Growth, Firm Scale, and the Energy Intensity of Production*

Kathryn McDonald
Columbia University

Noémie Pinardon-Touati
Columbia University

Conor Walsh
Columbia University

March 2026

Abstract

We uncover a new mechanism that links growth and a decline in the energy intensity of production, observed globally since 1990. Using microdata from India and a causal research design, we demonstrate that the expenditure share of energy declines steeply with firm scale, due both to physical scaling laws and technology investment. Given the fact that average firm size increases with growth, this scale dependence in energy demand implies that production endogenously becomes less energy-intensive with aggregate growth. We develop a model of this mechanism in general equilibrium, and quantify significant reductions in aggregate energy intensity as low- and middle-income countries (LMICs) like India grow. We conclude with a discussion of the future path of emissions in India.

Keywords: Growth, Energy Cost Shares, Firm Size, Energy Efficiency

JEL Codes: O44, Q43, D24, L25

*We are grateful to Francine Montecinos for excellent research assistance. We thank Ashley Langer (discussant), Severin Bornstein, Danial Lashkari, Kevin Kuruc, Ezra Oberfield and Joseph Shapiro for helpful comments. We thank conference participants at the NYU LAEF, OzMacro, Utah Winter Business Economics conferences, as well as seminar participants at Berkeley ARE, Columbia, Georgetown, INSEAD, Middlebury, NY Fed, Richmond Fed, SF Fed, and Stanford.

1. INTRODUCTION

The tension between economic growth and environmental damage is one of the central challenges of our times. As long as production requires energy, and energy usage emits CO2, there exists a trade-off between expanding production and reducing carbon emissions. A key determinant of this trade-off is the energy intensity of production, which measures the energy required to produce an additional unit of output. Figure 1 shows that the energy intensity of production has declined significantly since the 1990s at the global level. This almost entirely explains the reduction in CO2 intensity in recent decades, as the CO2-to-energy ratio has remained relatively stable up until 2020. Understanding the drivers and potential persistence of this trend is crucial for assessing the sustainability of future economic growth.

The decline in energy intensity could be driven by many factors: energy-saving technological improvement, a change in the relative price of energy, or changes in policy, for instance.[1] The core contribution of this paper is to uncover a mechanism that endogenously links economic growth to a decline in energy intensity. This mechanism combines two facts. First, it is well-known that average firm size increases as economies grow (Tybout 2000; Hsieh and Klenow 2014; Bento and Restuccia 2017). Second, we document that relative energy demand falls as firms grow. Together, these facts imply that aggregate production endogenously becomes less energy-intensive as economies grow.

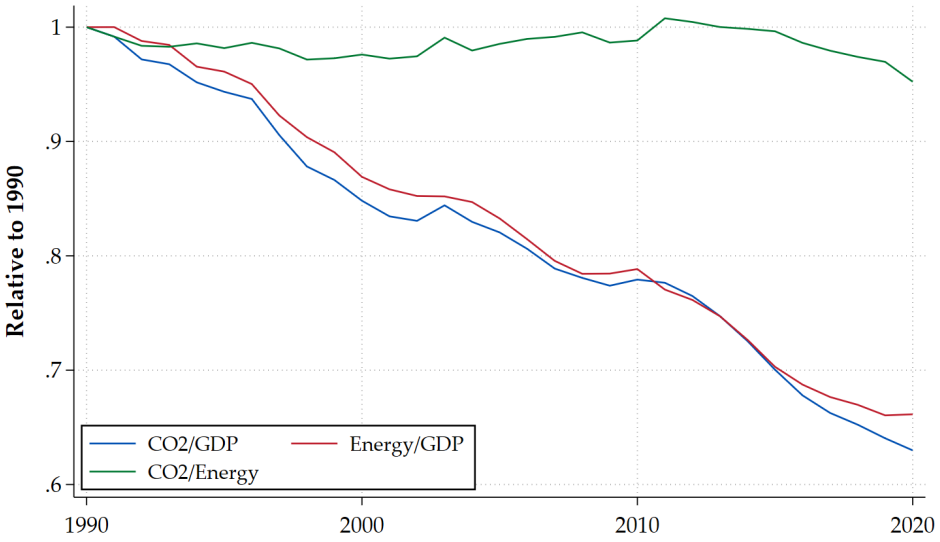
We conduct our investigation in the context of India. India is both representative of the broader group of low- and middle-income countries (LMICs) in its patterns of energy use, and is interesting in itself due to its size and projected growth trajectory. We proceed in two steps. First, we exploit firm-level data to empirically document that the energy expenditure share declines as firms grow. We show that this relationship is causal and is driven by a scale elasticity of energy demand (as opposed to substitution due to differences in relative prices). Second, we develop a model of heterogeneous firms with non-homothetic energy demand, where scale dependence arises from both physical scaling laws and endogenous technology choice. Disciplining the model using our causal estimates, we quantify the role of scale dependence in explaining the aggregate decline in energy intensity observed in India in previous decades.

For the first step, we use firm-level data from the Annual Survey of Industries, which provides detailed information on energy expenditures and usage across a representative sample of manufacturing establishments from 1998 to 2017. We document a robust negative relationship between firm size and their expenditure share on energy inputs, within narrowly defined industry-year cells. Firms in the tenth decile of the firm size distribution have an expenditure share on energy which is about half that of firms in the first decile.

To address concerns related to omitted variables that could be correlated with firm size in the cross-section, we develop a causal research design where we instrument for output growth using firm-specific demand shocks. Our instrument exploits the interaction of firm-level ex-ante product shares with the evolution of nationwide sales across products. With this design, we obtain a negative elasticity of approximately -0.4, which is stable at different time horizons.

[1] Appendix Figure B.1 shows that this decline is *not* explained by reallocation of economic activity across countries and sectors—the latter accounts for less than 10% of the total decline.

FIGURE 1: Global Energy and CO2 Intensity



Note: GDP is measured in constant 2021 USD PPP. CO2/GDP is tonnes of CO2 per dollar of output. Energy/GDP is computed as joules of energy per dollar of output. CO2/Energy is tonnes of CO2 emitted per joule of energy. Values are computed by summing values for a balanced panel of 134 countries which jointly account for 95% of the world's total energy consumption. All series are normalized to 1 in 1990.

