit-Moe 2018). A second strand develops models of endogenous, directed technical change with energy-saving technology (Hassler, Krusell and Olovsson 2021, 2022; Casey 2023). Hassler et al. (2021) aim at understanding the income share of energy

in the U.S. in the long-run. A similar model is used in Hassler et al. (2022) to study the use of oil at the global level and Casey (2023) who builds on this framework to study the long-run effects of climate policies on energy consumption.[1] We rely on the models of Hassler et al. (2021) and Casey (2023) to derive identifying restrictions which allow us to empirically isolate shocks to input-saving technologies for energy versus other inputs. Applying these restrictions in a quarterly structural VAR of the U.S. economy, we then study their dynamic effects on emissions and the macroeconomy.

The third strand is the literature that employs structural VARs to disentangle different types of technology shocks. Examples include Fisher (2006), Galí and Gambetti (2009), and Altig, Christiano, Eichenbaum and Lindé (2011). In line with these studies, we use a representation of the data in which two underlying technology shocks capture the bulk of the long-run variation in TFP. We then attribute these innovations to two orthogonal structural shocks with different implications for the efficient use of various production inputs, including energy. We find that the energy-saving technology shock, which is associated with a persistently lower fossil fuel income share, substantially contributes to the longer-run movements in TFP and real activity and implies impulse responses similar to TFP news shocks identified in the prior literature (e.g. Kurmann and Sims 2021). The orthogonal NEST shock, in turn, increases the energy income share in the medium-run and features mean-reverting dynamics of real output similar to the surprise TFP innovations identified by Barsky and Sims (2011) and Amir-Ahmadi and Drautzburg (2021).[2]

A related recent paper is Känzig and Williamson (2024). These authors use a max-share approach in the time domain to identify first a fossil energy price shock, second a residual shock that explains the maximal medium-run share of variation in capital/labor augmenting technology while being orthogonal to the energy price shock, and third a residual energy-saving shock that explains the most of the medium-run variation of energy-saving technology while being orthogonal to the other two shocks. They find that the energy-saving technology shock accounts for only around 10-20 percent of the business cycle variation in output and fossil energy consumption. Consistent with our results, Känzig and Williamson (2024) find that a large share of the variation in fossil energy consumption is unexplained by the energy-saving technology shock. In contrast to their findings, however, our *joint* identification of energy-saving and non-energy-saving technology shocks implies a substantially larger contribution of the EST shock for output and its components.

[1]The mechanism described in these models finds support in micro-empirical studies which document that higher energy prices cause an increase in energy-patenting intensity and raise R&D spending in polluting firms (Popp 2002; Aghion, Dechezleprêtre, Hémous, Martin and Van Reenen 2016; Brown, Martinsson and Thomann 2022).

[2]In a related paper, Khan, Metaxoglou, Knittel and Papineau (2019) estimate structural VARs on U.S. macroeconomic and emission data. They document that a significant fraction of the variation in per capita emissions is unexplained by news and surprise technology shocks. We explicitly separate between two input-saving technologies and find that both types of technology shocks combined explain only a moderate share of per capita carbon emissions. In addition, energy-saving technology innovations are associated with a hump-shaped rebound in aggregate carbon emissions in recent decades.

The fourth strand is the literature aiming to quantify the rebound effect in fossil energy use and carbon emissions after increases in energy efficiency. As documented in review articles by Gillingham, Rapson and Wagner (2016) and Brockway et al. (2021), while different theoretical mechanisms have been proposed to generate rebound effects, the literature is inconclusive about their macroeconomic importance. To the best of our knowledge, our study is one of the first to document an economy-wide rebound effect using state-of-art macroeconometric methods. An exception is Bruns, Moneta and Stern (2021). These authors explore the effects of an energy efficiency shock, which they identify as an innovation that generates a contemporaneous reduction in energy use and is orthogonal to innovations to energy prices and GDP. In contrast, we derive identifying restrictions from frontier models of directed technical change and provide evidence of a rebound effect on a macroeconomic scale following improvements in energy efficiency. Our results are also consistent with Bolton et al. (2023) who document using patent data that fuel efficiency innovations lower the emission intensity but at the same time lead to more sales and investment and – ultimately – higher emissions.

The remainder of this paper is organized as follows. Section 2 summarizes the econometric methodology used to identify the shocks. Section 3 describes the data and discusses the specification of the VAR model. In Section 4, we then present the results of our analysis and provide robustness checks. Section 5 concludes.

2 Identification of an Energy-Saving Technology Shock

This section presents our empirical approach to identifying an energy-saving technology shock in a structural VAR. We estimate the VAR using standard Bayesian methods and the Minnesota prior. Online Appendix A details the estimation approach. Given the posterior distribution of the VAR parameters, we recover two innovations that explain the bulk of low frequency variations in the intensity of energy use per unit of output and TFP. In a second step, we rotate these innovations into two orthogonal structural shocks. We achieve identification by means of sign restrictions. An energy-saving technology shock is a shock that lowers energy intensity and the income share of energy. The orthogonal non-energy-saving technology shock, in turn, increases TFP and the energy income share in the medium-run.

2.1 Assumptions Underlying the Identification Approach

We now discuss three assumptions that are sufficient to disentangle between the two shocks.

ASSUMPTION 1. Output is determined by a CES production function as considered e.g. in Hassler et al. (2021) and Casey (2023)

$$y_t = F(A_t k_t^{\alpha} l_t^{1-\alpha}, A_{et} e_t) = \left[(1-\gamma) \left(A_t k_t^{\alpha} l_t^{1-\alpha} \right)^{\frac{\varepsilon-1}{\varepsilon}} + \gamma (A_{et} e_t)^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}} \tag{1}$$

where y_t is output, k_t and l_t are capital and labor, and e_t is (fossil) energy use. The two variables A_t and A_{et} are the input-saving technology levels for capital/labor and energy, respectively. γ is a share parameter determining the relative importance of the two factors. Consistent with the models of Hassler et al. (2021) and Casey (2023), we assume that there is a low short-run substitutability between energy and the capital/labor input, i.e. $\varepsilon < 1$.

ASSUMPTION 2. The ratio of energy to the capital/labor composite (measured in efficiency units), $A_{et}e_t/A_t k_t^{\alpha} l_t^{1-\alpha}$, follows a stationary stochastic process.

This assumption is consistent with the balanced-growth path of standard macroeconomic models which features identical growth rates for output, consumption and capital. Since the production function is homogenous-of-degree-one, in steady state both factors must then grow at the rate of output.

By combining these two assumptions, we can write the following two expressions for energy intensity and the ratio of output over capital/labor composite

$$\frac{y_t}{e_t} = A_{et}\frac{y_t}{A_{et}e_t} = A_{et}F\left(\frac{A_t k_t^{\alpha} l_t^{1-\alpha}}{A_{et}e_t}, 1\right) = A_{et}\left[(1-\gamma)\left(\frac{A_t k_t^{\alpha} l_t^{1-\alpha}}{A_{et}e_t}\right)^{\frac{\varepsilon-1}{\varepsilon}} + \gamma\right]^{\frac{\varepsilon}{\varepsilon-1}} \tag{2}$$

$$\frac{y_t}{k_t^{\alpha} l_t^{1-\alpha}} = A_t\frac{y_t}{A_t k_t^{\alpha} l_t^{1-\alpha}} = A_t F\left(1, \frac{A_{et}e_t}{A_t k_t^{\alpha} l_t^{1-\alpha}}\right) = A_t\left[(1-\gamma) + \gamma\left(\frac{A_{et}e_t}{A_t k_t^{\alpha} l_t^{1-\alpha}}\right)^{\frac{\varepsilon-1}{\varepsilon}}\right]^{\frac{\varepsilon}{\varepsilon-1}} \tag{3}$$

or, in logs,

$$ln\left(\frac{y_t}{e_t}\right) = ln\left(A_{et}\right) + \xi_t \tag{4}$$

$$ln\left(\frac{y_t}{k_t^{\alpha} l_t^{1-\alpha}}\right) = ln\left(A_t\right) + \psi_t \tag{5}$$

where $\xi_t = ln\left[F(A_t k_t^{\alpha} l_t^{1-\alpha}/A_{et}e_t, 1)\right]$ and $\psi_t = ln\left[F(1, A_{et}e_t/A_t k_t^{\alpha} l_t^{1-\alpha})\right]$ are stationary under the above-mentioned assumptions. Equations (4)-(5) are key to our identification approach. They imply that the only sources of long-run variations in the output intensity of energy use and TFP are permanent shocks to A_{et} and A_t.

Notice that these two assumptions allow for A_e and A to be endogenous to one another– consistent with the notion that they are jointly determined in the longer-run. We rely on a two-shock representation of the long-run growth in A_e and A. To illustrate this, we follow Hassler et al. (2021) where directing R&D resources to increase the growth of one form of technology

comes at the expense of lowering growth for the other. As a result, there is a trade-off between the two forms of input-saving given by

$$G\left(\frac{1}{exp\left(\nu_{At}\right)}\frac{A_{t+1}}{A_t},\frac{1}{exp\left(\nu_{A_et}\right)}\frac{A_{et+1}}{A_{et}}\right)=0, \tag{6}$$

where ν_{A_et} is the energy-saving technology shock and ν_{At} is the capital/labor augmenting technology shock with mean zero and variance $\sigma^2_{A_e}$ and σ^2_A, respectively. Here, G is strictly increasing in both arguments. In a structural VAR, it would be impossible to separate these two shocks by simply constructing two innovations that explain the low frequency variations in energy intensity and TFP. The identification of the two shocks must therefore come from additional identifying restrictions. We rely on the following assumption to achieve identification.

ASSUMPTION 3. The medium-run income share of (fossil) energy falls following a shock to A_{et}, whereas it increases following a shock to A_t.

This assumption is consistent with the medium-run implications of energy-saving technological changes in the models of Hassler et al. (2021) and Casey (2023). Consistent with these models, an increase in A_{et} lowers the income share of energy in subsequent periods.[3] However, it also lowers the incentives to allocate R&D resources to improve the growth rate of A_{et}. Thus, the growth rate of A_{et} decreases and the growth rate of A_t increases. Both of these effects will then contribute to bringing the energy share back to its level on the balanced growth path. By a similar reasoning, an increase in A_t will lead to the opposite conclusion.

Taken together, Assumptions 1-3 provide sufficient restrictions to jointly identify the shocks to A_{et} and A_t. In a VAR including energy intensity, TFP and the energy share, we jointly identify the two shocks as linear combinations of two innovations that explain the bulk of the long-run variations in energy intensity and TFP, conditionally on satisfying sign restrictions on the impulse response functions. Given the link between the shocks and the long-run innovations to energy intensity and TFP, we naturally impose the sign of the response of energy intensity to be negative for an EST shock and the response of TFP to be positive for a NEST shock. Moreover, based on Assumption 3 we disentangle between the two shocks by imposing the sign restrictions that an EST shock lowers the energy income share while a NEST shock increases it. We impose these sign restrictions on the average of the impulse response functions from the moment when the shock hits up to some medium-run horizon, $0 \leq h \leq \bar{h}$. Table 1 summarizes the identifying restrictions.

The method we use to recover such long-run innovations closely follows the approach of Angeletos et al. (2020). These authors build on the approach of Uhlig (2003) that is widely

[3]The income share of energy is given by the following expression $e_t^{share} = \frac{p_t e_t}{y_t} = \left(\frac{(1-\gamma)}{\gamma}\left(\frac{A_{et}e_t}{A_t k_t^{\alpha} l_t^{1-\alpha}}\right)^{\frac{1-\varepsilon}{\varepsilon}}+1\right)^{-1}$.

Table 1: IDENTIFYING RESTRICTIONS IN THE BASELINE VAR

	SHOCK	
VARIABLE	EST SHOCK	NEST SHOCK
Fossil energy intensity	$-$ (avg. over horizons $0 \leq h \leq \bar{h}$)	n/a
TFP	n/a	$+$ (avg. over horizons $0 \leq h \leq \bar{h}$)
Fossil income share	$-$ (avg. over horizons $0 \leq h \leq \bar{h}$)	$+$ (avg. over horizons $0 \leq h \leq \bar{h}$)
Other variables	n/a	n/a

Notes: Restrictions on the average impulse responses of the variables over the horizons $0 \leq h \leq \bar{h}$ for the EST and NEST shocks. Respectively, $\pm$ and n/a denote the sign restrictions and unrestricted responses.

used for identification of long-run technology shocks in structural VARs (Barsky and Sims 2011; Kurmann and Sims 2021; Francis et al. 2014), and recover innovations that explain most of the variation of some target variable of interest in a particular frequency band. We then achieve identification by imposing restrictions on the sign of the impulse response functions following Uhlig (2005) and Rubio-Ramírez et al. (2010).

2.2 Econometric Approach

Let X_t denote an $n \times 1$ vector of quarterly time series which contains the energy intensity of U.S. output, TFP, the income share of energy, and a number of additional variables, which were chosen to represent different aspects of real activity, inflation and financial markets in the U.S. economy. We model these time series to have the VAR representation

$$A(L)X_t = \eta_t, \tag{7}$$

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 (7), one obtains the reduced-form moving average representation which expresses X_t in terms of current and past values of innovations:

$$X_t = A(L)^{-1}\eta_t. \tag{8}$$

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:

$$\eta_t = H\nu_t. \tag{9}$$

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

$$X_t = B(L)Q\nu_t, \text{ with } B(L) = A(L)^{-1}Chol(\Sigma_\eta), \tag{10}$$

where the impulse response function of X_t with respect to the ith shock is given by $B(L)Q_i$ with Q_i denoting 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.

