

NBER WORKING PAPER SERIES

UNRAVELING THE DRIVERS OF ENERGY-SAVING TECHNICAL CHANGE

Diego R. Känzig
Charles T. Williamson

Working Paper 34511
<http://www.nber.org/papers/w34511>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2025

We thank Christiane Baumeister, Adrien Bilal, André Kurmann, Andrea Lanteri, Givi Melkadze, Noémie Pinardon-Touati, Giorgio Primiceri, Martin Schneider, Conor Walsh, and an anonymous referee for the ECB working paper series for helpful comments and suggestions. The contribution by Diego Känzig has been prepared under the Lamfalussy Fellowship Programme sponsored by the ECB. Any views expressed are only those of the authors and do not necessarily represent the views of the ECB or the Eurosystem. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Diego R. Känzig and Charles T. Williamson. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Unraveling the Drivers of Energy-saving Technical Change
Diego R. Känzig and Charles T. Williamson
NBER Working Paper No. 34511
November 2025
JEL No. E0, O30, Q32, Q43, Q55

ABSTRACT

We explore the increasing divergence between economic growth and energy consumption through energy-saving technical progress. Proposing a new measure of energy-saving technology, we study the underlying drivers in a semi-structural model of the U.S. economy. Our analysis shows that energy price shocks reduce consumption and stimulate energy-saving innovation, but also cause economic downturns and crowd out other innovations. Only energy-saving technology shocks can explain the negative co-movement between output and energy use. These sudden efficiency gains emerge as the primary driver of energy-saving technical change. Our findings highlight the importance of fostering energy-saving innovations in transitioning to a low-carbon economy.

Diego R. Känzig
Northwestern University
Department of Economics
and NBER
dkaenzig@northwestern.edu

Charles T. Williamson
Northwestern University
cwilliamson@u.northwestern.edu

1. Introduction

Energy use has historically been a key driver of economic growth. The availability of cheap and accessible energy sources has been crucial in powering industries, facilitating transportation, and enabling technological advances. This importance has often made energy resources a focal point of geopolitical conflicts and economic competition.

The close relation between energy consumption and economic output is a salient feature of the data. Figure 1 shows that energy use and output in the United States evolved in lockstep for the first twenty years of the post-war period. Starting in the 1970s, however, there was a remarkable decoupling between energy consumption and economic growth. While real GDP continued to grow at an exponential rate, energy consumption suddenly grew much more slowly. Confronting the looming climate crisis requires drastic emission reductions across the globe. Against this backdrop, it is crucial to understand the drivers behind the decoupling of output and energy use.




Figure 1: U.S. real GDP and energy consumption since the 1950s

Notes: This figure shows the evolution of U.S. real GDP together with energy consumption, measured as coal, oil, and natural gas consumption, since the 1950s, both expressed as an index normalized to 100 in 1970Q1.

In this paper, we construct new measures of input-saving technical change and examine their underlying drivers based on an aggregate production function framework. We assume that output is produced using capital, labor and fossil energy according to a nested constant elasticity of substitution (CES) production function. From the cost minimization problem, we obtain expressions for capital/labor-saving and energy-saving technical change as a function of output and input growth, changes in utilization and changes in factor shares. A crucial parameter is the elasticity of substitution between en-

ergy and the other inputs. Our framework gives rise to a natural estimation equation for this elasticity, based on the energy demand schedule. A key identification challenge is that energy consumption and prices are jointly determined. Thus, we use energy price shocks as supply shifters to estimate the elasticity of substitution. We find that energy is highly complementary to the other inputs, with a quarterly elasticity of 0.026. This estimate is robust along a number of dimensions, including the use of different energy price shock instruments, the sample period, or the inclusion of additional controls.

With the estimated elasticity of substitution in-hand, we construct a new series of energy-saving technical change. The series displays a pronounced kink in the 1970s, marked by a significant acceleration in its growth rate, around the same time when output and energy use have started to diverge. We show that neither the structural transformation nor the increased use of renewables can account for this pattern; energy-saving technological improvements seem to be the key factor behind the decoupling of output and energy use.

To shed light on the driving forces behind energy-saving technical change, we develop a semi-structural model of the U.S. economy and identify three potential drivers: capital/labor-saving technology shocks, energy-saving technology shocks and energy price shocks. We identify these shocks in a structural VAR model, by maximizing their contributions to the forecast error variance in energy prices and input-saving technologies, respectively. Specifically, we assume that energy price shocks are the main driver of energy price fluctuations at short horizons. Next, we identify a capital/labor-saving technology shock as an innovation that is orthogonal to energy price shocks and accounts for the largest portion of the residual variation in capital/labor-saving technology in the longer run. Finally, we identify an energy-saving technology shock as the innovation that explains the maximum share of the residual variation in energy-saving technical change at longer horizons while being orthogonal to the energy price and the other input-saving technology shock.

We find that capital/labor-saving technology shocks lead to a significant increase in output and energy consumption and a fall in prices. Interestingly, increases in capital/labor-saving technologies do not seem to crowd out energy-saving innovations, as the response is mostly insignificant. Importantly, the shock leads to a significant increase in the energy intensity, defined as energy use over output, as energy consumption increases much more strongly than output. Overall, the responses are very similar to the effects documented in response to broad-based technology shocks such as total factor productivity (TFP) shocks in the literature.

Energy-saving technology shocks transmit to the economy very differently. They lead

to a significant fall in energy consumption and an increase in output. Energy prices and the GDP deflator do not respond significantly. This is an important finding because, in principle, a large decline in energy prices could offset some of the drop in energy consumption. We do not find evidence for such a rebound effect of energy efficiency. There is some evidence that other input-saving technologies are crowded out, at least in the short term, which may attenuate the positive output effects. Crucially, the decline in energy consumption is far larger than the increase in output, implying a substantial reduction in energy intensity.

Energy price shocks lead to a significant increase in energy prices and a decrease in energy consumption. This cost shock to an important input in production has severe economic consequences: output falls and the GDP deflator increases significantly. The effects are consistent with the responses to oil supply shocks in the literature. While the scarcity of energy creates an incentive to innovate in energy-saving technologies – we observe a substantial increase in these technologies after some lag – it also significantly crowds out innovations in other input-saving technologies. This helps explain the rather persistent effects on output.

We conclude that only energy-saving technology shocks can generate the negative co-movement between energy use and output consistent with the decoupling of these variables in recent decades. This is confirmed in a variance decomposition exercise, where energy-saving technology shocks emerge as the main driver of energy-saving technical change. Note that this result arises naturally, as we only require energy-saving technology shocks to explain the maximum share of variations in energy-saving technology conditional on the other shocks. Energy price shocks also contribute meaningfully to energy-saving technological change, especially at longer horizons. After eight years, they account for about a third of the variations. Thus, while endogenous, directed technical change is an important mechanism, it is not the major source of energy-saving innovation. The historical decomposition reveals that the drivers of energy-saving technology have varied across time. While the energy price shocks in the 1970s were a crucial impetus for developing energy-saving technologies, energy-saving technology shocks explain a large share of the variation in the 1990s and early 2000s, when the United States introduced a number of major programs to improve energy efficiency such as the Energy Star program.

How should we interpret energy-saving technology shocks? We think of such shocks as breakthroughs in the development of energy-saving technologies unrelated to energy prices. To corroborate this interpretation, we rely on patent data, which is commonly used as a proxy for innovation. The European Patent Office has devised a detailed classification scheme for patents in energy-saving technologies. We find that green patenting

and energy-saving technical change co-move meaningfully over our sample of interest. Interestingly, green patenting tends to lead energy-saving technical change, consistent with the notion that new green technologies that are patented become operational with a lag. We show that a shock to green patenting, identified by maximizing the explained variation in the patent share in energy-saving technologies, looks very much akin to an energy-saving technology shock – after we control for energy price shocks. Importantly, the green patenting shock also generates the negative co-movement between energy consumption and output, resulting in a significant fall in the energy intensity.

A comprehensive series of sensitivity checks indicate that the results are robust along a number of dimensions including the identification design, the model specification and the sample period.

Related literature and contribution. Our paper relates to a long-standing literature studying the relation between economic growth and energy use. A key object of interest is the elasticity of substitution between energy and other inputs, which is estimated to be relatively low in the short-run (Berndt and Wood, 1975) but considerably higher in the longer run (Griffin and Gregory, 1976). A number of early theoretical studies rationalized this stylized fact using vintage structures (Atkeson and Kehoe, 1999; Wei, 2003). Motivated by empirical evidence pointing to input-saving innovation in response to resource scarcity (Popp, 2002; Hanlon, 2015), a more recent theoretical literature studies directed technical change, which over time can be another source of factor substitutability (Hassler, Krusell, and Olovsson, 2012; Acemoglu et al., 2012; Aghion et al., 2016; Fried, 2018; Cakir Melek and Orak, 2023; Casey, 2024). In a closely related paper, Hassler, Krusell, and Olovsson (2021) investigate energy-saving technical change. Assuming an aggregate production function, they isolate series of input-saving technology given annual data on inputs and their prices. Motivated by the stark increase in energy-saving technical change after the big oil shocks in the 1970s, they develop a theory of directed technical change and emphasize energy price shocks as a key driver of energy saving.

We propose a refined approach to measure energy-saving technical change and provide novel measures of input-saving technologies at the *quarterly* frequency. This involves a significant data effort, collecting and constructing the relevant data on inputs and prices at the quarterly frequency for the entire post-WWII period. In particular, we also adjust our measures for changes in utilization, which has been found to be important to measure total factor productivity (Basu, Fernald, and Kimball, 2006; Basu et al., 2013; Fernald, 2014). Equipped with these new measures of input-saving technology, we investigate the drivers of energy-saving technical change. We allow for input-saving technology shocks

in addition to energy price shocks and find that exogenous energy-saving technology shocks are the main driver of energy saving. From a policy perspective, this points to an important role of R&D subsidies as in [Acemoglu et al. \(2012\)](#) complementing the price signal emphasized in [Hassler, Krusell, and Olovsson \(2021\)](#). Relative to these studies, we arrive at these conclusions based on a semi-structural approach, imposing a minimum set of assumptions about the exact structure of the economy.

