Energy and Production in US Manufacturing

Eva Lyubich Will Rafey Joseph S. Shapiro Reed Walker*

September 10, 2026

Abstract

Real output per unit of energy input in US manufacturing has more than doubled over the last half century, a pattern often attributed to successful energy efficiency policy. We assess whether this trend reflects targeted improvements in energy use or broader technical progress. To do so, we develop a model of factor-biased technical change and apply it to five decades of establishment-level data from the U.S. Census Bureau. Our analysis indicates that innovations in energy-augmenting productivity account for about one-third of establishment-level improvements in energy efficiency, with aggregate energy-augmenting productivity growth driven primarily by reallocation across existing firms. An implication is that the majority of historic improvements in industrial energy efficiency reflect broader technological progress, not uniquely innovative energy use.

JEL Classifications: D24, L11, L60, O4, O13, Q43

Keywords: Energy efficiency, production functions, productivity decompositions, factor-augmenting technology, bias of technical change, US manufacturing

*Lyubich: Anthropic; Rafey: Department of Economics, University of California, Los Angeles and NBER; Shapiro and Walker: University of California, Berkeley and NBER. Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product. This research began prior to when Lyubich joined the US Census Bureau, and was performed at a Federal Statistical Research Data Center under FSRDC Project Number 2211 (Disclosure approval CBDRB-FY25-CES005-019). We thank Max Olander and Yangkeun Yun for helpful research assistance.

1 Introduction

U.S. manufacturing has undergone a remarkable energy transformation over the past half century. Real output per unit energy consumed, a standard measure of energy efficiency, more than doubled between 1998 and 2022 (EIA, 2025). This improvement in energy efficiency has occurred alongside sustained growth in manufacturing output, suggesting that American industry has learned to produce more with less energy. Governments and policy organizations routinely cite these trends as evidence that improvements in energy efficiency can simultaneously enhance industrial competitiveness and environmental outcomes (e.g., CEA, 2017, pp. 449–474, Belzer *et al.*, 2020, EIA, 2021, and IEA, 2022). With industrial production generating more than one third of global greenhouse gas emissions (Worrell *et al.*, 2009), understanding the sources of these energy efficiency gains also matters for climate and energy policy discussions worldwide.

Our paper aims to better understand this progress, starting from the observation that many different economic changes could generate the same aggregate improvement in energy efficiency, each with distinct implications for productivity measurement, technological change, and policy. First, firms may invent or adopt existing technologies aimed to reduce energy consumption and make energy inputs more productive, such as more efficient boilers or combined heat and power systems, representing energy-augmenting technical change. Second, improvements may reflect general, Hicks (1932)-neutral productivity growth that raises output from all inputs proportionally, with energy intensity declining as a mechanical consequence of general efficiency gains. Third, advances such as automation and computerization may represent labor-augmenting technical change that primarily targets labor productivity but reduces energy use through input complementarities. Fourth, changes in input prices can shift factor demand in ways that could affect measures of output per unit energy. Finally, reallocation of production within and across industries may affect aggregate improvements.

These distinctions matter for policy and measurement, yet existing empirical evidence cannot distinguish between them. Standard productivity measurement combines all forms of technical change into a single total factor productivity (TFP) residual (Solow, 1957; Syverson, 2011). When energy intensity falls, conventional methods attribute the improvement to general productivity growth without identifying whether energy itself became more productive. This measurement approach can obscure critical information, leading to mismeasurement of sources of productivity growth. On the

other hand, if observed energy efficiency improvements primarily reflect Hicks-neutral productivity growth rather than energy-specific innovations, then energy-focused policies may be less effective than believed, and engineering-based predictions of future efficiency potential may be overly optimistic.

The inability to measure energy-augmenting technical change separately creates what we call the mirage of energy efficiency: aggregate statistics show dramatic improvements in energy efficiency, which casual observation may attribute to energy efficiency policy, but we cannot determine whether these reflect genuine energy-specific innovations or other dimensions of productivity growth that end up reducing energy use. We extend recent advances in production function estimation (Doraszelski and Jaumandreu, 2018; Demirer, forthcoming; Rubens *et al.*, 2025) to separately identify energy-augmenting productivity, labor-augmenting productivity, and Hicks-neutral productivity at the plant level. The key insight is that conditional on output, other inputs, relative prices, and total factor productivity, a plant's energy demand reveals its energy-augmenting productivity. If two plants produce the same output using similar capital, labor, and materials, but one uses substantially less energy despite facing the same energy and other input prices, that plant must have higher energy-augmenting productivity. Tracking how plants' energy use evolves over time relative to these other variables lets us infer changes in energy-specific technical change.

To implement this approach, we construct a novel dataset spanning five decades of confidential U.S. Census Bureau microdata. We link administrative records from the Census of Manufacturers, the Annual Survey of Manufacturers, and the Manufacturing Energy Consumption Surveys to create plant-level measures of physical energy use (kilowatt hours of electricity, cubic feet of natural gas, etc.) and fuel-specific prices. The resulting dataset, containing detailed energy information for hundreds of thousands of plant-years, is the most comprehensive data on manufacturing energy use in the U.S. that we know.

We find that energy-augmenting technical change accounts for only about one third of plant-level improvements in energy efficiency. This component represents true energy-specific innovation—improvements that make energy inputs specifically more productive. The remaining two-thirds of observed efficiency gains reflect other mechanisms. Hicks-neutral and labor-augmenting productivity growth together account for almost one third of average establishment-level improvements in energy efficiency. Thus, a meaningful portion of what appears as energy efficiency improvements is actually productivity growth unrelated to energy per se. Plants produce more output from

all inputs, and energy use per dollar of output falls as a mechanical consequence. This Hicks-neutral component shows up as declining energy intensity, or improving energy efficiency, in published statistics, but it does not represent energy-specific technical change. Similarly, as automation and computerization have increased labor productivity, the complementarity between labor and energy has caused energy use per unit output to decline even without energy-specific innovations. This channel has been largely overlooked in prior discussions of industrial energy efficiency. The final third of improvements results from price effects and substitution, i.e., establishments substituting away from energy as other inputs get relatively cheaper.

We then examine the mechanisms underlying aggregate changes in energy efficiency and the three dimensions of productivity with the unique feature of Census data that allows us to see entry and exit for each plant. Specifically, we decompose changes into within-plant technique innovations, across-plant but within-industry output reallocation, and shifts in industry composition. Our results indicate that the greatest aggregate energy-augmenting productivity gains arise from high-productivity plants increasing their market share over time, echoing findings in the broader reallocation literature (Baily *et al.*, 1992; Foster *et al.*, 2001, 2008; Autor *et al.*, 2020). Energy-augmenting productivity improvements were most pronounced in the period between 2007 and 2012, coinciding with the fracking boom, a period of significant innovation in energy extraction that led to much more abundant cheap natural gas. Energy productivity improvements were largest for the most energy-intensive industry in our sample, non-metallic mineral mining, consistent with strong incentives to reduce energy costs leading to energy-specific productivity innovations.

These findings revise our understanding of industrial energy efficiency improvements. Rather than a story of continuous innovation in energy-saving technologies, partly driven by energy efficiency policies, we find a more complex narrative where factor-neutral productivity growth, factor-specific technical change, and reallocation all play important roles. Relative magnitudes of each matter for policy: if most historical improvement comes from Hicks-neutral productivity, then policies specifically targeting energy efficiency, such as appliance standards, subsidized industrial energy audits, or information campaigns, may be less effective than policies promoting factor-neutral productivity growth or facilitating reallocation.

In addition to its specific empirical results, our work makes two more general contributions to the literature. First, in identifying and estimating plant-level productivity, we adapt recent methods that allow for multiple dimensions of unobserved and persis-

tent productivity to the case of energy. These methods extend the classic production function framework, which allows for heterogeneity across producers without imposing strong parametric assumptions on either the underlying differences across plants or the evolution of productivity within a plant over time (Olley and Pakes, 1996; Ackerberg *et al.*, 2015; De Loecker and Syverson, 2021), to multidimensional productivity, by adding auxiliary moments derived from factor demand (see, e.g., Klein, 1953; Christensen and Greene, 1976; León-Ledesma *et al.*, 2010; Doraszelski and Jaumandreu, 2018; Gandhi *et al.*, 2020; Rafey, 2023; Rubens *et al.*, 2025; Demirer, forthcoming) to the standard moments that enable identification of one-dimensional productivity models from producer-level panel data. In estimation, we apply these methods to energy and labor jointly, and show that they can be used for large administrative datasets without prohibitive computational difficulty through a sequential two-step algorithm that first recovers factor-augmenting productivities and estimates the elasticity of substitution between factors (using moments derived from the plant's factor demand equations and instruments for relative factor prices), then estimates the remaining production function parameters and Hicks-neutral productivity in a second step that is essentially the usual strategy used to estimate single-dimensional models of productivity.

In particular, our estimates of multidimensional plant-level productivity complement existing work exploring the role of energy efficiency and emissions intensity in the U.S. manufacturing sector, such as differences in energy efficiency across incumbent and entrant plants (Linn, 2008; Hawkins-Pierot and Wagner, 2021), electricity demand and prices (Davis *et al.*, 2013), and emissions across plants (Lyubich *et al.*, 2018; Ganapti *et al.*, 2020), which has focused on responsiveness of energy demand (often, energy divided by output) to energy prices, rather than its underlying determinants—such as the various components of productivity (i.e., the three dimensions discussed above); the evolution of productivity within firms over time; and the different roles of other input decisions, prices, and productivities, each of which can change energy efficiency, but with very different economic implications.[1]

Second, in showing how the evolution in productivity at the plant level aggregates across plants to determine aggregate energy intensity—and how these aggregate outcomes reflect entry, exit, and reallocation—we advance the study of firm dynamics

[1]Our work also complements analyses of energy and production outside of the United States, such as recent studies of EU firms in Butters *et al.* (2026) and Indian firms in McDonald *et al.* (2026), as well as the literature on endogenous innovation, such as the evidence of Newell *et al.* (1999), Popp (2002), and Aghion *et al.* (2016) on induced innovation in energy efficiency (air conditioners, patents, and cars, respectively). While we do not endogenize a plant's productivity advances, we assess the direction of energy-specific innovation relative to other sources of innovation.

and clean growth. Full entry-exit-reallocation decompositions, such as those of Hicks-neutral productivity (e.g., Olley and Pakes, 1996; Melitz and Polanec, 2015) and the aggregate labor share (e.g., Autor *et al.*, 2020), are uncommon because they require administrative data on the universe of plants that allows the researcher to observe entry and exit over time. In particular, our focus on aggregate energy intensity complements the longstanding interest in macro in the technological responses to energy price shocks (Atkeson and Kehoe, 1999; Hassler *et al.*, 2021), and more recently to climate policy (e.g., Golosov *et al.*, 2014; Acemoglu *et al.*, 2018; Casey, 2024). These macro studies motivate their analyses with models of heterogeneous firm dynamics, but rely primarily on aggregate time-series data to calibrate the quantitative exercises. We lack the ability to assess the general equilibrium channels present in such studies, but our work complements these efforts with aggregate results obtained from integrating over the empirical distributions of productivity across plants, in the spirit of Oberfield and Raval (2021); our results correspond to the empirical distribution of plants in the observed equilibrium, without imposing parametric assumptions about the underlying distribution of productivity across, or rate or direction of growth within, plants.

The remainder of the paper outlines the data used (Section 2), the empirical model we specify (Section 3), and how the two together deliver our main results (Section 4).

2 Data and descriptive evidence

2.1 Data

We use administrative, establishment-level data from the U.S. Census Bureau for our analysis. We construct a sample of U.S. manufacturing establishments from 1976 to 2020 using the Longitudinal Business Database (LBD), the Census of Manufacturers (CM) and the Annual Survey of Manufacturers (ASM).

The CM covers the universe of establishments in the manufacturing sector, conducted in years ending in 2 and 7. It measures establishment revenue, labor hours, payroll, capital expenditures, material expenditures, and energy expenditures. The ASM supplements the CM by surveying about 50,000 establishments annually (U.S. Census Bureau, 2024). To harmonize NAICS industry definitions and track plant entry and exit, we link our manufacturing plant panel to the LBD.

We impose several sample restrictions on the establishments in our data. First, we exclude plants with imputed, missing, or zero values for any of our main analysis vari-

ables. Second, we drop observations with labor, materials, or input cost shares below 1/10 of the first percentile or above 10x the 99th percentile of the within-industry-year distribution of outcomes. Finally, establishments sometimes switch industry (NAICS3) classifications between years. These could reflect real switches in changes to primary production output, or they could reflect changes in the quality of information reported to the Census. Either way, because we estimate production functions separately for each industry and want establishments that are representative of production processes in their assigned industries, we drop these (roughly 8%) of establishments from our estimating sample and retain only those assigned a single industry over all the years in which we observe them. Appendix B provides additional details on data cleaning and preparation.

Table 1: Sample Statistics

	Full Sample	CM	ASM	Energy Intense
Total Output ($1,000s)	29,070	27,760	29,400	47,150
Energy Share	0.02	0.02	0.02	0.04
Electricity Share of Energy	0.71	0.71	0.71	0.59
Materials Share	0.44	0.44	0.44	0.49
Labor Share	0.26	0.26	0.26	0.18
Obs. per Establishment	4.8	2.2	5.0	6.5
Entered	0.60	0.53	0.64	0.59
Exited	0.71	0.68	0.69	0.66
N obs. (1,000s)	2,103	781	1,323	454
N estabs. (1,000s)	435	353	263	70
N counties	3,100	3,100	3,100	2,800

Notes: This table presents summary statistics for the data for each of four samples. Column 1 presents results for the full sample, which includes establishments in CM and ASM years with non-missing and non-outlier analysis variables, as described in the text. Column 2 presents results for just establishments in CM years, and column 3 does the same for ASM years. Column 4 restricts the full sample to establishments in the five most energy-intensive industries. Establishment-by-year level statistics are weighted by ASM sample weights. Establishment level statistics are unweighted. *Source*: Authors' calculations from Census of Manufacturers (1977-2017) and Annual Survey of Manufacturers (1976-2020) data.

Table 1 presents statistics for our data. Column 1 consists of our full analysis sample. Columns 2 and 3 are observations from only CM years, and only ASM years, respectively. The sample moments come from establishment-by-year level observations weighted by ASM sample weights.[2] Establishment-level statistics are unweighted. Es-

[2]In CM years, all establishments are assigned an equal weight of one, and in each year we then rescale weights to sum to one by year.

tablishments surveyed in the ASM are slightly larger than those in the CM sample, but otherwise have similar input cost shares. Establishments in the five most energy-intensive industries (NAICS3) are larger than average. By definition, they have higher energy cost shares than average. A smaller share of their energy costs comes from electricity, and they have higher material costs and lower labor cost shares.