The contribution of the mth shock to the spectral density of X_{jt} over the frequency band $[\underline{\omega}, \overline{\omega}]$ is given by $\int_{\omega\in[\underline{\omega},\overline{\omega}]} \left(\overline{B^{[j]}\left(e^{-i\omega}\right)Q_m}B^{[j]}\left(e^{-i\omega}\right)Q_m\right)$, 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_m is the mth column of Q, and $\bar{x}$ denotes the conjugate transpose of x. The relative contribution of shock m to the variation of variable j is then

$$\begin{aligned}\Omega_{j,m}(\underline{\omega},\overline{\omega}) &= \frac{\int_{\omega\in[\underline{\omega},\overline{\omega}]}\left(\overline{B^{[j]}\left(e^{-i\omega}\right)Q_m}B^{[j]}\left(e^{-i\omega}\right)Q_m\right)}{\int_{\omega\in[\underline{\omega},\overline{\omega}]}\left(B^{[j]}\left(e^{-i\omega}\right)\overline{B^{[j]}\left(e^{-i\omega}\right)}\right)} \\ &= Q_m'\left(\frac{\int_{\omega\in[\underline{\omega},\overline{\omega}]}\left(\overline{B^{[j]}\left(e^{-i\omega}\right)}B^{[j]}\left(e^{-i\omega}\right)\right)}{\int_{\omega\in[\underline{\omega},\overline{\omega}]}\left(B^{[j]}\left(e^{-i\omega}\right)\overline{B^{[j]}\left(e^{-i\omega}\right)}\right)}\right)Q_m \qquad (11)\end{aligned}$$

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

We seek to identify two shocks that jointly explain the maximum shares of longer-run variation in two variables: TFP and the ratio of energy consumption over real output. We capture the longer-run by focusing on frequencies longer than 80 quarters ($0 \leq \omega \leq 2\pi/80$). To disentangle between an EST and NEST shock, we then rely on sign restrictions on the average of the impulse responses up to 80 quarters ahead. Let the EST shock be indexed by 1 and the NEST shock by 2. To achieve identification, we start in step (1) with an orthonormal matrix $Q = [Q_1\ Q_2\ Q_\bullet]$ with its first two columns selected by solving the following optimization problem

$$\begin{aligned}\underset{Q_1,Q_2}{\operatorname{argmax}}\ &\Omega_{\text{energy intensity},1}(0,2\pi/80) + \Omega_{\text{energy intensity},2}(0,2\pi/80) + \\ &\Omega_{\text{TFP},1}(0,2\pi/80) + \Omega_{\text{TFP},2}(0,2\pi/80), \qquad (12)\end{aligned}$$

subject to the restrictions $Q_1'Q_1 = 1$, $Q_2'Q_2 = 1$ and $Q_2'Q_1 = 0$. This implies that Q_1 and Q_2 are the eigenvectors associated with the first two largest eigenvalues of $S_{\text{energy intensity}}(0,2\pi/80) + S_{\text{TFP}}(0,2\pi/80)$. In step (2), we then post-multiply Q by a $n \times n$ matrix $W = \begin{pmatrix} W_{1:2} & 0 \\ 0 & I_{n-2}\end{pmatrix}$ where $W_{1:2}$ is a 2×2 orthonormal matrix that is obtained from the QR decomposition of a matrix of the same dimensions with elements drawn from the standard normal distribution. Letting $\mathrm{r}_{j,m}(h) = C_{h,jm}$ denote the resulting impulse response function for variable j for shock m where C_h is the hth lag matrix in $C(L) = B(L)QW$, we identify the two shocks by repeating step (2) multiple times until the resulting impulse response functions satisfy the sign restrictions in

Table 1. The restrictions are therefore given by

$$\begin{cases} \frac{1}{80}\sum_{h=0}^{79} \mathrm{r}_{\text{ energy intensity,1}}(h) \leq 0, \ \frac{1}{80}\sum_{h=0}^{79} \mathrm{r}_{\text{ energy share,1}}(h) \leq 0; \\ \\ \frac{1}{80}\sum_{h=0}^{79} \mathrm{r}_{\text{ TFP,2}}(h) \geq 0, \ \frac{1}{80}\sum_{h=0}^{79} \mathrm{r}_{\text{ energy share,2}}(h) \geq 0. \end{cases} \tag{13}$$

While our approach of disentangling energy-saving and non-energy-saving technology shocks shares some features with that in Känzig and Williamson (2024), there are at least three important differences. First, we *jointly* identify two shocks instead of identifying the EST shock as a residual shock that is orthogonal to other innovations. This allows us to disentangle energy-saving from non-energy-saving technological innovations without imposing a specific ordering. Second, we use the observed energy intensity and TFP directly as target variables rather than relying on series of input-saving technology series backed out from a production function. We document the robustness of our results to varying the target variables in Sections 4.2 and 4.4.2. Third, in line with Angeletos et al. 2020 we identify shocks in the frequency domain. This approach avoids relying on a specific target horizon in the time domain to construct the long-run technology innovations (Francis et al. 2014; Kurmann and Sims 2021). We document the robustness of our results when using a time-domain approach in Section 4.4.3.

3 Data and VAR Specification

The data we employ to estimate our baseline VAR consist of quarterly observations on 13 variables. Following the model of Hassler et al. (2021) and to account for the fact that U.S. energy consumption has been dominated by fossil fuels in our sample, our baseline specification relies on the intensity of fossil energy and its income share as target variables. In Section 4.4.2 we document the robustness of our main findings with respect to a broader measure of energy consumption, which additionally includes renewable and nuclear energy and excludes electrical system energy losses.

For our baseline analysis, we measure the energy intensity of U.S. output as the ratio of fossil energy use over real GDP, both expressed in per capita terms. We further construct a composite index of real fossil fuel prices that includes coal, natural gas and petroleum, closely following Hassler et al. (2021). The remaining ten variables are chosen to represent key macroeconomic and financial aggregates of the U.S. 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 (PCE), real private fixed investment, the change in private inventory investment, the federal funds rate, PCE inflation, the real S&P 500 index and a trade-weighted index of real exchange rates where the latter two are deflated by the GDP deflator.

We also entertain several extensions of the baseline VAR to study the effects of the identified shocks on the breakdown of CO2 emissions and fossil fuel consumption by end-use sectors. Specifically, we estimate larger VARs, each additionally including a variable capturing the per capita value of CO2 emissions and end-use fossil fuel consumption, and the ratio of the fossil fuel embodied in electricity use over the end-use of fossil fuel in each of the following sectors: Industrial, Residential, Commercial, and Transportation.

We corroborate our identification by showing that the estimated energy-saving technology shock series is highly correlated with future growth in several measures of energy-related technological innovation. Specifically, we consider indicators of energy-saving patents and government research, development and demonstration (RD&D) budget for energy-related technologies. We obtain the former as the number of *triadic* patents which are related to technologies related to mitigation and adaptation against climate change and improved power network operation. These are subsumed under the Y02/Y04S scheme in the standard patent classification (see Veefkind et al. 2012; Angelucci et al. 2018, Popp 2019). Our patent data are drawn from the February 2022 version of the OECD Triadic Patent Families database.[4] We also obtain two measures of government energy RD&D budget. The first measure involves RD&D for all energy technologies. The second captures only those technologies that are related to production, conversion and combustion of fossil fuels and CO2 capture and storage, which are classified as "Group 2–Fossil Fuels: Oil, Gas and Coal" (see IEA 2011). The government RD&D budget data are obtained from the International Energy Agency (IEA) Energy Technology RD&D Budgets database.[5] The data sources and the construction of these public energy RD&D measures are detailed in Online Appendix B.

Moreover, we obtain monthly emission and energy consumption data from the U.S. Energy Information Administration (EIA). We first seasonally adjust all of these series using the X-12 method and then convert them to quarterly values by summing over the quarter's three monthly values. Finally, the macroeconomic indicators were all obtained from the Federal Reserve Economic Database (FRED) and the Bureau of Economic Analysis (BEA). We include the PCE inflation rate, the Federal Funds rate, and the change in private inventory investment in percent and all of the remaining variables in natural logs.

Our baseline sample starts in 1973:I which is the earliest date for which the fossil fuel and emission data are available quarterly, and ends in 2019:IV, and thus before the COVID-19 pandemic started. The VARs are estimated with four lags using Bayesian methods subject to

[4]The dataset is made available through the OECD Intellectual property (IP) statistics and analysis website at https://www.oecd.org/sti/intellectual-property-statistics-and-analysis.htm#ip-data. We follow standard practice and use a count of all triadic patents by application date which are filed in each quarter in the U.S. Patent office (USPTO), the European Patent Office (EPO), and the Japanese Patent Office (JPO). While we consider patents that have been filed in all three major patent offices, we restrict our sample to patents granted by the USPTO. Until 2001, the USPTO has published only granted patent applications. To have a consistent series before and after 2001, we counted only those triadic patents which have been granted by the USPTO. For more detail on triadic patents and uses of patent statistics, see Dernis and Khan (2004), Griliches (1990), OECD (2009).

[5]The database is made available at https://www.iea.org/data-and-statistics/data-product/energy-technology-rd-and-d-budget-database-2.

a Minnesota prior. We generate 20,000 draws from the posterior distribution via the Gibbs sampler, where we discard the first 4,000 as burn-in and keep every fourth draw from the subsequent 16,000. In Online Appendix A, we provide details on the Bayesian estimation. In Section 4.4, we further show that our main findings are robust to using alternative values for the hyperparameters, and a sample period excluding the shale gas boom.

4 Results

In this section we document our empirical results. We provide impulse response functions (IRFs) and forecast error variance decompositions (FEVDs) for the energy-saving and non-energy-saving technology shock in Section 4.1. Section 4.2 documents that the estimated EST shock is indeed strongly correlated with measures of energy-saving innovations. In Section 4.3, we then shed light on the economic underpinnings of the rebound effect for fossil fuel consumption and carbon emissions that we document. Finally, section 4.4 shows the robustness of our baseline results to (i) orthogonalizing our shocks with respect to a fossil energy price shock; (ii) adding renewable and nuclear energy; (iii) the use of different hyperparameter values for the Bayesian estimation, different horizons for the sign-restrictions, different sample periods and an alternative approach to recovering long-run innovations.

4.1 Energy-Saving and Non-Energy-Saving Technology Shocks

The top panel of Figure 1 provides the posterior median IRF of the variables in our baseline VAR for the two shocks, along with the 16-84 percent posterior coverage intervals. The top-left chart shows that on impact the EST shock is associated with a strongly significant one percent decline of fossil energy intensity, measured as fossil fuel consumption per unit of output. While the fossil fuel intensity reverts some of the initial decline in the following quarters, it remains significantly compressed thereafter. In contrast, the NEST shock is associated with an *increase* of the fossil fuel intensity, which gradually returns to its initial value.

Let's turn to the response of TFP shown in the next chart to the right. Initially, the EST shock has no effect on TFP, but rises significantly within two years and remains strongly elevated even ten years after the shock. In sharp contrast, TFP jumps significantly on impact following an NEST shock, but then slowly reverts back to its pre-shock level. The negative response of the energy income share follows directly from the sign restrictions we impose. The response of the income share to an EST shock echoes the statistically significant decline of the fossil fuel price shown in the top-right chart of Figure 1. Importantly, the NEST shock is instead associated with a similarly persistent, protracted *rise* of energy prices.

Turning to the first chart in the second row, we see that fossil fuel consumption per capita strongly declines on impact following the EST shock, but then reverts back to its initial level within two years. Hence, we see a rebound effect of fossil fuel consumption: an energy-saving technology shock which persistently reduces the energy intensity of output does not lead to

persistently lower fossil fuel consumption. The response to the NEST shock is very different. It is associated with a pronounced initial increase of fossil fuel consumption which then slowly dissipates within five years of the shock. Consistent with the finding that the NEST shock is also associated with a protracted increase in fuel prices, fossil fuel consumption turns negative after some years.

The short-run response of fossil energy consumption for the two technology shocks is consistent with the directed technical change models of Hassler et al. (2021) and Casey (2023): energy-saving innovations allow economic agents to use less energy for operating capital and labor. In contrast, non-energy-saving innovations increase energy use, consistent with a low short-run substitutability between energy and the other production factors. In an alternative identification approach documented in Section 4.4.2, we back out two series representing energy-saving and capital/labor-augmenting technologies given a low short-run substitution elasticity between these input factors as suggested by Hassler et al. (2021). Identifying energy-saving and capital-labor-augmenting technology shocks from sign restrictions imposed on the IRFs of these two variables, we obtain essentially the same results.

Our finding that an EST shock only leads to a temporary decline of aggregate fossil fuel consumption followed by a hump-shaped increase represents novel evidence for a rebound effect in economy-wide fossil energy consumption.[6] Such rebound effects have first been discussed by Jevons (1865) related to the use of coal in England during the industrial revolution. He made the paradoxical observation 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, evidence for energy consumption at an aggregate level has been scarce thus far (Gillingham et al. 2016; Brockway et al. 2021). As we will see below, the rebound effect in fossil energy consumption that we document also gives rise to a rebound effect of carbon emissions. In Section 4.3, we explore this issue in more detail.

Why does fossil fuel consumption rebound although the fossil energy intensity of output is persistently reduced? The reason is that output features a pronounced hump-shaped increase in response to the EST shock, as shown in the second chart of the same row. While the initial response of output per capita is small, it rises gradually over the first five years and peaks at a level a little less than one percent above its initial value, before declining somewhat over the next several years. Ten years after the shock, output is still a strongly statistically significant 0.7 percent higher than its pre-shock value, highlighting the long-term impact of energy-saving technology 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 five years

[6] Bruns, Moneta and Stern (2021) use a reduced-form statistical approach to identify an energy-efficiency shock which is orthogonal to innovations to GDP and energy prices in a VAR for the U.S. economy. They find that overall energy consumption significantly drops on impact but then recovers in subsequent quarters following this shock.

after the shock and remains about 0.8 percent above its initial level at the ten year horizon. Fixed investment and hours worked show a pronounced hump-shaped response. Interestingly, the EST shock is associated with a significant negative response of inventory investment within two quarters. 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 to enhance efficiency. The inverse of the relative price of investment goods also rises persistently. This adds support to the EST shock being associated with gradual improvements in the quality of newly produced investment goods.

Let us now turn to the IRFs for the NEST shock. In sharp contrast to the results above, this shock only has a transitory impact on output and its components. Output, consumption and fixed investment peak within ten quarters before reverting to their pre-shock values, consistent with the persistent but transitory effect on TFP. Moreover, hours worked briefly decline on impact and then experience a short-lived hump-shaped increase. Similarly as for an EST shock, the inverse of the relative price of investment goods rises persistently.

How quantitatively important are the two technology shocks for business cycle variation? The variance decompositions provided in the bottom panel of Figure 1 paint quite different pictures. The EST shock explains more than 30 percent of the variation of fossil fuel consumption per unit of output at all horizons, highlighting its importance for energy efficiency. It is also an important driver of TFP. While it captures less than ten percent of TFP variation during the first five years, this share gradually increases to around 40 percent after ten years. Conversely, the NEST shock explains only small fractions of fossil fuel intensity, but initially a somewhat larger share of TFP variation. Combined, both technology shocks capture more than 80 percent of the longer-run variation in productivity and output. We also find that both shocks account for increasing fractions of the fossil energy income share and fuel prices at longer horizons, albeit with a larger contribution by the EST shock. The FEVDs for the remaining variables echo our previous findings that the EST shock is associated with pronounced and persistent macroeconomic dynamics, while the NEST shock contributes more at shorter horizons.

A key contribution of our analysis is to document that the EST and NEST shocks give rise to substantially different macroeconomic dynamics. Moreover, as evidenced in Figure 2, combined they capture large shares of variation of key macroeconomic aggregates. Much of the previous literature has focused on TFP news shocks as a key driver of longer-run variation in productivity and output (Beaudry and Portier (2006); Barsky and Sims (2011)). As a point of reference, we also construct a TFP news shock as an innovation explaining a maximum share of TFP variation over low frequencies of 80 quarters and beyond. Figure 2 shows the variance shares explained by this TFP news shock, along with those implied by the EST and NEST shocks, individually and combined. While the overall shares of variation captured by the two shocks together are similar to those explained by the TFP news shock for output, consumption and fixed investment, the EST and NEST shocks combined account for substantially larger variance shares for TFP itself and some other variables, most importantly the energy income share of output and energy prices. Hence, ignoring the different sources of productivity growth

one would underestimate the importance of technological change for longer-run productivity.