Methodologically, we build on the literature on identifying technology shocks using VAR techniques (see [Ramey, 2016](#), Chapter 5). A commonly used identification strategy are medium-run restrictions ([Barsky and Sims, 2011](#); [Kurmman and Otrok, 2013](#); [Kurmman and Sims, 2021](#)). We use this approach to identify input-saving technology shocks. In that sense, we also relate to an influential literature studying different facets of technical change. For instance, [Fisher \(2006\)](#) shows how to identify investment-specific technology shocks à la [Greenwood, Hercowitz, and Krusell \(1997\)](#) using the relative price of investment. We contribute to this literature by shedding light on the role of energy-saving technology shocks, documenting meaningful impacts on output and energy use. Relatedly, [Khan et al. \(2019\)](#) analyze how different technology shocks affect emissions and find that emissions generally rise following such shocks. Consistent with this evidence, we show that only energy-saving technology shocks can generate the negative co-movement between output and emissions.

Our paper also relates to a large environmental economics literature on the rebound effect of energy efficiency – the idea that improvements in energy efficiency may trigger behavioral responses leading to an increase in energy use. The micro-level evidence on the rebound effect is mixed, and model-based estimates vary widely depending on structural assumptions and calibration targets ([Turner, 2013](#); [Gillingham, Rapson, and Wagner, 2016](#)). There is little direct evidence on the rebound effect at the macroeconomic level. We contribute to this literature by showing that an energy-saving technology shock leads to a persistent fall in energy consumption, pointing to weak rebound effects in the aggregate. This result contrasts with two recent studies that analyze the macroeconomic effects of energy-efficiency shocks ([Bruns, Moneta, and Stern, 2021](#); [Soofi-Siavash and Moench, 2023](#)), both of which find strong dynamic rebound effects that essentially offset the initial decline in energy use. Importantly, however, neither of these studies account for the potentially confounding effects of *both* energy price and other input-saving technology shocks, which turns out to be crucial in this context.

Outline. The paper proceeds as follows. In the next section, we explain our approach to measure energy-saving technical change. Section 3 provides new estimates on the short-

term elasticity of substitution between energy and other inputs. In Section 4, we introduce our semi-structural model of the U.S. economy and study the underlying drivers of energy-saving technical change. Section 5 concludes.

2. Measuring Energy-saving Technical Change

How have energy-saving technologies evolved over the last several decades and what are the underlying drivers? To answer these questions, we need to obtain measures of input-saving technologies. In this section, we discuss the conceptual framework that forms the basis for our measurement of input-saving technical change.

2.1. Conceptual framework

Our starting point is an aggregate production function, which relates aggregate domestic output to the aggregate inputs of the production factors. We assume that aggregate output is produced using three factors: aggregate capital K , aggregate labor L and aggregate fossil energy E . Further, we posit that capital is utilized at the rate Z , and that the labor effort exerted per unit of labor is N . Specifically, our production function takes the following form

$$Y_t = F \left(A_t (Z_t K_t)^\alpha (N_t L_t)^{1-\alpha}, A_{et} E_t \right), \quad (1)$$

where A and A_e are input-saving technology levels for capital/labor and energy, respectively, and $\alpha \in (0, 1)$ is a share parameter. Since fossil energy is both an intermediate and final good, Y will not exactly equal domestic business sector production, rather it will be business sector output less the value of energy use outside of domestic production (i.e. energy consumed by consumers and net energy exports).

In the short run, energy is thought to be strongly complementary to capital and labor (see e.g. [Pindyck and Rotemberg, 1983](#)). We formalize this by using a nested CES production function as in [Hassler, Krusell, and Olovsson \(2021\)](#):

$$Y_t = \left[(1 - \gamma) \left(A_t (Z_t K_t)^\alpha (N_t L_t)^{1-\alpha} \right)^{(\varepsilon-1)/\varepsilon} + \gamma (A_{et} E_t)^{(\varepsilon-1)/\varepsilon} \right]^{\varepsilon/(\varepsilon-1)}, \quad (2)$$

where ε is the elasticity of substitution between energy and the other inputs and γ is a share parameter. Note that A_t and A_{et} are well defined only when ε is not equal to one.

From the cost minimization problem, we obtain the factor demands for labor and en-

ergy. Assuming competitive factor markets, these are given by:

$$(1 - \alpha)(1 - \gamma) \left(\frac{A_t(Z_t K_t)^\alpha (N_t L_t)^{1-\alpha}}{Y_t} \right)^{(\varepsilon-1)/\varepsilon} = \frac{W_t N_t L_t}{Y_t} \equiv L_t^{\text{share}}$$

$$\gamma \left(\frac{A_{et} E_t}{Y_t} \right)^{(\varepsilon-1)/\varepsilon} = \frac{P_{et} E_t}{Y_t} \equiv E_t^{\text{share}},$$

where W_t is the real wage and P_{et} is the real price for fossil fuel energy.

Log-linearizing and rearranging these equations, we obtain the following expressions for the two input-saving technology series:

$$\Delta \ln A_t = \Delta \ln Y_t - (\alpha \Delta \ln K_t + (1 - \alpha) \Delta \ln L_t) - \Delta \ln U_t - \frac{\varepsilon}{\varepsilon - 1} \Delta \ln L_t^{\text{share}} \quad (3)$$

$$\Delta \ln A_{et} = \Delta \ln Y_t - \Delta \ln E_t - \frac{\varepsilon}{\varepsilon - 1} \Delta \ln E_t^{\text{share}}, \quad (4)$$

where $\Delta \ln U_t = \alpha \Delta \ln Z_t + (1 - \alpha) \Delta \ln N_t$ is (the change in) aggregate utilization. Under our assumptions, we can back out how input-saving technology has evolved over time, given data on output, capital, labor, energy use, the labor share, the energy share, and utilization.

The key unknown parameter is ε , which we estimate in Section 3. The share parameter α is set to match the relative capital and labor shares.¹ Note that the share parameter γ drops out when we consider log changes.²

2.2. Data

Since our production function is gross in nature with energy being an intermediate input, we can interpret Y as domestic business output minus the value of energy consumption outside of domestic production. Thus, we have to net out net energy exports from business output. For business output, we use gross value added for total business from the National Income and Product Accounts (NIPA) tables.

For energy, we focus on the three major types of fossil fuels: coal, oil and natural gas. From the Energy Information Administration (EIA), we are able to source data on consumption, imports, exports, and prices of these fossil fuels. For energy consumption, we construct a fossil fuel composite, summing over the consumption of coal, oil and natu-

¹From the relative factor shares we have $\alpha/(1 - \alpha) = K^{\text{share}}/L^{\text{share}}$, from which we can solve for α .

²The same is true if we relax the assumption of competitive factor markets. As long as the markup is constant, it will drop out when considering log changes.

ral gas expressed in British thermal units (Btu), $E = E^c + E^o + E^g$. To measure energy consumption in the business sector, we focus on consumption of each fossil fuel in the commercial and industrial sectors as well as the business share of fossil fuel consumption in the transportation and electric power generation sectors. Similarly, we construct composites of fossil energy exports and imports to obtain a measure of net energy exports. Finally, we construct a corresponding price index as a weighted average of the coal, oil and natural gas price, $P_e = (P^c E^c + P^o E^o + P^g E^g) / E$. The energy share is then given by business energy expenditure relative to business output. A key challenge with the energy data is that quarterly information is only available from the mid-1970s. For the period before, we temporally disaggregate the annual variables using the Chow-Lin approach, see Appendix A for details.

For capital K , we use capital input calculated from disaggregated quarterly NIPA investment data using the perpetual inventory method. The labor input L is constructed based on data for hours in the business sector and labor composition and quality data from the Bureau of Labor Statistics (BLS). Utilization is constructed as a weighted average of industry utilization rates using industry weights from [Basu, Fernald, and Kimball \(2006\)](#). Industry utilization rates are constructed as $\Delta \ln U_i = \beta_i \Delta \ln (H^i / N_w^i)$ where β_i are estimates from [Basu et al. \(2013\)](#) and H^i / N_w^i is hours-per-worker by industry from the monthly employment reports of the BLS. The labor share of income is constructed from NIPA data for the corporate sector. All these series are sourced from [Fernald \(2014\)](#).³

We discuss the sources and the construction of all the variables used in the measurement of the input-saving technology measures in more detail in Appendix A.

3. Estimating the Energy Complementarity

As discussed in the previous section, the elasticity of substitution between energy and other inputs is a crucial parameter in our context. While energy is often thought to be strongly complementary to other inputs, at least at shorter horizons, the empirical estimates vary largely (see e.g. [Koetse, De Groot, and Florax, 2008](#), for a meta analysis).

Our aggregate production function framework gives rise to a straightforward estimation equation for the elasticity of substitution. By rearranging the energy demand

³We use a recent vintage of the TFP data where hours worked are detrended using a biweight filter as discussed in [Fernald \(2015\)](#).

schedule (4), we obtain

$$\Delta \ln Y_t - \Delta \ln E_t = \beta \Delta \ln E_t^{\text{share}} + \mathbf{X}_t' \boldsymbol{\phi} + u_t, \quad (5)$$

where β is a function of the elasticity of substitution and u_t captures unobserved variations in energy-saving technical change $\Delta \ln A_{et}$. To generalize this specification we allow for the inclusion of additional controls, $\mathbf{X}_t$.

To estimate the slope of the demand schedule, β , we need a supply shifter different from energy-saving technology shocks, which are captured in u_t . A natural candidate for such a supply shifter are oil supply shocks. Such shocks tend to explain a meaningful share of the contemporaneous variations in the energy share, satisfying the relevance requirement. Importantly, they will likely be uncorrelated with energy-saving technology, at least contemporaneously, given the substantial lags in directed energy-saving innovation. This suffices for our purposes as we are estimating the short-run elasticity of substitution. Specifically, we use the oil supply shocks by [Baumeister and Hamilton \(2019\)](#) and the oil supply news shocks from [Känzig \(2021\)](#) as instruments for the energy share. To improve precision, we also control for four lags of the dependent variable and the energy share. However, our results are robust to the selection of controls.⁴ Given an estimate of $\hat{\beta}$ we can then back out the elasticity of substitution as $\hat{\varepsilon} = \hat{\beta} / (\hat{\beta} - 1)$.

Table 1 shows our estimates for the elasticity of substitution between energy and other inputs for four different sample periods: our baseline sample from 1976Q2 onward for which we have quarterly data on energy prices and consumption and consistently observe the oil supply shock instruments, the full sample going back to 1949Q1 for which we use interpolated data for energy prices and consumption in the earlier part of the sample, a sample that starts in 1983Q1 after the big oil shocks of the 1970s, and a sample that stops before the global financial crisis in 2008Q3.