The negative relationship between size and the energy expenditure share is robust to measuring energy demand relative to either material inputs, material and labor inputs, total inputs including capital, or output. The decreasing relationship between firm size and the expenditure share on energy inputs translates into a decreasing relationship between firm size and energy intensity (measured both as physical quantity of energy per unit of output and as physical quantity of energy per dollar of revenue).

We then turn to the mechanisms that can explain this result. First, we explore *how* firms reduce their energy expenditure shares when scaling up. The reduction in the energy expenditure share is fully attributable to a decline in the use of energy relative to other material inputs—energy expenditures relative to capital and labor expenditures are unchanged. That is, as scale increases, firms are using less energy to process a given amount of material input.[2] The relationship between energy demand and scale is not driven by changes in relative input prices as firms grow. That is, our estimated coefficients reflect a scale elasticity, not a substitution elasticity.

Second, we explore the reasons *why* firms disproportionately reduce their energy usage when scaling up. Our main hypothesis is that energy as an input is characterized by a large fraction of the energy purchased being wasted, and these energy losses scale less than proportionately with firm output. We show that this is due to two distinct forces: (i) for a given technology, operating at a higher scale is more energy efficient because physical scaling laws imply smaller relative losses as scale increases; (ii) as they grow, firms are able to undertake larger fixed-cost investments that improve energy efficiency.

[2]One potential concern is that the growing material share reflects higher indirect energy use as firms grow. We estimate total energy use—the sum of energy used by the firm and of the energy embedded in its material inputs—and show that total energy use does decline as firms grow.

How does the scale dependence of energy demand affect aggregate energy intensity along the growth path? To answer this question, we develop a heterogeneous firms model with non-homothetic energy demand. We consider a nested non-homothetic CES (nhCES) firm-level production function that combines energy and non-energy inputs (Sato 1974, 1977; Comin, Lashkari, and Mestieri 2021; Lashkari, Bauer, and Boussard 2024). This specification assumes that the output elasticity of energy demand relative to other factors is a constant structural parameter. This parameter is identified by our reduced-form evidence on the causal relationship between firm scale and the energy expenditure share. Calibrating the model requires another key elasticity: the elasticity of substitution between energy and other factors. We obtain this elasticity by estimating the response of relative energy demand to shocks to the relative price of energy inputs. We find that energy is a strong complement to other inputs.

We find that the scaling up channel has economically meaningful effects on the aggregate energy expenditure share. That channel generates a reduction in the energy expenditure share of about 45% over 1990 to 2023. This corresponds to almost all of the decline in the energy expenditure share from 1990 to 2023. In the data, the relative price of energy to the final good has increased over time. When feeding this time series for the relative price to the model, we find that the model predicts a 32% decline in the energy expenditure share, corresponding to roughly 70% of the decline in the expenditure share. There are two conclusions from this analysis. First, our theory explains most of the change in the energy expenditure share. The residual not accounted for by our theory is any energy-specific aggregate productivity shifts. Second, the decline in the aggregate energy share in India is attributable to the scaling up channel—as opposed to a substitution channel—since the change in the relative price of energy works in the opposite direction.

With the calibrated model, we perform two further exercises. First, we use the model to project the path of aggregate energy intensity in India in future decades. Compared to a benchmark without scale dependency of energy demand, our analysis implies a 40% reduction in the growth of total energy demanded between 2023 and 2050. Second, we use our model to examine the paths of CO2/Energy intensity that would be required to meet India's Paris Climate Agreement goals to reduce the emissions intensity of its GDP by 45 percent by 2030. We find that the CO2/Energy ratio would be required to fall significantly more, and thus require significantly more policy-led decarbonization, without accounting for the mechanism we uncover.

In summary, our study shows that there exists a causal mechanism linking economic growth and lower relative energy usage. This has two key policy implications. First, it suggests that the trade-off between economic outcomes and environmental degradation endogenously improves as LMICs grow. Second, while this paper explores the implications of the scale elasticity of energy demand for growth, another key implication is for policies affecting the firm size distribution: industrial policies favoring large firms will reduce the aggregate energy intensity of production. Finally, our findings encourage the developers of Integrated Assessment Models (IAMs) to consider incorporating energy scaling laws with development. Doing so could be a fruitful avenue for future research.

Related Literature. Our work most closely relates to a growing literature investigating the consequences of economic growth for energy use and CO2 emissions. A large body of work has focused on the endogenous relationship between growth and the adoption of renewable energy sources (Acemoglu, Aghion, Bursztyn, and Hemous 2012; Acemoglu, Aghion, and Hémous 2014; Arkolakis and Walsh 2023). Meanwhile, we focus on the endogenous relationship between growth and the energy intensity of production. Previous work has documented the importance of directed technical change for the energy transition (Shanker and Stern 2018; Hassler, Krusell, and Olovsson 2021; Acemoglu, Aghion, Barrage, and Hémous 2023; Casey 2024). We highlight technological improvements in response to scale, as opposed to a response to the relative price of energy. Gertler, Shelef, Wolfram, and Fuchs (2016), Caron and Fally (2022) and Aghion, Boppart, Peters, Schwartzman, and Zilibotti (2024) explore the role of non-homothetic consumption for the relationship between growth and the energy/CO2 intensity of consumption. We document a meaningful non-homotheticity in production.

Second, our paper also relates to the literature that decomposes the time series of CO2 or pollution intensity into the composition of products, the composition of firms within products, and within-firm changes (Levinson 2015; Shapiro and Walker 2018; Martin 2011; Barrows and Ollivier 2018). Compared to this literature, our work provides an endogenous mechanism that links growth and energy intensity. In their work, Barrows and Ollivier (2018) document that CO2 intensity is falling in the cross-section of firm size in their sample of 2,500 Indian firms. While their sample only covers the extreme right tail of the firm size distribution, we show that this pattern holds in a sample approximately 100 times larger. In addition, we go beyond correlations to document a causal effect of firm scale on relative energy demand. Finally, we quantify how the scale dependence of energy demand shapes aggregate energy intensity along the growth path.