One might be concerned that the two identified technology shocks to some degree capture persistent demand-driven innovations. However, if that was the case, they would likely give rise to substantial business cycle variations in the price of fossil energy. As the top-right chart in the bottom panel of Figure 1 shows, this is not the case. Both shocks at best explain about 15 percent of the variation of energy prices at horizons one to eight years ahead. In Section 4.4, we explicitly orthogonalize the two shocks with respect to a fossil energy price shock and show that our results don't change.

Note also that our finding that energy-saving technology shocks – when identified *jointly* with capital/labor augmenting technology shocks – account for a large fraction of macroeconomic variation stands in contrast to the results in Känzig and Williamson (2024) who find that residual energy-saving technology shocks capture only about 10-20 percent of the medium-run variation in output. That said, these authors' finding that the two technology shocks cannot account for the bulk of the variation in fossil energy consumption is consistent with our results.

4.2 EST Shocks and Measures of Energy-Saving Technology

Our identification assumptions were informed by models of directed technical change and energy efficiency where the elasticity of substitution between energy and capital/labor is low. We now corroborate the interpretation of the identified EST shock as capturing energy-saving technological innovations. We do so in two complementary ways. First, we contrast the estimated shock series from our baseline VAR with measures of energy-related technological change. Specifically, we construct a measure of clean energy patenting intensity as the share of the number of Y02/Y04S patents to the overall number of patents using the OECD Triadic Patent Families database. As emphasized by Popp (2019), the Y02/Y04S classification offers the most reliable scheme for identifying clean energy technologies. Those patents cover technology subclasses related to adaptation to climate change (Y02A), buildings (Y02B), capture and storage of greenhouse gases (Y02C), ICT aiming at the reduction of own energy use (Y02D), production, distribution and transport of energy (Y02E), industry and agriculture (Y02P), transportation (Y02T), waste and wastewater (Y02W) and smart grids (Y04S). The analysis below focuses on the total number of Y02/Y04S patents, as well as the number of the patents covered in these subclasses. We also use measures of government RD&D in different energy technologies as a share of the NIPA government R&D investment from the IEA Energy Technology RD&D Budgets database.

To compare the dynamics of our identified EST shock series with these measures of innovation, we perform the following exercise. First, since the government RD&D budget series is only available at the annual frequency, we aggregate the quarterly median estimated shock series to the annual frequency by averaging over four quarters in each year. We then run the annual median shock through an AR(1) filter with autoregressive coefficient of 0.9 and use an OLS regression to describe the growth in the energy innovation indicators associated with the

Table 2: Identifying Restrictions in the Extended VAR

	SHOCK	
VARIABLE	EST SHOCK	NEST SHOCK
Energy-saving technology (A_e)	+ (avg. over horizons $0 \leq h \leq \bar{h}$)	n/a
Capital/labor augmenting technology (A)	n/a	+ (avg. over horizons $0 \leq h \leq \bar{h}$)
Fossil income share	− (avg. over horizons $0 \leq h \leq \bar{h}$)	+ (avg. over horizons $0 \leq h \leq \bar{h}$)
Other variables	n/a	n/a

Notes: Restrictions on the average impulse responses of the variables over the horizons $0 \leq h \leq \bar{h}$ for the EST and NEST shocks. Respectively, $\pm$ and n/a denote the sign restrictions and unrestricted responses.

smoothed shock series. To be precise, we regress k-year ahead growth of each energy innovation indicator on each of the smoothed series for the EST and NEST shocks with k equal to 2, 4, 6, 8 and 10 years. Figure 3 shows the estimated slope obtained from those regressions along with robust standard errors. The results are provided for those patent subclasses which have the largest correlation with our shock series. They show that the energy-saving technology shock is strongly associated with accelerated growth in patenting intensity in most of the clean energy technology subclasses, as well as with future growth in the public energy RD&D shares. In contrast, the NEST shock exhibits little correlation with these measures, except for the Y04S subclass which covers the mitigation technologies related to smart grids.

We complement this exercise with an alternative structural VAR analysis that explicitly includes measures of input-saving technologies. These technology indicators are backed out from the production function (1) by setting the substitution elasticity equal to $\varepsilon = 0.02$, in line with Hassler et al. (2021). We then replace TFP and energy intensity with the capital/labor augmenting and energy-saving technology series in the VAR model. Online Appendix C describes the details on the construction of these technology series. Figure 4 depicts the resulting energy-saving technology level (A_e) together with the capital/labor augmenting technology level (A).[7]

We use this model to identify shocks to A_e and A and contrast the resulting impulse responses with those estimated from our baseline specification. Specifically, we first recover two innovations that jointly target the variations in the two input-saving technology series over frequencies of 80 quarters and beyond. Second, we rotate them into an energy-saving and capital/labor augmenting technology shock by imposing sign restrictions such that (i) the energy-saving technology shock increases the level of energy-saving technology and lowers the fossil energy share, and (ii) a capital/labor augmenting technology shock increases the level of capital/labor augmenting technology and raises the fossil energy share. For both we maintain the sign-restricted horizons at $0 \leq h \leq 79$ quarters as in our baseline identification. These alternative identifying restrictions are summarized in Table 2.

Figure 5 provides scatter plots of the two sets of identified technology shocks across the

[7]This can be compared to the right panel in Figure 3 in Hassler et al. (2021) which is based on annual data for the period 1949- 2018. Instead, we use quarterly data on inputs and output for our sample from 1973:I to 2019:IV.

baseline and the alternative identification. Both approaches yield highly correlated identified shocks. Figure 6 shows the resulting median estimate of the impulse responses and the corresponding variance contributions for the alternative shock identification. We superimpose the corresponding results for the shocks identified in our baseline VAR as solid lines. The dashed lines in the first two charts show that the EST shock is associated with a marked and persistent increase in the energy-saving technology level and a delayed but long-lived increase in the capital/labor augmenting technology level. The properties of the remaining IRFs and FEVDs echo our previous findings. We interpret the fact that the IRFs and FEVDs across the two identification approaches are essentially identical as corroborating evidence that we capture the drivers of the two types of input-saving technologies implied by the production function in Equation (1).

4.3 EST Shocks and the Rebound Effect in Carbon Emissions

We have shown in the previous section that the use of fossil fuels initially declines following an energy-saving technology shock, but then quickly rebounds. In this section, we seek to answer two follow-up questions. First, what role did changes in the energy mix of the U.S. economy play in explaining this rebound effect? Second, what does this rebound of fossil fuel use imply for carbon emissions? To answer these questions, we first propose a decomposition of the carbon emission intensity and then estimate auxiliary VAR models which add carbon emissions at the aggregate and sectoral level in Sections 4.3.1 and 4.3.2.

4.3.1 A Decomposition of U.S. Fossil Energy Mix and Carbon Emissions

Consider the following decomposition of the carbon emission intensity of output, denoted by $\frac{\text{CO2}_t}{\text{Y}_t}$:

$$\frac{\text{CO2}_t}{\text{Y}_t} = \left(\frac{\text{CO2}_t}{\text{Fuel}_t}\right) \times \left(\frac{\text{Fuel}_t}{\text{Y}_t}\right). \tag{14}$$

The emission intensity equals emissions per total fossil fuel use, $\frac{\text{CO2}_t}{\text{Fuel}_t}$, multiplied by the fossil energy intensity of output, $\frac{\text{Fuel}_t}{\text{Y}_t}$. We can further decompose the total fossil fuel consumption into the end-use of fossil fuels and fossil fuels consumed by the electric power sector, i.e., $\text{Fuel}_t = \text{Fuel}_t^{End-Use} + \text{Fuel}_t^{EPowerSector}$. We can then write the fossil fuel intensity as equal to the ratio of total to end-use fossil fuel use, $\frac{\text{Fuel}_t}{\text{Fuel}_t^{End-Use}}$, times the end-use fossil fuel per unit of output $\frac{\text{Fuel}_t^{End-Use}}{Y_t}$. This yields

$$\frac{\text{CO2}_t}{\text{Y}_t} = \left(\frac{\text{CO2}_t}{\text{Fuel}_t}\right) \times \left(\frac{\text{Fuel}_t}{\text{Fuel}_t^{End-Use}}\right) \times \left(\frac{\text{Fuel}_t^{End-Use}}{\text{Y}_t}\right), \tag{15}$$

$$\text{where} \quad \frac{\text{Fuel}_t}{\text{Fuel}_t^{End-Use}} = 1 + \frac{\text{Fuel}_t^{EPowerSector}}{\text{Fuel}_t^{End-Use}}.$$

The first term on the right-hand side captures the carbon emission intensity of fossil fuel consumption. A decline in this component could be due e.g. to a compositional change between coal and natural gas where the latter produces fewer emissions per unit of energy consumed. The second term measures the fossil fuels embodied in electricity relative to fossil fuels directly consumed by end-use sectors. Finally, the third term captures the fossil fuel end-use intensity of output. The advantage of this decomposition is that it represents the carbon intensity of output in terms of fossil fuels consumed directly as well as indirectly through electricity.[8]

The top panel of Figure 7 depicts the evolution of the emission intensity of output along with its three components relative to their initial levels in 1973, the beginning of our sample. The blue and purple lines show the emission intensity and end-use fossil intensity of output, respectively. Both have seen a strong reduction of about 70% over the past five decades. That said, for much of the sample there has been a wedge between the two series, which is explained by the other two components in the decomposition. Specifically, the yellow line documents an increase in the fossil fuel consumption embodied in electricity relative to the end-use of fossil fuels. This shows that while the U.S. economy has relied less and less on direct fossil fuels over the past several decades, it has increased its reliance on fossil fuels via the electric power sector. Finally, the emission intensity of fossil fuel use (orange line) has been fairly stable over much of the sample, but has experienced a decline in recent years that is likely primarily associated with the shale gas revolution.

We now study the impulse responses of these components to an EST and a NEST shock. We obtain these by re-estimating our baseline VAR, adding one at a time each of the measures in logs. Hence, the sum of the three components equals the log emission intensity. Since the third component measuring the emission intensity per unit of total fossil fuel consumption does not move much over the sample, we drop it here for brevity. The results are shown in the bottom panel of Figure 7. The impulse responses for the emission intensity closely echo the result obtained for fossil fuel intensity documented above. Both experience a highly persistent decline of about one percent in response to the EST shock. Looking at the components of emission intensity, we see that the sharp initial decline of emission intensity is mainly driven by a persistent reduction of the end-use fossil fuel intensity of output. At the same time, the fossil fuel embodied in electricity sees a marked and persistent increase. Hence, the energy-saving technology shock is associated with a substitution from direct fossil fuel end-use to fossil fuels used in electricity.

In addition to the emission intensity of output, it is also instructive to study the impulse responses of emissions per capita. These are provided in the last column of Figure 7. The EST shock is associated with a sharp initial decline, followed by a strong hump-shaped rebound. It takes ten years before emissions per capita are back to their initial level after the shock. Not

[8]In 2020, fossil fuels accounted for approximately 60% of the total energy consumed by the electric power sector, with coal and natural gas being the primary inputs. The share of coal used to generate electricity relative to overall coal consumption by the U.S. economy amounted to 90%. The share of natural gas used by the electric power sector accounted for 40% of overall natural gas consumption. In turn, only about 1% of overall petroleum consumption in the U.S. was used for electricity generation.

surprisingly, this result closely mimics the corresponding dynamics of fossil fuel consumption.

The impulse responses for the NEST shock are quite different. The emission intensity rises somewhat on impact and then declines, but the response is not statistically significant. This increase is driven by a rise of both the end-use fossil fuel intensity as well as the fossil fuel use embodied in electricity. Moreover, per capita emissions increase immediately and stay elevated for about five years.

The top panel of Figure A.2 in the Online Appendix provides the corresponding FEVDs. Two comments are in order. First, the EST shock accounts for sizable fractions of the variance of both the end-use fossil fuel intensity and the emission intensity of U.S. output. At the ten-year ahead forecast horizon, about 50 percent of their variation is explained by the shock. Second, the bulk of variation of emissions per capita is left unexplained by the two shocks. Combined, they explain less then 50 percent of the variation of per capita emissions. This is consistent with Khan et al. (2019) who explore the effects of different news and surprise technology shocks and document that none explains more than one third of the variation of carbon emissions. It also suggests that other factors than technology must have caused the observed decline of per capita emissions in the U.S. in recent decades.

4.3.2 Rebound Effect in Sector-Level Data

We have documented that emissions per capita mimic the strong rebound effect of fossil fuel consumption. In this section, we further decompose emissions per capita at the sectoral level. Specifically, we write

$$\frac{\text{CO2}_t^i}{\text{Pop}_t} = \left(\frac{\text{CO2}_t^i}{\text{Fuel}_t^{End-Use,i} + \text{Electricity}_t^i \times \text{S}_t^{EPowerSector}} \right) \times \left(\frac{\text{Fuel}_t^{End-Use,i} + \text{Electricity}_t^i \times \text{S}_t^{EPowerSector}}{\text{Fuel}_t^{End-Use,i}} \right) \times \left(\frac{\text{Fuel}_t^{End-Use,i}}{\text{Pop}_t} \right), \tag{16}$$

for sector $i \in \{\text{transportation, industrial, residential, and commercial}\}$. Here, $\text{Fuel}_t^{End-Use,i}$ denotes the fossil fuel consumed directly by the respective end-use sector, and $\text{Electricity}_t^i \times \text{S}_t^{EPowerSector}$ captures the fossil fuel embodied in electricity sales as well as electrical grid losses associated with sector i, where

$$\text{S}_t^{EPowerSector} = \frac{\text{Fuel}_t^{EPowerSector}}{\text{Fuel}_t^{EPowerSector} + \text{Renewable}_t^{EPowerSector} + \text{Nuclear}_t^{EPowerSector}}, \tag{17}$$

is the energy share of fossil fuel in the electric power sector.

The top panel of Figure 8 shows this decomposition for aggregate emissions per capita as well as for the four end-use sectors. Several points are worth making. First, the blue lines show a decline of emissions per capita at both the aggregate level and in all end-use sectors. That

said, the magnitudes of these reductions are quite different across sectors. While emissions per capita from the industrial and the residential sectors have declined by almost 60 and 50 percent from 1973 until 2019, the reduction has been much less pronounced in the transportation and commercial sector. Second, much of this reduction has been accounted for by a secular decline of the direct end-use fossil fuel consumption per capita in all sectors (shown as purple lines). Third, with the exception of the transportation sector which has historically used petroleum as the main energy source, all four end-use sectors have increasingly relied on electricity over the past several decades (as shown by the yellow lines). That said, the fossil fuels embodied in electricity have declined in all sectors but the transportation sector since around 2008, consistent with the shale gas boom. As a result, carbon emissions per unit of fossil fuel consumption (shown as orange lines) have also seen a decline in recent years in almost all sectors.