Oil supply shocks are relatively strong instruments for the energy share. In our baseline sample, the robust F-statistic is 15.63, which is safely above the commonly used threshold of 10. The estimate for the elasticity of substitution between energy and other inputs is 0.026. Note that while we cannot reject the null hypothesis that the elasticity is zero, we can safely reject the null that the elasticity is greater than or equal to 1. Thus, we find stark evidence for a strong complementarity between energy and other inputs. Remarkably, the estimated value aligns closely with the findings of [Hassler, Krusell, and](#)

⁴To mitigate the role of influential observations, we drop values that are below the 1 or above the 99 percentile of the distribution of the predicted energy share values from the first stage. In our baseline sample, there are two such values, 2008Q4 and 2020Q2, which are clear outliers.

Table 1: The elasticity of substitution between energy and other inputs

	Baseline	Post-WWII	Post-70s	Pre-GFC
Elasticity of substitution (ϵ)	0.026 (0.023)	0.030 (0.029)	0.023 (0.022)	0.021 (0.037)
First-stage F statistic	15.63	10.38	18.06	13.17
Observations	181	287	154	229

Notes: This table shows the estimates for the elasticity of substitution between energy and the other inputs estimated from equation (5), instrumenting the energy share using the oil supply shocks from [Baumeister and Hamilton \(2019\)](#) and [Känzig \(2021\)](#). Our baseline estimates, reported in the first column, are from the sample 1976Q2 to 2022Q4, for which we have non-interpolated data for energy consumption and prices, and consistently observe the oil supply shocks. The second column reports estimates from the full sample, 1949Q1 to 2022Q4, using interpolated data for energy consumption and censoring the missing oil supply shocks to zero. The third column contains estimates from a sample excluding the big oil shocks in the 1970s, 1983Q1 to 2022Q4. The last column shows estimates from a sample excluding the global financial crisis, 1949Q1 to 2008Q2. Robust standard errors are displayed in parentheses, together with the robust first-stage F-statistic below.

[Olovsson \(2021\)](#), who structurally estimate a model of endogenous resource-saving technical change using Bayesian maximum likelihood. Their posterior mean of the elasticity of substitution is 0.021 and thus well within the confidence bands of our estimates.

Our finding of an elasticity of substitution close to zero is highly robust. As can be seen from the second and third columns in Table 1, the estimated elasticity is consistent across different sample periods. Excluding the big oil shocks from the 1970s or stopping the sample before the global financial crisis yields very similar results. In Appendix B.1, we perform a number of additional sensitivity checks. In particular, we obtain similar estimates if we just use one of the two oil supply shock instruments or if we use the global oil demand shocks from [Baumeister and Hamilton \(2019\)](#) instead. Furthermore, we show that our results are robust with respect to the inclusion of additional controls. While there is some model uncertainty – our estimates range from 0.006 to 0.045 – these results strongly support the notion that energy is highly complementary to other inputs, at least in the short run.

4. The Drivers of Energy-saving Technology

In this section, we study the drivers of energy-saving technical change. We first analyze the historical evolution of input-saving technical progress before investigating into its underlying driving forces through the lens of our semi-structural model, allowing for

technology and energy price shocks.

4.1. The evolution of input-saving technological change

Equipped with our estimate of the elasticity of substitution, $\hat{\varepsilon} = 0.026$, we can construct time series of input-saving technology, A_t and A_{et} , based on equations (3)-(4). We first construct a series for the growth rates of the input-saving technologies and then cumulate these to get the corresponding indices.




Figure 2: Input-saving technologies

Notes: This figure displays our measures of input-saving technologies, constructed based on equations (3)-(4). Panel A shows the indices for capital/labor-saving technology and energy-saving technology, normalized to 1 in 1949Q1. Panel B shows the corresponding growth rates, smoothed using a simple 4-period moving average.

Figure 2 shows the evolution of capital/labor-input saving technology on the left and energy-saving technology on the right. We can see that capital/labor-saving technologies have improved substantially over our sample but the trend has somewhat slowed

from the 1980s onward. This is also visible from the lower panel, which shows a marked slowdown in capital/labor saving technology growth in the second part of the sample. The measure turns out to be fairly highly correlated with total factor productivity from [Fernald \(2014\)](#), with a correlation coefficient for the growth rates of above 0.9. Energy-saving technologies increased only marginally from 1950 to 1970, but have experienced a markedly higher growth following the big oil shocks of the 1970s. This trend has been sustained through the present. The kink in energy-saving technology is consistent with the evidence in [Hassler, Krusell, and Olovsson \(2021\)](#), even though it is somewhat less pronounced in our measure. The consistent growth in energy-saving technical change over this period cannot be accounted for by a transition from fossil to renewable energy, rather it is driven by a more efficient use of fossil energy in the production process (as can be seen in Appendix Figures [B.1](#) and [B.2](#)). Importantly, there were significant variations in the rate of energy-saving technological change throughout our sample.

Energy-saving technological change turns out to be only weakly correlated with capital/labor-saving technological change, and in some instances the two even move in opposite directions. This weak correlation between the respective growth rates is not a mere consequence of our low estimated elasticity of substitution. In fact, the finding is robust to varying the estimated elasticity of substitution from 0.005 to 0.1.

4.2. A semi-structural model of the U.S. economy

We now introduce our econometric framework to study the drivers of energy-saving technology. We follow a semi-structural approach, imposing a minimum set of restrictions to identify potential drivers of energy-saving technological change.

We are interested in modeling the U.S. economy. Let $\mathbf{Y}_t$ denote a $k \times 1$ vector of quarterly time series, consisting of our input-saving technological change measures, real GDP, total energy consumption, energy prices, and the GDP deflator. We assume that the dynamics of $\mathbf{Y}_t$ can be characterized by the following structural vector moving-average representation:

$$\mathbf{Y}_t = \mathbf{B}(L)\mathbf{S}\varepsilon_t, \quad (6)$$

where ε_t is a vector of structural shocks driving the economy with $\mathbb{E}[\varepsilon_t \varepsilon_t'] = \mathbf{I}$, $\mathbf{B}(L) \equiv \mathbf{I} + \mathbf{B}_1 L + \mathbf{B}_2 L^2 + \dots$ is a matrix lag polynomial, and $\mathbf{S}$ is the structural impact matrix.

Assuming that the vector-moving average process (6) is invertible, it admits the fol-

lowing VAR representation:

$$\mathbf{A}(L)\mathbf{Y}_t = \mathbf{S}\boldsymbol{\varepsilon}_t = \mathbf{u}_t,$$

where $\mathbf{u}_t$ is a $k \times 1$ vector of reduced-form innovations with variance-covariance matrix $\mathbb{E}[\mathbf{u}_t\mathbf{u}_t'] = \boldsymbol{\Sigma}_u$ and $\mathbf{A}(L) \equiv \mathbf{I} - \mathbf{A}_1L - \dots$ is a matrix lag polynomial. Truncating the VAR to order p , we can estimate the model using standard techniques and recover an estimate of $\mathbf{B}(L)$.

The main identification problem is then to find the structural impact matrix $\mathbf{S}$. From the linear relation between the structural shocks and the reduced-form innovations, we obtain the following covariance restrictions $\mathbf{S}\mathbf{S}' = \boldsymbol{\Sigma}_u$. The Cholesky decomposition of $\boldsymbol{\Sigma}_u$, denoted by $\tilde{\mathbf{S}}$, is one potential solution satisfying these restrictions. The entire set of permissible values of $\mathbf{S}$ consistent with $\mathbf{S}\mathbf{S}' = \boldsymbol{\Sigma}_u$ can then be described by $\tilde{\mathbf{S}}\mathbf{Q}$, where $\mathbf{Q}$ is an orthonormal rotation matrix. The identification problem thus amounts to finding the appropriate rotation $\mathbf{Q}$.

Max-share approach. Our goal is to identify three distinct potential drivers of energy-saving technology: energy price shocks, capital/labor-saving technology shocks, and energy-saving technology shocks. We remain agnostic about the other drivers of the economy. We identify these shocks using the max-share approach of Uhlig (2004). The idea is to identify shocks that explain the maximum amount of the forecast error variance (FEV) in a target variable at a certain horizon. We detail the approach below.

The h -step ahead forecast error of $\mathbf{Y}_t$ can be written as

$$\mathbf{Y}_{t+h} - \mathbb{E}_{t-1} \mathbf{Y}_{t+h} = \sum_{l=0}^h \mathbf{B}_l \tilde{\mathbf{S}}\mathbf{Q}\boldsymbol{\varepsilon}_{t+h-l}.$$

The FEV share of variable i attributable to shock j at horizon h is then

$$\Omega_{i,j}(h) = \frac{\sum_{l=0}^h \mathbf{B}_{i,l} \tilde{\mathbf{S}}\mathbf{q}_j \mathbf{q}_j' \tilde{\mathbf{S}}' \mathbf{B}_{i,l}'}{\sum_{l=0}^h \mathbf{B}_{i,l} \boldsymbol{\Sigma}_u \mathbf{B}_{i,l}'},$$

where $\mathbf{B}_{i,l}$ is the i th row of lag polynomial at lag l and $\mathbf{q}_j$ is the j th column of $\mathbf{Q}$.

The idea is then to choose $\mathbf{q}_j$ to maximize the FEV share of a given variable i at a certain horizon H :

$$\max_{\mathbf{q}_j} \Omega_{i,j}(H) \quad \text{s.t. } \mathbf{q}_j' \mathbf{q}_j = 1,$$

as in [Kurmman and Sims \(2021\)](#). The constraint ensures that $\mathbf{q}_j$ belongs to an orthonormal matrix. The relevant variable and horizon(s) are chosen based on economic theory with the aim to motivate a certain economic interpretation of the shock.⁵

Identifying assumptions. Given that we are interested in identifying three separate shocks, we proceed sequentially. We first identify an energy price shock as the shock that explains the maximum FEV share at the four year horizon, obtaining $\mathbf{q}_{pe}$. This is motivated by the fact that energy price shocks seem to have the largest impact on energy prices at horizons up to four years (see e.g. [Kilian, 2009](#); [Baumeister and Hamilton, 2019](#); [Känzig, 2021](#)). As a robustness check, we alternatively identify this shock using energy supply disruptions as an external instrument for the energy price shock.

Next, we identify an orthogonalized shock to capital/labor-saving technical change as the shock that is orthogonal to the energy price shock and maximizes the FEV share in capital/labor-saving technology ten years out. This strategy builds on a large literature identifying technology shocks using the max-share approach at medium-run horizons ([Barsky and Sims, 2011](#); [Kurmman and Otrok, 2013](#); [Kurmman and Sims, 2021](#)). In this way, one can leverage the uncontroversial assumption that technical change is driven by technology shocks at longer horizons without facing the challenges of estimating the long-run variance.