Within the literature on firms and energy or CO2 intensity, it has also been shown that CO2 intensity is negatively correlated with firms' revenue productivity (Klenow, Pastén, and Ruane 2024; Gao 2025; Kim 2025; Kabir and Tan 2025) with implications for carbon taxation policies and the effects of green financing on aggregate production.

Third, we contribute to the literature examining the nexus between economic development, energy usage, and CO2 emissions in the context of LMICs. Closest to the spirit of our exercise, several papers (e.g., Bruckner, Hubacek, Shan, Zhong, and Feng 2022; Wollburg, Hallegatte, and Mahler 2023) quantify the effect of growth in low-income countries on global CO2 emissions. However, these papers build projections using historical levels of CO2 intensity. Our work shows that energy intensity endogenously evolves as these countries grow. Several papers investigate the role of the electricity sector in the process of growth (Abeberese 2017; Lee, Miguel, and Wolfram 2020; Singer 2024; Colmer, Lagakos, and Shu 2024). Abeberese (2017) and Singer (2024) investigate how shocks to the relative price of electricity affect firm growth and electricity intensity in India. We study how growth affects relative energy demand.

Finally, a distinct literature investigates the effects of trade on CO2 and energy intensity. This literature has documented that exporting firms—which tend to be larger—are cleaner (Batrakova and Davies 2012; Forslid, Okubo, and Ulltveit-Moe 2018; Barrows and Ollivier 2021). Barrows

and Ollivier (2021) show that, among Indian exporters, exposure to foreign demand shocks leads to reductions in the CO2 intensity of output. Also in India, Martin (2011) documents improvements in firms' energy efficiency following the 1991 trade liberalization. It is critical to evaluate the relationship between scale and energy intensity beyond the sample of exporters; exporting firms could for instance be cleaner because they are more exposed to foreign demand for clean products or foreign climate policy.

The remainder of the paper is structured as follows. Section 2 describes the data sources and presents aggregate facts on energy intensity. Section 3 presents our main results on firm scale and energy intensity. Section 4 proposes a theory of aggregate energy intensity. Section 5 quantifies the importance of this mechanism to explain the historical trends in India, and section 6 explores the implications for Indian energy use and CO2 emissions going forward. Finally, Section 7 concludes by discussing policy implications and avenues for future research.

2. Data and empirical context

2.1 Data

Annual Survey of Industries. The main data source is the Annual Survey of Industries (ASI), India's mandatory annual establishment-level manufacturing survey. Its long history since 1953 makes it a relatively reliable data source in the developing country context. The formal firms in the ASI represent approximately 80% of output, with the remaining 20% made up of informal firms or firms with less than 10 employees. We use an annual panel from 1998 to 2017. A detailed description of the data and variables we use is contained in Appendix A.1.

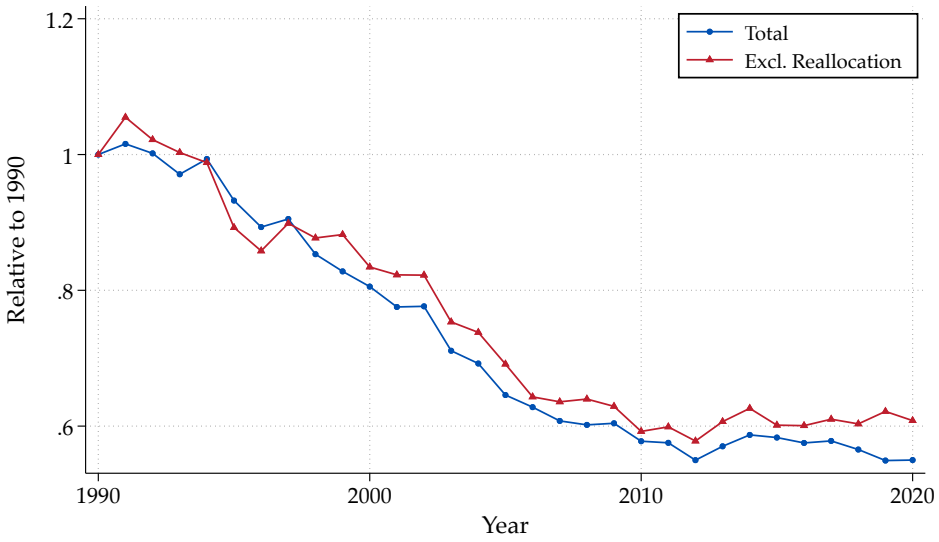
The ASI contains detailed data on energy usage by manufacturing plants: we know total expenditures on energy inputs, broken down into electricity, coal, oil, and (since 2008) gas. For each energy source except oil, expenditures are broken down into price and quantity. For oil, we approximate quantity consumed by dividing expenditures by the price of oil from the Petroleum Planning & Analysis Cell (PPAC). This allows us to construct the physical quantity of energy used by each firm.

At several points in the paper, we exploit the fact that, for both the products that manufacturing plants produce and the inputs they buy, we observe information on sales, quantities, and unit values, at the product-code level. This allows us to construct firm-level changes in prices and quantities, for both products and material inputs (see details in Appendix A.1).

We drop observations in non-manufacturing industries, winsorize the lowest and highest percentile of each variable within each year to reduce sensitivity to outliers, and deflate all monetary values in the ASI to the base year of 2011 throughout the paper. All statistics are weighted by the sampling weight.

Other data sources. We exploit several additional data sources, including: (i) International Energy Agency data, (ii) Indian national accounts, (iii) World Bank Enterprise Survey, (iv) United Nations Industrial Development Organization (UNIDO) Energy Efficiency and Renewable Energy Technology Compendiums, and (v) India's National Sample Survey on Unorganised Manufacturing Enterprises (NSS). These data sources are described in Appendix A.3.

FIGURE 2: Energy Intensity Components in India



Note: This figure displays the change in energy intensity in India. It is defined as energy used in the production of goods and services divided by GDP (measured in constant 2021 USD PPP), and is normalized to 1 in 1990. The change in energy intensity is decomposed into a within-sector component, a reallocation component, and a cross term (see equation (B.1.1) for details). The sectors are: Agriculture and Fishing, Industry (mining, manufacturing, construction), Transport and Communications, and Other Services.