We now study the impulse responses of these components to the EST and NEST shocks. Specifically, we estimate auxiliary VARs adding them one at a time in log levels. Panel B of Figure 8 provides the results. The top row shows that the emissions per capita from all end-use sectors feature dynamics similar to the aggregate emissions per capita shown in the top-right chart of Figure 7. Per capita emissions initially decline sharply, but then recover in subsequent quarters in all sectors except for the Commercial sector. As can be seen in the second row of Panel B of Figure 8, this rebound of per capita emissions is largely driven by the fossil fuels directly consumed by these sectors. Finally, the fossil fuels embodied in electricity consumption persistently increase in all end-use sectors but Transportation. Hence, energy-saving technology shocks are associated with a broad-based substitution away from fossil fuels towards electricity. However, in our sample period from 1973-2019 a sizable fraction of this additional electricity has been produced using fossil fuels.

Turning to the IRFs for the NEST shock, we make the following observations. Emissions per capita significantly increase in all end-use sectors, before reverting back to their initial levels. These dynamics closely track those of the direct fossil fuel consumption in each of the sectors, as shown in the middle panel of the figure. According to the decomposition in Equation 16, the difference between the IRFs for per capita emissions and per capita end-use of fossil fuels is closely related to the IRFs for fossil fuels embodied in electricity. These show that in contrast to the EST shock, the NEST shock is not associated with a significant substitution away from direct fossil fuel consumption towards electricity.

The FEVDs associated with these impulse responses are provided in the second panel of Figure A.2 in the Online Appendix. The main takeaway from these charts is that the bulk of the variation of per capita emissions is unexplained by the two identified technology shocks, despite the fact that particularly the EST shock explains a sizable fraction of the emission intensity of output.

4.4 Robustness Checks

4.4.1 The EST and NEST Shocks versus Fossil Energy Price Shocks

A potential concern is that our identified EST shock might be contaminated by shocks to the price of fossil fuels. Particularly the large oil price shocks in the 1970s have been shown to have triggered substantial technological progress in the use of fossil fuels (Hassler et al. 2021). Moreover, a growing literature documents microeconomic evidence for price increases in carbon intensive inputs acting as an important catalyst for energy-saving technologies.[9] We have shown in Section 4.1 that the two technology shocks in our baseline analysis account for only a small fraction of the business cycle variation in the real fossil fuel price. This suggests little overlapping information with energy price shocks, as the existing literature often attributes a sizable share of the variance of real energy prices at short and business cycle frequencies to energy price shocks (see, e.g., Känzig and Williamson 2024).

Still, to rule out that our results are driven by energy price variations, here we orthogonalize our two technology shocks with respect to fossil energy price shocks. Specifically, we rerun our baseline VAR and recover a fossil energy price shock as an innovation that explains the maximum share of the variation in the real price of fossil fuels over business cycle frequencies from six to 32 quarters. We then proceed by identifying the EST and NEST shocks by solving the optimization problem (12), but with the additional restriction that both shocks are orthogonal to the fossil energy price shock.

The results are presented in Online Appendix Figure A.3. Orthogonalizing the two technology shocks to a fossil fuel price shock, the IRFs and FEVDs associated with the two shocks remain essentially unchanged. This shows that we are indeed picking up shocks to energy-saving and other input-saving technologies which are not themselves driven by fossil fuel price dynamics.

4.4.2 Using a Broad Measure of Energy Consumption

In our baseline analysis documented thus far, we have identified the EST and NEST shocks via sign restrictions imposed on the intensity of fossil energy and its income share. This choice follows the model of Hassler et al. (2021) and reflects the fact that the U.S. economy has heavily relied on fossil fuels over the past decades. In this section, we provide a comparison to a specification using a broad measure of end-use energy instead of fossil energy.

Specifically, we consider two changes to our baseline VAR. First, we include the intensity of end-use energy, defined as the per output value of the sum of fossil, renewable and nuclear energy minus electrical energy system losses. Second, we construct the income share using the annual “Total end-use energy average price (TETXD)", retrieved from the State Energy Data System (SEDS) database. We interpolate this annual end-use energy price to construct a quarterly series using a random walk interpolator as in Stock and Watson (2020) and express it

[9] See, for example, Popp (2002), Aghion et al. (2016) and Brown et al. (2022).

in real terms using the GDP deflator.[10]

We then repeat our identification exercise using these alternative variable definitions. We first identify the two technology shocks as jointly explaining the bulk of the variation in end-use energy intensity and TFP over frequencies of 80 quarters and longer. Second, we rotate these shocks to satisfy the following sign restrictions: an EST shock is required to lower the intensity of end-use energy as well as its income share for horizons $0 \leq h \leq 80$, while a NEST shock is required to raise TFP and the end-use energy income share over the same horizons.

Online Appendix Figure A.4 documents that the two identified shock series are highly correlated with the ones from our baseline identification. The IRFs and FEVDs are provided in Figure A.5. They are essentially identical to those in our baseline analysis which we superimpose. If anything, the differences between the EST and NEST shocks are somewhat more pronounced when considering the income-share of end-use energy, the real end-use price of energy, and end-use energy consumption instead of their fossil fuel counterparts. In sum, the effects of energy-saving technology shocks that we highlight in our baseline results are not specific to our reliance on fossil rather than overall end-use energy.

4.4.3 Alternative Specifications

We now present results from a battery of alternative robustness checks. We consider several departures from the baseline analysis in terms of the values for the hyperparameters used in the Bayesian estimation of the VAR, the horizon used for imposing sign restrictions, achieving identification by alternatively targeting a specific horizon in the time domain, and sample periods.

Figure 9 shows the impulse responses and the variance contributions for the five main variables. As described in Section 3 and Online Appendix A, our baseline results make use of the Minnesota prior with the hyperparameters set to $\gamma_1 = 0.2$, $\gamma_2 = 0.5$, $\gamma_3 = 2$ and $\gamma_4 = 10^5$ following Canova (2007). The first three robustness checks we compare our baseline results to those obtained using alternative values for γ_1, γ_2 and γ_3. The first assumes a looser γ_1 ($\gamma_1 = 0.3$, $\gamma_2 = 0.5$, $\gamma_3 = 2$). The second takes a looser γ_2, which represents the relative tightness of the prior distribution for other variables ($\gamma_1 = 0.2$, $\gamma_2 = 1$, $\gamma_3 = 2$). The third imposes $\gamma_3 = 1$ and thus reduces the relative tightness of the prior standard deviation for lags beyond one ($\gamma_1 = 0.2$, $\gamma_2 = 0.5$, $\gamma_3 = 1$).

The forth robustness analysis varies the horizons over which we impose sign restrictions according to (13). Specifically, we decrease the number of quarters from $0 \leq h \leq 79$ to $0 \leq h \leq 39$.

The fifth contrasts the shocks with those identified by maximizing the long-run variation of fossil energy intensity and TFP in the time instead of the frequency domain. Here, we follow Francis et al. (2014) and Kurmann and Sims (2021) and construct the innovations that

[10] In doing so, the annual observations are assumed to be averages of the four quarterly values, and quarterly values are modeled as following a random walk with drift and are estimated using the Kalman smoother.

explain the maximal share of the forecast error variances of the target variables at a long, but finite horizon.[11] Given the resulting long-run innovations, identification is then also achieved by imposing the sign restrictions (13). We maximize the forecast error variance shares at a horizon of 80 quarters as in Kurmann and Sims (2021) and impose the sign restrictions for horizons of $0 \leq h \leq 39$ quarters.

Our sixth robustness check studies whether our results are driven by the shale gas boom. Specifically, we rerun our analysis for the subsample 1973:I-2005:IV. Acemoglu, Aghion, Barrage and Hémous (2023), among others, document that natural gas increasingly replaced coal in electricity production in the mid 2000s when the development of fracturing and horizontal drilling led to a boom in shale gas production.

Ideally, we would also like to study the implications of EST and NEST shocks in the post-2005 sample. However, since we identify innovations that maximize *long-run* variance shares, we cannot meaningfully replicate our analysis for such a short sample. We therefore assess the dynamic effects of the two shocks over different subsamples adopting the IV regression approach of Stock and Watson (2012, 2018). We first recover the shock series for the full sample from our baseline analysis and then use them as instruments for different subsamples keeping the VAR coefficients constant. The seventh and eighth robustness checks perform this analysis for the two subsamples 1973:I-2005:IV and 2006:I-2019:IV, respectively.

Figure 9 superimposes the IRFs (top panel) and FEVDs (bottom panel) for the eight different robustness analyses and also provides the posterior coverage intervals for the baseline specification. Evidently, the basic character of the IRFs and FEVDs is robust to all of these modifications. This gives us confidence in our main result: an energy-saving technology shock is associated with a marked and persistent response of real economic activity as well as a pronounced hump-shaped rebound of fossil fuel consumption and emissions.[12]

[11] This involves rewriting the optimization problem (12) to capture the variation of the variables at a fixed horizon $\bar{k}$. More concretely, let $B_k^{[j]}$ denote the jth row of the kth lag matrix in $B(L)$ such that $B_k^{[j]}Q_i$ is the effect of shock i on variable j after k periods, and let $\Theta_{j,m}\left(\bar{k}\right) = \frac{\Sigma_{k=0}^{\bar{k}}\left(B_k^{[j]}Q_m\right)^2}{\Sigma_{k=0}^{\bar{k}}B_k^{[j]}B_k^{[j]\prime}}$ represent the share of forecast error variance of variable j explained by shock m. The innovations are then constructed by solving $\underset{Q_1,Q_2}{\text{argmax}}\,\Theta_{\text{energy intensity},1}\left(\bar{k}\right) + \Theta_{\text{energy intensity},2}\left(\bar{k}\right) + \Theta_{\text{TFP},1}\left(\bar{k}\right) + \Theta_{\text{TFP},2}\left(\bar{k}\right)$ subject to the restrictions $Q_1'Q_1 = 1$, $Q_2'Q_2 = 1$ and $Q_2'Q_1 = 0$. This implies that Q_1 and Q_2 are the eigenvectors associated with the two largest eigenvalues of the matrix $\frac{\Sigma_{k=0}^{\bar{k}}B_k^{[\text{energy intensity}]\prime}B_k^{[\text{energy intensity}]}}{\Sigma_{k=0}^{\bar{k}}B_k^{[\text{energy intensity}]}B_k^{[\text{energy intensity}]\prime}} + \frac{\Sigma_{k=0}^{\bar{k}}B_k^{[\text{TFP}]\prime}B_k^{[\text{TFP}]}}{\Sigma_{k=0}^{\bar{k}}B_k^{[\text{TFP}]}B_k^{[\text{TFP}]\prime}}$.

[12] To conserve space, we only show the main five variables in this figure. We obtained similarly robust results for the remaining eight variables in the VAR.

5 Conclusion

In this paper, we disentangle energy-saving and non-energy-saving technology shocks using identifying restrictions derived from state-of-the-art models of directed technical change. We show that the identified EST shock series is correlated with subsequent growth in key measures of energy innovations based on patents and public energy RD&D spending, corroborating our identifying restrictions. Although the EST shock is associated with a persistent reduction of fossil fuel intensity and thus also the carbon intensity of output, it leads to a rebound in fossil fuel consumption and per capita emissions. The reason for this rebound effect is that output and its components strongly and persistently increase following energy-saving technological innovations.

We explore the economic underpinnings of the rebound effect by studying the breakdown of U.S. emissions by source and sector. The EST shock is associated with a compositional change from the direct consumption of fossil fuels towards electricity. Yet, the bulk of electricity has been produced with fossil fuels in our sample from 1973 through 2019. As such, the substitution of end-use fossil fuels by electricity did not come with a substantial reduction of emissions in the short to medium-run. The rebound effect in emissions becomes smaller after the mid 2000s when the shale gas boom increasingly replaced coal by natural gas in electricity production.

While our results provide a coherent account of the dynamics of carbon emissions and output in the U.S. over the past few decades, only a general equilibrium analysis can deliver robust policy conclusions. That said, our findings can inform the debate about how the transition to a net-zero carbon economy might be achieved. We have shown that in the past decades technological advances lowering the fossil fuel and carbon emission intensity of output have led to a rebound in fossil fuel consumption and per capita carbon emissions. We also provide evidence that this rebound is largely explained by a compositional change from end-use of fossil fuels towards fossil fuels embodied in electricity production.

This evidence corroborates a key point emphasized in the IEA (2020)'s Sustainable Development Scenario: low-carbon electricity is likely to be the largest contributor to reaching net-zero carbon emissions. While the electric power sector has heavily relied on fossil fuels over the past decades, achieving net-zero would require a much higher share of low-carbon electricity generation. Going forward, improvements in technology will thus likely have to go hand in hand with policies or social norms affecting the supply of and demand for fossil fuels to achieve a meaningful reduction of emissions. Our results also suggest that technological innovations leading to a more efficient use of fossil energy may have substantial positive effects on economic growth.

References

Acemoglu, D., Aghion, P., Barrage, L., Hémous, D., 2023. Climate Change, Directed Innovation, and Energy Transition: The Long-Run Consequences of the Shale Gas Revolution. Working Paper 31657. National Bureau of Economic Research.

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.

Altig, D., Christiano, L.J., Eichenbaum, M., Lindé, J., 2011. Firm-Specific Capital, Nominal Rigidities and the Business Cycle. Review of Economic Dynamics 14, 225–247.

Amir-Ahmadi, P., Drautzburg, T., 2021. Identification and Inference with Ranking Restrictions. Quantitative Economics 12, 1–39.

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

Angelucci, S., Hurtado-Albir, F.J., Volpe, A., 2018. Supporting Global Initiatives on Climate Change: The EPO's "Y02-Y04S" Tagging Scheme. World Patent Information 54, S85–S92.

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.

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.

Bolton, P., Kacperczyk, M.T., Wiedemann, M., 2023. The CO2 Question: Technical Progress and the Climate Crisis. Available at SSRN 4212567 .

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.

Brown, J.R., Martinsson, G., Thomann, C., 2022. Can Environmental Policy Encourage Technical Change? Emissions Taxes and R&D Investment in Polluting Firms. The Review of Financial Studies 35, 4518–4560.

Bruns, S.B., Moneta, A., Stern, D.I., 2021. Estimating the Economy-Wide Rebound Effect Using Empirically Identified Structural Vector Autoregressions. Energy Economics 97, 105158.

Canova, F., 2007. Methods for Applied Macroeconomic Research. volume 13. Princeton University Press.

Casey, G., 2023. Energy Efficiency and Directed Technical Change: Implications for Climate Change Mitigation. The Review of Economic Studies 91, 192–228.