The maximization problem in this case reads:

$$\max_{\mathbf{q}_a} \Omega_{a,a}(40) \quad \text{s.t. } \mathbf{q}'_a \mathbf{q}_a = 1 \text{ and } \mathbf{q}'_a \mathbf{q}_{pe} = 0.$$

Finally, we identify an orthogonalized energy-saving technology shock as the shock that is orthogonal to the energy price and the capital/labor-saving technology shocks and maximizes the FEV share in energy-saving technological change ten years out:

$$\max_{\mathbf{q}_{ae}} \Omega_{ae,ae}(40) \quad \text{s.t. } \mathbf{q}'_{ae} \mathbf{q}_{ae} = 1, \mathbf{q}'_{ae} \mathbf{q}_a = 0 \text{ and } \mathbf{q}'_{ae} \mathbf{q}_{pe} = 0.$$

The ordering of the energy price shock and the capital/labor-saving technology shock does not turn out to be consequential (see Section 4.6). However, for identifying energy-saving technology shocks, it is crucial to control for these other driving forces. In Appendix C, we provide details on the implementation of our identification strategy.

At this point, a side note on the role of technology news shocks is in order (see [Beaudry](#)

⁵Alternatively, one can also choose the shock to maximize the FEV of a given variable up to a certain horizon H : $\max_{\mathbf{q}_j} \sum_{h=0}^H \Omega_{i,j}(h)$ s.t. $\mathbf{q}'_j \mathbf{q}_j = 1$, as in [Uhlig \(2003\)](#) or [Barsky and Sims \(2011\)](#).

and Portier, 2014, for an overview). It is important to note that we do not aim to isolate technology news from surprise shocks – in general our identified technology shocks may capture both news and surprise shocks. Our key goal is to disentangle energy-saving technology from capital/labor-saving technology shocks.

Empirical specification. We estimate the model on quarterly data from 1949Q1 to 2022Q4. The data on real GDP and the GDP deflator is from FRED. We construct a measure of total domestic energy consumption, akin to our business energy consumption measure using data from the EIA. The input-saving technology measures and energy prices are based on our measurement as described in Section 2. All variables enter in log-levels. We set the lag order of the VAR to 4, as is customary with quarterly data.

4.3. The impact of technology and energy price shocks

How do input-saving technology and energy price shocks transmit to the economy? And which shocks are the most promising to explain the decoupling of output and energy consumption? Figure 3 shows the impulse responses to one-standard deviation input-saving technology shocks. The responses to capital/labor-saving technology shocks are depicted in blue, the responses to energy-saving technology shocks are shown in green. The solid lines are the point estimates and the shaded areas are 68 and 90 percent confidence bands based on 10,000 bootstrap replications.

A shock to capital/labor-saving technology leads to a persistent increase in input-saving technologies for capital and labor. This in turn generates a substantial increase in output while prices fall. Since energy is complementary to capital and labor, energy consumption also increases significantly and persistently. Interestingly, energy prices are not very responsive on impact but tend to increase at longer horizons, likely in response to higher energy demand. This may help explain the response of the GDP deflator, which falls but not very persistently. Importantly, capital/labor-saving technology shocks do not seem to crowd out energy-saving technological change. Energy-saving technology falls slightly after about one year but the response is not statistically significant.

Capital/labor-saving technology shocks have a difficult time explaining the disconnect between output and energy use, as the two variables co-move strongly positively conditional on these shocks. This can also be seen from the response of the energy intensity, defined as energy use per unit of output. As shown in Figure 5, the energy intensity even tends to increase slightly in response to capital/labor-saving technology shocks, as energy consumption increases somewhat more strongly than output.




Figure 3: The responses to input-saving technology shocks

Notes: This figure shows the impulse responses to a one standard deviation shock in capital/labor-saving technology (in blue) and energy-saving technology (in green), respectively. The point estimates are depicted as solid lines, and the 68 and 90 percent confidence bands are the dark and light shaded areas.

Energy-saving technology shocks propagate through the economy in a very different manner. As expected, an energy-saving technology shock leads to a significant and persistent increase in input-saving technologies in energy. This leads to a fall in energy consumption; however, given that energy can be used more efficiently, output still increases, albeit the response is weaker than for capital/labor-saving technology shocks. The responses of prices are less clear. Energy prices fall on impact but rebound quickly. However, the response is rather imprecisely estimated. The GDP deflator even tends to increase slightly but again, the response features a considerable degree of estimation uncertainty. The energy intensity, defined as energy use over output, falls substantially and the response persists even ten years out.

These results speak directly to concerns about rebound effects of energy efficiency improvements ([Gillingham, Rapson, and Wagner, 2016](#)). We find that an energy-saving technology shock that raises energy efficiency does *not* lead to increased energy use. Instead, these shocks lead to a reduction in aggregate energy consumption that is significant and persists through the end of our impulse horizon. This stands in contrast to recent evidence on the macroeconomic effects of energy-efficiency shocks ([Bruns, Moneta, and Stern, 2021](#); [Soofi-Siavash and Moench, 2023](#)). These studies find that energy consumption declines on impact after such shocks but then reverts to its initial level over time. [Bruns, Moneta, and Stern \(2021\)](#) use independent component analysis – a more statistical approach to identification – in a simple trivariate model of energy use, output and energy prices. However, their approach does not explicitly account for other input-saving technology shocks. [Soofi-Siavash and Moench \(2023\)](#) combine medium-run and sign-restrictions on energy intensity and aggregate productivity derived from economic theory, but their implementation does not account for potentially confounding energy price shocks. In contrast, our approach relies on direct measures of input-saving technological change and crucially, controls for other potentially important drivers of energy-saving technical change.

The muted response of energy prices is consistent with our finding of limited rebound effects. One channel through which rebound effects can arise is a decline in energy prices, which would weaken incentives to reduce energy consumption. The absence of such economy-wide rebound effects has important implications for the effectiveness of energy-efficiency policies.

There is also little evidence that energy-saving technology shocks crowd out other input-saving technologies. We observe a short-lived fall in capital/labor-saving technology, which may explain the somewhat weaker output response but overall, these technology spillovers do not seem to be very substantial. Thus the much discussed produc-

tivity slowdown does not seem to be driven by an uptick in energy-saving technological change, at least not conditional on energy-saving technology shocks.

Given the negative co-movement between output and energy use, resulting in a substantial fall in the energy intensity, energy-saving technology shocks are a promising candidate to account for the divergence between output and energy use.




Figure 4: The responses to energy price shocks

Notes: This figure shows the impulse responses to a one standard deviation energy price shock, identified using the max-share approach. The point estimates are depicted as solid lines, and the 68 and 90 percent confidence bands are the dark and light shaded areas.

What about energy price shocks? [Hassler, Krusell, and Olovsson \(2021\)](#) argue that such shocks are a crucial driver of energy-saving technological change. Figure 4 shows the impulse responses to a one-standard deviation energy price shock. Energy prices increase by close to 10 percent on impact and the response persists even five years out.

Output and energy consumption do not respond on impact but start to fall significantly and persistently after about a year, given the increase in cost. The increase in energy prices also passes through other prices, as reflected in the persistent increase in the GDP deflator. Overall, these responses are very similar to the effects of oil price shocks documented in the literature (see e.g. [Kilian, 2009](#); [Känzig, 2021](#)).

Interestingly, energy-saving technologies increase significantly after some lag, in line with the evidence in [Hassler, Krusell, and Olovsson \(2021\)](#). However, these advances come at a cost of a slowdown in capital/labor-saving technical change, which helps explain the quite persistent, adverse effects of energy price shocks on output. Energy consumption falls by more than output and the response is also more persistent, resulting in a substantial decline in the energy intensity at longer horizons (see Figure 5). While energy price shocks may help explain the increasing disconnect between output and energy use, they are likely not the main driver as they imply a positive co-movement between output and energy use.




Figure 5: The impact on emissions intensity

Notes: This figure shows the impulse responses of the energy intensity, defined as energy use over output, to a one standard deviation shock in capital/labor-saving technology (in blue), energy-saving technology (in green), and energy price shocks (in orange). The point estimates are depicted as solid lines, and the 68 and 90 percent confidence bands are the dark and light shaded areas.

They will, however, likely be a key driver of energy use, especially at longer horizons, as the persistent fall in economic activity coupled with a more efficient use of energy causes a sustained decline in energy consumption.

An interesting question is whether it is necessary to separate between different input-saving technologies. Could aggregate TFP shocks account for the decoupling of output and energy consumption? To shed light on this question, we identify an aggregate TFP




Figure 6: The responses to a TFP shock

Notes: This figure shows the impulse responses to a one standard deviation TFP shock, identified using the max-share approach. The point estimates are depicted as solid lines, and the 68 and 90 percent confidence bands are the dark and light shaded areas.

shock as the shock maximizing the FEV share in the utility-adjusted TFP measure by [Fernald \(2014\)](#) ten years out.⁶

Figure 6 shows the corresponding impulse responses. A TFP shock leads to a significant increase in output and a fall in prices. Importantly, however, output and energy consumption also increase, and energy consumption increases by even more than output, as can be seen from the increase in the energy intensity. This illustrates the importance of distinguishing between energy-saving and other technical change. Not all technological progress is energy-saving, and in fact innovations in TFP seem to be driven by the non-energy saving part. In fact, the estimated responses are quite similar to the identified capital/labor-saving technology shock.⁷

4.4. Historical importance

How important are energy-saving technology and energy price shocks in explaining variations in the variables of interest? To answer this question, we perform a variance and historical decomposition. Figure 7 shows the variance decomposition of the capital/labor-saving and energy-saving technology measures, as well as output and energy consumption. Depicted is the fraction of forecast error variance that is explained by the three identified shocks: capital/labor-saving technology shocks (in blue), energy-saving technology shocks (in green), and energy price shocks (in orange).

Energy price shocks and the associated directed technical change turns out to be an important driver of energy saving, as emphasized in [Hassler, Krusell, and Olovsson \(2021\)](#). At longer horizons, energy price shocks explain about a third of the variation in energy-saving technical change. However, the bulk of the variation in energy-saving technology, especially at shorter horizons, is explained by energy-saving technology shocks, which turn out to be the main force behind the rapid growth in energy-saving technology over the past half century. On the other hand, capital/labor-saving technology shocks barely explain any of the variation in energy-saving technology. Importantly, these re-

⁶Note that this shock differs from the capital/labor-saving technology shock in two respects. First, we use the original Fernald TFP measure as the relevant technology series. Furthermore, we do not orthogonalize this shock to the energy price shock.

⁷Our approach seems to identify a convolution of news and surprise shocks, resulting in a technology response that displays a more immediate impact that only slightly builds up over time. This is consistent with the findings of [Kilian, Plante, and Richter \(2023\)](#) who argue that using medium-run restrictions only will in general recover a combination of TFP surprise and news shocks. We do not attempt to disentangle the TFP news from the surprise shock, as our main goal is to disentangle capital/labor-saving technology from energy-saving technology shocks. In Appendix B.7, we discuss the relationship of our approach to the TFP news literature in more detail.