2.2 Energy Intensity and the Energy Expenditure Share in India

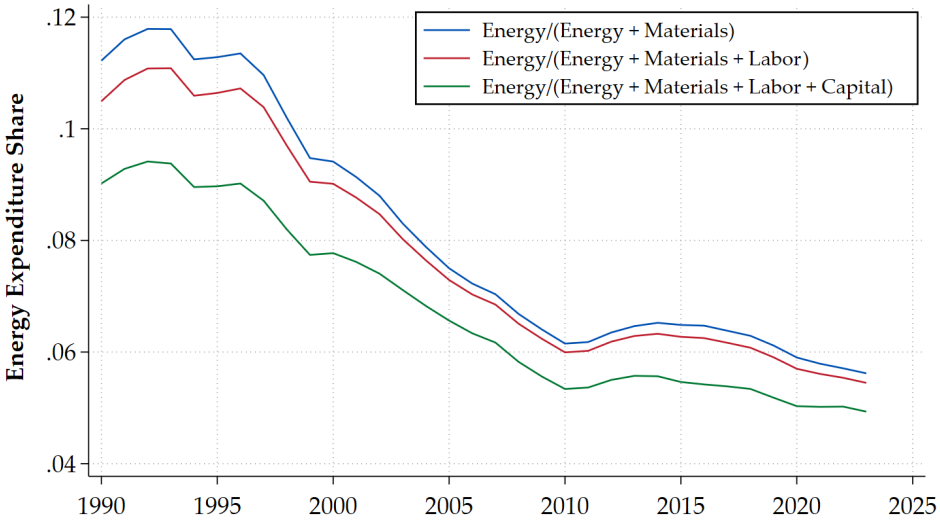
As in the rest of the world, energy intensity has declined markedly in India since 1990. Figure 2 shows that aggregate energy intensity declined by close to 50% between 1990 and 2020—slightly more than the global average decline.

Our theory speaks to declines in sectoral energy intensities as India (and its sectors grow). Another source of decline in aggregate energy intensity may be reallocation across sectors, for instance if production reallocates from energy-intensive manufacturing to less energy-intensive services. To evaluate the magnitude of this channel, we implement a within-across decomposition of the change in aggregate energy intensity and report the time series excluding the reallocation component (see details in Appendix B.1.1). The sectors are agriculture, industry (mining, manufacturing and construction), transport, and services. Reallocation accounts for little of the decline. The industry share remains roughly constant, and the decline in moderately energy-intensive agriculture is offset by a rise in both low-energy services and highly energy-intensive transport.

For the remainder of the paper, we focus on energy use in the manufacturing sector. Manufacturing accounts for roughly 70% of energy used in the production of goods and services in India. Focusing on manufacturing allows us to use the ASI data to look beyond energy intensity and investigate the expenditure share of energy inputs. Figure 3 shows the cost share of energy relative to the cost share of other inputs over time. It reveals that the cost share of energy has declined *relative to that* of other inputs. Taking the green line which uses total costs as a denominator, the energy expenditure share has declined by 45% from 1990 to 2023.[3]

[3]Figure B.3 displays the energy expenditure share in ASI purged of reallocation across 2-digit industries and

FIGURE 3: Energy Expenditure Share in Indian Manufacturing



Note: This figure plots the aggregate cost share of energy inputs relative to other inputs in the ASI data.

3. EMPIRICAL EVIDENCE ON FIRM SCALE AND ENERGY DEMAND

One simple way to think about the relationship between energy intensity and scale is to imagine a firm i with four basic inputs $\mathcal{X} \subseteq \{E, M, L, K\}$ where E is energy, M is material inputs, L is labor and K is capital. They have a cost function C_i for output that depends on these inputs and the choice of output y_i, as in

$$C_i(y_i) = \mathcal{C}_i(y_i, \mathbf{w}),$$

where $\mathbf{w}$ is the vector of input prices. Then, if we write energy intensity as energy inputs in joules per dollar value of output, we have

$$\Theta_i \equiv \frac{e_i}{p_i y_i},$$

where p_i is the unit price.[4] Now, define the expenditure share of energy in inputs as

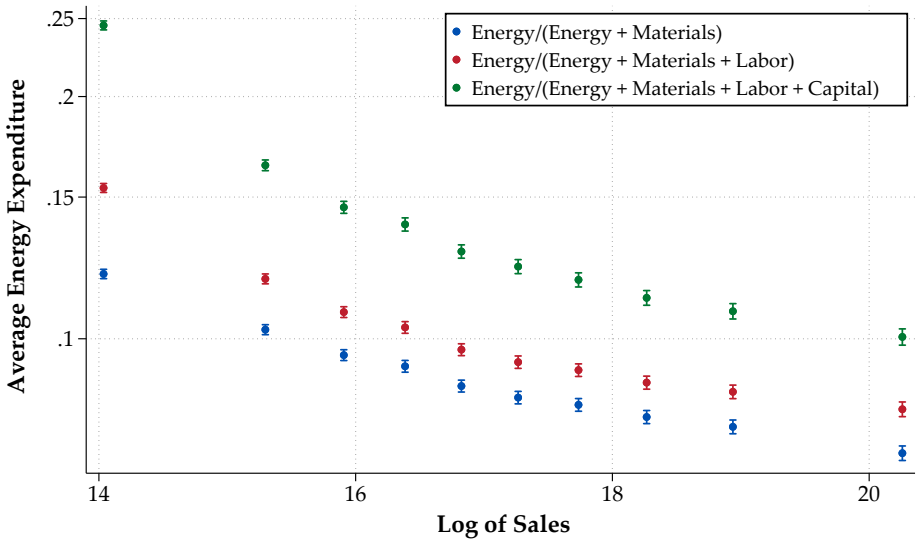
$$\Omega_i^E \equiv \frac{w_i^E e_i}{\sum_{X \in \mathcal{X}} w_i^X x_i} = \frac{w_i^E e_i}{\mathcal{C}_i(y_i, \mathbf{w}_i)}, \tag{1}$$

where w_i^X is the (possibly firm-specific) price of input X and x_i is the input choice. We can then write

shows that reallocation across manufacturing industries plays no role in the decline in the energy share.