Chay, K.Y., Greenstone, M., 2005. Does Air Quality Matter? Evidence from the Housing Market. Journal of Political Economy 113, 376–424.

Copeland, B.R., Shapiro, J.S., Taylor, M.S., 2022. Globalization and the Environment, in: Gopinath, G., Helpman, E., Rogoff, K. (Eds.), Handbook of International Economics: International Trade. volume 5, pp. 61–146.

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.

Dernis, H., Khan, M., 2004. Triadic Patent Families Methodology OECD Science, Technology and Industry Working Paper 2004/2, OECD Publishing, Paris. http://dx.doi.org/10.1787/443844125004.

Edenhofer, O., Pichs-Madruga, R., Sokona, Y., Farahani, E., Kadner, S., Seyboth, K., Adler, A., Baum, I., Brunner, S., Eickemeier, P., Kriemann, B., Savolainen, J., Schlömer, S., von Stechow, C., Zwickel, T., Minx, J.C., 2014. Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change .

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.

Galí, J., Gambetti, L., 2009. On the Sources of the Great Moderation. American Economic Journal: Macroeconomics 1, 26–57.

Gillingham, K., Rapson, D., Wagner, G., 2016. The Rebound Effect and Energy Efficiency Policy. Review of Environmental Economics and Policy 10, 68–88.

Griliches, Z., 1990. Patent Statistics as Economic Indicators: A Survey. Journal of Economic Literature 28, 1661–1707.

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

Hassler, J., Krusell, P., Olovsson, C., 2022. Finite Resources and the World Economy. Journal of International Economics 136, 103592.

Henderson, J., 1996. Effects of Air Quality Regulation. American Economic Review 86.

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

IEA, 2011. IEA Guide to Reporting Energy RD&D Budget/Expenditure Statistics. Https://www.iea.org/reports/iea-guide-to-reporting-energy-rdd-budget-expenditure-statistics.

IEA, 2020. Energy Technology Perspectives 2020. Technical Report. Https://www.iea.org/reports/energy-technology-perspectives-2020.

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.

Känzig, D.R., Williamson, C., 2024. Unraveling the Drivers of Energy-Saving Technical Change. ECB Working Paper 2984.

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

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

Masson-Delmotte, V., Zhai, P., Pörtner, H.O., Roberts, D., Skea, J., Shukla, P.R., Pirani, A., Moufouma-Okia, W., Péan, C., Pidcock, R., Others, 2018. Global Warming of 1.5°C. An IPCC Special Report on the Impacts of Global Warming of 1.5°C above Pre-Industrial Levels and Related Global Greenhouse Gas Emission Pathways, in the Context of Strengthening the Global Response to the Threat of Climate Change, Sustainable Development, and Efforts to Eradicate Poverty. IPCC.

OECD, 2009. OECD Patent Statistics Manual. Technical Report. Https://www.oecd.org/sti/inno/oecdpatentstatisticsmanual.htm.

Popp, D., 2002. Induced Innovation and Energy Prices. American Economic Review 92, 160–180.

Popp, D., 2019. Environmental Policy and Innovation: A Decade of Research. Working Paper 25631. National Bureau of Economic Research.

Rubio-Ramírez, J.F., Waggoner, D.F., Zha, T., 2010. Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference. The Review of Economic Studies 77, 665–696.

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.

Stock, J.H., Watson, M.W., 2012. Disentangling the Channels of the 2007-09 Recession. Brookings Papers on Economic Activity , 120–157.

Stock, J.H., Watson, M.W., 2018. Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments. The Economic Journal 128, 917–948.

Stock, J.H., Watson, M.W., 2020. Trend, Seasonal, and Sectorial Inflation in the Euro Area, Central Bank of Chile. number 27 in Changing Inflation Dynamics, Evolving Monetary Policy, Series on Central Banking, Analysis and Economic Policies, pp. 317 – 344.

Uhlig, H., 2003. What Moves Real GNP?, Unpublished.

Uhlig, H., 2005. What are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure. Journal of Monetary Economics 52, 381–419.

Veefkind, V., Hurtado-Albir, J., Angelucci, S., Karachalios, K., Thumm, N., 2012. A New EPO Classification Scheme for Climate Change Mitigation Technologies. World Patent Information 34, 106–111.

Figures

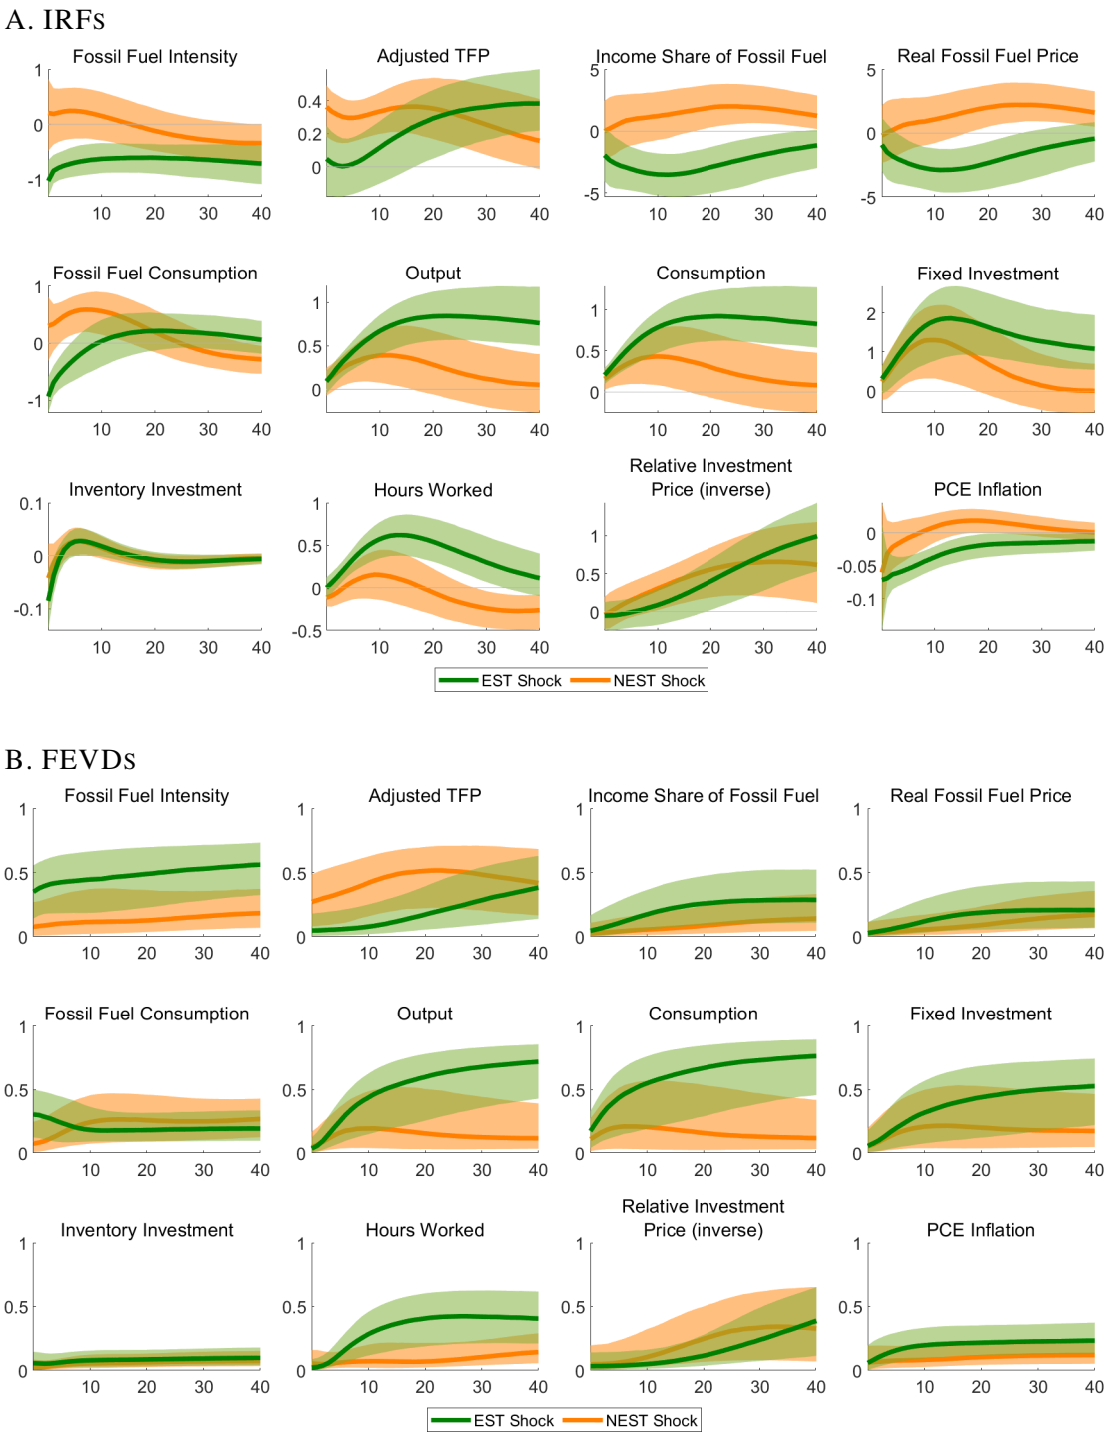

Figure 1: IRFs and FEVDs for EST and NEST Shocks

Notes: The top panel shows the IRFs for the EST shock (green) and the NEST shock (brown) from the structural VAR. The shocks are reported as one-standard-deviation impulses. The bottom panel displays the corresponding FEVDs. The shaded bands correspond to the 16 to 84 percent posterior coverage intervals.

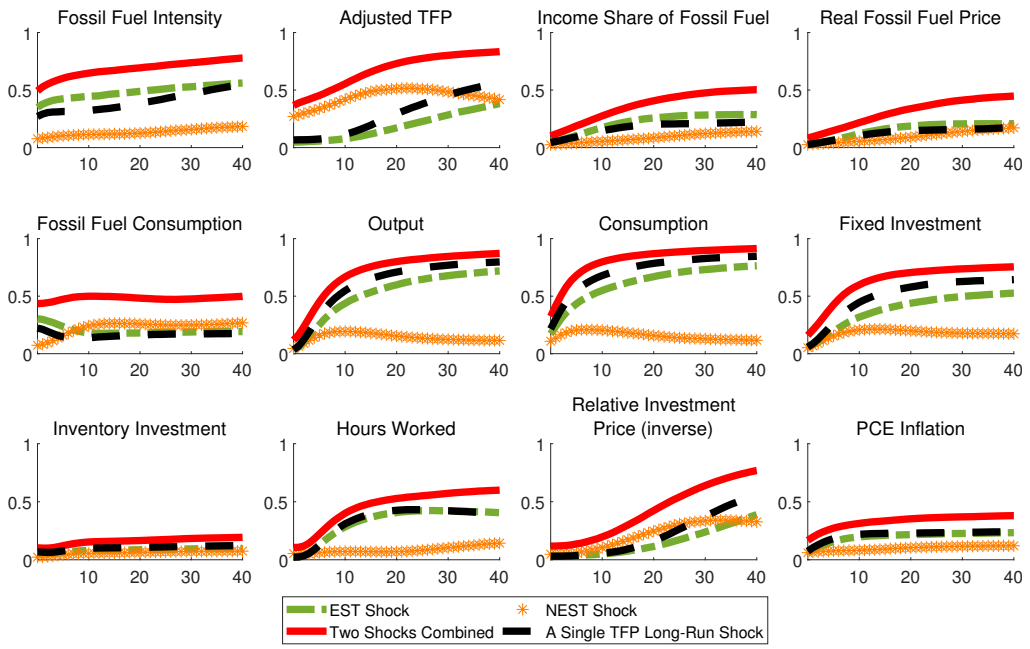

Figure 2: FEVDs for EST and NEST Shocks Combined versus One Single Long-Run Shock to TFP

Notes: This figure shows the posterior median forecast error variance shares explained by the EST shock (green dash-dot), the NEST shock (brown stars), the EST and NEST shocks combined (red solid), and a single long-run shock to TFP (black dashed) from the structural VAR.

Figure 3: EST and NEST shock series versus k-year ahead growth of energy innovation indicators

Notes: The top-left chart shows the median estimate of the shock series. Each bar represents the aggregate of the quarterly median estimated shock series in the annual frequency, computed by averaging over four quarters in each year. The remaining charts estimate the regression coefficient of k-year ahead growth of different energy innovation indicators as the dependent variable (for k=2, 4, 6, 8 and 10 years) on a smoothed version of our shock series which we constructed by running through an AR(1) filter with autoregressive coefficient of 0.9. The charts also show the robust standard error bands for these estimates.

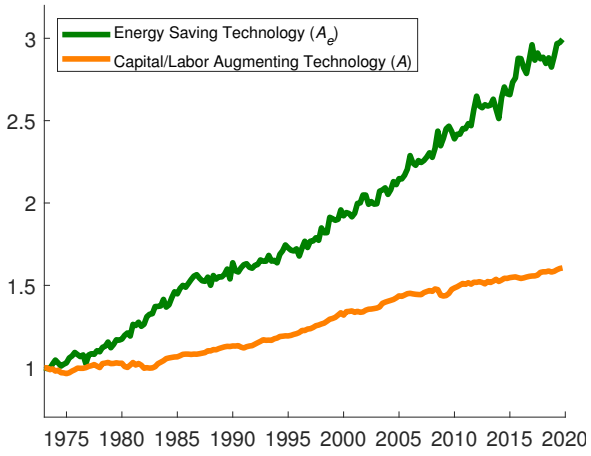

Figure 4: Input-Saving Technologies, $\varepsilon = 0.02$

Notes: The green and brown lines depict the energy saving and capital/labor augmenting technology levels backed out from the production function (1) closely following Hassler et al. (2021) (see Online Appendix B for details on the construction of these technology series). Each series is normalized to 1 in 1973:I.