Figure 7: Variance decomposition

Notes: This figure shows the fraction of forecast error variance in input-saving technologies, output and energy consumption that can be accounted for by input-saving technology shocks and energy price shocks, at horizons ranging from one quarter to ten years.

sults emerge naturally as we do not impose that energy-saving technology shocks explain the largest share of the variation in energy-saving technology. This is because we identify energy-saving technology shocks by maximizing the FEV share conditional on the identified energy price and capital/labor-saving technology shocks.

Looking at other variables, we can see that energy-saving technology shocks explain very little of the variation in other input-saving technologies but account for a considerable share of the variations in output and energy consumption. Thus, energy-saving innovation is a facet of technical change that is relevant for understanding aggregate fluctuations.

As expected, capital/labor-saving technology shocks explain the bulk of the variation in capital/labor-saving technology. These shocks are also the most important factor explaining variations in output and explain a large share of the variation in energy consumption. Energy price shocks on the other hand are the dominant driver of energy price

fluctuations (see Figure B.3 in the Appendix), and they also explain up to 20 percent of the variation in output and close to 40 percent of the variation in energy consumption at longer horizons.

The finding that energy price shocks crowd out non-energy-saving technological change seems to be a quantitatively important one. At horizons up to ten years, energy price shocks explain over 15 percent of the variations in capital/labor-saving technology.




Figure 8: Historical decomposition of energy-saving technological change

Notes: This figure shows the historical contribution of input-saving technology shocks and energy price shocks to energy-saving technological change over the post-WWII period. The solid blue line is the historical contribution of the respective shock and the black, dashed line are the actual, annualized changes in energy-saving technology.

Another interesting question is how input-saving technology and energy price shocks have contributed to changes in energy-saving technology during specific historical episodes. To this end, we construct a historical decomposition of energy-saving tech-

nical change. Figure 8 shows the historical contribution of the shocks together with the realized change in energy-saving technology.

We can see that both energy-saving technology and energy price shocks have contributed meaningfully to historical variations in energy-saving technical change. In contrast, the contributions of capital/labor-saving technology shocks have been minimal throughout our sample. Energy-saving technology shocks turn out to be the most important driver. Particularly in the 1990s and early 2000s, they have contributed substantially to changes in energy-saving technical change. Interestingly, this period coincides with major government initiatives to promote energy efficiency, such as the Energy Star program, established in 1992. However, energy-saving technology shocks have also contributed to energy-saving technical change in many other episodes.

Energy price shocks are the second most important driver of energy-saving technical change. They were particularly important in explaining changes in energy-saving technology after the big oil shocks in the 1970s, where we see a significant positive contribution of energy price shocks. Similarly, they have contributed negatively to energy-saving technological change after the oil price collapse in 2014.

4.5. Interpreting the drivers of energy-saving technology

We have seen that energy-saving technology shocks and energy price shocks are the main drivers of energy-saving technical change. How should we think about energy-saving technology shocks? We interpret these shocks as sudden breakthroughs in the development of energy-saving technologies unrelated to energy prices. As such they should be visible in patenting data.

Green patents to proxy energy-saving technical change. The Cooperative Patent Classification (CPC) features a classification scheme for patents in climate change mitigation technologies. We focus on the subset of patents within this scheme that is related to energy-saving technologies, and compute the share of energy-saving patents filed at the US patent office relative to all patents filed at a given point in time (see Appendix A.4 for details).

Our measure of energy-saving technical change co-moves meaningfully with energy-saving patents (see Figure B.5 in the Appendix). Starting in the 1970s, both series display a strong uptick, followed by a subsequent slowdown. Interestingly, patenting in energy-saving technologies seems to lead energy-saving technical change, consistent with the notion that technologies in patents may only become operational with some lag.




Figure 9: The responses to an energy-saving patenting shock

Notes: This figure shows the impulse responses to a one standard deviation shock in energy-saving patenting, identified using the max-share approach. The point estimates are depicted as solid lines, and the 68 and 90 percent confidence bands are the dark and light shaded areas.

To further corroborate this interpretation, we identify an energy-saving patenting shock as the shock that is orthogonal to the energy price and the capital/labor-saving technology shock and explains the maximum FEV in the share of energy-saving patents 15 years out. We pick here a slightly longer horizon than for input-saving technology to account the leading behavior discussed above.

The results are shown in Figure 9. The green patenting shock looks akin to our baseline energy-saving technology shock, identified based on our energy-saving technical change measure. Both lead to a significant increase in energy-saving technology, a fall in energy consumption, an increase in output and a fall in the energy intensity. The only qualitative difference is the response of the GDP deflator, but the difference turns out to be not statistically significant. Importantly, the shock is not associated with an increase in the energy price. Overall, this confirms the notion that our identified energy-saving technology shocks correspond to breakthroughs in the development of green technologies unrelated to energy price shocks that allow for substantial energy saving in production. Our results suggest that these shocks have been instrumental in decoupling economic growth from energy use.

Alternative identification of energy price shocks. Even though our identification of energy price shocks is grounded in the stylized fact that such shocks are typically persistent but not permanent and dissipate after about four to five years (see e.g. [Känzig, 2021](#)), it is somewhat non-standard. An alternative approach to identify energy price shocks that has become increasingly popular is to use external instruments. Oil price shocks are the prime example for energy price shocks, and several shock measures have been put forward in the literature ([Kilian, 2008](#); [Baumeister and Hamilton, 2019](#); [Känzig, 2021](#)). One challenge is that oil price shocks only became significant in the 1970s, as the U.S. grew increasingly dependent on imported oil. However, as discussed in [Gertler and Karadi \(2015\)](#), the external instrument approach can also be employed if the instrument is only available for a subsample. Therefore, we alternatively use the oil supply shock from [Baumeister and Hamilton \(2019\)](#) and the oil supply news shock from [Känzig \(2021\)](#) as an instrument for the period they are observed, to identify the energy price shock.

This, however, likely reduces the relevance of the instrument. While oil supply news shocks turn out to be a relatively strong instrument for the energy price shock, with a robust F-statistic of 10.7 the F-statistic for oil supply shocks is considerably weaker and below the commonly used threshold of 10.

Figure 10 shows the impulse responses to a one standard deviation energy price shock, identified using the oil shock instruments, together with the energy price shock identified




Figure 10: Energy price shocks identified using external instruments

Notes: This figure shows the impulse responses to a one standard deviation energy price shock, identified using the oil supply and oil supply news shocks from [Baumeister and Hamilton \(2019\)](#) and [Känzig \(2021\)](#) as an external instrument, together with the energy price shock identified using the max-share approach. The point estimates are depicted as solid lines, and the 68 and 90 percent confidence bands are the dark and light shaded areas.

using the max-share approach. We can see that the point estimates are very similar, particularly for the oil supply news shock. The oil supply shock is associated with a much weaker response of the GDP deflator and weaker increase in energy-saving technologies. However, none of the responses are statistically significantly different from each other. Overall, these results suggest that our approach to identifying energy price shocks is robust in the sense that other more commonly used approaches yield comparable results. Because the responses based on the external instruments are somewhat less precisely estimated, likely because the oil shock instruments are only available for a subsample, we focus on the responses from the max-share approach as the baseline.

Structural transformation as a driver of energy-saving technical change? An alternative hypothesis for the decoupling of output and energy consumption is the structural transformation. Over the last decades, the service sector has become increasingly important for the U.S. economy. Therefore, one might hypothesize that this shift from a manufacturing-oriented to a service sector-oriented economy could be driving some of the key results.

In Figure B.6, we show the evolution of the service sector share in the U.S. together with our measure of energy-saving technical change. Note that the service sector share has been increasing throughout our sample, and there is no apparent break in the 1970s. If at all, the increase in the service sector share seems to have slowed in the 1970s before picking up again in the 1980s. This contrasts with our measure of energy-saving technical change, which starts to increase more rapidly around the first oil shock.

To mitigate concerns that our results may be driven by the structural transformation more formally, we include the service share as an additional control variable in our VAR. The results are shown in Figure 11. We can see that controlling for the service share leaves the responses to an energy-saving technology shock virtually unchanged.

As an additional check, we identify an innovation to the service share, by maximizing the share of forecast error variance in the service sector share at horizons up to 10 years. The results are shown in Figure B.7 in the appendix. Reassuringly, an innovation to the service share looks nothing like an energy-saving technology shock. In fact, the shock leads to a slight fall in energy-saving technology, a significant fall in the energy price, and tends to increase energy consumption and its intensity. Thus, such shocks are unlikely to be a key driver of energy-saving technical change.




Figure 11: Controlling for the service share

Notes: This figure shows the impulse responses to a one standard deviation energy-saving technology shock, identified using the max-share approach. We depict our baseline results in green; the estimates from the model controlling for the service sector share are shown in blue. The point estimates are depicted as solid lines, and the 68 and 90 percent confidence bands are the dark and light shaded areas.

4.6. Sensitivity Analysis

In this section, we perform a comprehensive series of robustness checks. The corresponding figures and tables are shown in Appendix B.6.

Degree of energy complementarity. First, we assess the robustness of our results with respect to the elasticity of substitution between energy and non-energy inputs. While our empirical estimates from Section 3 consistently point to low elasticities that are close to zero, there is quite a bit of model and estimation uncertainty around these estimates. Therefore, we study how our results are affected by the value of that elasticity. Specifically, we vary the elasticity from 0.005 to 0.1. Note that the upper bound corresponds to the upper confidence bound at the 95 percent level of the specification with the highest estimated elasticity of substitution (0.045). It turns out that our results are robust to assuming a higher elasticity of between energy and non-energy inputs (see Figure B.8). The only substantive difference is that energy price shocks have a more immediate effect on energy-saving technology when the elasticity is higher.

Orthogonalization of shocks. Next, we study the role of the orthogonalization of energy price, capital/labor-saving and energy-saving technology shocks. Energy price shocks are sometimes thought of as instances of negative technology shocks. However, we think of them as cost push shocks, for instance periods where OPEC unexpectedly constrained their supplies. Our finding that identifying the energy price shock using external instruments for oil supply shocks from the literature produces very similar results to the baseline responses from the max-share approach supports this notion. Nevertheless, it is important to analyze the sensitivity of our results with respect to the order of the orthogonalization. In Figure B.9, we show the responses when we identify the capital/labor-saving technology shock before the energy price shock. The results turn out to be consistent with our baseline responses.