[4]We want to begin from revenue intensity at the micro level, not output intensity e_i/y_i, since we want a way to compare units across firms with heterogeneous outputs (ceramics and casts, for example). This is also the correct notion we need in order to disaggregate economy-wide energy intensity E/GDP into firm-level components.

FIGURE 4: Energy Expenditure Shares and Firm Size in the Cross-Section



Note: This figure shows average energy expenditure shares by deciles of firm size as measured by revenue (defined within each 3-digit industry×year) in the ASI data. The dot is the point estimate and the bars indicate the 95% confidence interval.

$$\Theta_i = \frac{\Omega_i^E(y)C_i}{w_i^E p_i y_i} = \Omega_i^E(y)(1-\Pi_i)(w_i^E)^{-1}, \tag{2}$$

where $\Pi_i \equiv (p_i y_i - C_i)/(p_i y_i)$ is the profit rate of the firm. As such, energy intensity can vary with scale for three reasons. First, expenditure shares can vary with scale, as we will show below. Firms scaling up may choose to alter their input mix, buying relatively more or less energy. Second, the profit rate may vary with firm size. This could occur because of economies of scale, or because of variable markups with firm size. Third, prices of energy could change with firm scale. In this paper, we focus primarily on the first channel. In addition, there is evidence in both our data and the broader literature for the second channel that large firms have higher profit rates, and we explore this below. The third channel, differential input prices by size, does not appear to be present in our data.

3.1 Firm Scale and Relative Energy Demand

3.1.1 Energy Expenditure Shares in the Cross-Section of Firms

Figure 4 shows the relationship between the expenditure share on energy relative to total inputs and firm size in the cross-section of firms. It plots the average energy expenditure share by deciles of firm size (defined within each industry×year). The denominator of the energy expenditure share is defined as either the cost of energy and materials ($X = \{E, M\}$), energy, materials and labor ($X = \{E, M, L\}$), or energy, materials, labor and capital ($X = \{E, M, L, K\}$). It shows a strongly negative relationship between firm size and energy expenditure share.

A key threat to the interpretation of these results as a scale elasticity of energy demand is if larger firms systematically experience different energy-specific productivities or supply shocks.

One potential concern is that size correlates with age, and firms in different cohorts use capital of different vintages, which may be more or less energy-efficient. Additionally, it may be the case that smaller firms systematically differ in their probabilities to experience electricity shortages, or differ in the probability that their energy-using machinery breaks down, or can only access energy inputs at higher prices or of lower quality. Finally, it may be the case that size is not only correlated with factor-neutral productivity, but also factor-biased productivity, e.g., large firms are large because they are particularly labor- or energy-productive. Appendix B.3.2 formalizes these concerns for the identification of the scale elasticity of energy demand through the lens of the production function used in Section 5. To address these identification concerns, we turn to an instrumental variable strategy.

3.1.2 Causal Effects of Firm Growth on Relative Energy Demand

To examine the causal effect of firm scale on the intensity of energy demand, we construct an instrument for firm-level output growth and estimate the effect of a change in output on the change in the energy expenditure share across h years. We estimate regressions with the following general form:

$$\Delta^h \log\left(\frac{\Omega^E_{it}}{1-\Omega^E_{it}}\right) = \alpha_{st} + \eta^h \, \Delta^h \log Output_{it} + \varepsilon_{it} \tag{3}$$

Δ^h is the difference operator across h years: $\Delta^h x_{it} = x_{it} - x_{it-h}$. $Output_{it}$ is the value of firm sales (in constant 2011 INR). The outcome variable is the log of energy expenditures relative to expenditures on other inputs. We choose this outcome variable because the estimated elasticity directly maps to the structural parameter of the model presented in Section 5, but show the robustness of our results to choosing other outcome variables below. α_{st} are industry$\times$time fixed effects, where industries are defined by 3-digit NIC codes. In some specifications, we further interact these fixed effects with cohort fixed effects (we group year of initial production in bins of four years back to 1974, and define two bins for the earlier years). For each horizon h, η^h is the elasticity of the energy expenditure share with respect to firm scale. We present results for $h = \{1, 3, 5\}$, allowing long-run elasticities to differ from short-run elasticities.[5]

In order to address potential endogeneity concerns associated with estimating equation (3), we proceed with an instrumental variable strategy. We construct a shift-share instrument for firm-specific demand shocks.[6] We interact ex-ante firm-level product shares with the time-series evolution of aggregate sales at the product level. The logic is that a firm will grow if it is specialized in products that are subject to positive product-specific demand shocks. More

[5]The ASI is a representative sample where large units are sampled every year, and small units are sampled with a probability of 20%. This implies that looking at h-horizon differences will imply dropping many observations for small firms. The best populated difference is the first-difference $h = 1$.

[6]Our instrument need not be a demand shifter and could instead be a factor-neutral productivity shock. The key requirement is that it is orthogonal to factor-specific productivity shocks.

precisely, defining $\mathcal{J}$ as the set of all products, we let

$$\mathcal{D}_{it}^{h} = \sum_{j \in \mathcal{J}} \omega_{i,j,t-h} \, \Delta^{h} \log Output_{jt} \tag{4}$$

$\Delta^{h} \log Output_{jt}$ is the aggregate growth rate of sales of product j. $\omega_{i,j,t-h}$ is the share of product j in the sales of firm i at time $t-h$. We use $\mathcal{D}_{it}^{h}$ as an instrument for $\Delta^{h} \log Output_{it}$.

The instrument is relevant: Appendix Table B.3 reports first-stage results and shows that the F-stat varies between 297 and 502 across horizons and specifications.

The instrument will be valid as long as firm exposure to product-level growth is orthogonal to the unobserved determinants of energy shares ε_{it}, in particular any firm-specific energy supply or productivity shock. That is, firms must not sort across products such that firms with high (low) energy supply shocks systematically have high shares in high (low) growth products (Borusyak, Hull, and Jaravel 2022). In particular, it must be that products whose demand is expanding in the aggregate are not also experiencing energy-specific technical change in their production. Figure B.6 shows that the instrument is uncorrelated with firm characteristics that are likely correlates of energy supply or productivity, in particular the levels and past trends in factor shares. We also provide evidence that our instrument captures a demand shift by showing that it is uncorrelated with firm-level changes in TFPQ.[7] Appendix B.3.2 provides further details on identification with the shift-share design following the "shifters-based" approach in Borusyak et al. (2022).