We deflate plant-level revenues and expenditures using industry-by-year (NAICS6) price indices from the NBER's manufacturing productivity database, except where plant-specific price deflators or plant-specific quantity measures are available. For example, we observe total hours of labor in a year and the total wage bill. We divide the latter by the former to come up with a measure of the average wage in a plant-year. The CM and ASM also contain plant-level physical electricity (mmBtu) inputs and expenditures, which we use to construct average electricity prices.[3]

The CM and ASM also record plant-level fuel expenditures, but they do not record physical quantities or distinguish individual fuels besides electricity. Hence, we merge fuel price microdata from the Manufacturing Energy Consumption Survey (MECS) and state-by-year data on industrial sector fuel prices from EIA's State Energy Data System (SEDS). MECS is conducted every 4 years, and collects detailed information on prices and quantities of plant-level energy consumption across several energy sources. We use fuel quantities (in mmBtu) for natural gas, coal, and oil to construct average fuel shares by industry and year. We then use the SEDS state-by-year industrial fuel pricesto construct a fuel price index that we use as a price deflator for fuel expenditures. Appendix Figure A.1 displays the trend in these energy prices over time.

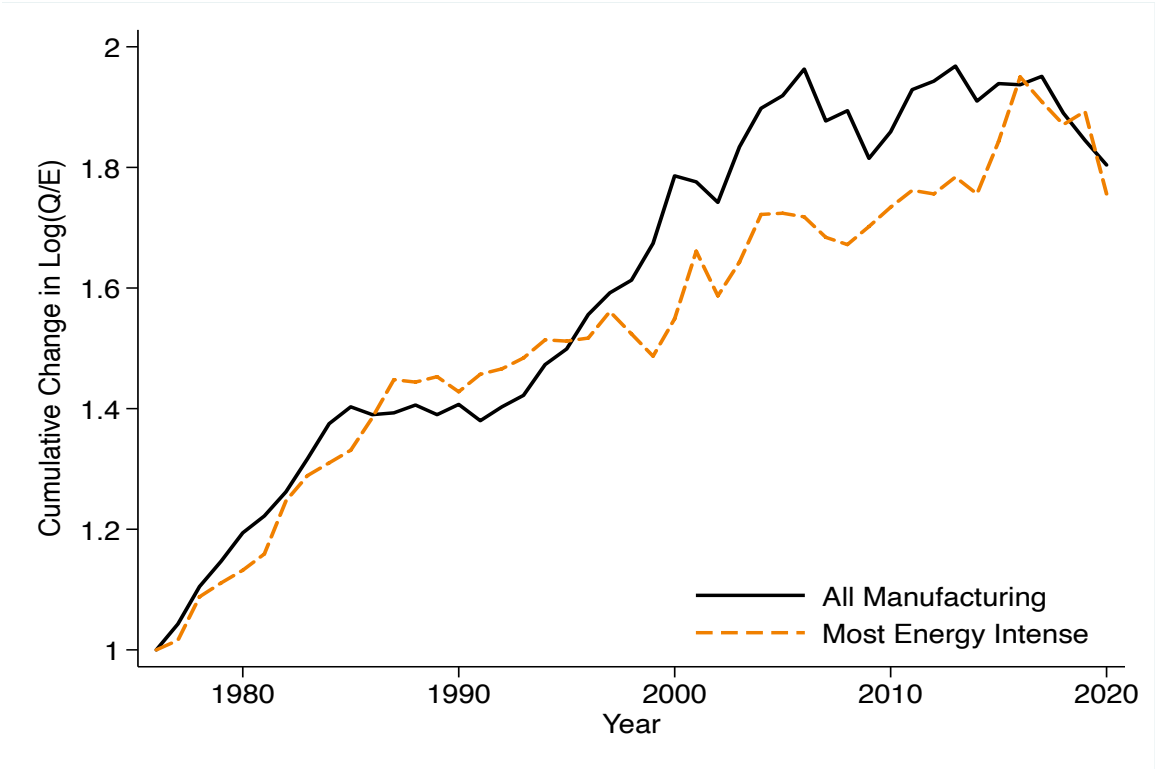
2.2 Descriptive Results

We measure energy efficiency for every establishment i as log output per unit energy, $\log(Q_i/E_i)$, where Q_i denotes physical output and E_i represents energy inputs in mmBtu. Figure 1 displays the cumulative change in aggregate energy efficiency over our sample period for all manufacturing and for the most energy-intensive industries. Aggregate energy efficiency nearly doubled over the past half-century, rising by 0.8 in logs. While growth in energy efficiency was somewhat slower among energy-intensive industries in the early 2000s, it converges to the level of the full sample by the end of

[3]We explore constructing plant-level measures of materials with the Census' Material Trailer; these are theoretically more attractive than industry-level materials price indices, and have been used successfully in narrower case studies of specific industries (e.g., Collard-Wexler and De Loecker (2015))—but we found the data quality unable to support an exercise aimed to capture the majority of the U.S. manufacturing sector.

the period, with an overall gain of 0.76 in logs relative to baseline. These efficiency improvements correspond to declining energy cost shares from 4.8% to 3.0% in energy-intensive industries and from 2.4% to 1.8% sector-wide, with the larger absolute decline in energy-intensive sectors reflecting their higher baseline energy use.

Figure 1: Energy Efficiency Over Time



Notes: This figure shows cumulative growth in aggregate log(Q/E) for the entire manufacturing sector (black line) and the five most energy-intensive 4-digit NAICS industries in the manufacturing sector. *Source*: Authors' calculations from U.S. Census of Manufacturers (1977-2017) and Annual Survey of Manufacturers (1976-2020) data.

3 An Empirical Model of Energy Productivity

To better understand the key drivers of these patterns, we now develop a flexible empirical model of plant production building on Doraszelski and Jaumandreu (2018) to separately identify energy-augmenting technical change from other sources of technological improvements and facilitate subsequent decompositions.

3.1 Production Technology

Plant i produces output Q_{it} in year t using capital K_{it}, energy E_{it}, labor L_{it}, and materials M_{it}:

$$\ln Q_{it} = \ln F_j(K_{it}, e^{\omega_{it}^E} E_{it}, e^{\omega_{it}^L} L_{it}, M_{it}) + \omega_{it}^H + \varepsilon_{it} \tag{1}$$

The function F_j represents components of production common across plants within industry j. Within each industry and over time, plants also have three dimensions of productivity, $\omega_{it} \equiv (\omega_{it}^H, \omega_{it}^L, \omega_{it}^E)$, corresponding to Hicks-neutral, labor-augmenting, and energy-augmenting productivities. Energy and labor augmenting productivity scale their respective input uses to deliver a measure of effective energy $(exp(\omega_{it}^E)E_{it})$ or effective labor $(exp(\omega_{it}^L)L_{it})$. Distinguishing between these productivities lets plants improve their efficiency of energy or labor use without necessarily improving their productivity in using all factors together. The distribution of productivity across plants within a period captures cross-sectional technical differences, while its evolution over time captures technological differences within a given plant over time. Measurement error ε_{it} also generates differences in observed output across plants, but—unlike productivity—is not observed by the plant.

For parsimony, empirical tractability, and consistency with some prior literature, we parameterize F_j using the simplest production function that allows for factor-augmenting technical change, the constant elasticity of substitution (CES) production function:

$$\ln Q_{it} = \frac{\nu_j \sigma_j}{\sigma_j - 1} \cdot \ln\left(\beta_{jK} K^{\frac{\sigma_j - 1}{\sigma_j}} + (e^{\omega_{it}^L} L)^{\frac{\sigma_j - 1}{\sigma_j}} + (e^{\omega_{it}^E} E)^{\frac{\sigma_j - 1}{\sigma_j}} + \beta_{jM} M^{\frac{\sigma_j - 1}{\sigma_j}}\right) + \omega_{it}^H + \varepsilon_{it} \tag{2}$$

Here ν_j denotes returns to scale, σ_j denotes the elasticity of substitution among factors, and β_{jK} and βjM are distributional parameters, common to plants in industry j.[4] Our approach, described below, does not rely on the assumption that the production function be CES. However, the function F_j must be flexible enough to allow us to identify factor-augmenting productivity; for example, in the limit where $\sigma_j \to 1$, the function F_j becomes Cobb-Douglas, which by assumption rules out the separate identification of factor-augmenting from Hicks-neutral productivity.

We assume that plants take as given energy, labor, and materials prices, P_{it}^E, P_{it}^L, and P_{it}^M. We also assume that energy, materials, and labor are static (variable) inputs

[4] Similar parameters for energy and labor are not identified from the fully nonparametric distributions of $\{\omega_{it}^E\}$ and $\{\omega_{it}^L\}$.

that plants can adjust in each period t. The assumption that producers adjust energy, labor, and at least one other input in response to changing prices is indispensable because, as the next section shows, we use these adjustments to identify differences in the factor-augmenting productivities across plants and over time.[5,6]

We leave output markets mostly unspecified. Output prices $P_j(Q_{it}; \vartheta_{it})$ are continuously differentiable functions of the establishment's own output Q_{it} and some unobservable shifter, ϑ_{it}. Full control function approaches require strong restrictions on conduct in output markets (Doraszelski and Jaumandreu, 2021 provide a discussion). We can accommodate greater flexibility because we instead use general dynamic panel methods. We do not need to impose the stringent requirements that we can express all aspects of output market structure as a function of various time-t observables, for reasons similar to those described in De Loecker *et al.* (2016).

3.2 Identification

Our approach to identify the model above with panel data on output, physical inputs, and input prices involves two sets of assumptions. First, given factor prices, we provide assumptions to recover factor-augmenting technical change from establishment-level physical inputs. Second, we provide assumptions to identify the full production function from data where we replace energy and labor inputs with effective inputs.

3.2.1 Identifying Factor-Augmenting Technical Change

For factor-augmenting technical change, our identification strategy uses observed factor prices and auxiliary assumptions on factor demand. Under the assumptions in the previous section, plant i's energy demand at t is continuous in its energy-augmenting

[5]Treating labor as flexible has been used extensively by prior literature (Levinsohn and Petrin, 2003; De Loecker and Warzynski, 2012; De Loecker *et al.*, 2016; Doraszelski and Jaumandreu, 2018; Demirer, forthcoming), though it has been recently questioned, particularly for annual data (Ackerberg *et al.*, 2015; Foster *et al.*, 2022). Our approach still applies if some inframarginal labor, energy, and/or materials inputs are predetermined or involve adjustment costs, for the same reason that in-house labor and in-house materials can involve costly adjustment as long as (observable) outsourced labor and imported materials are true variable inputs in the analysis of labor-augmenting technical change in Doraszelski and Jaumandreu (2018).

[6]Energy, materials, and labor are also static inputs of production, in the sense that their use at time t does not affect productivity in subsequent periods $t' > t$. This rules out forms of innovation that arise directly from learning-by-doing, for example Benkard (2000).

productivity and characterized by

$$\left[\frac{P_t'(Q_{it};\vartheta_{it})Q_{it}\mathbb{E}[e^{-\varepsilon_{it}}]}{P_t}+1\right]P_t(Q_{it};\vartheta_{it})e^{\omega_{it}^H+\omega_{it}^E}\frac{\partial F_j(K_{it},e^{\omega_{it}^E}E_{it},L_{it})}{\partial E_{it}}=P_{it}^E. \quad (3)$$

Dividing (3) by a similar first-order condition for materials eliminates markups, the output price, and total factor productivity, $\left[\frac{P_t'(Q_{it};\vartheta_{it})Q_{it}\mathbb{E}[e^{-\varepsilon_{it}}]}{P_t}+1\right]P_t(Q_{it};\vartheta_{it})e^{\omega_{it}^H}$. We can obtain energy-augmenting technical change via

$$\omega_{it}^E=p_{it}^E-p_{it}^M-\ln\frac{\partial F_j}{\partial E}+\ln\frac{\partial F_j}{\partial M} \quad (4)$$

Equation (4) depends only on the production function parameters for F_j, the observed current inputs, and the ratio of input prices, P^E/P^M. We emphasize that these objects identify factor-augmenting productivities without imposing restrictions on output market structure or Hicks-neutral productivity.

With a candidate vector of production function parameters θ_j that characterize $F_j(\cdot)$, it is straightforward to solve for some $\tilde{\omega}^E$ satisfying (4).

For CES, equation (4) simplifies to a straightforward expression for energy-augmenting productivity, ω_{it}^E, in terms of the physical energy-materials input ratio, $e_{it}-m_{it}$, relative factor prices, $p_{it}^M-p_{it}^E$, and the elasticity of substitution, σ_j, as

$$\omega_{it}^E=(1-\sigma_j)^{-1}\left[e_{it}-m_{it}+\sigma_j\ln\beta_{jM}-\sigma_j(p_{it}^M-p_{it}^E)\right] \quad (5)$$

We can recover an analogous equation for labor-augmenting productivity, the labor-materials input ratio, and relative factor prices. Note that, in (5), if contemporary factor prices are co-determined with factor productivity, then the first-order conditions cannot be used directly to consistently estimate σ_j. Consequently, only together with instruments for factor prices do these conditions on firm behavior serve as the basis to identify the first stage of our empirical strategy.

In short, our model identifies unobserved factor-augmenting productivity by attributing all residual differences in physical inputs across plants—i.e., those that cannot be explained by differences in plant-specific factor prices or sector-level factor substitution parameters—to factor-augmenting productivity.

3.2.2 Identifying Hicks-neutral technical change

After recovering a plant's effective energy input $\tilde{E}_{it} = e^{\omega_{it}^E} E_{it}$, we still face the standard identification issues that arise because producers choose inputs based on unobserved productivity that also determines output (Marschak and Andrews, 1944). Our solution relies on a combination of assumptions on how productivity and input prices evolve relative to the timing of input decisions, which allow us to identify productivity using instruments whose validity is implied by the model. Specifically, we employ a dynamic panel approach for similar reasons as discussed in (Ackerberg *et al.*, 2015, §.4.3.3). In particular, our identification assumptions resemble the baseline strategies of Ackerberg *et al.* (2015) and Doraszelski and Jaumandreu (2018), except that we do not use control functions. This means we cannot "purge" our productivity estimates of measurement error, nor can we allow productivity to evolve according to more complicated nonlinear functions of past productivities beyond AR(1). However, unlike these papers, we do not need to assume strict monotonicity in a proxy variable or address finite-sample bias in the nonparametric estimator for these control functions.