To further shed light on the role of shock orthogonalization, we study the responses without imposing any orthogonalization in Figure B.10. These results are obtained by estimating three separate VARs, identifying one shock at a time. While the orthogonalization does not turn out to be that important for energy price and capital/labor-saving technology shocks, as discussed above, the orthogonalization turns out to be crucial for energy-saving technology shocks. When we do not impose the orthogonalization, the identified shocks to energy-saving technology seem to be picking up other shocks as well. Of particular concern is that these shocks also pick up energy price shocks, which also

explain a non-negligible portion of the variance in energy-saving technology at longer horizons. This intuition is confirmed by the strong and significantly positive response of energy prices in the model with no orthogonalization.

Max-share horizon. In Figure B.11, we further study how sensitive our results are with respect to the choice of the horizon in the max-share approach. We show that identifying the energy price shock as the shock that explains the maximum share of variations in energy prices at the two year horizon produces very similar results. Likewise, the results are robust to setting the horizon for technology shocks to 15 or even 20 years. Finally, the results are also robust to maximizing the sum of the FEV shares up to the given horizon, as in Barsky and Sims (2011).

Sample period. A striking fact that Hassler, Krusell, and Olovsson (2021) document and we confirm in this paper is the stark break in energy-saving technical change in the 1970s. While this break is very interesting as it coincides with the big oil shocks, it may pose some problems for our shock identification approach based on a fixed-parameter model. To mitigate concerns that the structural break in the 1970s may confound our results, we also estimate our VAR on the more recent sample starting in 1974. From Figure B.12, we can see that the results turn out to be very similar. The main difference is that the capital/labor-saving technology shock is associated with an increase in energy-saving technology. Furthermore, energy-saving technology shocks seem to be somewhat more inflationary. Qualitatively, however, the findings are in line with our baseline responses based on the full sample.

In Figure B.13, we further assess the sensitivity of our results with respect to the sample period. It turns out that excluding the Covid-19 pandemic or stopping the sample before the global financial crisis yields consistent results. The results are also robust to starting the sample in the 1960.

Informational sufficiency. We estimate our responses based on a medium-scale VAR. A crucial assumption underlying the VAR approach is that the model spans all the relevant information to recover the structural shock of interest. This is referred to as invertibility. To assess whether our results are subject to non-invertibility concerns, we augment our model with factors estimated based on a large set of macroeconomic and financial variables from the FRED-MD database (McCracken and Ng, 2016). Unfortunately, this database is not available for our full sample. Therefore, we perform this exercise for the shorter sample starting in 1974. Specifically, we augment our VAR with the first four es-

timated factors, which explain a large part of the variation in the database. Our results turn out to be very robust to the inclusion of these additional factors, suggesting that non-invertibility is likely less of a concern in our application (see Figure B.14).

Estimation method. In our baseline specification, we estimate the impulse response functions of our identified shocks based on the VAR model. To mitigate concerns about lag truncation bias, we show in Figure B.15 that our results are robust to estimating the impulse response functions via local projections. Specifically, we impose the same identifying assumptions as described in Section 4.2 to estimate the structural shocks, and then estimate the impulse response functions to these measured shocks using local projections. While the estimated responses are somewhat less stable and precise, many of the qualitative patterns remain the same.

Other specification choices. Finally, our results are robust to the specification of the VAR model. Focusing on business-sector output and energy consumption produces similar results (Figure B.16). Thus, energy-saving technology shocks help to improve energy efficiency beyond the business sector. The results are also robust to increasing the lag order (Figure B.17).

5. Conclusion

The once close relation between economic growth and energy use has weakened substantially in recent decades. While the structural transformation and in particular the expansion of the service sector may have contributed to this phenomenon, the key factor is energy-saving technical change. To confront the climate crisis, it is crucial to understand how to best foster innovation in energy-saving technologies. Should we increase the price of emitting, for instance using carbon taxes, or subsidize green innovation?

We shed light on the historical drivers of energy-saving innovation. Energy price shocks lead to a significant fall in energy consumption and a persistent increase in energy-saving technology: as energy becomes more scarce – reflected in higher prices – firms have stronger incentives to direct innovation toward energy-saving technologies. While this mechanism is operative, we find that the bulk of the historical variations in energy-saving technical change is explained by energy-saving technology shocks. These shocks are the only ones that generate a negative co-movement between output and energy use, and the associated decline in energy consumption is both sizeable and persistent, pointing to weak economy-wide rebound effects. However, because energy-saving technology

shocks do not create a price signal in the form of higher energy prices – in fact theoretically prices could even decrease and in the data we find that the response is insignificant – they do not account for a major share of the variations in energy use. Our results are thus consistent with the findings in [Acemoglu et al. \(2012\)](#): to confront the climate challenge a combination of policies that increase the price of emitting and policies that support green innovation is necessary.

References

- Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous** (2012). "The environment and directed technical change". *American Economic Review* 102.1, pp. 131–166.
- Aghion, Philippe, Antoine Dechezleprêtre, David Hemous, Ralf Martin, and John Van Reenen** (2016). "Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry". *Journal of Political Economy* 124.1, pp. 1–51.
- Atkeson, Andrew and Patrick J. Kehoe** (1999). "Models of energy use: Putty-putty versus putty-clay". *American Economic Review* 89.4, pp. 1028–1043.
- Barsky, Robert B. and Eric R. Sims** (2011). "News shocks and business cycles". *Journal of Monetary Economics* 58.3, pp. 273–289.
- Basu, Susanto, John Fernald, Jonas Fisher, and Miles Kimball** (2013). "Sector-specific technical change". *Manuscript, Federal Reserve Bank of San Francisco*.
- Basu, Susanto, John G. Fernald, and Miles S. Kimball** (2006). "Are technology improvements contractionary?" *American Economic Review* 96.5, pp. 1418–1448.
- Baumeister, Christiane and James D. Hamilton** (2019). "Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks". *American Economic Review* 109.5, pp. 1873–1910.
- Beaudry, Paul and Franck Portier** (2014). "News-driven business cycles: Insights and challenges". *Journal of Economic Literature* 52.4, pp. 993–1074.
- Berndt, Ernst R. and David O. Wood** (1975). "Technology, prices, and the derived demand for energy". *The Review of Economics and Statistics*, pp. 259–268.
- Bruns, Stephan B, Alessio Moneta, and David I Stern** (2021). "Estimating the economy-wide rebound effect using empirically identified structural vector autoregressions". *Energy Economics* 97, p. 105158.
- Cakir Melek, Nida and Musa Orak** (2023). "The Role of Technology and Energy Substitution in Climate Change Mitigation". *Federal Reserve Bank of Kansas City Working Paper* 23-15.
- Casey, Gregory** (2024). "Energy efficiency and directed technical change: implications for climate change mitigation". *Review of Economic Studies* 91.1, pp. 192–228.
- Fernald, John G.** (2014). "A quarterly, utilization-adjusted series on total factor productivity". Federal Reserve Bank of San Francisco.
- (2015). "Productivity and potential output before, during, and after the great recession". *NBER Macroeconomics Annual 2014*. Ed. by J. A. Parker and M. Woodford. Vol. 29. University of Chicago Press, pp. 1–51.

- Fisher, Jonas D. M.** (2006). “The dynamic effects of neutral and investment-specific technology shocks”. *Journal of Political Economy* 114.3, pp. 413–451.
- Fried, Stephie** (2018). “Climate policy and innovation: A quantitative macroeconomic analysis”. *American Economic Journal: Macroeconomics* 10.1, pp. 90–118.
- Gertler, Mark and Peter Karadi** (2015). “Monetary policy surprises, credit costs, and economic activity”. *American Economic Journal: Macroeconomics* 7.1, pp. 44–76.
- Gillingham, Kenneth, David Rapson, and Gernot Wagner** (2016). “The rebound effect and energy efficiency policy”. *Review of environmental economics and policy*.
- Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell** (1997). “Long-run implications of investment-specific technological change”. *The American Economic Review*, pp. 342–362.
- Griffin, James M. and Paul R. Gregory** (1976). “An intercountry translog model of energy substitution responses”. *American Economic Review* 66.5, pp. 845–857.
- Hanlon, Walker W.** (2015). “Necessity is the mother of invention: Input supplies and Directed Technical Change”. *Econometrica* 83.1, pp. 67–100.
- Hassler, John, Per Krusell, and Conny Olovsson** (2012). *Energy-saving technical change*. Tech. rep. National Bureau of Economic Research.
- (2021). “Directed technical change as a response to natural resource scarcity”. *Journal of Political Economy* 129.11, pp. 3039–3072.
- Känzig, Diego R.** (2021). “The macroeconomic effects of oil supply news: Evidence from OPEC announcements”. *American Economic Review* 111.4, pp. 1092–1125.
- Khan, Hashmat, Konstantinos Metaxoglou, Christopher R. Knittel, and Maya Papineau** (2019). “Carbon emissions and business cycles”. *Journal of Macroeconomics* 60, pp. 1–19.
- Kilian, Lutz** (2008). “The economic effects of energy price shocks”. *Journal of economic literature* 46.4, pp. 871–909.
- (2009). “Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market”. *American Economic Review* 99.3, pp. 1053–1069.
- Kilian, Lutz, Michael D. Plante, and Alexander W. Richter** (Apr. 2023). *Estimating macroeconomic news and surprise shocks*. Working Paper 2304. Revised November 2023. Federal Reserve Bank of Dallas.
- Koetse, Mark J., Henri L. F. De Groot, and Raymond J. G. M. Florax** (2008). “Capital-energy substitution and shifts in factor demand: A meta-analysis”. *Energy Economics* 30.5, pp. 2236–2251.
- Kurmann, André and Christopher Otrok** (2013). “News shocks and the slope of the term structure of interest rates”. *American Economic Review* 103.6, pp. 2612–2632.
- Kurmann, André and Eric Sims** (2021). “Revisions in utilization-adjusted TFP and robust identification of news shocks”. *Review of Economics and Statistics* 103.2, pp. 216–235.