Main results: Cost shares. In Table 1 we report the results of this estimation strategy. As a baseline, we report results where the energy share is defined relative to energy, material and labor ($X = \{E, M, L\}$). Panel A reports OLS results for $h = 1$, $h = 3$, and $h = 5$. The elasticities are systematically negative and approximately equal to -0.3 at all horizons. Panel B reports our IV estimates. The IV elasticities are similar to the corresponding OLS ones, with an elasticity equal to approximately -0.4. That is, a 10% increase in output leads to a 4% drop in the relative energy expenditure share. In Appendix Tables B.5 and B.6, we show that these results hold for any choice of denominators: $X = \{E, M\}$, $X = \{E, M, L, K\}$, as well as gross output or value added.

Figure 5 visualizes the relationship using binned scatter plots (of variables residualized on industry$\times$year fixed effects). The relationship is approximately linear, providing support for the choice of our specification.

How large are the differences in the energy share across firms implied by this scale elasticity? Consider a firm that grows from the 10th to the 90th percentile of the size distribution: this corresponds to $\Delta \log Output_{it} = 5.2$ (output is multiplied by 184). From our estimated elasticity, the log relative energy expenditure share will fall by 2.1 (be divided by 8). This corresponds to approximately 75% of the p10-p90 difference in the log energy share distribution.

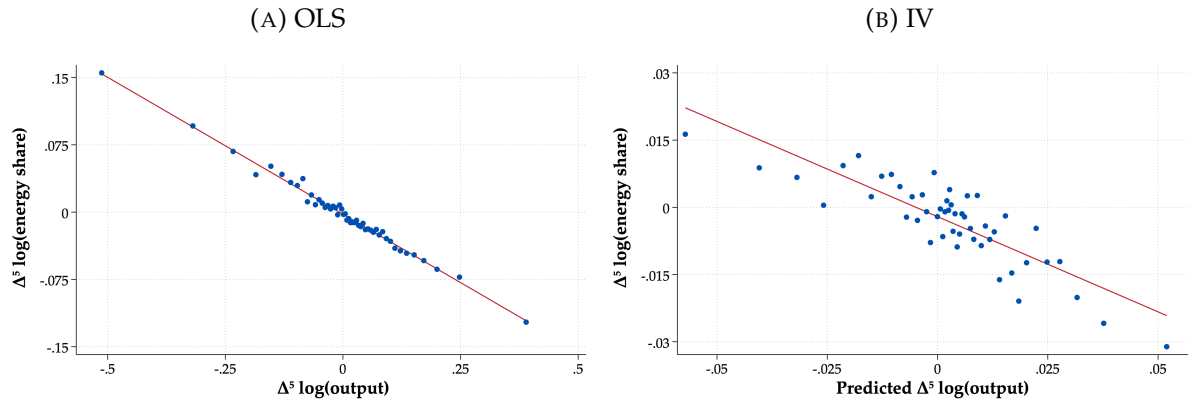
[7] Remember that it is *not* a problem for identification if scale is shifted by a factor-neutral productivity shock. However, it is plausible that if our instrument were correlated with energy-specific productivity, then it would be correlated with overall productivity too, so that this test is informative.

TABLE 1: Energy Expenditure Shares and Firm Growth

	OLS			IV		
	$h=1$	$h=3$	$h=5$	$h=1$	$h=3$	$h=5$
	(1)	(2)	(3)	(4)	(5)	(6)
Δ^h log(output)	-0.315***	-0.319***	-0.313***	-0.481***	-0.409***	-0.451***
	(0.008)	(0.008)	(0.008)	(0.033)	(0.029)	(0.034)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓
F-Stat				382.8	393.8	295.0
R-squared	0.074	0.103	0.114	0.055	0.093	0.088
Observations	392,889	282,570	205,479	370,868	262,870	187,690

Note: This table presents the results of estimating equation (3). Columns 1-3 present OLS results. Columns 4-6 present results where output growth $\Delta^h \log Output_{it}$ is instrumented by the firm-level demand shock defined in (4). Standard errors are clustered at the firm level. ***, ** and * indicate significance at the 1%, 5% and 10% levels, respectively.

FIGURE 5: Energy Expenditure Share and Firm Size: Binned Scatterplots



Note: This figure presents binned scatter plots corresponding to equation (3) for $h=5$. Panel A presents OLS results. Panel B presents results where output growth $\Delta^h \log Output_{it}$ is instrumented by the firm-level demand shock defined in (4).

Main results: Energy intensity. The decreasing relationship between firm size and the relative expenditure share on energy inputs translates directly into a decreasing relationship between firm size and the energy intensity of production. To document this, we employ the same specification (3), but the outcome variable is now defined as $\log(e_{it}/x_{it})$. e_{it} is the firm's total energy usage measured in megajoules. Table 2 presents the results. We let the denominator x_{it} be, respectively, revenue py, value-added va, or physical output y (see Appendix A.1 for details on the construction of firm-level physical output). In columns (1) to (3), the explanatory variable is sales growth (as above), and in columns (4) to (6) we present the same results where the explanatory variable is physical output growth. The coefficient in column (1) is again negative and statistically significant, and greater in magnitude than the relationship estimated for the cost share in Table 1. The magnitudes are instead similar to the cost share results when the denominator is total output (Table B.6). As implied by equation (2), if energy intensity (over revenue) falls faster than the cost share with scale, this either implies rising prices with scale or an increasing profit rate. We show in section 3.2.1 below that there is no evidence for the former

in our data—suggesting rising profit rates with scale.