3.3 Estimation procedure

The primitives to estimate are the production technology parameters $\theta_j \equiv (\beta_{jM}, \beta_{jK}, \sigma_j, \nu_j)$ for each sub-sector j, productivities $\{\omega_{it}^E, \omega_{it}^L, \omega_{it}^H\}$ for each establishment i and time t, and Markov transitions ρ_j^k for productivity $k \in \{E, L, M\}$. The estimator involves minimizing a generalized method of moments (GMM) objective function that we obtain through the following loop for each sector j:

(a) Estimate a common $\hat{\sigma}_j$ across the energy-materials and labor-materials ratios using moments constructed as

$$\mathbb{E}[\xi_{it}^k(\sigma_j) Z_{i,t-1}^k | X_{it}] = 0 \quad \text{for } k \in \{E, L\},$$

where $\xi_{it}^k(\sigma_j) \equiv \omega_{it}^k(\sigma_j) - \rho_j^k \omega_{i,t-1}^k(\sigma_j)$, with ω_{it}^k for $k \in \{E, L\}$ constructed from (5) and $Z_{i,t-1}^k$ consisting of lagged electricity prices (for $k = E$), lagged wages (for $k = L$), and lagged materials prices (for $k \in \{L, H\}$), and X_{it} includes industry fixed effects and a time trend.

(b) Use $\hat{\sigma}_j$, observed inputs, and factor prices to solve (5) for the fixed points $(\tilde{\omega}^E, \tilde{\omega}^L)$, up to the constant $\ln \beta_{jM}$, for each i and t.

(c) Conditional on $\hat{\sigma}_j$, $\tilde{\omega}^E$, and $\tilde{\omega}^L$, estimate the remaining parameters θ_j of the production function and the evolution of productivity, with instruments for input

decisions, using moments

$$\mathbb{E}[\xi_{it}^H(\theta_j) Z_{i,t-1}^H | X_{it}, \tilde{\omega}_{it}^E, \tilde{\omega}_{it}^L] = 0,$$

for $\xi_{it}^H(\theta_j) \equiv \omega_{it}^H(\theta_j) - \rho_j^H \omega_{i,t-1}^H(\theta_j) + \varepsilon_{it} - \rho_j^H \varepsilon_{i,t-1}$, with $\omega_{it}^H + \varepsilon_{it}$ constructed from (2) and $Z_{i,t-1}^H$ consisting of lagged effective energy and labor inputs, lagged materials inputs, and current and lagged capital, and X_{it} includes industry fixed effects and a time trend.

Valid instruments for the endogenous variables—factor prices in step (a) and input decisions in step (c)—must be uncorrelated with productivity innovations, ξ_{it}^k for $k \in \{E, L\}$ in (a) and ξ_{it}^H in (c). We instrument for factor prices in (a) using lagged electricity prices for current energy prices, lagged wages for current wages, and lagged materials prices for current materials prices, under the assumption that the innovation in factor-augmenting productivity is uncorrelated with these historical factor prices. We use lagged *electricity* prices as shifters for a plant's current own energy price because those electricity rates may be less endogenous than the effective energy price that the plant faces, which also reflects, in part, how the plant chooses to use fuel versus electricity. In (c), we instrument for current inputs using lagged input decisions, under the assumption that a plant's historical (but not current) flexible factor decisions and capital are uncorrelated with that plant's productivity innovation. Such lags can be excluded but still relevant shifters of those decisions, through, e.g., inframarginal adjustment costs and autocorrelation in factor prices. Note that past productivities can influence current factor prices, and future factor prices can influence historical productivity; what we rule out is that dependence of factor prices on the within-period innovation in productivity.

One justification of our choice to use a standard set of internal instruments[7] is to take fullest advantage of the breadth and near-universal coverage of the administrative Census data. Ideally, we would construct instruments for input prices in step (a) and input demand in (c) using more specific sources of quasi-experimental variation, such as tax or other policy changes (Griliches and Mairesse, 1998). In our study, the major limitation to this conceptually more attractive approach is that such instruments must be narrowly-tailored to the specific market and period of study. Our decomposition analysis of aggregate economic changes, however, require that we recover productivi-

[7] "Internal" in the sense of their validity being implied by the model; see, e.g., Ackerberg *et al.* (2015) for a discussion, and Doraszelski and Jaumandreu (2018), Rubens *et al.* (2025), and Demirer (forthcoming) for some examples.

ties for the near universe of US manufacturing establishments over a half-century of economic development.

4 Results

Here we report aspects of our estimates of the production functions and productivities (Section 4.1), implications for aggregate energy efficiency (4.2), and the decomposition of aggregate changes into incumbent, entering, and exiting plants (4.3). Sections 4.4–4.5 study additional heterogeneity over time and observable plant characteristics.

4.1 Production Function Estimates

Appendix Table A.1 reports elasticities of substitution, distributional parameters, and output elasticities for each industry.

Before examining the distribution of productivity, we note two aspects of these production function estimates. First, unlike a naive OLS estimator that does not correct for endogeneity between factor price ratios and unobserved factor-augmenting technical change (i.e., where $\hat{\sigma}_j > 1$ for nearly all j), our dynamic panel estimator for $\hat{\sigma}_j$ delivers elasticities that are largely below 1.[8] Elasticities below 1 align with prior estimates for US manufacturing (Raval, 2019; Oberfield and Raval, 2021), and the endogeneity problem with the naive estimator for the elasticity of substitution also aligns with prior literature (see, e.g., Doraszelski and Jaumandreu (2018)).

Second, because factor-augmenting productivity is not separately identified from Hicks-neutral productivity when the production function is additively separable in log inputs, we cannot implement our productivity analysis for the four industries for which $\hat{\sigma}_j > 1$. We note that none of these industries comprise nontrivial shares of energy use. In addition, for one industry (primary metals), the estimated $\hat{\sigma}_j < 0$, implying that we cannot reject non-substitutability, i.e., a Leontief production function. For this industry, we estimate a Leontief production function directly.[9] Consequently, for the remainder of our analysis, we restrict our sample to the 15 industries with elasticities

[8] There are five industries, where we cannot reject that the production technology is Cobb-Douglas: Wood products, printing and related support activities, furniture and related products, machinery, and miscellaneous.

[9] Primary metals manufacturing is NAICS 331. For the steel industry (NAICS 3311), we can also construct plant-specific materials price indices, and when we use these price indices and estimate separate production functions for minimills versus vertically integrated steel plants, we obtain well-behaved elasticity estimates. See e.g. Collard-Wexler and De Loecker, 2015.

of substitution estimated to lie between 0.05 and 0.95 and the 1 industry we estimate to be Leontief.

4.2 Productivity Types and Energy Efficiency

Given production function estimates, we can recover productivities for each establishment in each year in our sample, then decompose improvements in energy efficiency. Our goal is to understand the extent to which the change in establishment-level energy efficiency, $\log(Q_i/E_i)$, reflects energy-specific technological progress versus other sources. This is similar to the analysis undertaken in Doraszelski and Jaumandreu (2018) to understand the output effects of labor-augmenting productivity; here, we care not about output effects per se, but output per energy input (or energy efficiency) effects. The total change in $\log(Q/E)$ can be expressed as:

$$\begin{aligned}\Delta \log(Q_i/E_i) \approx \frac{\partial \log(Q_i/E_i)}{\partial \omega_i^H}\Delta\omega_i^H + \frac{\partial \log(Q_i/E_i)}{\partial \omega_i^E}\Delta\omega_i^E + \frac{\partial \log(Q_i/E_i)}{\partial \omega_i^L}\Delta\omega_i^L \\ + \sum_{x\in\{K_i,L_i,E_i,M_i\}} \frac{\partial \log(Q_i/E_i)}{\partial p_x}\Delta p_x\end{aligned} \tag{6}$$

This structural decomposition attributes changes in energy efficiency to its ultimate technological and economic drivers.

The complete analytical characterization of these partial derivatives is complicated by the fact that the firm's optimal choice of inputs, including energy (E), is a function of all technology parameters and prices. For example, a Hicks-neutral shock ω_H not only increases output directly but also induces a scale effect, leading the firm to adjust its usage of all inputs. To make the problem tractable, we adopt a simplified, partial-equilibrium approach. We analyze the effect of each technology shock by isolating its most direct impacts: the direct effect on output (holding physical inputs constant) and the direct effect on input composition (via substitution). This approach abstracts from more complex equilibrium scale effects, but delivers a clear and interpretable first-order decomposition.

4.2.1 Hicks-neutral productivity

Differentiating the production function (1) with respect to ω_H gives $\partial log Q_i/\partial\omega_i^H = 1$. In other words, holding inputs fixed, increasing Hicks-neutral productivity increases log output one-for-one. While this shock also induces scale effects that change energy

use, our simplified framework isolates the direct output effect. We therefore define the contribution of Hicks-neutral productivity to energy efficiency improvements as

$$\text{HN Effect} = \Delta\omega_i^H \tag{7}$$

By raising the efficiency of all inputs, including energy, Hicks-neutral productivity increases observed energy efficiency, Q_i/E_i.

4.2.2 Energy-augmenting productivity

Energy-augmenting productivity ω_i^E increases energy efficiency Q_i/E_i through two distinct channels. First, it increases effective energy $e^{\omega_{it}^E E_{it}}$ through a pure technological improvement that boosts output even when physical energy use remains constant. We derive this direct output effect from the CES production function (2). by differentiating log output with respect to ω^E while holding physical inputs constant:

$$\begin{aligned}
\left.\frac{\partial \log(Q_i)}{\partial \omega_i^E}\right|_{E,M,L,K\text{const.}} &= -\frac{\nu\sigma}{1-\sigma}\frac{1}{X_i}\frac{\partial X_i}{\partial \omega_i^E} \\
&= -\frac{\nu\sigma}{1-\sigma}\frac{1}{X_i}\left[-\frac{1-\sigma}{\sigma}(\exp(\omega_i^E)E_i)^{-\frac{1-\sigma}{\sigma}}\right] \\
&= \frac{\nu}{X_i}(\exp(\omega_i^E)E_i)^{-\frac{1-\sigma}{\sigma}} \\
&= \varepsilon_i^E
\end{aligned}$$

where X represents the CES aggregator and ε_i^E denotes the output elasticity with respect to effective energy.[10]

Second, by making energy more productive, the shock induces input substitution. The relative demand of energy and materials in equation 5 lets us derive how physical energy use responds to energy-augmenting technical change. Differentiating $e_i = m_i - \sigma_j \ln(\beta^M) + \sigma(p_M - p_E) + (1-\sigma_j)\omega^E$ with respect to ω^E while holding materials and prices constant gives $\partial e_i/\partial \omega_{Ei} = \sigma_j - 1$. Thus, the substitution effect reduces a plant's physical energy inputs when $\sigma_j < 1$. Combining both channels, the total contribution

[10]Empirically, we estimate an industry-wide average output elasticity ε_j^E and plug this into the computation of the output effect of energy-augmenting productivity.

of energy-augmenting productivity to energy efficiency is

$$\begin{aligned}\text{EA Effect} &= \underbrace{\varepsilon_j^E \cdot \Delta\omega_i^E}_{\text{output effect}} - \underbrace{(\sigma_j - 1) \cdot \Delta\omega_i^E}_{\text{substitution effect}} \\ &= (\varepsilon_j^E + 1 - \sigma_j)\Delta\omega_i^E.\end{aligned}$$

4.2.3 Labor-augmenting productivity

Labor productivity also affects energy efficiency through both output and substitution channels. Analogous to energy, the direct output effect equals $\varepsilon_j^L \cdot \Delta\omega_i^L$, where ε_j^L is the industry-average output elasticity of effective labor. The relative demand between energy and labor shows that the substitution effect is $(1-\sigma_j)\Delta\omega_i^L$. The total effect on energy efficiency is therefore:

$$\begin{aligned}\text{LA Effect} &= \underbrace{\varepsilon_j^L \cdot \Delta\omega_i^L}_{\text{output effect}} - \underbrace{(1 - \sigma_j) \cdot \Delta\omega_i^L}_{\text{substitution effect}} \\ &= (\varepsilon_j^L - 1 + \sigma_j)\Delta\omega_i^L\end{aligned}$$

When inputs are complements ($\sigma_j < 1$), labor-augmenting technical change has an ambiguous effect on energy efficiency Q_i/E_i: it raises output per unit input, but also increases relative energy use.

4.2.4 Price effects and residual

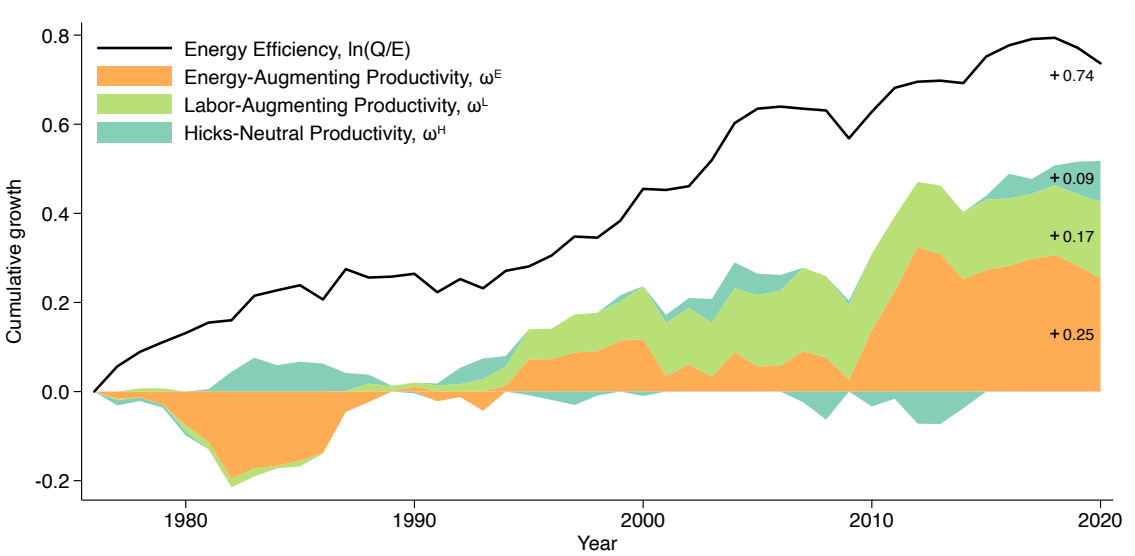
The remaining variation in energy efficiency captures price-induced substitution and approximation errors in our first-order decomposition:

$$\text{Price Effect} = \Delta \log(Q/E) - (\text{HN Effect}) - (\text{EA Effect}) - (\text{LA Effect})$$

We apply this decomposition by calculating the year-by-year contribution of each productivity component to growth in energy efficiency. For establishments with more than one year between observations, we compute average annual growth between t_{-1} and t for each ω_i as $(\omega_{i,t} - \omega_{i,t_{-1}})/(t - t_{-1})$. We aggregate across establishments by averaging weights across t and $t-1$. We construct weights by multiplying ASM sampling weights by their output share to obtain representative industry/sector-level aggregates.[11]

[11] Our results are somewhat sensitive to our time choice of weights. As Doraszelski and Jaumandreu

Figure 2: Components of Growth in Establishment Energy Efficiency, 1977-2022



Notes: This figure shows observed cumulative growth in aggregate energy efficiency (black line) and the contributions to cumulative growth from energy-augmenting productivity (orange area), labor-augmenting productivity (green area), and Hicks-neutral productivity (blue area). Annual growth rates are estimated at the establishment level and aggregated using ASM × q weights averaged over t and t-1. *Source*: Authors' calculations using data from U.S. Census of Manufacturers (1977-2017), Annual Survey of Manufacturers (1976-2020), and Manufacturing Energy Consumption Survey (1985-2018) data.