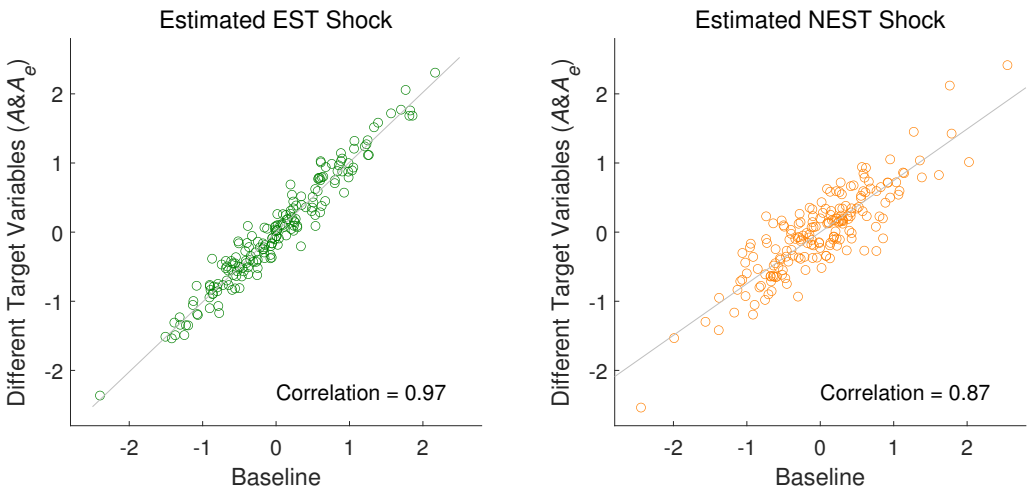

Figure 5: Correlation among median estimate of shock series obtained from baseline analysis and alternative approach using different target variables ($A\&A_e$)

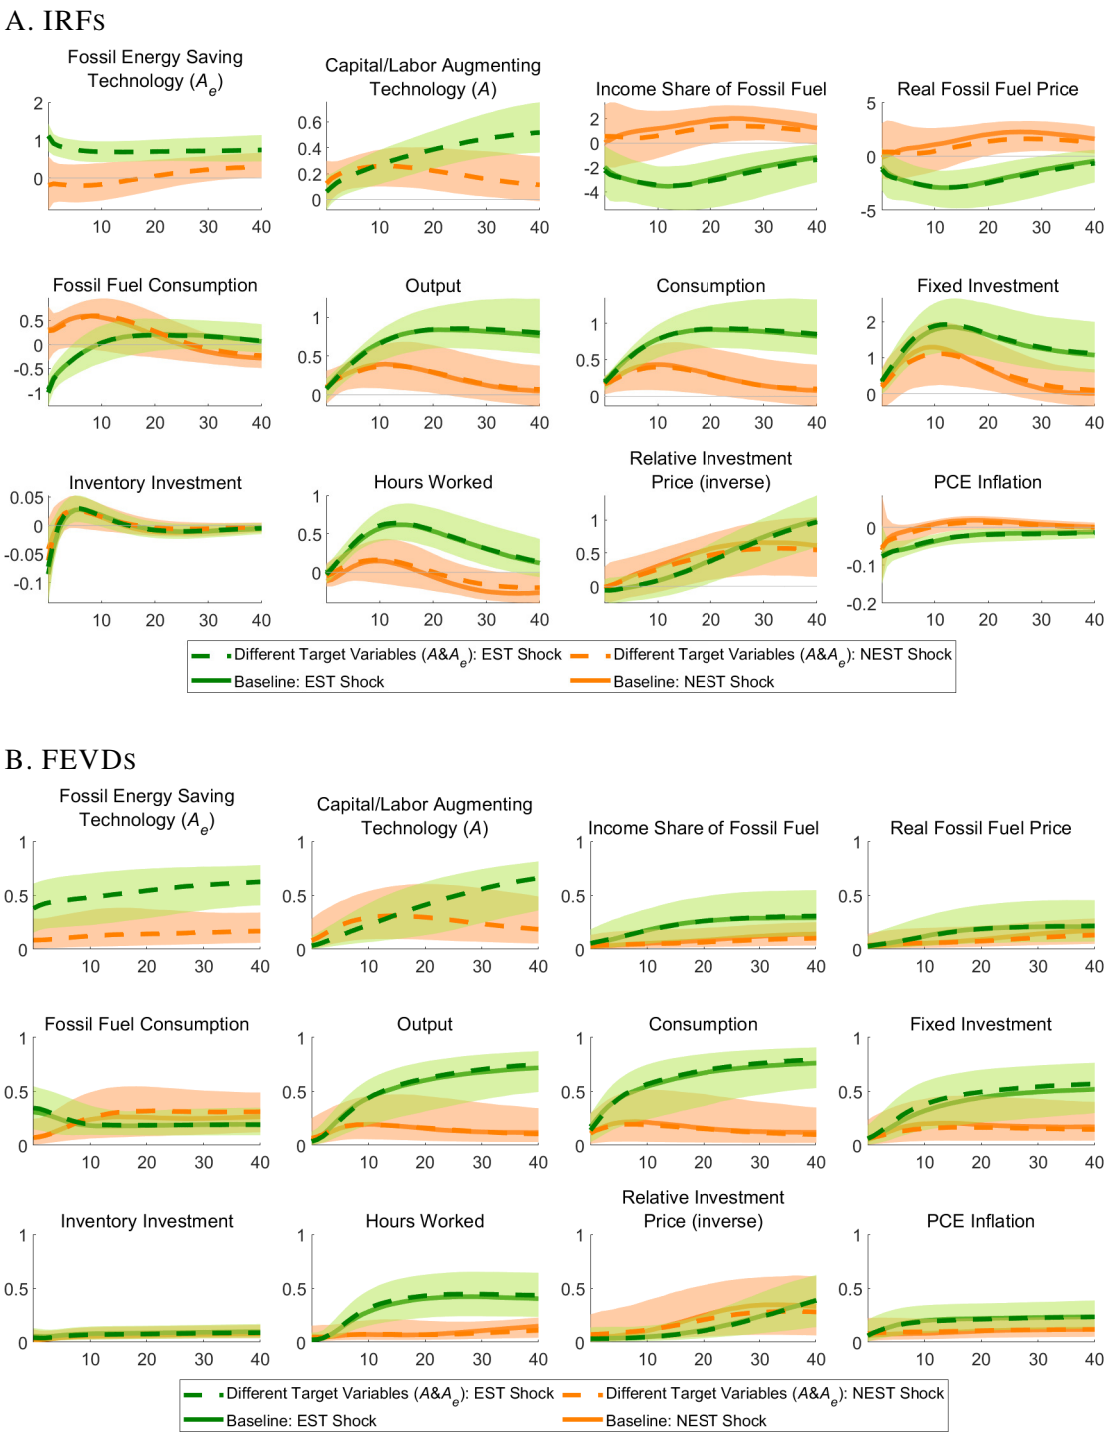

Figure 6: Different Target Variables (A&A_e): IRFs and FEVDs for EST and NEST Shocks

Notes: The top panel shows the IRFs for the EST shock (green) and the NEST shock (brown) from the structural VAR. The shocks are reported as one-standard-deviation impulses. The bottom panel displays the corresponding FEVDs. The shaded bands correspond to the 16 to 84 percent posterior coverage intervals.

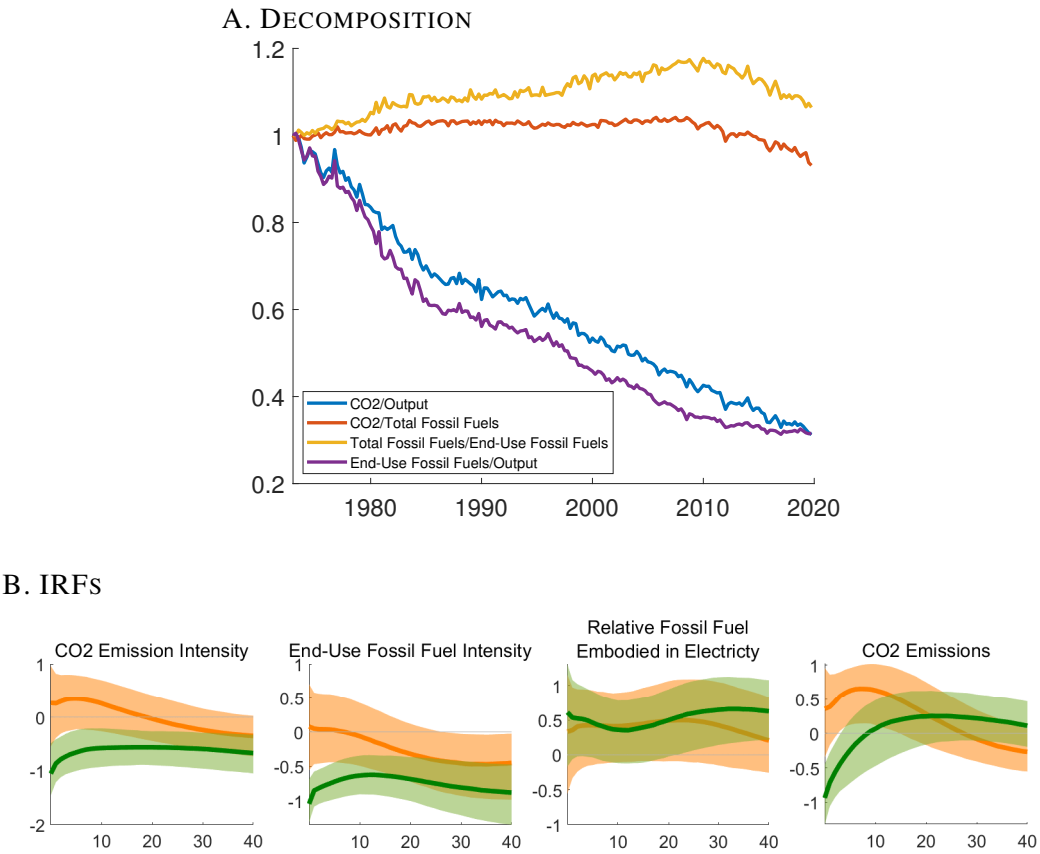

Figure 7: Decomposition of Aggregate CO2 Emission Intensity, and IRFs for Corresponding Components

Notes: The blue and purple lines depict the intensity of carbon emissions and end-use fossil fuel consumption per unit of U.S. output in the top panel. The gap between these lines are captured by the product of the brown and yellow lines which, respectively, reflect changes in the ratio of carbon emissions to total fossil fuel consumption (which consists of fossil fuels consumed in end-use and electric power sectors) and the ratio of total fossil fuel consumption to end use of fossil fuel. Each series is normalized to 1 in 1973:I. The bottom panel shows the IRFs of the components for the EST shock (green) and the NEST shock (brown) from the structural VAR. The shocks are reported as one-standard-deviation impulses. The shaded bands correspond to the 16 to 84 percent posterior coverage intervals.

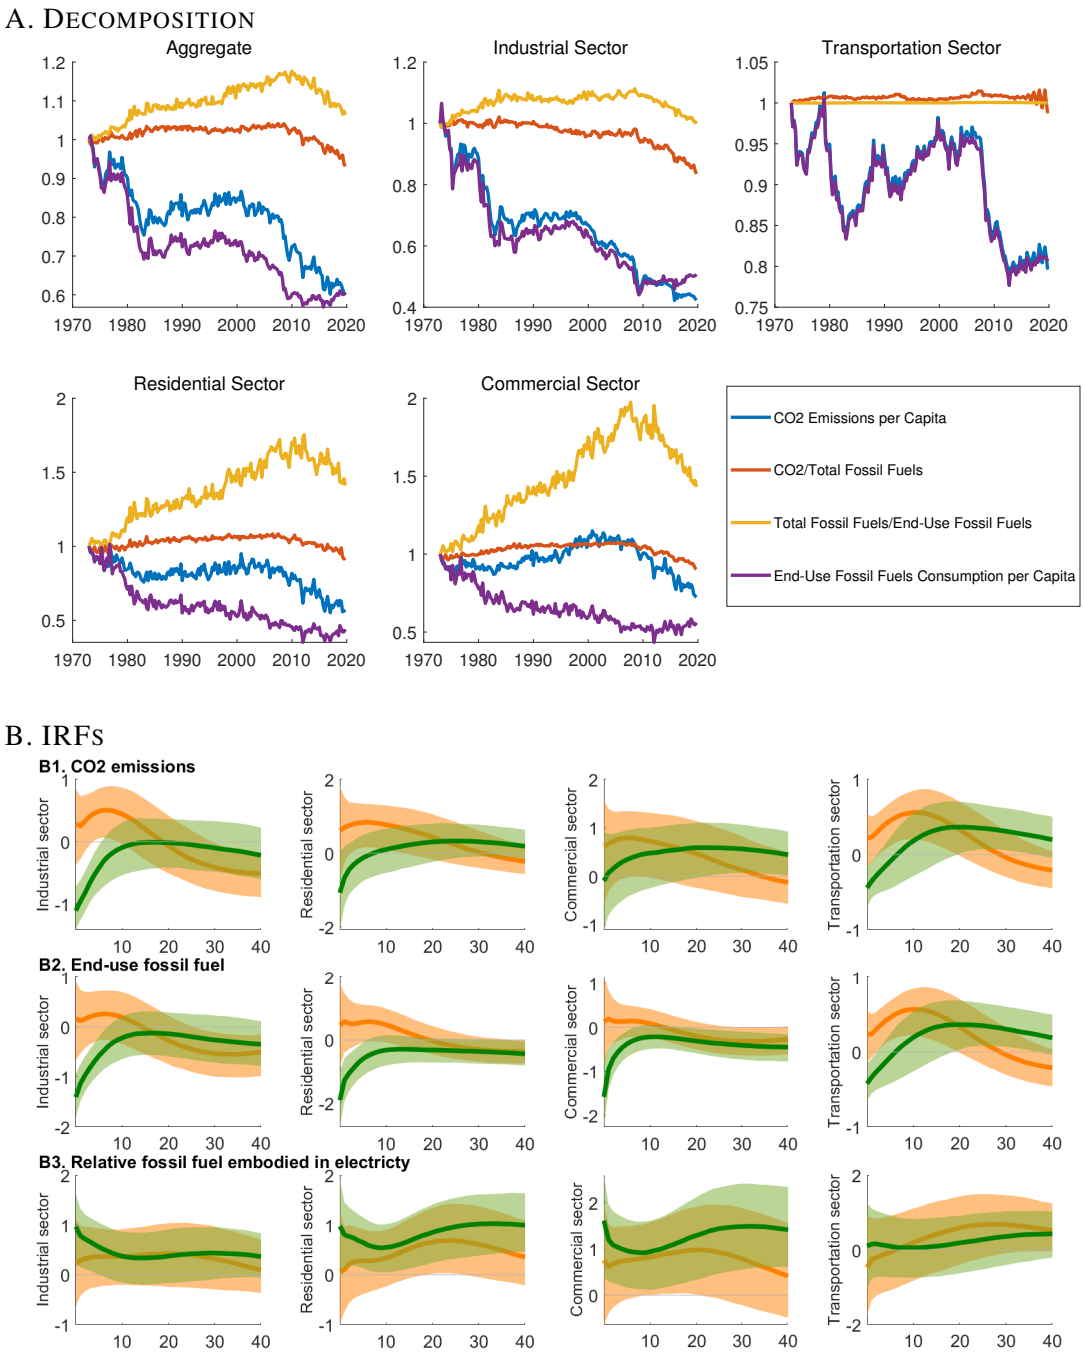

Figure 8: Decomposition of CO2 Emission per Capita by Sectors, and IRFs for Corresponding Components

Notes: The blue and purple lines depict per capita values of carbon emissions and end-use fossil fuel consumption both in aggregate and sector-level in the top panel. The gap between these lines are captured by the product of the brown and yellow lines which, respectively, reflect changes in the ratio of carbon emissions to total fossil fuel consumption (which consists of fossil fuels consumed in end-use and electric power sectors) and the ratio of total fossil fuel consumption to end use of fossil fuel. Each series is normalized to 1 in 1973:I. The bottom panel shows the IRFs of the components at the sectoral level for the EST shock (green) and the NEST shock (brown) from the structural VAR. The shocks are reported as one-standard-deviation impulses. The shaded bands correspond to the 16 to 84 percent posterior coverage intervals.