TABLE 2: Energy-output Ratio and Firm Size

	Δ^5 log(outcome)					
	E/PY (1)	E/VA (2)	E/Y (3)	E/PY (4)	E/VA (5)	E/Y (6)
Δ^5 log(PY)	-0.517***	-0.504***	-0.340***			
	(0.039)	(0.065)	(0.111)			
Δ^5 log(Y)				-0.641***	-0.789***	-0.422***
				(0.102)	(0.201)	(0.087)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓
F-Stat	298.1	261.3	332.0	41.9	19.2	41.9
Observations	182,910	158,062	137,119	137,119	118,490	137,119

Note: This table presents the results of estimating equation (3), where output growth is instrumented by the firm-level demand shock defined in (4). In columns (1)-(3), the endogenous variable is the log change in sales value. In columns (4)-(6), the endogenous variable is the log change in quantity produced. Standard errors are clustered at the firm level. ***, ** and * indicate significance at the 1%, 5% and 10% levels, respectively.

Robustness. We provide a number of robustness exercises and complementary tests for these findings.

Alternative specifications. The results above are robust to looking at energy expenditure shares in logs or in levels, to dropping any industry or year, and to using alternative measures of firm size in place of output.

Composition effects. Our findings could be explained by firms changing the composition of their product portfolio as they grow, tilting it towards less energy-intensive products. Likewise, multi-establishment plants could be reallocating their production across plants. In Table B.10, we repeat our main test when splitting the sample by single- vs. multi-product or single- vs. multi-plant firms, and find highly similar results across subsamples. Hence, our results are not driven by such composition effects. Similarly, Table B.11 shows that our results are highly similar for firms using the exact same material inputs at t and $t+h$, as opposed to firms changing their material input bundle.

Choice of denominator. Appendix Tables B.5 and B.6 show that our results hold for any choice of denominators. One alternative interpretation of the ratio $\frac{w_{it}^E e_{it}}{p_{it} y_{it}}$ being decreasing in firm size would be that larger firms have larger markups or larger input wedges in the language of the misallocation literature. By taking the ratio of energy expenditures over the expenditures on other inputs, we rule out this concern.

Scale elasticity along the firm size distribution. Table B.12 shows the results from estimating equation (3) by bins of ex-ante firm size. These results show that the scale elasticity is significantly negative for both ex-ante smaller and larger firms, and does not systematically vary across bins.

Capacity utilization. One hypothesis is that cost shares changing as firms grow just reflects a change in the utilization of some fixed inputs. Our main results alleviate this concern by looking

at the effects of firm growth on cost shares at relatively long horizons, where it is unlikely that firms can adjust only by changing capacity utilization. In addition, in Table B.13, we proxy for capacity utilization using the share of days worked, and show that our results are similar for firms with and without changes in capacity utilization.

Representativeness of continuing firms. One potential concern with our estimation of η is that our identification strategy exploiting within-firm changes restricts the sample to firms present in two consecutive periods (i.e., continuing firms). Figure B.4 shows that changes in the energy share of continuing firms are representative of the dynamics of the energy share in the whole sample.

Informal sector. We evaluate the robustness of our finding to the inclusion of the informal sector by combining the ASI data with the National Sample Survey on Unorganised Manufacturing Enterprises (NSS). The NSS does not have a panel dimension, hence we report cross-sectional results similar to Figure 4. These results are presented in Table B.18. We document a similarly negative elasticity focusing on the NSS sample, as well as when we run the regression in the pooled sample. It is worth noting that while formal firms represent only 25% of employment, they represent approximately 80% of output and energy expenditures, hence, our results are representative in a value-weighted sense.

US data. These results demonstrate that as firm size increases, the energy expenditure share on energy falls. We discuss the channels by which it does below. Before we proceed, we note that we suspect this relationship between firm scale and energy use holds outside the Indian context, and is representative of a general feature of production functions. In Appendix B.5, we plot energy intensity by firm size for the cross-section of firms in the United States. While we are unable to employ our causal research design for this data, energy intensity falls with scale in a similar manner there.

3.2 Mechanisms

We now dive deeper into the underlying mechanisms behind the results above. We first present descriptive results that provide further detail about how relative energy demand is falling with size. Specifically, we rule out that the results above are coming from falling relative energy prices with scale, or differential input wedges. We then discuss the role of changing intermediate input shares. Finally, we move to discussing evidence about the underlying physical and economic mechanisms.

3.2.1 How Does Relative Energy Demand Fall with Size?

The first channel that we examine is whether the price of energy relative to other inputs evolves endogenously with firm scale. For instance, suppose that firms face firm-specific supply curves in input markets, and that the supply curve for energy is more elastic than for other inputs. Then, one causal effect of firm growth is that the relative price of energy faced by the firm declines. This could introduce both a mechanical change, if the price of energy declines relative to other inputs, and could induce factor substitution towards energy. If energy and other inputs are complements, this leads to a decline in the energy expenditure share. One could also

TABLE 3: Relative Input Price Changes and Firm Growth

	$\Delta^5 \log(w_{it})$			$\Delta^5 \log(\Omega^E_{it}/(1-\Omega^E_{it}))$		
	w^E_{it}	w^M_{it}	w^L_{it}	E/EM	E/EML	E/EMLK
	(1)	(2)	(3)	(4)	(5)	(6)
Δ^5 log(output)	0.009	0.153	0.091***	-0.449***	-0.383***	-0.336***
	(0.009)	(0.103)	(0.018)	(0.044)	(0.042)	(0.042)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓
$\Delta^5 \log(w_{it})$ controls				✓	✓	✓
Observations	183,536	149,343	185,191	147,693	147,693	147,041

Note: Columns (1)-(3) regress changes of log input prices at the firm level on changes in log output, using the instrumental variable strategy discussed in equation (4). Energy prices are constructed as the weighted average of unit prices for electricity and fuel, weighted by expenditure shares. Material prices are constructed as the weighted average of material unit prices. Labor costs are calculated as the total wage bill divided by number of employees. In columns (4)-(6), we regress changes in relative expenditures on changes in log output.

imagine the opposite scenario, where energy is substitutable with other inputs, and the price rises with firm scale.

This suggests a horse race: one can regress the change in the energy expenditure share on the log change in output and the log change in the relative price of energy. In the ASI, we observe input prices at the firm level, and so can construct firm-level indices of the price of energy relative to other inputs. We perform this analysis in Table 3. In columns (1)-(3), we show what happens to input prices at the firm level as firms grow, with firm growth instrumented for using the strategy above. We do not find evidence that energy or material prices rise with growth. We do find that firm-level wages rise with firm growth, consistent with evidence from Mertens and Schoefer (2024) . However, in columns (4)-(6), we show what happens to relative energy expenditures with firm growth, but including as controls the change in firm-level input prices. Controlling for the change in input prices leaves our results essentially unchanged. This suggests that the results documented above are not driven by changes in relative prices as firms grow.