Figure 2 shows the contribution of each dimension of productivity to energy efficiency growth, stacked on one another. We estimate that growth in energy-augmenting productivity explains about one-third of overall energy efficiency improvements (25/74). Enegy-augmenting productivity has some interesting cyclical patterns. The two clearest episodes of energy-augmenting innovation occurred during the 1980s during the oil glut that followed the Energy Crisis of the 1970s, and from 2009 to 2012, following the Great Recession and the onset of the fracking boom (Bartik *et al.*, 2019). While these episodes differ, both represent periods of energy market transformation that appear to have catalyzed innovation in energy-using technologies (e.g., Linn, 2008; Aghion *et al.*, 2016).

Labor-augmenting productivity and Hicks-neutral productivity together account

(2018) highlight, we expect changes in output shares to correlate with changes in productivities, with more productive firms growing and less productive firms shrinking. This correlation leads estimates using lagged weights to more heavily weight shrinking firms, while those using weights in period t to more heavily weight growing firms. Our approach balances these competing biases, and is also consistent with how we weight decomposition terms across industries in the next section of the paper, following Autor *et al.* (2020).

for another third of growth in energy efficiency ((17+9)/74). The rest reflects input price effects and higher-order interactions. Notably, labor-augmenting productivity has increased energy efficiency more than Hicks-neutral productivity has. One interpretation is that establishment-level innovation to economize on labor, such as automation, process controls, and organization improvements, simultaneously reduce energy requirements.

Finally, price effects are non-negligible. When energy prices are relatively high, establishments respond by substituting to other inputs. While persistently high prices can lead to price-induced productivity innovation, these market forces can also lead to energy efficiency improvements even without longer-term changes in production technologies.

Appendix Figure A.2 shows growth in energy efficiency and the energy efficiency effect of each of the three dimensions of productivity, aggregated to industry (NAICS3) rather than the entire manufacturing sector. The figures reveal substantial heterogeneity across industries, underscoring the distinct channels through which efficiency improvements occur. Despite overall gains in aggregate energy efficiency, five of fifteen industries experienced declines in energy-augmenting productivity. Energy efficiency grew fastest in Computer and Electronic Product Manufacturing; interestingly, this industry has the largest effect of labor-augmenting productivity on energy efficiency. This demonstrates how labor-augmenting technical change can meaningfully drive energy efficiency through production complementarities. The heterogeneity across industries, and contrast between industries where energy-augmenting productivity drives efficiency gains and where labor-augmenting or Hicks-neutral productivity do instead, shows that pathways for improving energy efficiency differ by industry.

4.3 Drivers of Aggregate Energy Efficiency

To clarify how aggregate energy efficiency and productivity gains emerge from establishment-level technological improvements, reallocation across plants, and changes in the sectoral composition of industries, we implement the decomposition framework of Melitz and Polanec (2015) (which extends Olley and Pakes (1996)) at the industry level. Following Autor *et al.* (2020), we augment our decomposition to compare within-industry and between-industry components.

For each outcome ϕ of interest (energy efficiency, energy-augmenting productivity, labor-augmenting productivity, and Hicks-neutral productivity) the aggregate value of

ϕ in industry j is

$$\Phi_j \equiv \sum_{i \in j} s_i \phi_i = \bar{\phi}_j + \sum_{i \in j} (s_i - \bar{s}_j)(\phi_i - \bar{\phi}_j) \tag{8}$$

$$= \bar{\phi}_j + \text{cov}_j(s_i, \phi_i) \tag{9}$$

where s_i is establishment i's share of output in industry j, and $\bar{s}_j$ and $\bar{\phi}_j$ are unweighted industry-level mean values. We can therefore write the Olley and Pakes (1996) decomposition of within-industry *change* in aggregate productivity $\Delta\Phi_j$ between period $t-1$ and t for a balanced panel of establishments:

$$\Delta\Phi_j = \Delta\bar{\phi}_j + \Delta\text{cov}_j(s_i, \phi_i). \tag{10}$$

Melitz and Polanec (2015) extend this decomposition to account for changes from plant entry and exit:

$$\Delta\Phi_j = \underbrace{\Delta\bar{\Phi}_j^S}_{\text{technique}} + \underbrace{\Delta\text{cov}_j^S(s_i, \phi_i)}_{\text{survivor reallocation}} + \underbrace{s_{jt}^E(\Phi_{jt}^E - \Phi_{jt}^S)}_{\text{entry}} + \underbrace{s_{j,t-1}^X(\Phi_{j,t-1}^S - \Phi_{j,t-1}^X)}_{\text{exit}} \tag{11}$$

where S, E, and X index the establishments that survive from period $t-1$ to period t, enter in period t, or exit after period $t-1$, respectively. s_{jt}^E represents the output share of entrants in period t, $s_{j,t-1}^X$ represents the output share of exiters in period $t-1$, and Φ^G represents the share-weighted mean value of ϕ_i in group G.

Equation (11) decomposes within-industry aggregate changes into four components. The technique component $\Delta\bar{\Phi}_j^S$ reflects the change in unweighted mean productivity among survivors, which provides a direct analog to the establishment-level gains that the previous section measured.[12] Positive values imply that establishments are innovating over time, e.g. through adopting new productivity-enhancing technologies. The survivor reallocation component Δcov_j^j measures the changes in covariance between establishment market shares and productivity. When $\Delta\text{cov}_j^j > 0$, more productive survivors gain market share. Entry and exit contribute through selection. If $\Phi_{jt}^E > \Phi_{jt}^S$, entry raises aggregate productivity because entrants have greater productivity than incumbents; if $\Phi_{j,t-1}^X < \Phi_{j,t-1}^S$, exit improves productivity because less productive es-

[12]Note however that results are aggregated with different weights. In the previous analysis, establishment changes between t and $t-1$ were weighted by average establishment market shares over those two periods. In this analysis, establishments are weighted equally within industry, and then industry-specific within components are aggregated across industries using average market share.

tablishments exit. The sum of survivor reallocation, entry, and exit equals the total within-industry reallocation effect, i.e., the aggregate change in productivity due to only changing market shares of differentially productive establishments, holding establishment-level improvements in productivity constant.

Finally, to measure manufacturing-wide structural transformation, we decompose manufacturing-wide changes in average productivity, $\Delta\Phi = \Delta \sum_j s_j \Phi_j$, into between-industry and within-industry components:

$$\Delta\Phi = \underbrace{\sum_j \tilde{s}_j \Delta\Phi_j}_{\text{Within-industry}} + \underbrace{\sum_j \Delta s_j \tilde{\Phi}_j}_{\text{Industry Composition}} \quad , \tag{12}$$

Here $\tilde{s_j}$ denotes the average output share in industry j over times $t-1$ and t and $\tilde{\Phi}_j$ denotes the average of the output-share weighted mean productivity in industry j over times $t-1$ and t. The within-industry term aggregates the industry-specific changes from equation (11), weighted by average industry shares. The between-industry term captures shifts in output composition across industries with different productivity levels.

Table 2 presents decomposition results for energy efficiency and the three dimensions of productivity, scaled by the relevant elasticities to express their contribution to energy efficiency. We estimate decompositions over five-year intervals between CM years and sum across periods to obtain cumulative changes over our sample period. We classify establishments as entrants or exiters based on the LBD, so that temporary data gaps do not affect classification.[13] Columns (1)-(4) present the four within-industry Melitz-Polanec decomposition components from equation (11) – technique, survivor reallocation, entry, and exit – weighted across industries by industry time-average market share as in (12). Columns (5) and (6) present the industry shift-share decomposition terms from (12)—the first is the weighted mean within-industry change (i.e., the sum of columns (1)-(4)), and the second is the composition term, indicating whether industries with higher Φ are gaining or losing market share. The final column shows the aggregate change in the manufacturing sector overall.[14]

[13]This approach yields smaller entry and exit contributions than a purely sample-based definition would generate. The five-year interval structure similarly attenuates entry and exit effects relative to a single long-difference decomposition, as fewer establishments enter or exit within shorter windows.

[14]While equation (6) expresses changes in establishment level energy efficiency as a sum of changes in each of the three productivity parameters (rescaled by relevant elasticities) and a component driven by price effects, the components in the last three rows of Table 2 do not sum to the first row. The reason is that the Melitz-Polanec decomposition operates on *levels*. For example, the exit term

Table 2: Aggregate Decomposition

	Within-Industry Decomposition				Industry Shift-Share		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Φ	Survivor Technique	Survivor Reallocation	Entry	Exit	Within Industry	Composition	Total Change
q-e	0.60	0.35	0.09	-0.07	0.97	-0.08	0.89
ω^E	0.07	0.36	0.11	-0.08	0.46	0.01	0.46
ω^L	-0.00	0.02	0.10	-0.10	0.01	-0.18	-0.17
ω^H	0.43	-0.33	-0.07	0.06	0.09	-0.05	0.04

Notes: This table presents results from the within-industry Melitz-Polanec and cross-industry shift-share decompositions shown in equations (11) and (12). We present decompositions for energy efficiency and each estimated productivity component. Decompositions are estimated over adjacent 5-year CM periods, and cumulated over our sample period. For each outcome Φ, column (1) shows the change in Φ due to within-establishment changes in technique, column (2) captures the change in Φ that results from reallocation of market share across surviving establishments, and columns (3) and (4) capture the change in Φ that results from reallocation of market share as a result of entry and exit, respectively, of more or less productive establishments. Column (5) presents the sum of these four components, which together make up the within-industry change in Φ. Column (6) presents how Φ changes as a result of changing industry compositions across the manufacturing sector. Column (7) presents the total change in Φ, equal to (5) + (6). *Source*: Authors' calculations using Census of Manufacturers (1977-2017) and Manufacturing Energy Consumption Survey (1984-2018) data.

Each row of Table 2 presents results for one of our outcomes. Our estimates suggest that aggregate energy efficiency improvements primarily reflect technique rather than reallocation. Row 1 shows that technique accounts for 67% (0.60/0.89) of total energy efficiency improvements; reallocation contributes 42% (37/89). Entry and exit approximately offset, with both entrants and exiters exhibiting higher average efficiency than incumbents. Industry composition works against energy efficiency gains: the negative term in column (6) indicates that manufacturing output has shifted towards industries with lower average energy efficiency.

We find the opposite pattern for energy-augmenting productivity. Row 2 of Table 2

$s^X_{j,t-1}(\Phi^S_{j,t-1} - \Phi^X_{j,t-1})$ compares survivors' and exiters' aggregate productivity levels in period $t-1$. In contrast, equation (6) relates *changes* in energy efficiency to changes in the three productivity parameters through a first-order approximation. Because the production function mapping from productivities to energy efficiency in (2) is nonlinear, the composition components across the four rows of Table 2 do not have a simple additive relationship (except the within survivor component). Each (other) component of the energy efficiency decomposition reflects a complex weighted combination of the corresponding components from the three productivity decompositions, where weights depend on the terms in the production function.

shows that reallocation (the sum of survivor reallocation, entry, and exit) accounts for nearly 85% (.39/.46) of the aggregate gains while technique contributes only 15% (.07/.46). Thus aggregate energy-augmenting productivity improvements result primarily from establishments with higher energy-augmenting productivity gaining market share, rather than survivors' adoption of new energy-augmenting technologies over time. Again, energy-augmenting productivity of entering and exiting firms exceeds that of incumbents. Shifts in industry composition hardly affect energy-augmenting productivity. Thus, while row 1 shows that growing industries are energy intensive, row 2 shows that growing industries do not have systematically different energy-augmenting productivity.

It is worth discussing the point from both the first and second findings which indicate that exiting establishments have higher energy efficiency and higher energy-augmenting productivity than incumbents. In many models of firm dynamics with one dimension of productivity, incumbents (survivors) are more productive on average, because the least productive firms are more likely to exit (e.g., Hopenhayn, 1992, Ericson and Pakes, 1995). The divergence between that common assumption and our results likely reflects the vintage of capital. New entrants adopt the most energy-efficient equipment and have elevated turnover, while older incumbents face higher costs of upgrading legacy capital. Indeed, younger plants in a given year have higher energy efficiency and energy productivity than older plants.

Row 3 of Table 2 shows that labor-augmenting productivity has decreased energy efficiency, with industry composition accounting for nearly the entire decline. Figure 2 showed that at the establishment level, labor-augmenting productivity has increased energy efficiency. How can we reconcile these contrasting findings? Appendix Figure A.2 highlights heterogeneity across industries in how labor-augmenting productivity affects energy efficiency. The negative composition term in this table could reflect increased size of industries where labor productivity, combined with factor substitutability, increases relative energy use (i.e., negative energy efficiency effects). Another possible mechanism is if the growing energy-intensive industries have also economized on labor. Again, for entrants and exiters more than incumbents, labor-augmenting productivity increases energy efficiency.

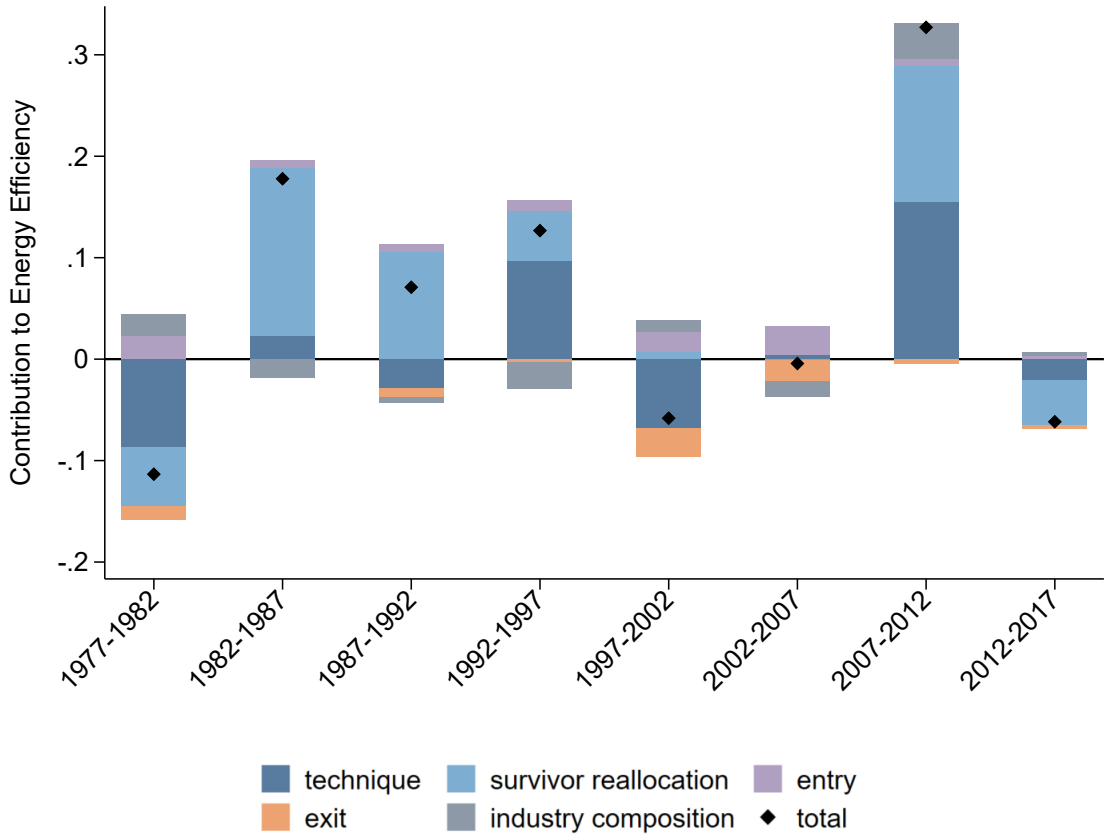
Row 4 of Table 2 indicates that Hicks-neutral productivity improvements hardly affect overall energy efficiency. Despite substantial technique effects, reallocation within industries has shifted market share away from establishments with greater Hicks Neutral productivity (perhaps in favor of establishments with high factor-augmenting pro-

ductivity).