- McCracken, Michael W. and Serena Ng** (2016). "FRED-MD: A monthly database for macroeconomic research". *Journal of Business & Economic Statistics* 34.4, pp. 574–589.
- Pindyck, Robert S. and Julio J. Rotemberg** (1983). "Dynamic factor demands and the effects of energy price shocks". *American Economic Review* 73.5, pp. 1066–1079.
- Popp, David** (2002). "Induced innovation and energy prices". *American Economic Review* 92.1, pp. 160–180.
- Ramey, Valerie A** (2016). "Macroeconomic shocks and their propagation". *Handbook of macroeconomics* 2, pp. 71–162.
- Soofi-Siavash, Soroosh and Emanuel Moench** (2023). *Carbon Intensity, Productivity, and Growth*. Tech. rep. Bank of Lithuania.
- Turner, Karen** (2013). "'Rebound' effects from increased energy efficiency: a time to pause and reflect". *The Energy Journal* 34.4, pp. 25–42.
- Uhlig, Harald** (2003). "What moves real GNP?" *Unpublished manuscript*.
- (2004). "Do technology shocks lead to a fall in total hours worked?" *Journal of the European Economic Association* 2.2-3, pp. 361–371.
- Wei, Chao** (2003). "Energy, the stock market, and the putty-clay investment model". *American Economic Review* 93.1, pp. 311–323.

Online Appendix

Unraveling the Drivers of Energy-saving Technical Change

Diego R. Känzig[†]

Charles Williamson[‡]

Contents

A. Data Sources and Construction	39
A.1. Output, capital and labor	39
A.2. Energy consumption and prices	40
A.3. Seasonal adjustment and temporal disaggregation	41
A.4. Other data	41
A.5. Summary	42
B. Additional Results	44
B.1. Estimating the energy complementarity	44
B.2. Role of renewable energy	45
B.3. Variance decomposition	46
B.4. Interpreting the drivers of energy-saving technology	48
B.5. Identifying innovations to service share	49
B.6. Sensitivity analysis	51
B.7. Relation to TFP news literature	62
C. Max-Share Identification	64
References Appendix	65

[†]Northwestern University, CEPR and NBER. E-mail: dkaenzig@northwestern.edu. Web: www.diegokaenzig.com.

[‡]Northwestern University. E-mail: cwilliamson@u.northwestern.edu.

A. Data Sources and Construction

In this appendix, we provide details about the data we use to construct our measures of input-saving technology. As discussed, our approach requires data on output, input quantities and their prices. Given that our goal is to construct measures of technical change at the quarterly frequency, we need quarterly data for all the relevant series. This is a challenge because for many variables, quarterly data is only available from the mid-1970s. However, the 70s have been crucial for energy-saving technical change and ideally we want to have series of technology that span the entire post-WWII period. To this end, we have to temporally disaggregate some of the relevant series, at least for the earlier part of the sample.

A.1. Output, capital and labor

For business output, we use gross value added for total business from the NIPA tables. Following the BLS, we use the expenditure-side measure. As our focus is on a closed economy setting, we have to adjust business output for energy use outside of domestic production. To this end, we subtract net energy exports from our measure of business output.

Since our production function is gross in nature with energy being an intermediate input, we can interpret Y as domestic business output minus the value of energy use outside of domestic production. Thus, we have to net out net energy exports from business output. For business output, we use gross value added for total business from the National Income and Product Accounts (NIPA) tables.

For capital K , we use capital input calculated from disaggregated quarterly NIPA investment data using the perpetual inventory method. L is constructed based on data for hours in the business sector and labor composition and quality data from the Bureau of Labor Statistics (BLS). Utilization is constructed as a weighted average of industry utilization rates using industry weights from [Basu, Fernald, and Kimball \(2006\)](#). Industry utilization rates are constructed as $\Delta \ln U_i = \beta_i \Delta \ln (H^i / N_w^i)$ where β_i are estimates from [Basu et al. \(2013\)](#) and H^i / N_w^i is hours-per-worker by industry from the monthly employment reports of the BLS. The labor share of income is constructed from NIPA data for the corporate sector. All these series are sourced from [Fernald \(2014\)](#).

A.2. Energy consumption and prices

For energy, we focus on fossil fuels. From the Energy Information Administration (EIA), we are able to source data on consumption, imports, exports, and prices. We focus on three fossil fuel types: coal, oil and natural gas. Data on energy consumption by source and sector is provided by U.S. Energy Information Administration in Table 2.1 of their Monthly and Annual Energy Reviews. The data is available for five sectors: residential, commercial, industrial, transportation, and electric power generation. To obtain energy consumption by the business sector, we sum over the consumption of each fossil fuel in the commercial and industrial sectors and the share of the fossil fuel consumption in the transportation and electric power generation sectors consumed by businesses. For the transportation sector, we attribute 40% of the consumption to the business sector for each fossil fuel. We obtain this share from the Annual Energy Outlook – looking at the Outlooks going back to 1979 suggests that the business share of transportation energy consumption has been roughly constant over time. For the electric power generation sector, the EIA provides data on electric power consumption by sector and so we are able to subtract off fossil fuels used by the electric power generation sector for the purpose of providing electricity to non-business uses.¹

Data on imports and exports of energy by source are also provided by the EIA in Table 1.4 of their Monthly and Annual Energy Reviews. As for energy consumption, we construct a series of net energy exports by summing over the exports and imports of each fuel type.

The construction of petroleum and natural gas price data follows [Hassler, Krusell, and Olovsson \(2021\)](#). For coal prices, we deviate from their data sources because the EIA does not provide monthly data on coal prices. For petroleum prices, we use data from the EIA on the U.S. Crude Oil First Purchase Price. The units of this data are in dollars per barrel, we convert this to dollars per BTU by dividing by 5,551,365.23. To construct a series of natural gas prices for the entire sample period, we have to combine two sources. The first series is the U.S. natural gas wellhead price in dollars per thousand cubic feet. The second series is the U.S. natural gas electric power price. The electric power price series first becomes available in 2002 and the wellhead price is available only through 2012. To construct a consistent series, we scale the electric power price series to match

¹To construct our measure of business energy consumption, we simply sum over the consumption of each fuel type. By contrast, [Hassler, Krusell, and Olovsson \(2021\)](#) compute the energy composite as a weighted sum, multiplying oil consumption by 3.82 and gas consumption by 1.63 because the average price per Btu of coal is on average 3.82 times higher for oil and 1.63 times higher for natural gas over our sample of interest. We find that using a weighted or unweighted energy composite produces very similar results.

the wellhead price series in 2002 and then use the scaled electric power price from then on. To convert from dollars per thousand cubic feet to dollars per BTU, we divide by 1,037,000. For coal prices we use the “Producer Price Index by Commodity: Fuels and Related Products and Power: Coal” series since the EIA does not provide coal price data at the monthly or quarterly level. We scale this index to match EIA coal price per BTU in 1949. We confirm that, at the annual frequency, the correlation between the scaled producer price index and the EIA series is 0.97. Based on these series, we construct an energy price index, as weighted average of oil, coal and natural gas prices, weighted by the respective share in energy consumption.

Based on the price and quantity data, we can obtain a measure of energy expenditure, and compute the energy share as business energy expenditure relative to business output.

A.3. Seasonal adjustment and temporal disaggregation

A key challenge with the energy data is that from 1949 through the mid-1970s, the EIA provides data on energy consumption and prices only at the annual frequency. Only after, the EIA started to publish monthly figures. When monthly data is available, we construct quarterly measures by aggregating over the respective months. For quantities, we take the sum. For prices, we take an average. To account for seasonality in the data, we seasonally adjust the energy consumption and the energy share series using the X-13ARIMA-SEATS procedure.

For the data before 1976, we temporally disaggregate the annual data using the Chow-Lin disaggregation procedure with indicators by [Quilis, 2024](#). Specifically, we use maximum-likelihood estimation with an intercept model and use a grid of $[0, 1)$ with 1,000 points for ρ .

A.4. Other data

For the construction of input-saving technology, we focus on variables for the U.S. business sector. In our VAR model, however, we focus on aggregate U.S. data. Thus, we construct an energy consumption measure as described above but for all sectors. Data on real GDP and the GDP deflator is from FRED.

To construct our measures of green patenting, we use data on patent applications from the Google Patents Public Dataset. Specifically, we compute the share of patents filed in energy-saving technologies (CPC codes Y02E, Y02T, Y04S, and the parts of Y02P relevant for energy-saving), relative to the overall number of patents filed in a given period.

A.5. Summary

An overview of all the data, its sources and coverage is shown in Table A.1.

Table A.1: Data Description, Sources, and Coverage

Variable	Description	Source	Sample
Business sector variables			
Y_prod	Business output, expenditure side	Fernald TFP data	1949Q1-2022Q4
NETEEXP	Net energy exports (coal, oil, natural gas)	EIA/own calculations	1949Q1-2022Q4
K	Capital input	Fernald TFP data	1949Q1-2022Q4
L	Labor input (business sector hours $\times$ labor composition/quality)	Fernald TFP data	1949Q1-2022Q4
UTIL	Utilization of capital and labor	Fernald TFP data	1949Q1-2022Q4
L_share	Labor share (computed as 1 - capital share)	Fernald TFP data	1949Q1-2022Q4
ECONS_bus	Energy consumption (coal, oil, natural gas) in business sector	EIA/own calculations	1949Q1-2022Q4
E_share	Energy expenditure share (business energy consumption over business output)	EIA/own calculations	1949Q1-2022Q4
Ae	Energy-saving technical change	Own calculations	1949Q1-2022Q4
A	Capital/labor-saving technical change	Own calculations	1949Q1-2022Q4
Macro variables			
GDPC1	U.S. Real Gross Domestic Product	FRED	1949Q1-2022Q4
GDPDEF	Gross Domestic Product: Implicit Price Deflator	FRED	1949Q1-2022Q4
ECONS	Aggregate energy consumption (coal, oil, natural gas)	EIA/own calculations	1949Q1-2022Q4
Pe	Energy price (weighted average of coal, oil, natural gas price)	EIA/own calculations	1949Q1-2022Q4
TFP_util	Utilization-adjusted TFP	Fernald TFP data	1949Q1-2022Q4
EPATENT_share	Share of energy-saving patents	Google Patents Public Data /own calculations	1949Q1-2022Q4

The transformed series used in our structural VAR model are depicted in Figure A.1.




Figure A.1: The series in our baseline VAR

B. Additional Results

B.1. Estimating the energy complementarity

In this appendix, we assess the sensitivity of our estimates for the elasticity of substitution between energy and other inputs. We can see that our estimates are robust to using different energy price shocks as the instrumental variable, as well as to the inclusion of controls.

Table B.1: Using alternative energy price shock instruments

	Baseline	Känzig supply	BH supply	BH demand
Elasticity of substitution (ϵ)	0.026 (0.023)	0.012 (0.031)	0.045 (0.028)	0.020 (0.025)
First-stage F statistic	15.63	13.5	10.93	8.20
Observations	181	181	181	181

Notes: This table shows the estimates for the elasticity of substitution between energy and the other inputs estimated from equation (5) using different instruments for energy price shocks. The first column is our baseline, where we instrument the energy share using the oil supply shocks from [Baumeister and Hamilton \(2019\)](#) and [Känzig \(2021\)](#). In the second column, we just use the oil supply news shocks from [Känzig \(2021\)](#) as an instrument and in the third column, we just use the oil supply shock from [Baumeister and Hamilton \(2019\)](#). In the fourth column, we use the oil inventory and consumption demand shocks from [Baumeister and Hamilton \(2019\)](#) instead. Robust standard errors are displayed in parentheses, together with the robust first-stage F-statistic below.