A. ROBUSTNESS: IRFS

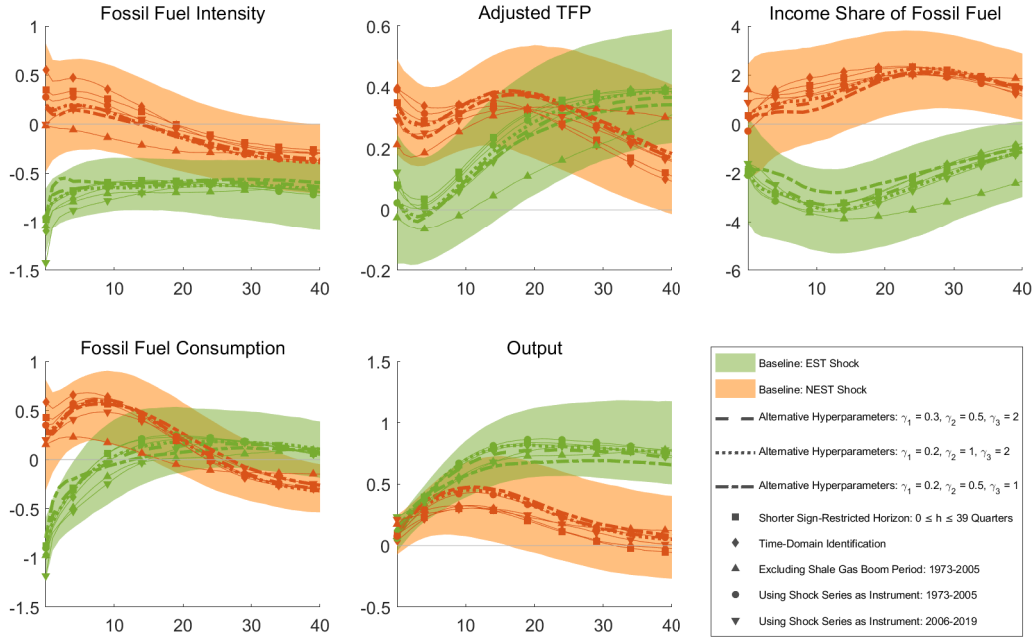

B. ROBUSTNESS: FEVDS

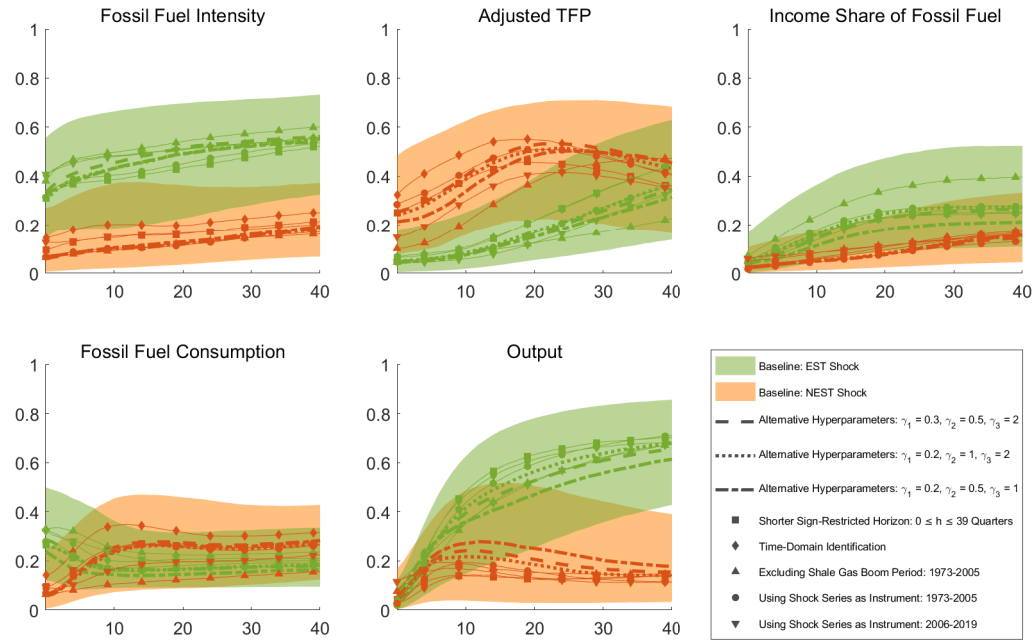

Figure 9: Robustness: IRFs and FEVDs for EST and NEST Shocks

Notes: The baseline case uses the same sample period, Bayesian VAR specification and identification as in Figure 1.The other cases are different from the baseline as indicated. In three cases, we use alternative values for the hyperparameters (see Online Appendix A for details on the Bayesian estimation). For the time-domain identification, we alternatively identify two shocks explaining the maximum forecast error variance share of fossil energy intensity and TFP at the 80-quarter horizon and satisfy sign restrictions up to 40 quarters ahead. For IV regressions IRFs, we treat the shock series estimated over the full sample as instrument over two subsamples while keeping the VAR coefficients unchanged.

Online Appendix

In this appendix, we first describe the details of the Minnesota prior we used for Bayesian estimation and inference in the VARs we estimate. We then describe the details on the public energy RD&D budget measures and energy-saving technology series.

A Priors

Following a common convention in the literature on Bayesian VARs, we make use of the Minnesota prior, which is based on the belief that the univariate behavior of each time series variable included in the VAR is well described by a random walk model. In particular, for a VAR model of the form

$$X_t = c + A_1 X_{t-1} + \cdots + A_p X_{t-p} + \eta_t,$$

where X_t denotes an $n \times 1$ vector of quarterly time series, we use a representation for the prior information that sets $c = 0$, $A_1 = I_n$ and $A_2 = A_3 = ... = A_{p-1} = 0$.

Moreover, the Minnesota prior takes the following standard deviation for the prior distribution of $a_{ij}^{(s)}$

$$\frac{\gamma_1}{s^{\gamma_3}}$$

when $i = j$, and

$$\frac{\gamma_1 \gamma_2 \widehat{\sigma}_i}{s^{\gamma_3} \widehat{\sigma}_j}$$

when $i \neq j$, and also takes the standard deviation $\gamma_4 \widehat{\sigma}_i$ for the constant term c_i, where $\widehat{\sigma}_i$ is estimated by the standard deviation of the residuals from the OLS regression of x_{it} on a constant and p of its own lags. Following Canova (2007), we set $\gamma_1 = 0.2$, $\gamma_2 = 0.5$, $\gamma_3 = 2$ and $\gamma_4 = 10^5$.

B Construction of Public Energy RD&D Budget Shares

The shares we use are defined as

$$\text{GovTotalEnergyRDD\_Share}_t = \frac{\text{GovTotalEnergyRDD}_t}{\text{GovRD}_t},$$

$$\text{GovFossilEnergyRDD\_Share}_t = \frac{\text{GovFossilEnergyRDD}_t}{\text{GovRD}_t},$$

where $\text{GovTotalEnergyRDD}_t$, $\text{GovFossilEnergyRDD}_t$ and GovRD_t respectively denote the government RD&D budget for all energy-related technologies, the government RD&D budget for

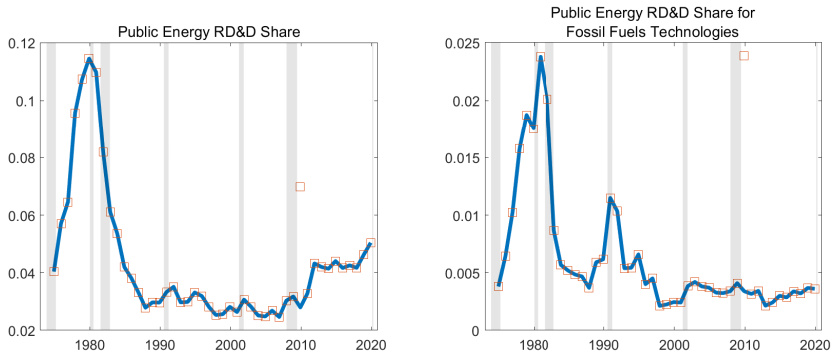

Figure A.1: Annual Series of Public Energy RD&D Budget

the fossil fuel technologies, and the NIPA government R&D investment. GovTotalEnergyRDD$_t$ is taken from the IEA Energy Technology RD&D Budgets database. We use "Group 2–Fossil Fuels: Oil, Gas and Coal" from the same dataset for GovFossilEnergyRDD$_t$. For GovRD$_t$, we use "Government Gross Investment: Intellectual Property Products: Research and Development (Y057RC1A027NBEA)", retrieved from FRED, Federal Reserve Bank of St. Louis. As described in the IEA's database documentation, there has been a large increase in RD&D spending associated with the American Recovery and Reinvestment Act of 2009. Since this is a one-year appropriation and the following year sees a substantial decrease, we treat the 2009 observations as outliers and replace them with the median of the five preceding observations. The series are plotted in Figure A.1.

C Construction of Input-Saving Technology Variables

Following Hassler et al. (2021), we use the production function (1) and data on output, inputs and their prices to back out the energy-saving technology level (A_{et}) and the capital/labor augmenting technology level (A_t). To this end, under the assumption of perfect competition in input markets, we solve for the two technology trends given a value of the substitution elasticity. Letting $l_t^{share} = w_t l_t / y_t$ and $e_t^{share} = p_t e_t / y_t$, this yields

$$A_t = \frac{y_t}{k_t^{\alpha} l_t^{1-\alpha}} \left[\frac{l_t^{share}}{(1-\alpha)(1-\gamma)} \right]^{\frac{\varepsilon}{\varepsilon-1}},$$

and

$$A_{et} = \frac{y_t}{e_t} \left[\frac{e_t^{share}}{\gamma} \right]^{\frac{\varepsilon}{\varepsilon-1}}.$$

We keep the substitution elasticity equal to $\varepsilon = 0.02$ and set $\gamma = 0.05$ and $\alpha = 0.25/0.95 = 0.2632$ as in Hassler et al. (2021). Here, e_t and p_t correspond to fossil fuel consumption and a composite index of real fossil fuel prices defined as

$$e_t = e_t^{\text{coal}} + e_t^{\text{petroleum}} + e_t^{\text{gas}},$$
$$p_t = \frac{p_t^{\text{coal}} e_t^{\text{coal}} + p_t^{\text{petroleum}} e_t^{\text{petroleum}} + p_t^{\text{gas}} e_t^{\text{gas}}}{e_t^{\text{coal}} + e_t^{\text{petroleum}} + e_t^{\text{gas}}}.$$

Monthly data on e_t^i and p_t^i for $i \in \{\text{coal}, \text{petroleum}, \text{gas}\}$ for January 1973 to December 2019 are taken from the EIA. Specifically, we use consumption of coal, petroleum and natural gas from Table 1.3 "Primary energy consumption by source". We first seasonally adjust these series using the X-12 method and then convert them to quarterly values by adding the monthly values.[13] We further use Table 9.9 "Cost of fossil-fuel receipts at electric generating plants" to retrieve p_t^i. These price series, measured in Dollars per million Btu (including taxes), are deflated by the GDP deflator.

Quarterly data on total compensation of employees are taken from the BEA. This variable is deflated by the GDP deflator. It is then used to computed the labor share of income as

$$l_t^{share} = \frac{\text{TotalEmployeeCompensation}_t}{y_t}.$$

For l_t, we use "Employment Level (CE16OV)" , retrieved from FRED.

For k_t, we use annual data on "Capital Stock at Constant National Prices for United States (RKNANPUSA666NRUG)", also obtained from FRED. We interpolate the annual capital stock to construct a quarterly capital stock series using a random walk interpolator (see footnote 10).

Finally, we follow Hassler et al. (2021) and define output as

$$y_t = \text{GDP}_t - \text{NetExport}_t^{\text{Fuel}},$$

where

$$\begin{aligned}\text{NetExport}_t^{\text{Fuel}} = {} & p_t^{\text{coal}} \left(X_t^{\text{coal}} - M_t^{\text{coal}} \right) \\ & + p_t^{\text{petroleum}} \left(X_t^{\text{petroleum}} - M_t^{\text{petroleum}} \right) \\ & + p_t^{\text{gas}} \left(X_t^{\text{gas}} - M_t^{\text{gas}} \right).\end{aligned}$$

We use coal, petroleum and natural gas imports and exports from Tables 1.4a-b "Primary Energy Imports by Source" and "Primary Energy Export by Source" from the EIA. These series are originally available monthly. Similarly to the data on fossil fuel consumption, we first seasonally adjust them and then convert them to the quarterly frequency by adding the monthly values. This output series is in turn used in constructing the fossil energy income share. None

[13] The corresponding sectoral data used in Section 4.3.2 are taken from Tables 2.1a-b "Energy consumption: Residential, commercial, and industrial sectors" and "Energy consumption: Transportation sector, total end-use sectors, and electric power sector".

of the results change, however, if we use real GDP instead.