4.4 Dynamics of Energy Productivity

We now focus our attention on the dynamics of energy-augmenting productivity. Figure 3 decomposes energy-augmenting productivity into its constituent components for each five-year interval between CM years. Stacked bars show the contribution of each component, and black diamonds show total changes.

Figure 3: Decomposition of ω_E



Notes: This figure presents results from the within-industry Melitz-Polanec and cross-industry shift-share decompositions shown in equations (11) and (12). The figure shows results for energy productivity ω_E in each 5-year CM interval. Components are stacked on one another, and the total change is expressed by the black diamonds. *Source*: Authors' calculations using Census of Manufacturers (1977-2017) and Manufacturing Energy Consumption Survey (1984-2018) data.

Several patterns emerge. Energy-augmenting productivity actually declines over the initial period 1977–1982, driven primarily by negative technique and survivor reallocation terms. This period coincides with the oil crisis of 1979, which more than

doubled oil prices that year and triggered the 1980–1982 recession. Perhaps this price shock and macroeconomic turbulence disrupted energy-specific innovation. At the same time, this period saw the passage of the first major CAA amendments, which made regulation binding for many polluting establishments (Greenstone, 2002), perhaps diverting resources away from energy-augmenting productivity investments and towards pollution abatement investments.

Figure 3 shows that following this initial decline, energy-augmenting productivity grew steadily from 1982 through 1997. Reallocation among survivors primarily drives this, i.e., establishments with better energy-specific technologies gained market share. Aggregate energy-augmenting productivity then stagnates between 1997–2007, the ten-year period with the largest negative contribution from exit ($-.05$ in total).

Strikingly, Figure 3 shows a 2007-2012 surge in energy-augmenting productivity. This episode spans both the Great Recession and the fracking boom. In contrast to the 1977–1982 period, which also experienced a severe recession but had energy scarcity and high prices following oil shocks, the 2007-2012 period had energy abundance and rapidly falling natural gas prices. This suggests that falling natural gas prices could have spurred aggregate improvements in energy-augmenting productivity. The 2007-2012 period has a large technique component (.16), suggesting that establishments responded to cheap natural gas by making investments in energy-augmenting technologies, for instance by converting production processes from oil to natural gas, or upgrading to more efficient natural gas-powered equipment. Simultaneously, the large survivor reallocation component (0.13) indicates that establishments with higher energy-augmenting productivity, which were perhaps best positioned to exploit low natural gas prices, gained market share. We see a compositional effect across industries (0.03), wherein industries with higher energy-augmenting productivity gained market share.

4.5 Heterogeneity by Regulatory and Market Exposure

The preceding results indicate that the level and determinants of productivity growth vary by year. The five decades we study also had changing environmental regulation (Clean Air Act), trade competition (China shock), macroeconomic conditions (1980 and 2008 recessions), and energy price shocks (a rise after 1979 and fall after 2008). To explore how these changes relate to energy efficiency and its components, we construct sub-samples of establishments with varying exposure to these shocks. Our goal is to capture a range of conditions that we expect to impact our three dimensions of

productivity differentially.

4.5.1 Clean Air Act

Our first sub-sample examines the set of establishments exposed to Clean Air Act (CAA) regulation, a landmark U.S. environmental policy.[15] Following Walker (2013), we create an indicator for whether an establishment is regulated under the Clean Air Act based on the interaction between its county and industry. Specifically, an establishment appears in our "CAA" sample if for any of the six criteria pollutants, it is in a 6-digit NAICS industry that is responsible for at least .05% of national emissions for that pollutant, and if it is in a county ever designated as in non-attainment of a specific National Ambient Air Quality Standard (NAAQS) for that pollutant.

4.5.2 China Shock

Our second sub-sample, "China," includes establishments exposed to higher competition from trade due to the "China Shock," i.e., the integration of China into the global trading system, a transformation shock to U.S. manufacturing with significant labor market impacts (Autor *et al.*, 2013). Following Acemoglu *et al.* (2016) we define a measure of import penetration based on industry and commuting zone of the establishment. Our sample includes establishments in the top quartile of this measure.

4.5.3 High Natural Gas

Our third sub-sample, "High NG," consists of establishments in one of the five industries with the highest pre-2005 average shares of natural gas usage. These industries were poised to capture the largest gains from the fracking boom that transformed U.S. energy markets beginning in the late-2000s. We construct these using the MECS data.[16]

[15] The Clean Air Act was first passed in 1963 and amended in 1970 to establish a list of six criteria pollutants (sulfur dioxide, carbon monoxide, nitrogen dioxide, lead, particulate matter, and ozone) and create National Ambient Air Quality Standards (NAAQS) which set maximum allowable ambient pollution levels for each criteria pollutant. Amendments in 1977 created non-attainment designations, determined at the county level annually. The EPA sets industry-specific emission standards in non-attainment counties, requiring new facilities to adopt "lowest achievable emissions rate" technologies and existing facilities to adopt "reasonably available control technologies." Additional amendments in 1990 and in the 2000s tightened standards for criteria pollutants. See (Currie and Walker, 2019) for details.

[16] This is also the set of five industries with the highest share of cogeneration capacity.

4.5.4 Highest Energy Intensity

Our final sub-sample, "top EI," comprises establishments in the most energy intense NAICS3 industry in our sample: nonmetallic mineral products mining, which includes cement and concrete product manufacturing. This industry faces particularly strong incentives to improve energy efficiency, as energy costs represent a substantial share of total costs. It also happens to be the only one of the top five most energy-intensive industries that is not also in the "High NG" sample.

Table 3: Aggregate Decomposition of ω_E – Heterogeneity

	Within-Industry Decomposition				Industry Shift-Share			
Sample	(1) Survivor Technique	(2) Survivor Realloc.	(3) Entry	(4) Exit	(5) Within Industry	(6) Industry Comp.	(7) Total Change	(8) Δ Sample Share
All	0.07	0.36	0.11	-0.08	0.46	0.01	0.46	.
CAA	0.29	0.26	0.07	-0.02	0.61	-0.33	0.28	-0.03
China	-0.01	0.44	0.07	-0.05	0.45	0.05	0.50	0.06
High NG	0.06	0.46	0.03	-0.10	0.45	0.03	0.48	0.06
Top EI	0.39	0.12	0.70	-0.30	0.91	.	0.91	-0.01

Notes: This table presents results from the within-industry Melitz-Polanec and cross-industry shift-share decompositions shown in equations (11) and (12). We decompose energy productivity ω_E for each sub-sample described above. Decompositions are estimated over adjacent 5-year CM periods, and cumulated over our sample period. For each sub-sample, column (1) shows the change in ω_E due to within-establishment changes in technique, column (2) captures the change in ω_E that results from reallocation of market share across surviving establishments, and columns (3) and (4) capture the change in ω_E that results from reallocation of market share as a result of entry and exit, respectively, of more or less productive establishments. Column (5) presents the sum of these four components, which together make up the within-industry change in ω_E, and column (6) presents how ω_E changes as a result of changing industry compositions across the manufacturing sector. Column (7) presents the total change in ω_E (i.e. (5) + (6)). Column (8) presents the change in sector market share of the relevant sub-sample. *Source*: Authors' calculations using Census of Manufacturers (1977-2017) and Manufacturing Energy Consumption Survey (1984-2018) data.

Table 3 reports the resulting decompositions of energy-augmenting productivity ω^E over the period 1977-2028, analogous to row 2 of Table 2. Row 2 shows that among the four sub-samples, establishments facing Clean Air Act regulation had the slowest growth in energy-augmenting productivity. A large negative industry composition term depresses the total change, despite above-average within-industry improvements, and in particular, high innovations in technique. In other words, CAA-regulated industries decreased energy intensity through technique improvements within plants, but increased it through compositional shifts to more energy-intensive industries. Column (8) shows a 3.3 percentage point decline in the sample share of regulated establish-

ments (or a 14% decline from a base of 24%), consistent with environmental regulation decreasing competitiveness.

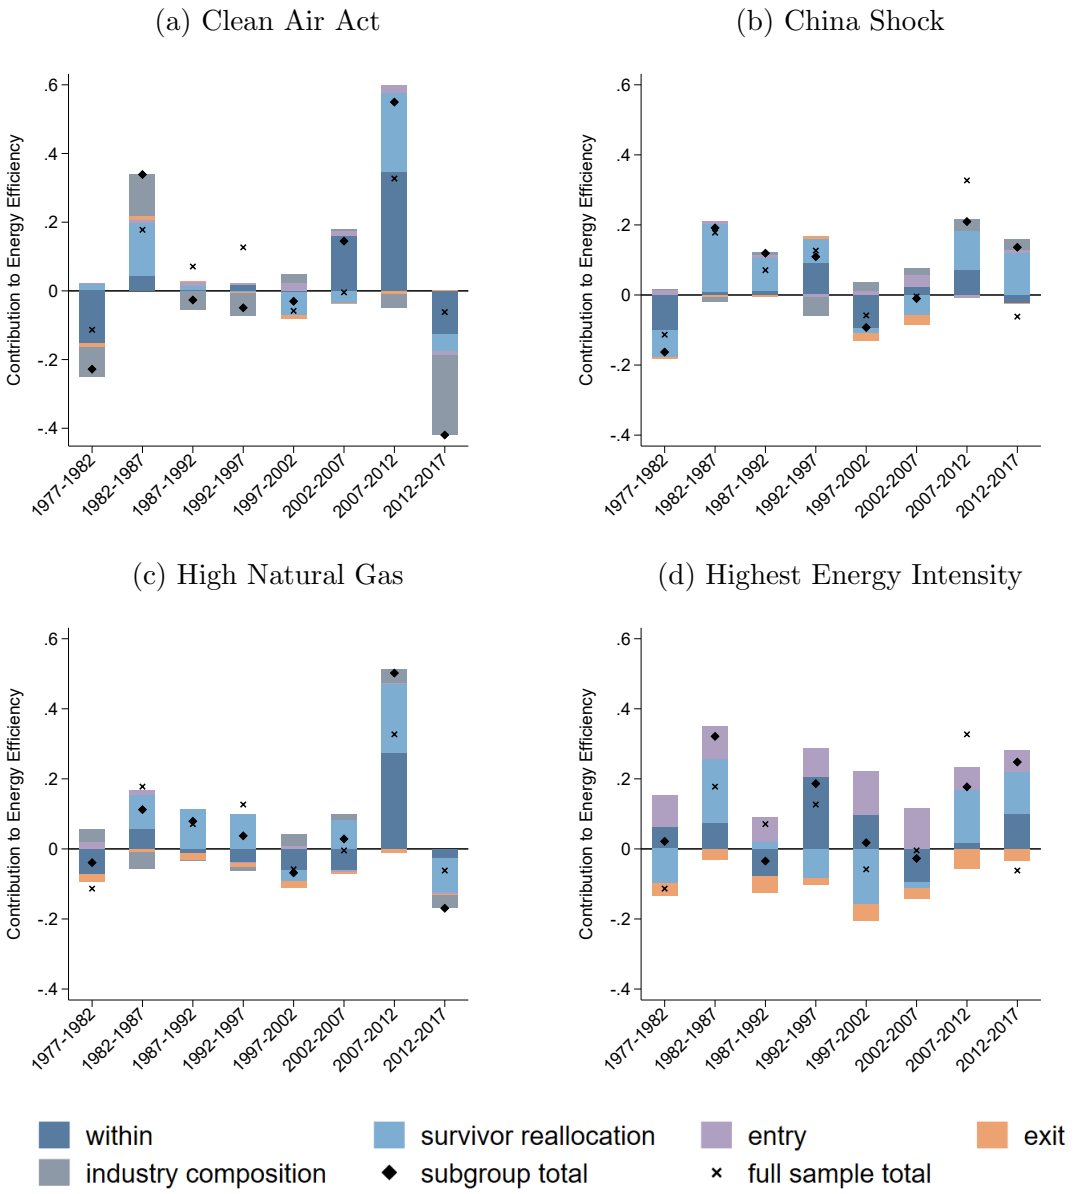
In contrast, row 3 of Table 3 shows that establishments exposed to the China shock demonstrate close to average energy-augmenting productivity growth. Survivor reallocation primarily drives this improvement, and survivor technique improvements contribute essentially nothing. China-shock-exposed establishments, which initially represented 21 percent of manufacturing output, expanded their share of the sample by 5.6 percentage points (27%) despite elevated import competition, suggesting their productivity gains may have increased domestic competitiveness.

Row 4 of Table 3 shows that industries highly dependent on natural gas had energy-augmenting productivity growth similar to manufacturing overall, including similar total change and within-survivor reallocation dominating technique. The group expanded its output share by 12% relative to baseline (six percentage points from an initial share of .49). Thus, manufacturing energy efficiency improved despite growing output from industries dependent on natural gas that initially have high energy intensity.

Finally, row 5 of Table 3 shows that the most energy-intensive industry—nonmetallic mineral products—has dramatically different dynamics from other samples. Aggregate growth in energy-augmenting productivity was twice that of the manufacturing sector overall. Entry is especially important, contributing to 77% (0.70/0.91) of the total change, compared with 24% (0.11/0.46) in manufacturing overall. Exit also has large and negative effects on energy efficiency, indicating that exiting establishments in this industry are substantially more efficient than survivors. Paired with the large entry effects, this suggests a pattern of high failure rates among young (though more efficient) entrants. This could reflect capital vintages within the industry, where new establishments can incorporate the latest energy efficient technologies without facing adjustment costs, while older establishments with sunk capital investments face adjustment frictions (Atkeson and Kehoe, 1999). Entry is not the only driver of innovation in the nonmetallic mineral products industry, however. Technique makes a substantial contribution, indicating that continuing establishments change production processes to make their energy use more productive.