Table B.2: Sensitivity with respect to controls included

	No lags	2 lags	4 lags	6 lags
Elasticity of substitution (ϵ)	0.006 (0.029)	0.016 (0.026)	0.026 (0.023)	0.026 (0.022)
First-stage F statistic	12.41	13.63	15.63	19.21
Observations	185	183	181	179

Notes: This table shows the estimates for the elasticity of substitution between energy and the other inputs estimated from equation (5) using different sets of additional controls. As in our baseline, we instrument the energy share using the oil supply shocks from [Baumeister and Hamilton \(2019\)](#) and [Känzig \(2021\)](#). In the first column, we omit any controls than a constant. In the second to fourth columns, we control for 2, 4, and 6 lags of the dependent variable and the energy share. Robust standard errors are displayed in parentheses, together with the robust first-stage F-statistic below.

B.2. Role of renewable energy




Figure B.1: U.S. business sector output and energy use by source

Notes: This figure shows the evolution of U.S. business output together with energy consumption from all sources, including renewables, and fossil energy consumption measured as coal, oil, and natural gas consumption. All three series are expressed as an index normalized to 100 in 1970Q1.




Figure B.2: U.S. business sector output and energy use by source

Notes: This figure shows the share of business sector energy consumption coming from renewable energy sources.

B.3. Variance decomposition




Figure B.3: Variance decomposition

Notes: This figure shows the fraction of forecast error variance in energy prices and the GDP deflator that can be accounted for by input-saving technology shocks and energy price shocks, at horizons ranging from one quarter to ten years.




Figure B.4: Historical decomposition of energy consumption

Notes: This figure shows the historical contribution of input-saving technology shocks and energy price shocks to energy consumption over our sample of interest. The solid green line is the historical contribution of the respective shock and the black, dashed line are the actual, annualized changes in energy consumption.

B.4. Interpreting the drivers of energy-saving technology




Figure B.5: Energy-saving technical change and green patenting

Notes: This figure shows the evolution of our measure of energy-saving technical change (in logs) together with the share of filed patent applications in energy-saving technologies (Y02E, Y02T, Y04S, and the relevant parts of Y02P).




Figure B.6: The structural transformation

Notes: This figure shows the evolution of our measure of energy-saving technical change (in logs) together with the U.S. service sector share.

B.5. Identifying innovations to service share




Figure B.7: The responses to a service-share shock

Notes: This figure shows the impulse responses to a one standard deviation service share shock, identified using the max-share approach. The point estimates are depicted as solid lines, and the 68 and 90 percent confidence bands are the dark and light shaded areas.

B.6. Sensitivity analysis

In this section, we include the corresponding graphs of the sensitivity analysis.




Figure B.8: Robustness with respect to the elasticity of substitution between energy and non-energy inputs




Figure B.9: Robustness with respect to the order of the orthogonalization





Figure B.10: The role of orthogonalizing the shocks




Figure B.11: Robustness with respect to max-share horizon




Figure B.12: Robustness with respect to sample period




Figure B.13: Robustness with respect to sample period





Figure B.14: Informational sufficiency




Figure B.15: Estimation method




Figure B.16: Responses for overall economy versus business sector




Figure B.17: Robustness with respect to lag order

B.7. Relation to TFP news literature

How do our identified TFP and capital/labor-saving technology shocks relate to the TFP shocks identified in the literature? Since we build on the max-share approach that is commonly used in the literature (Barsky and Sims, 2011; Kurmann and Otrok, 2013; Kurmann and Sims, 2021), the question arises why our responses look somewhat different from a TFP news shock. In particular, we find more of an immediate response of technology, suggesting that the capital-labor saving technology shock we recover with the Kurmann and Sims (2021) identification strategy is a combination of a surprise and a news shock. While this is to be expected given the findings of Kilian, Plante, and Richter (2023), it is still useful to understand the source of the discrepancies with Kurmann and Sims (2021).

To study where this difference comes from, we start by reproducing the results in Kurmann and Sims (2021). They estimate a VAR in utility-adjusted TFP, consumption, inflation and hours on a sample from 1960-2007. The impulse responses to the identified technology shock are shown in Figure B.18.




Figure B.18: The responses to a Kurmann and Sims (2021) TFP shock

The purple responses basically reproduce the results in Kurmann and Sims (2021).

Interestingly, however, when we estimate the same model over our longer sample and identify the technology shock using the same approach, the resulting shock is more similar to what we find in our baseline model: TFP displays a more immediate response that builds up less over time. One possible reason for this is that in this longer sample, there may have been some surprise TFP shock, which had very persistent effects, and thus get also picked up by the medium-run restrictions implemented using the max-share approach – resulting in a shock that captures both TFP news and surprise shocks.

Recall, that this does not pose a challenge for our results, since our goal is not to distinguish between technology news and surprise shock but rather disentangle capital/labor-saving technology from energy-saving technology shocks.

C. Max-Share Identification

Consider the reduced-form VAR introduced in Section 4.2 of the paper. Let k be the number of variables in the VAR, let Φ_t denote the $k \times k$ matrix of coefficients on reduced form errors in period t and Σ_u be the $k \times k$ variance-covariance matrix of the reduced form errors. Let $\tilde{\mathbf{S}}$ be the Cholesky decomposition of Σ_u . Let h_i be the horizon at which we are maximizing the sum of the forecast-error variance explained of that variable. Our max-share identification approach is then implemented as follows. As we identify three shocks, we follow this procedure for $i = 0, 1, 2$.

Step 1. Find an orthonormal matrix $\mathbf{Z}$ satisfying orthogonality constraints. If $i = 0$, let $\mathbf{Z} = \mathbf{I}_k$. If $i = 1$, then $\mathbf{Z}$ needs to be a $k \times (k - 1)$ orthonormal matrix and satisfy $\mathbf{Z}'\mathbf{Q}_1 = 0$. If $i = 2$, then $\mathbf{Z}$ needs to be a $k \times (k - 2)$ orthonormal matrix and satisfy $\mathbf{Z}'\mathbf{Q}_1 = 0$ and $\mathbf{Z}'\mathbf{Q}_2 = 0$.

- (a) Let $\hat{\mathbf{Z}} = \mathbf{I}_k[:, i :]$ (e.g. if $i = 2$, then $\hat{\mathbf{Z}}$ is the last $k - 2$ columns of the $k \times k$ identity matrix).
- (b) We then concatenate any vectors ($\mathbf{Q}_x$) we've already estimated to $\hat{\mathbf{Z}}$ and apply a QR decomposition to the resulting matrix to obtain orthonormal matrix $\mathbf{J}$. Let $\mathbf{Z}$ be the last $k - i$ columns of $\mathbf{J}$ such that $\mathbf{Z}$ is an $k \times k - i$ matrix satisfying the orthogonality constraints.

Step 2. Let $\mathbf{E}_{ii}$ be a matrix filled with zeros and a 1 in position i, i .

$$\text{Define } \mathbf{M} = \sum_{j=0}^{h_i} \tilde{\mathbf{S}}' \Phi_j' \mathbf{E}_{ii} \Phi_j \tilde{\mathbf{S}}.$$

Step 3. Let γ be the eigenvector corresponding to the largest eigenvalue of $\mathbf{Z}'\mathbf{M}\mathbf{Z}$, then $\mathbf{Q}_{i+1} = \mathbf{Z}\gamma$.

Then, let $\mathbf{Q}$ be the matrix with columns $\mathbf{Q}_1, \mathbf{Q}_2, \mathbf{Q}_3$ (if applicable) and $\mathbf{S} = \tilde{\mathbf{S}}\mathbf{Q}$.

In case we estimate the structural impact vector of the energy price shock using an external instrument, we proceed as follows. First, obtain an estimate of the structural impact vector $\mathbf{s}$ using the approach explained in [Stock and Watson \(2018\)](#). Given an estimate of $\mathbf{s}$, we have $\mathbf{Q}_1 = \tilde{\mathbf{S}}^{-1}\mathbf{s}$ and we can follow the procedure outlined above for $i = 1, 2$.

References Appendix

- Barsky, Robert B. and Eric R. Sims** (2011). “News shocks and business cycles”. *Journal of Monetary Economics* 58.3, pp. 273–289.
- Basu, Susanto, John Fernald, Jonas Fisher, and Miles Kimball** (2013). “Sector-specific technical change”. *Manuscript, Federal Reserve Bank of San Francisco*.
- Basu, Susanto, John G. Fernald, and Miles S. Kimball** (2006). “Are technology improvements contractionary?” *American Economic Review* 96.5, pp. 1418–1448.
- Baumeister, Christiane and James D. Hamilton** (2019). “Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks”. *American Economic Review* 109.5, pp. 1873–1910.
- Fernald, John G.** (2014). “A quarterly, utilization-adjusted series on total factor productivity”. Federal Reserve Bank of San Francisco.
- Hassler, John, Per Krusell, and Conny Olovsson** (2021). “Directed technical change as a response to natural resource scarcity”. *Journal of Political Economy* 129.11, pp. 3039–3072.
- Känzig, Diego R.** (2021). “The macroeconomic effects of oil supply news: Evidence from OPEC announcements”. *American Economic Review* 111.4, pp. 1092–1125.
- Kilian, Lutz, Michael D. Plante, and Alexander W. Richter** (Apr. 2023). *Estimating macroeconomic news and surprise shocks*. Working Paper 2304. Revised November 2023. Federal Reserve Bank of Dallas.
- Kurmann, André and Christopher Otrok** (2013). “News shocks and the slope of the term structure of interest rates”. *American Economic Review* 103.6, pp. 2612–2632.
- Kurmann, André and Eric Sims** (2021). “Revisions in utilization-adjusted TFP and robust identification of news shocks”. *Review of Economics and Statistics* 103.2, pp. 216–235.
- Quilis, Enrique M.** (2024). *Temporal Disaggregation*. <https://www.mathworks.com/matlabcentral/fileexchange/69800-temporal-disaggregation>. MATLAB Central File Exchange. Retrieved January 2, 2024.
- Stock, James H. and Mark W. Watson** (2018). “Identification and estimation of dynamic causal effects in macroeconomics using external instruments”. *The Economic Journal* 128.610, pp. 917–948.