D Additional Figures

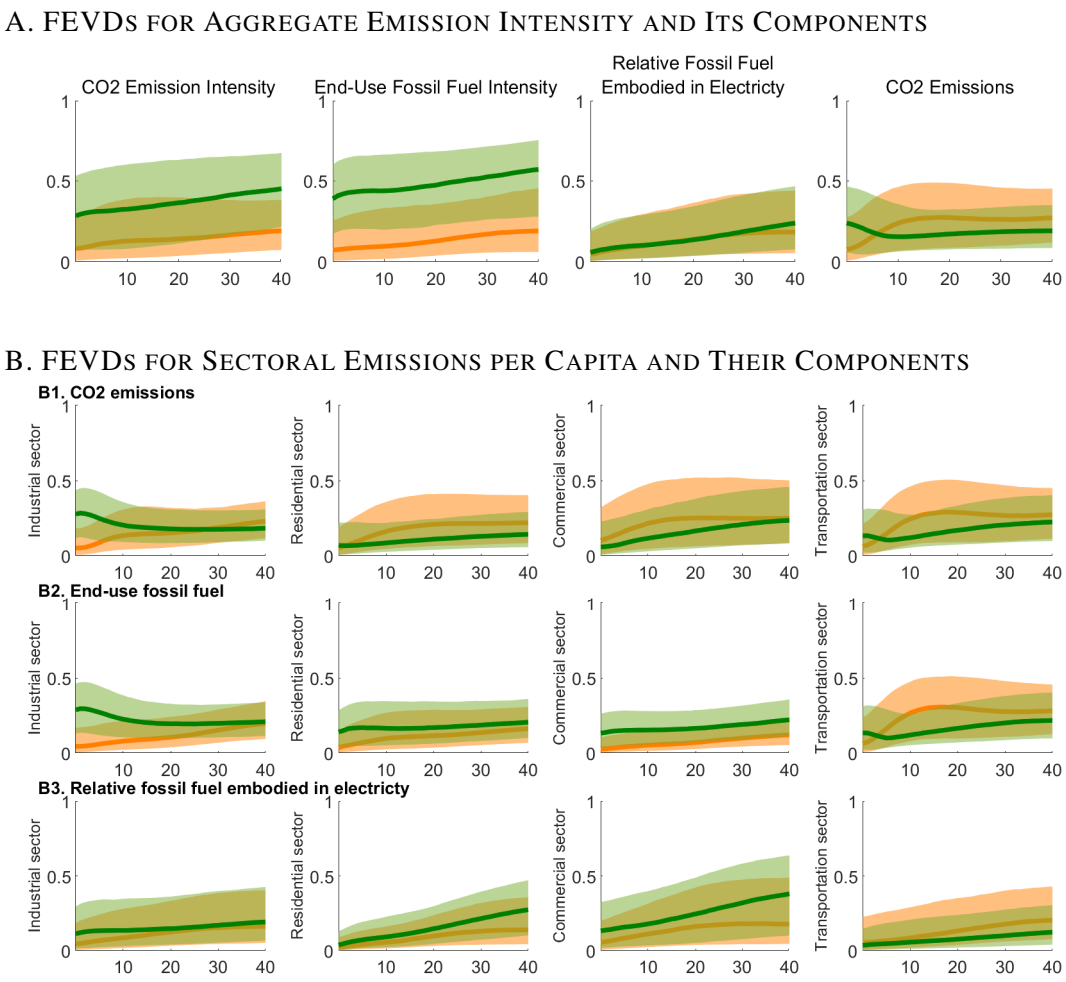

Figure A.2: FEVDs for Carbon Emission Intensity, per Capita Emissions, and Corresponding Components at Aggregate and Sector-Level

*Notes:*This figure shows the FEVDs for the EST shock (green) and the NEST shock (brown) from the structural VAR. The shaded bands correspond to the 16 to 84 percent posterior coverage intervals.

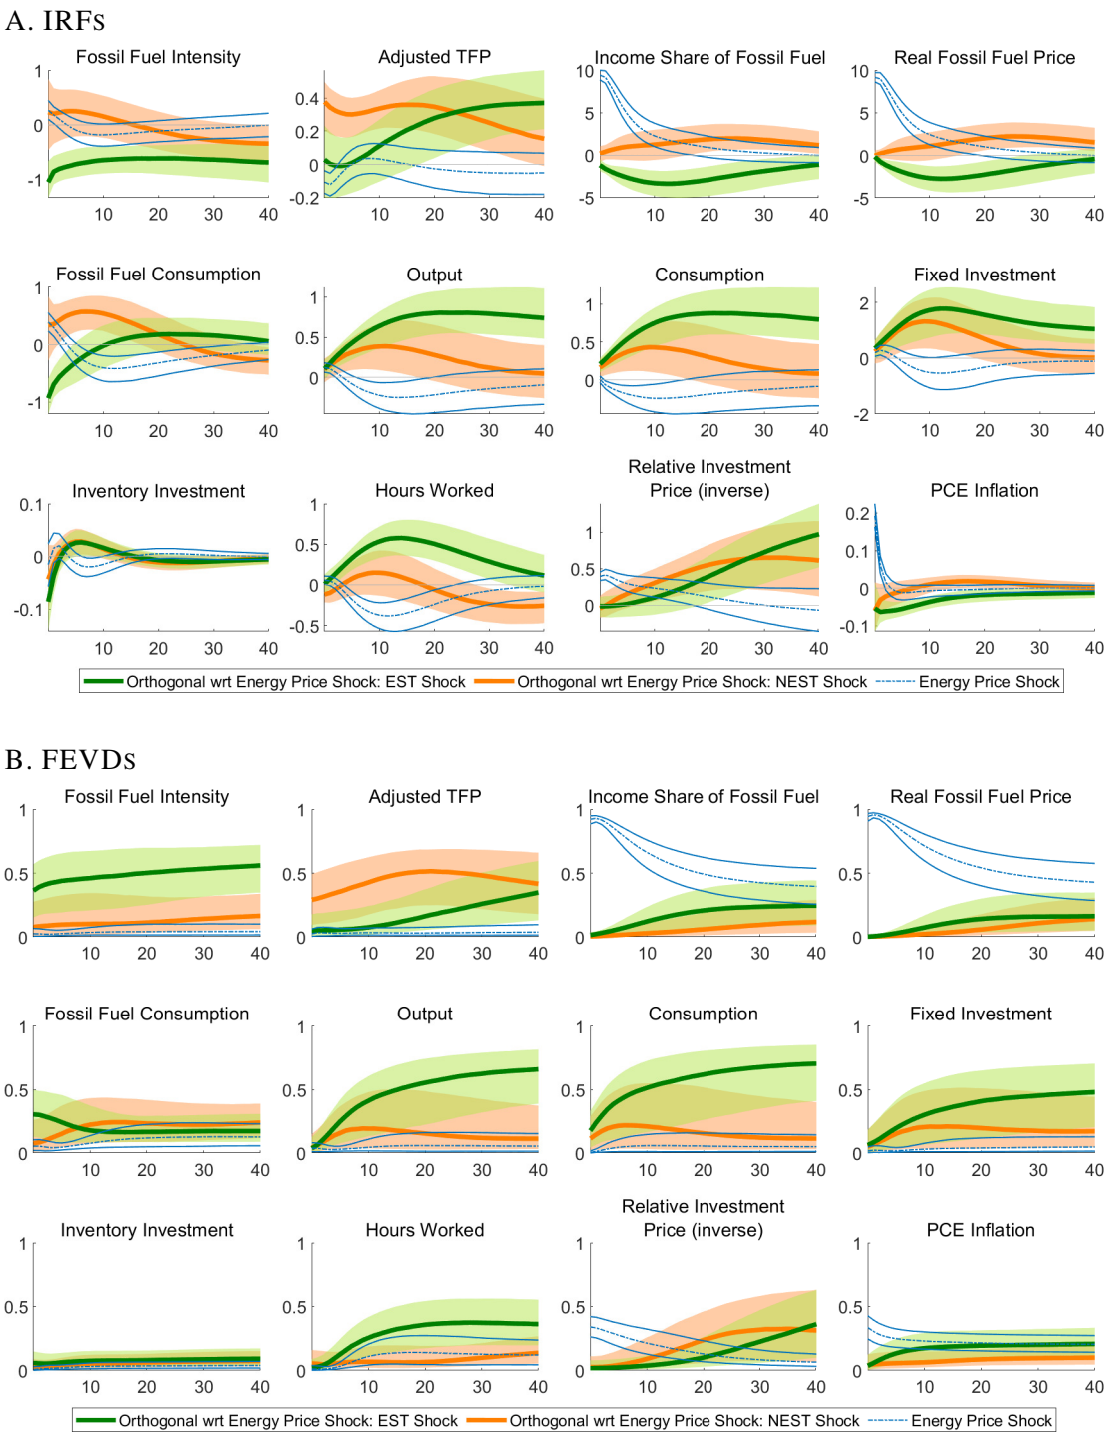

Figure A.3: Orthogonalizing w.r.t. to an Energy Price Shock: IRFs and FEVDs for EST and NEST Shocks

Notes: The top panel shows the IRFs for the EST shock (green), the NEST shock (brown), and the energy price shock (blue) from the structural VAR. The energy price shock is identified as a shock that explains the bulk of the volatility in the real fossil fuel price over the business-cycle frequencies (6-32 quarters). The shocks are reported as one-standard-deviation impulses. The bottom panel displays the corresponding FEVDs. The shaded bands correspond to the 16 to 84 percent posterior coverage intervals.

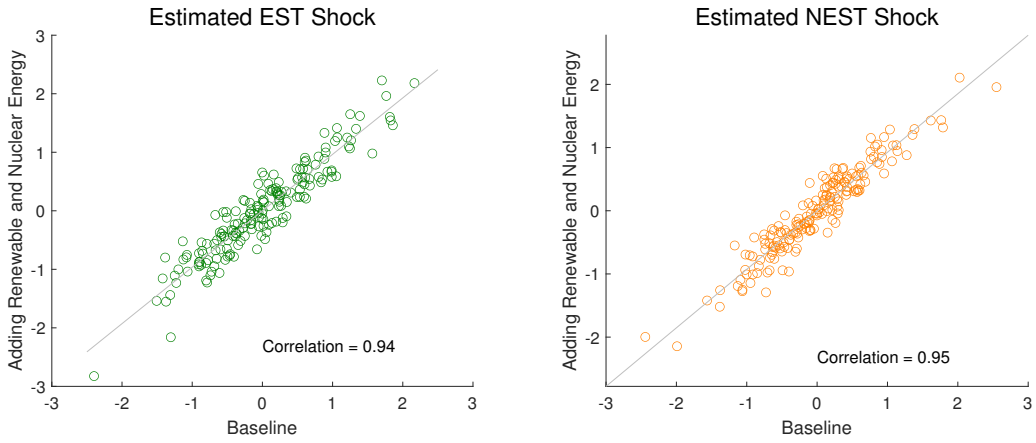

Figure A.4: Correlation among median estimate of shock series obtained from baseline analysis using prices and consumption of fossil energy and extended analysis adding additionally renewable and nuclear energy

A. IRFs

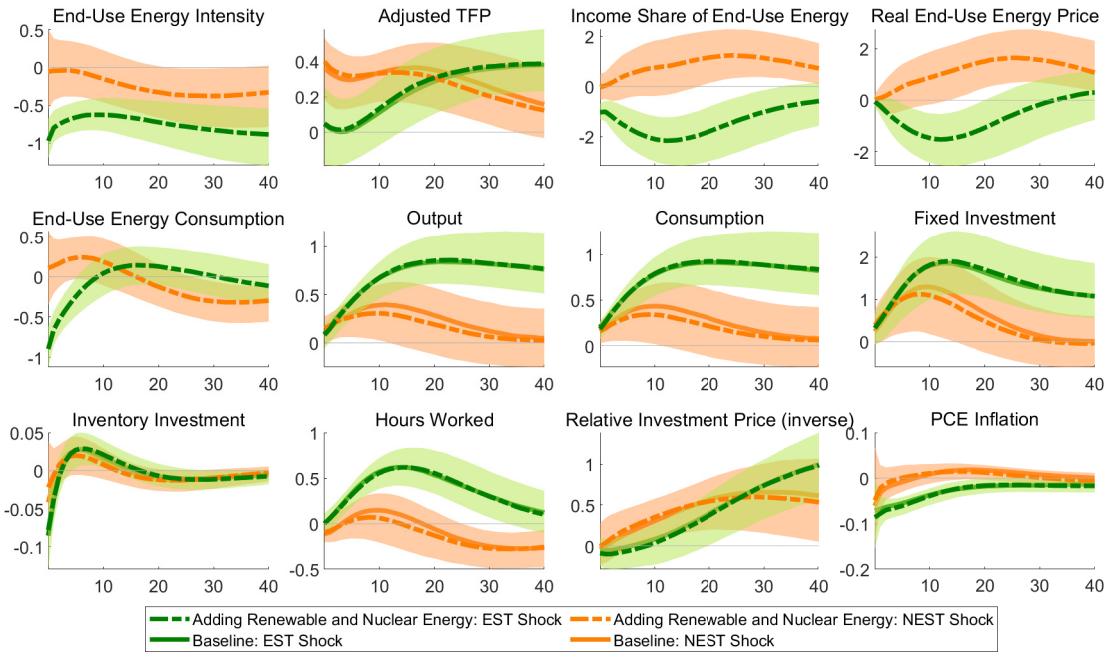

B. FEVDs

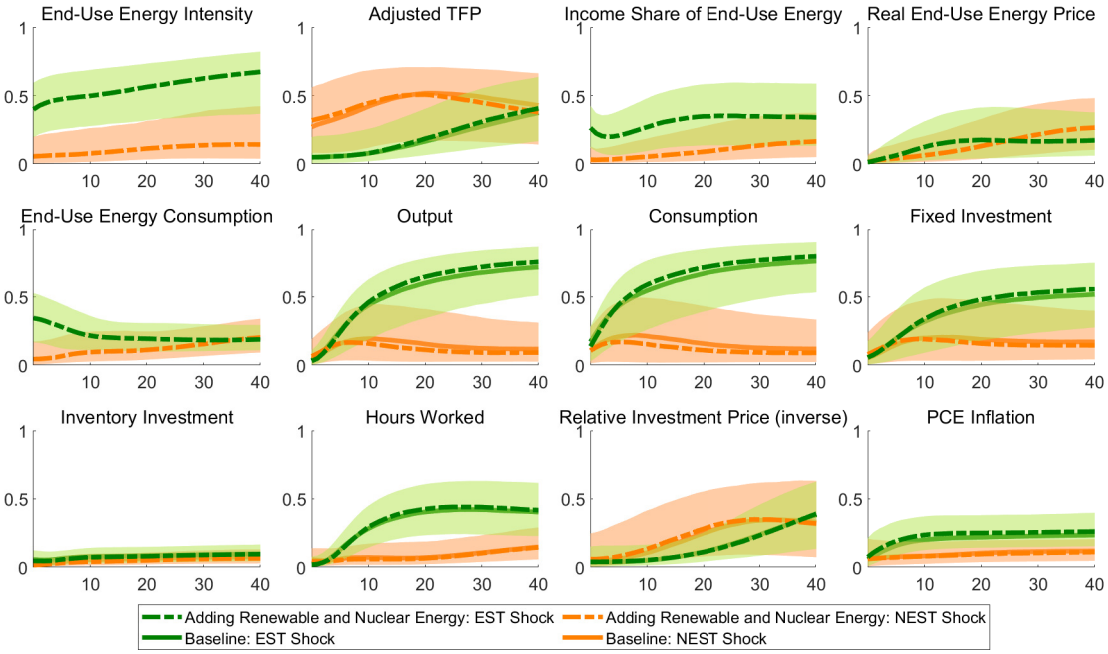

Figure A.5: Adding Renewable and Nuclear Energy: IRFs and FEVDs for EST and NEST Shocks

Notes: The top panel shows the IRFs for the EST shock (green) and the NEST shock (brown) from the structural VAR. The shocks are reported as one-standard-deviation impulses. The bottom panel displays the corresponding FEVDs. The shaded bands correspond to the 16 to 84 percent posterior coverage intervals.