To clarify transitional dynamics in these sub-samples, Figure 4 reports decomposition results for each five-year interval. All sub-samples except the most energy-intensive experienced the highest growth in energy-augmenting productivity ω^E 2007–2012, with especially large spikes for natural gas industries and CAA-regulated establishments. For all sub-samples, this jump reflects within-industry innovation, with large contribu-

Figure 4: Decomposition of ω_E by Subgroup



Notes: This figure presents results from the within-industry Melitz-Polanec and cross-industry shift-share decompositions shown in equations (11) and (12). The figure shows results for energy productivity ω^E for each 5-year CM interval. Each subfigure corresponds to a different sub-sample. Panel (a) shows results for the CAA sample, Panel (b) shows results for the China sample, Panel (c) shows results for the High NG sample. Panel (d) shows results for the highest energy intensity industry. Components are stacked. Black diamonds show the total change. For comparison, black x markers show totals for the entire sector. *Source*: Authors' calculations using Census of Manufacturers (1977-2017) and Manufacturing Energy Consumption Survey (1984-2018) data.

tions from both technique and survivor reallocation for CAA, China, and High NG, and almost entirely survivor reallocation in high EI. One interpretation of these results is

that, while many plants improved energy productivity during the fracking boom, fewer such innovations occurred in the most capital- and energy-intensive industry, where innovative establishments that could take advantage of the new natural gas abundance gained greater market share.

Figure 4 also shows that the CAA subsample had a negative contribution from industry composition in many of the intervals with important CAA amendments (1977, 1990, 2008, 2012; the exception is the period containing the 2005 amendments). This pattern is consistent with a story of environmental regulation shifting activity toward less energy-innovative industries within the regulated subsample, even as regulated establishments themselves may increase innovation efforts.

Finally, the China sub-sample exhibits a decline in ω^E coinciding with the China Shock, mostly driven by decreases in technique and a small but noticeable negative contribution from exit. This is consistent with the hypothesis that establishments facing increased import competition and automation pressures from China focused technology adoption or innovation on labor-augmenting or Hicks-neutral productivity to compete on dimensions other than energy efficiency during this period. Indeed, as we show in Appendix Tables B.2-B.5, while this subsample saw above-average energy efficiency improvements in 1997-2002, improvements in energy efficiency entirely reflected Hicks-neutral productivity, and any labor augmenting productivity innovations decreased aggregate energy efficiency.

5 Conclusion

Our results show that the dramatic improvements in U.S. manufacturing energy efficiency over the past half-century have reflected not only changes in how establishments use energy-specific technologies, but a complex interplay between innovations in energy-augmenting productivity, labor-augmenting productivity, Hicks-neutral productivity, price effects, and reallocation of market share between establishments. We find that only about one-third of establishment-level gains in energy efficiency reflect improvements in energy-augmenting productivity. On aggregate, a large share of economy-wide improvement in energy productivity reflects shifts in market share from less- to more-productive establishments. Our descriptive analysis of aggregate changes reveals substantial heterogeneity across industries and time periods, with an especially striking period of growth in energy productivity occurring during 2007-2012, a period of dramatic innovation in natural gas extraction. This period of energy-specific

innovation had positive energy-augmenting productivity innovations across the manufacturing sector, but especially among establishments in industries with high baseline usage of natural gas, which were poised to take most advantage of the technological shock. Our findings suggest that interventions narrowly focused on promoting energy-augmenting innovations may leave significant portions of the opportunity to improve energy efficiency on the table.

References

Acemoglu, D., Akcigit, U., Alp, H., Bloom, N. and Kerr, W. (2018). Innovation, reallocation and growth. *American Economic Review*, **108** (11), 3450–3491.

—, Autor, D., Dorn, D., Hanson, G. and Price, B. (2016). Import competition and the great us employment sag of the 2000s. *Journal of Labor Economics*, **34** (S1), S141–S198.

Ackerberg, D., Caves, K. and Frazer, G. (2015). Identification properties of recent production function estimators. *Econometrica*, **83** (6), 2411–2451.

Aghion, P., Dechezleprêtre, A., Hémous, D., Martin, R. and Reenen, J. V. (2016). Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry. *Journal of Political Economy*, **124** (1), 1–51.

Atkeson, A. and Kehoe, P. (1999). Models of energy use: Putty-putty versus putty-clay. *American Economic Review*, **89** (4), 1028–1043.

Autor, D., Dorn, D. and Hanson, G. (2013). The China syndrome: Local labor market effects of import competition in the United States. *American Economic Review*, **103** (6), 2121–2168.

—, —, Katz, L., Patterson, C. and Reenan, J. (2020). The fall of the labor share and the rise of superstar firms. *Quarterly Journal of Economics*, **135** (2), 645–709.

Baily, M. N., Hulten, C. and Campbell, D. (1992). Productivity dynamics in manufacturing plants. *Brookings Papers on Economic Activity: Microeconomics*, **1992**, 187–267.

Bartik, A., Currie, J., Greenstone, M. and Knittel, C. (2019). The local economic and welfare consequences of hydraulic fracturing. *American Economic Journal: Applied Economics*, **11** (4), 105–155.

Belzer, D., Cort, K. and Ganguli, S. (2020). *A comprehensive system of energy intensity indicators for the U.S.: Methods, data and key trends*. Tech. rep., OSTI, https://www.osti.gov/biblio/1633858.

Benkard, C. L. (2000). Learning and forgetting: The dynamics of aircraft production. *American Economic Review*, **90** (4), 1034–1054.

Butters, R. A., Dorsey, J. and Rudik, I. (2026). Pollution pricing in equilibrium: Production, reallocation, and aggregate impacts. *NBER Working Paper*, (35682).

Casey, G. (2024). Energy efficiency and directed technical change: Implications for climate change mitigation. *Review of Economic Studies*, **91** (1), 192–228.

CEA (2017). *Economic Report of the President*. Tech. rep., The White House Council of Economic Advisors.

Christensen, L. and Greene, W. (1976). Economies of scale in U.S. electric power generation. *Journal of Political Economy*, **84** (4), 655–676.

Collard-Wexler, A. and De Loecker, J. (2015). Reallocation and technology: Evidence from the U.S. steel industry. *American Economic Review*, **105** (1), 131–171.

Currie, J. and Walker, R. (2019). What do economists have to say about the clean air act 50 years after the establishment of the environmental protection agency? *Journal of Economic Perspectives*, **33** (4), 3–26.

Davis, S., Grim, C., Haltiwanger, J. and Streitwieser, M. (2013). Electricity unit value prices and purchase quantities: U.S. manufacturing plants, 1963–2000. *The Review of Economics and Statistics*, **95** (4), 1150–1165.

De Loecker, J., Goldberg, P., Khandelwal, A. and Pavcnik, N. (2016). Prices, markups, and trade reform. *Econometrica*, **84** (2), 445–510.

— and Syverson, C. (2021). An industrial organization perspective on productivity. In *Handbook of Industrial Organization*, vol. 4, *3*, North-Holland.

— and Warzynski, F. (2012). Markups and firm-level export status. *American Economic Review*, **102** (6), 2437–2471.

Demirer, M. (forthcoming). Production function estimation with factor-augmenting technology: An application to markups, econometrica.

Doraszelski, U. and Jaumandreu, J. (2018). Measuring the bias of technological change. *Journal of Political Economy*, **126** (3), 1027–1084.

— and — (2021). Reexamining the De Loecker and Warzynski (2012) method for estimating markups, working Paper.

EIA (2021). *U.S. energy intensity has dropped by half since 1983, varying greatly by state*. Tech. rep., U.S. Energy Information Administration.

EIA (2025). *2022 Manufacturing Energy Consumption Survey: Highlights from Data Release 1*. Tech. rep.

Ericson, R. and Pakes, A. (1995). Markov-perfect industry dynamics: A framework for empirical work. *Review of Economic Studies*, **62** (1), 53–82.

Foster, L., Haltiwanger, J. and Krizan, C. J. (2001). Aggregate productivity growth: Lessons from microeconomic evidence. In C. R. Hulten, E. R. Dean and M. J. Harper (eds.), *New Developments in Productivity Analysis*, Chicago: University of Chicago Press, pp. 303–372.

—, — and Syverson, C. (2008). Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *American Economic Review*, **98** (1), 394–425.

—, — and Tuttle, C. (2022). Rising markups or changing technology?, nBER Working Paper 30491.

Ganapti, S., Shapiro, J. and Walker, R. (2020). Energy cost pass-through in US manufacturing: Estimates and implications for carbon taxes. *American Economic Journal: Applied Economics*, **12** (2), 303–342.

Gandhi, A., Navarro, S. and Rivers, D. (2020). On the identification of gross output production functions. *Journal of Political Economy*, **128** (8).

Golosov, M., Hassler, J., Krusell, P. and Tsyvinski, A. (2014). Optimal taxes on

fossil fuel in general equilibrium. *Econometrica*, **82** (1), 41–88.

Greenstone, M. (2002). The impacts of environmental regulations on industrial activity: Evidence from the 1970 and 1977 Clean Air Act Amendments and the census of manufactures. *Journal of Political Economy*, **110** (6), 1175–1219.

Griliches, Z. and Mairesse, J. (1998). Production functions: The search for identification. In *Econometrics and economic theory in the twentieth century: The Ragnar Frisch Centennial Symposium*, Cambridge University Press.

Hassler, J., Krusell, P. and Olovsson, C. (2021). Directed technical change as a response to natural resource scarcity. *Journal of Political Economy*, **129** (11).

Hawkins-Pierot, J. and Wagner, K. (2021). Technology lock-in and optimal carbon pricing, working paper.

Hicks, J. R. (1932). *The Theory of Wages*. London: Macmillan, 2nd edn., 2nd edition published 1963.

Hopenhayn, H. (1992). Entry, exit, and firm dynamics in long run equilibrirum. *Econometrica*, **60** (5), 1127–1150.

IEA (2022). *SDG7: Data and Projections*. Tech. rep., International Energy Agency, Paris, France.

Klein, L. (1953). *A Textbook of Econometrics*. Evanston, IL: Row, Peterson and Co.

León-Ledesma, M., McAdam, P. and Willman, A. (2010). Identifying the elasticity of substitution with biased technical change. *American Economic Review*, **100**, 1330–1357.

Levinsohn, J. and Petrin, A. (2003). Estimating production functions using inputs to control for unobservables. *Review of Economic Studies*, **70**, 317–341.

Linn, J. (2008). Energy prices and the adoption of energy-saving technology. *The Economic Journal*, **118**, 1986–2012.

Lyubich, E., Shapiro, J. and Walker, R. (2018). Regulating mismeasured pollution: Implications of firm heterogeneity for environmental policy. *AEA Papers and Proceedings*, **108**, 136–42.

Marschak, J. and Andrews, W. (1944). Random simultaneous equations and the theory of production. *Econometrica*, **12** (3–4), 143–205.

McDonald, K., Pinardon-Touati, N. and Walsh, C. (2026). Growth, firm scale, and the energy intensity of production. *NBER Working Paper*, (35405).

Melitz, M. and Polanec, S. (2015). Dynamic Olley-Pakes productivity decomposition with entry and exit. *RAND Journal of Economics*, **46** (2), 362–375.

Newell, R., Jaffe, A. and Stavins, R. (1999). The induced innovation hypothesis and energy-saving technological change. *Quarterly Journal of Economics*, **114** (3), 941–975.

Oberfield, E. and Raval, D. (2021). Micro data and macro technology. *Econometrica*, **89** (2), 703–732.

Olley, G. S. and Pakes, A. (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, **64** (6), 1263–1297.

Popp, D. (2002). Induced innovation and energy prices. *American Economic Review*, **92** (1), 160–180.

Rafey, W. (2023). Droughts, deluges, and (river) diversions: Valuing market-based water reallocation. *American Economic Review*, **113** (2).

Raval, D. (2019). The micro elasticity of substitution and non-neutral technology. *RAND*

Journal of Economics, **50** (1), 147–167.

Rubens, M., Wu, Y. and Xu, M. (2025). Exploiting or augmenting labor? *American Economic Review: Insights*, **forthcoming**.

Solow, R. M. (1957). Technical change and the aggregate production function. *The Review of Economics and Statistics*, **39** (3), 312–320.

Syverson, C. (2011). What determines productivity? *Journal of Economic Literature*, **49** (2), 326–365.

U.S. Census Bureau (2024). About the annual survey of manufactures.

Walker, R. (2013). The transitional costs of sectoral reallocation: Evidence from the Clean Air Act and the workforce. *The Quarterly Journal of Economics*, pp. 1787—1835.

Worrell, E. *et al.* (2009). Industrial energy efficiency and climate change mitigation. *Energy Efficiency*, **2**, 109–123.

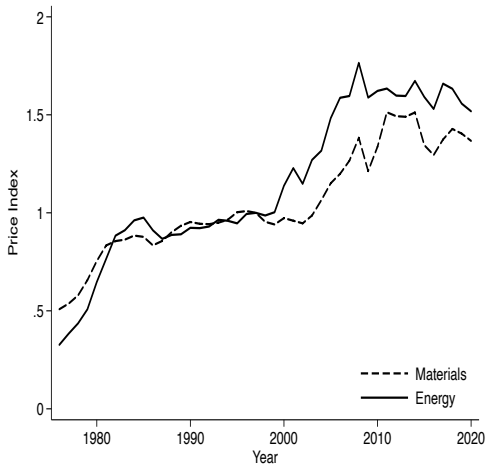
A Additional Tables and Figures

Table A.1: Production Function Parameter Estimates

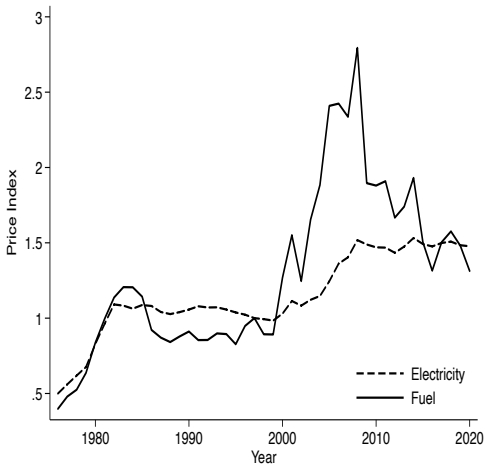
NAICS3	Description	σ	σ flag	β_M	β_K	ε_E	ε_L
327	Nonmetallic Mineral Products	0.78	0	1.70	0.20	0.056	0.30
331	Primary Metals	-0.07	1	2.50	0.00	0.037	0.18
322	Paper	0.69	0	2.50	0.20	0.041	0.23
325	Chemicals	0.72	0	1	0.30	0.057	0.19
324	Petroleum and Coal Products	0.61	0	1.90	0.10	0.057	0.10
313	Textile Mills	0.20	0	0.05	0.00	0.056	0.29
326	Plastic and Rubber Products	0.81	0	1.50	0.20	0.037	0.30
321	Wood Products	0.97	1	1.10	0.10	0.030	0.27
311	Food	0.32	0	0.90	0.01	0.28	0.19
332	Fabricated Metal Products	0.90	0	1.70	0.30	0.030	0.42
314	Textile Products	0.67	0	1.30	0.05	0.018	0.37
323	Printing and Related Support Activities	1.02	1	0.90	0.10	0.019	0.47
312	Beverage and Tobacco Products	0.91	0	1.30	0.30	0.030	0.23
335	Electrical Equipment, Appliance, and Components	0.70	0	0.60	0.05	0.020	0.33
316	Leather and Applied Products	0.65	0	1.60	0.10	0.019	0.34
337	Furniture and Related Products	1.18	1	1.90	0.20	0.018	0.38
315	Apparel	0.68	0	0.90	0.10	0.019	0.48
333	Machinery	1.01	1	1.90	0.30	0.017	0.40
336	Transportation Equipment	0.79	0	1.60	0.10	0.017	0.30
339	Miscellaneous	1.11	1	0.60	0.10	0.017	0.40
334	Computer and Electronic Products	0.87	0	1.50	0.40	0.017	0.40

Figure A.1: Price Indices

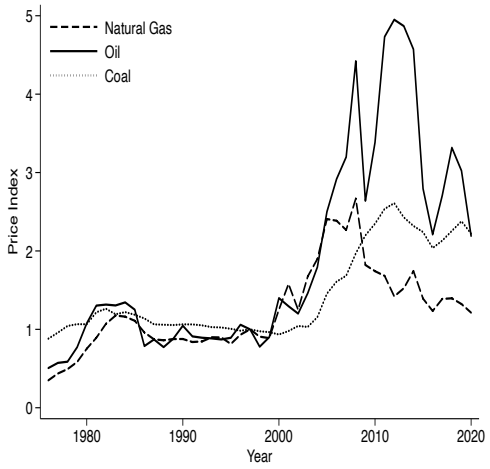
(a) Energy vs. Materials



(b) Electricity vs. Fuels



(c) Fuels by Type



(d) Electricity

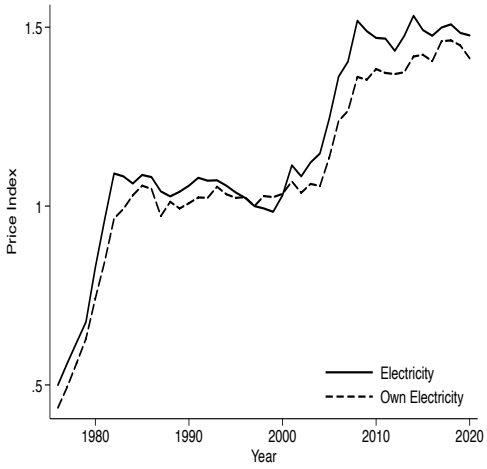
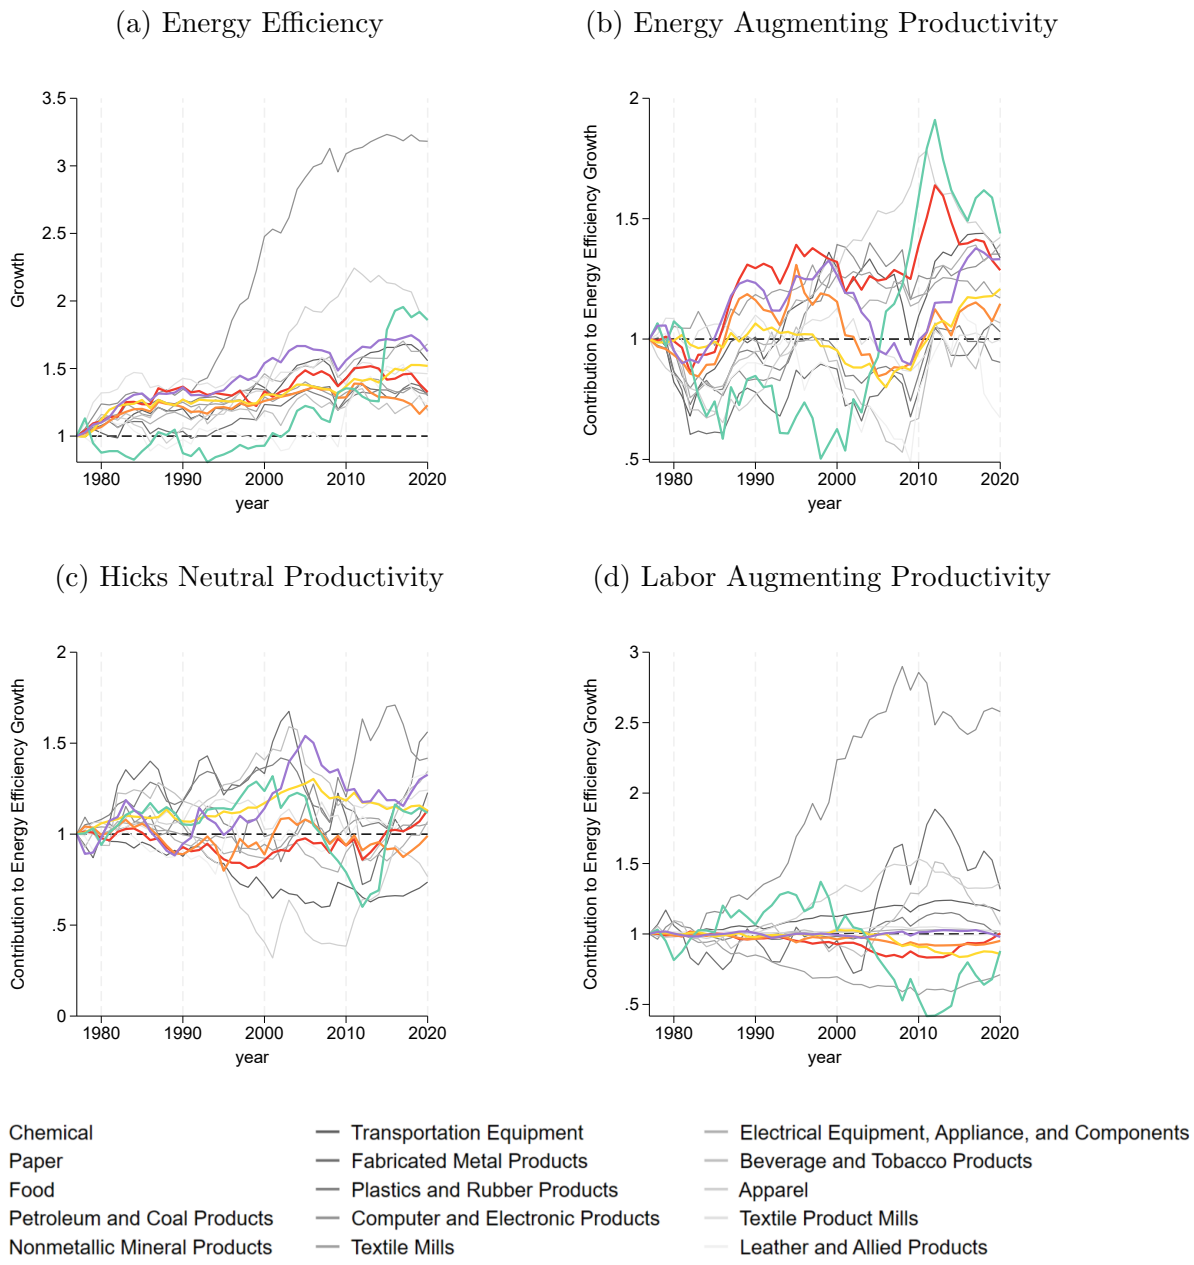
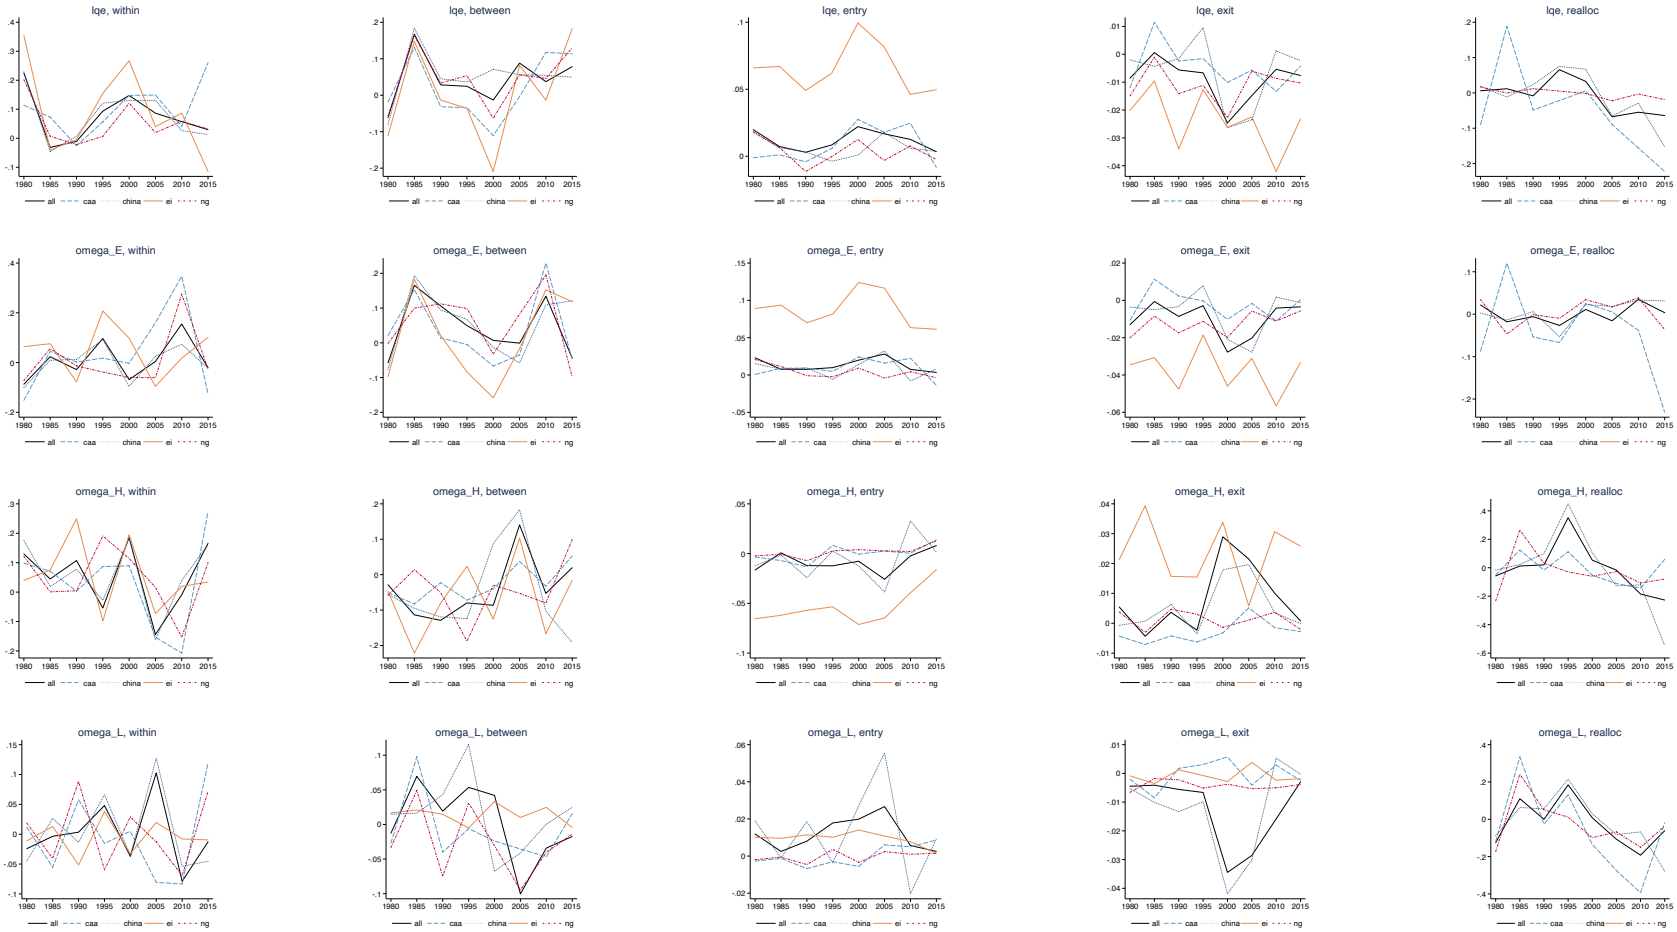


Figure A.2: Energy Efficiency Growth



Notes: This figure shows establishment-level growth of energy efficiency (panel a), energy-augmenting productivity (panel b), hicks-neutral productivity (panel c), and labor augmenting productivity (panel d) by 3-digit NAICS industry. The five most energy intense industries are highlighted in color, while remaining idustries are shown in scales of gray. *Source*: Authors' calculations using Census of Manufacturers (1977-2017), Annual Survey of Manufacturers (1976-2020), and Manufacturing Energy Consumption Survey (1984-2018) data.

Figure A.3: Decomposition Components



B Data details

Table B.1: Data Sources

var	cmf	units	details
t	1972–2019	calendar year	from ASM, CMF
Z_i	`fipsstate` × `cou`	county-state	
j	`naics_new`	six-digit industry	
Q_{it}	`q`		inventory-adjusted + deflated `tvs`
Y_{it}	`piship` · `q`	'1000 USD nominal	
	`tvs`	total value shipped	
	`fib`	finished good invtry, beginning of yr	
	`fie`	finished good invtry, end of yr	
	`wib`	works-in-progress invtry, b.o.y.	
	`wip`	works-in-progress invtry, e.o.y.	
P_{it}	`piship`	$1997 \equiv 1$	NBER-CES
E_{it}^{elec}	`pe`	kWh	
E_{it}^{fuel}	`cf`	'1000 USD nominal	
P_{it}^{elec}	`ee`/`pe`	'1000 USD nominal	average elec price
P_{it}^{fuel}		constructed	using MECS
	`pien`		NBER-CES
L_{it}	`te`	total employment hours	imputed from wage bill
	`ph`	total production worker hours, '1000	
P_{it}^{L}	`sw`/`pe`	USD nominal / hour	average wage
	`ww`/`ph`	USD nominal / hour	average wage, production workers
K_{it}	`kst_st` + `kst_eq`	'1000 USD nominal	structures + equipment
K_{i0}	`bae`	'1000 USD nominal	total buildings/structures assets, eoy
	`mae`	'1000 USD nominal	total machines/equipment assets, eoy
I_{it}^{K}	`nb`	'1000 USD nominal	new buildings/structures investment, eoy
	`nm`	'1000 USD nominal	new machines/equipment investment, eoy
Γ_{it}^{K}	`piinvs`	$1997 \equiv 1$	deflator, buildings/structures inv (BLS)
	`piinve`	$1997 \equiv 1$	deflator, machines/equipment inv (BLS)
M_{it}	`cp` + `cpc` + `cr` + `cw`	materials/parts + comms + resales + contract work	
P_{it}^{M}	`pimat`	$1997 \equiv 1$	materials price index (NBER)
	`pimat_bea`	$1997 \equiv 1$	materials price index (BEA)

Table B.2: Aggregate Decomposition of Q/E – Heterogeneity

	Within-Industry Decomposition				Industry Shift-Share			
Year Interval	(1) Survivor Technique	(2) Survivor Realloc.	(3) Entry	(4) Exit	(5) Within Industry	(6) Industry Comp.	(7) Total Change	(8) Δ Sample Share
CAA								
1977-1982	0.11	-0.02	-0.00	-0.01	0.08	-0.09	-0.01	-0.02
1982-1987	0.07	0.13	0.00	0.01	0.22	0.19	0.41	0.03
1987-1992	-0.03	-0.03	-0.00	-0.00	-0.06	-0.05	-0.11	-0.01
1992-1997	0.06	-0.03	0.01	-0.00	0.03	-0.02	0.00	-0.03
1997-2002	0.15	-0.11	0.03	-0.01	0.05	0.01	0.06	0.00
2002-2007	0.15	-0.00	0.02	-0.01	0.16	-0.09	0.07	0.02
2007-2012	0.04	0.12	0.02	-0.01	0.17	-0.16	0.01	-0.01
2012-2017	0.26	0.11	-0.01	-0.00	0.36	-0.22	0.14	-0.02
China								
1977-1982	0.23	-0.08	0.02	-0.00	0.17	0.02	0.18	0.01
1982-1987	-0.05	0.18	0.01	-0.00	0.14	-0.01	0.13	0.01
1987-1992	0.01	0.05	0.00	-0.00	0.05	0.02	0.08	0.02
1992-1997	0.12	0.04	-0.00	0.01	0.16	0.08	0.24	0.02
1997-2002	0.13	0.07	0.00	-0.03	0.18	0.07	0.24	0.01
2002-2007	0.13	0.06	0.02	-0.02	0.18	-0.07	0.12	-0.02
2007-2012	0.03	0.05	0.01	0.00	0.09	-0.03	0.06	0.00
2012-2017	0.01	0.05	0.00	-0.00	0.06	-0.15	-0.09	0.01
High NG								
1977-1982	0.20	-0.06	0.02	-0.01	0.15	0.02	0.17	0.02
1982-1987	0.01	0.16	0.01	-0.00	0.18	-0.00	0.17	0.00
1987-1992	-0.02	0.03	-0.01	-0.01	-0.01	0.01	-0.00	0.01
1992-1997	0.01	0.05	-0.00	-0.01	0.05	0.01	0.05	-0.05
1997-2002	0.12	-0.06	0.01	-0.02	0.05	-0.00	0.05	-0.01
2002-2007	0.02	0.06	-0.00	-0.01	0.07	-0.02	0.04	0.03
2007-2012	0.06	0.05	0.01	-0.01	0.10	-0.00	0.10	0.04
2012-2017	0.03	0.13	-0.00	-0.01	0.15	-0.02	0.13	0.01
Top EI								
1977-1982	0.35	-0.11	0.07	-0.02	0.29	.	0.29	-0.00
1982-1987	-0.04	0.15	0.07	-0.01	0.16	.	0.16	-0.00
1987-1992	-0.00	-0.01	0.05	-0.03	-0.00	.	-0.00	-0.00
1992-1997	0.15	-0.04	0.06	-0.01	0.17	.	0.17	0.00
1997-2002	0.27	-0.21	0.10	-0.03	0.13	.	0.13	-0.00
2002-2007	0.04	0.08	0.08	-0.02	0.18	.	0.18	0.00
2007-2012	0.09	-0.01	0.05	-0.04	0.08	.	0.08	-0.01
2012-2017	-0.11	0.18	0.05	-0.02	0.09	.	0.09	0.00

Table B.3: Aggregate Decomposition of ω_E – Heterogeneity

	Within-Industry Decomposition				Industry Shift-Share			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Year Interval	Survivor Technique	Survivor Realloc.	Entry	Exit	Within Industry	Industry Comp.	Total Change	Δ Sample Share
CAA								
1977-1982	-0.15	0.02	0.00	-0.01	-0.14	-0.09	-0.23	-0.02
1982-1987	0.05	0.15	0.01	0.01	0.22	0.12	0.34	0.03
1987-1992	0.00	0.01	0.01	0.00	0.03	-0.05	-0.03	-0.01
1992-1997	0.02	-0.00	0.00	-0.00	0.02	-0.07	-0.05	-0.03
1997-2002	-0.00	-0.07	0.02	-0.01	-0.06	0.03	-0.03	0.00
2002-2007	0.16	-0.03	0.02	-0.00	0.14	0.01	0.14	0.02
2007-2012	0.35	0.23	0.02	-0.01	0.59	-0.04	0.55	-0.01
2012-2017	-0.13	-0.05	-0.01	0.00	-0.19	-0.23	-0.42	-0.02
China								
1977-1982	-0.10	-0.08	0.02	-0.00	-0.17	0.00	-0.16	0.01
1982-1987	0.01	0.19	0.01	-0.00	0.21	-0.01	0.19	0.01
1987-1992	0.01	0.09	0.01	-0.00	0.11	0.01	0.12	0.02
1992-1997	0.09	0.07	-0.01	0.01	0.16	-0.05	0.11	0.02
1997-2002	-0.10	-0.01	0.01	-0.02	-0.12	0.02	-0.09	0.01
2002-2007	0.02	-0.06	0.03	-0.03	-0.03	0.02	-0.01	-0.02
2007-2012	0.07	0.11	-0.01	0.00	0.18	0.03	0.21	0.00
2012-2017	-0.02	0.12	0.01	-0.00	0.11	0.03	0.14	0.01
High NG								
1977-1982	-0.07	-0.00	0.02	-0.02	-0.07	0.03	-0.04	0.02
1982-1987	0.06	0.10	0.01	-0.01	0.16	-0.05	0.11	0.00
1987-1992	-0.01	0.11	-0.00	-0.02	0.08	-0.00	0.08	0.01
1992-1997	-0.04	0.10	-0.00	-0.01	0.05	-0.01	0.04	-0.05
1997-2002	-0.06	-0.03	0.01	-0.02	-0.10	0.03	-0.07	-0.01
2002-2007	-0.06	0.08	-0.00	-0.01	0.01	0.02	0.03	0.03
2007-2012	0.28	0.20	0.00	-0.01	0.46	0.04	0.50	0.04
2012-2017	-0.03	-0.10	-0.00	-0.01	-0.13	-0.04	-0.17	0.01
Top EI								
1977-1982	0.06	-0.10	0.09	-0.03	0.02	.	0.02	-0.00
1982-1987	0.08	0.18	0.09	-0.03	0.32	.	0.32	-0.00
1987-1992	-0.08	0.02	0.07	-0.05	-0.03	.	-0.03	-0.00
1992-1997	0.21	-0.08	0.08	-0.02	0.19	.	0.19	0.00
1997-2002	0.10	-0.16	0.12	-0.05	0.02	.	0.02	-0.00
2002-2007	-0.10	-0.02	0.12	-0.03	-0.03	.	-0.03	0.00
2007-2012	0.02	0.15	0.06	-0.06	0.18	.	0.18	-0.01
2012-2017	0.10	0.12	0.06	-0.03	0.25	.	0.25	0.00

Table B.4: Aggregate Decomposition of ω_L – Heterogeneity

	Within-Industry Decomposition				Industry Shift-Share			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Year Interval	Survivor Technique	Survivor Realloc.	Entry	Exit	Within Industry	Industry Comp.	Total Change	Δ Sample Share
CAA								
1977-1982	0.01	-0.03	-0.00	-0.00	-0.02	-0.12	-0.14	-0.02
1982-1987	-0.06	0.10	-0.00	-0.01	0.03	0.33	0.37	0.03
1987-1992	0.06	-0.04	-0.01	0.00	0.01	-0.03	-0.01	-0.01
1992-1997	-0.02	-0.01	-0.00	0.00	-0.02	0.13	0.11	-0.03
1997-2002	0.00	-0.02	-0.01	0.01	-0.02	-0.13	-0.15	0.00
2002-2007	-0.08	-0.04	0.01	-0.00	-0.11	-0.27	-0.39	0.02
2007-2012	-0.08	-0.05	0.00	0.00	-0.12	-0.39	-0.52	-0.01
2012-2017	0.12	0.02	0.01	-0.00	0.14	-0.02	0.12	-0.02
China								
1977-1982	-0.04	0.01	0.02	-0.01	-0.02	-0.09	-0.10	0.01
1982-1987	0.03	0.02	-0.00	-0.01	0.03	0.06	0.09	0.01
1987-1992	-0.01	0.04	0.02	-0.01	0.03	0.06	0.09	0.02
1992-1997	0.07	0.12	-0.00	-0.01	0.17	0.21	0.38	0.02
1997-2002	-0.04	-0.07	0.03	-0.04	-0.12	0.03	-0.09	0.01
2002-2007	0.13	-0.04	0.06	-0.03	0.11	-0.08	0.03	-0.02
2007-2012	-0.05	0.00	-0.02	0.01	-0.07	-0.07	-0.14	0.00
2012-2017	-0.05	0.02	0.01	-0.00	-0.01	-0.28	-0.29	0.01
High NG								
1977-1982	0.02	-0.03	-0.00	-0.01	-0.02	-0.17	-0.20	0.02
1982-1987	-0.04	0.05	-0.00	-0.00	0.01	0.24	0.25	0.00
1987-1992	0.09	-0.07	-0.00	-0.00	0.01	0.05	0.06	0.01
1992-1997	-0.06	0.03	0.00	-0.01	-0.03	0.01	-0.02	-0.05
1997-2002	0.03	-0.03	-0.00	-0.00	-0.01	-0.10	-0.11	-0.01
2002-2007	-0.01	-0.09	0.00	-0.01	-0.11	-0.07	-0.18	0.03
2007-2012	-0.07	-0.04	0.00	-0.01	-0.11	-0.15	-0.26	0.04
2012-2017	0.07	-0.01	0.00	-0.00	0.05	-0.04	0.02	0.01
Top EI								
1977-1982	-0.01	0.02	0.01	-0.00	0.01	.	0.01	-0.00
1982-1987	0.01	0.02	0.01	-0.00	0.04	.	0.04	-0.00
1987-1992	-0.05	0.01	0.01	0.00	-0.02	.	-0.02	-0.00
1992-1997	0.04	-0.00	0.01	-0.00	0.04	.	0.04	0.00
1997-2002	-0.03	0.03	0.01	-0.00	0.01	.	0.01	-0.00
2002-2007	0.02	0.01	0.01	0.00	0.04	.	0.04	0.00
2007-2012	-0.01	0.02	0.01	-0.00	0.02	.	0.02	-0.01
2012-2017	-0.01	-0.00	0.00	-0.00	-0.01	.	-0.01	0.00

Table B.5: Aggregate Decomposition of ω_H – Heterogeneity

	Within-Industry Decomposition				Industry Shift-Share			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Year Interval	Survivor Technique	Survivor Realloc.	Entry	Exit	Within Industry	Industry Comp.	Total Change	Δ Sample Share
CAA								
1977-1982	0.10	-0.05	-0.00	-0.00	0.04	-0.05	-0.01	-0.02
1982-1987	0.07	-0.08	-0.01	-0.01	-0.03	0.12	0.10	0.03
1987-1992	0.00	-0.02	-0.01	-0.00	-0.04	-0.02	-0.05	-0.01
1992-1997	0.09	-0.07	0.01	-0.01	0.02	0.11	0.13	-0.03
1997-2002	0.09	-0.04	-0.00	-0.00	0.05	-0.05	-0.01	0.00
2002-2007	-0.15	0.04	0.00	0.01	-0.11	-0.11	-0.22	0.02
2007-2012	-0.21	-0.03	0.00	-0.00	-0.24	-0.14	-0.38	-0.01
2012-2017	0.27	0.05	0.01	-0.00	0.34	0.06	0.40	-0.02
China								
1977-1982	0.18	-0.06	-0.01	-0.00	0.11	-0.02	0.09	0.01
1982-1987	0.02	-0.10	-0.00	0.00	-0.08	0.02	-0.05	0.01
1987-1992	0.08	-0.12	-0.02	0.01	-0.06	0.10	0.04	0.02
1992-1997	-0.03	-0.12	0.00	-0.00	-0.15	0.45	0.30	0.02
1997-2002	0.18	0.09	-0.01	0.02	0.28	0.11	0.38	0.01
2002-2007	-0.16	0.18	-0.04	0.02	0.00	-0.13	-0.13	-0.02
2007-2012	0.04	-0.10	0.03	0.00	-0.03	-0.12	-0.15	0.00
2012-2017	0.16	-0.19	0.00	-0.00	-0.03	-0.55	-0.58	0.01
High NG								
1977-1982	0.12	-0.06	-0.00	0.00	0.06	-0.23	-0.17	0.02
1982-1987	0.00	0.01	-0.00	-0.00	0.01	0.27	0.28	0.00
1987-1992	0.00	-0.05	-0.01	0.00	-0.05	0.03	-0.02	0.01
1992-1997	0.19	-0.19	0.00	0.00	0.01	-0.03	-0.02	-0.05
1997-2002	0.11	-0.03	0.00	-0.00	0.09	-0.06	0.03	-0.01
2002-2007	0.02	-0.05	0.00	0.00	-0.03	-0.03	-0.06	0.03
2007-2012	-0.15	-0.08	0.00	0.00	-0.23	-0.11	-0.33	0.04
2012-2017	0.10	0.10	0.01	-0.00	0.21	-0.08	0.13	0.01
Top EI								
1977-1982	0.04	-0.05	-0.07	0.02	-0.05	.	-0.05	-0.00
1982-1987	0.08	-0.22	-0.06	0.04	-0.17	.	-0.17	-0.00
1987-1992	0.25	-0.08	-0.06	0.02	0.13	.	0.13	-0.00
1992-1997	-0.10	0.02	-0.05	0.02	-0.11	.	-0.11	0.00
1997-2002	0.19	-0.13	-0.07	0.03	0.03	.	0.03	-0.00
2002-2007	-0.07	0.10	-0.06	0.01	-0.03	.	-0.03	0.00
2007-2012	0.02	-0.17	-0.04	0.03	-0.16	.	-0.16	-0.01
2012-2017	0.03	-0.02	-0.02	0.03	0.03	.	0.03	0.00