We also find that as firms grow, their energy *mix* from various sources does not change. Specifically, in Table B.16 in the Appendix we show that there is no significant change on average in the share of coal, gas, oil or electricity in the total energy expenditures with firm growth, properly instrumented.

We now investigate which inputs firms use more of when their relative energy demand declines. We show the results of estimating equation (3) where the outcome variable is the energy-to-materials expenditures ratio in Table 4. We find that the energy-to-materials expenditures ratio is strongly decreasing with firm growth. By contrast, we find that the ratio of energy expenditures to wage expenditures is essentially stable as firms grow. These results show that when firms grow their demand for energy falls relative to their demand for material inputs.

One potential explanation for the rise in the material share as firms grow is that firms grow by outsourcing more steps of the production process (Boehm and Oberfield 2023). This could lead to a fall in the energy share via two channels: (i) growing firms may outsource the most

TABLE 4: Size and the Energy-to-Materials Ratio

	E/M			E/L		
	$h = 1$	$h = 3$	$h = 5$	$h = 1$	$h = 3$	$h = 5$
	(1)	(2)	(3)	(4)	(5)	(6)
Δ^h log(output)	-0.527***	-0.453***	-0.486***	-0.025	0.012	-0.052
	(0.042)	(0.037)	(0.041)	(0.040)	(0.035)	(0.040)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓
F-Stat	507.3	426.3	297.5	506.6	425.9	297.5
R-squared	0.078	0.123	0.124	-0.002	0.001	-0.006
Observations	366,569	259,125	184,285	366,422	259,013	184,193

energy-intensive part of the production process; (ii) the production process is split in equally energy-intensive steps, but increasing fragmentation mechanically leads to a decline in the non-material share (and so in the energy share). This would imply that the firm-level decline in the energy share masks a reallocation in energy usage across firms, which would be partly or fully undone in general equilibrium.

To investigate these hypotheses, we perform two tests. To investigate point (i), we construct the average energy intensity of a firm's material inputs. For each product that a firm uses as a material input, we construct an average (sales-weighted) expenditure share on energy at the economy-wide level. We then average these values at the firm level to construct the energy intensity of the firm's material inputs. The results are presented in Table B.14. Energy intensity of inputs does not change as firms grow. This suggests that while firms may be fragmenting their production as they grow, partially accounting for the rise in the relative material share of expenditure, they are not doing it in a way that differentially takes out energy-intensive steps in the production process.

To investigate point (ii), we construct the indirect energy share of a firm as the energy purchases embodied in the material purchases of that firm. Using the sum of the direct and indirect energy shares as an outcome, our coefficient remains negative and statistically significant. This reveals that the energy content of materials purchased does not fully offset the reduction in the firm's own energy share. Contrasting our results using the direct energy share as an outcome with results using the sum of the direct and indirect energy shares as an outcome, we see that considering the indirect share reduces the size of our coefficient by 30-50% depending on the chosen specification. In the theoretical model and the quantitative analysis below, we consider the general equilibrium effects of relative increases in material intensity with aggregate growth and how this shapes the aggregate energy share.

3.2.2 Why Does Relative Energy Demand Fall with Size?

In this section, we present two mechanisms that jointly can explain much of the quantitative magnitude of the results presented above.

The first is the presence of fundamental scaling laws for energy losses that come from physical mechanisms. Second, there exist many energy-saving technologies, but those typically require

TABLE 5: Modeled Scale Elasticity of Energy Intensity Within Technology

Category	Examples of technology	Share of Total (%)	Scale elasticity
Process heating - fired	Furnace	51	−0.15
Process heating - steam	Boiler, dryer	9	−0.03
Process mechanical work	Motor, pump	17	−0.06
Other process		12	0
Non-process	Building heat, refrigeration	11	−0.10
Total		100	−0.10

Note: This table provides estimated scaling laws for energy intensity by the category of usage in manufacturing. All computations are detailed in Appendix E.

largely fixed costs, implying that they tend to be adopted by larger firms. This section provides empirical evidence for these two mechanisms. In section 4, we show that through the lens of our theoretical model, these two forces can quantitatively rationalize the scale elasticity documented above.

To begin with, it is useful to understand how energy is used in manufacturing. Table B.2 shows the share of energy by end-uses in manufacturing. The majority of energy (60%) is used to generate heat, mostly in fired systems like furnaces or kilns (51%). The remainder (9%) represents steam systems (e.g. a steam dryer). The second key use is mechanical work, representing 17% of energy, such as a motor or pump.

A characteristic of energy used as an input is that only a fraction of the energy input goes to serve its specific purpose. In 2018, the U.S. Energy Information Administration estimates that only 52% of the energy entering an industrial plant is applied to its desired use (U.S. Department of Energy 2022). These losses occur in on-site energy generation (steam boilers, on-site electricity generators), in energy distribution (pipes, transmission lines, etc.), in energy conversion (electrical to mechanical, heat exchangers...), and in process energy use (waste heat, flared gases, by-products...). We think of the share of energy lost in the United States as a lower bound for a country like India.

Physical Scaling Laws. Larger production units can achieve lower energy use per unit of output simply by virtue of physical scaling laws. As firms expand capacity, they benefit from geometric relationships that favor reduced losses at higher scales. These effects arise when scaling up a given technology, for instance, by increasing the capacity of a furnace while keeping its design unchanged. We use engineering models of key technologies within each energy usage class to develop estimates of these geometric relationships. Table 5 reports the estimated scale elasticities for representative technologies for each industry end-use. These are read as the percent decline in energy required per unit of output for a given percent increase in output. We provide some intuition on the sources of scale dependence in energy intensity in the main text and leave all derivations to Appendix E.

Heating. One key source of energy losses in heat generation is surface losses, where air (or other fluids) contact the hot surface, absorb heat, and carry it away through circulation or radiation. These losses scale with the surface area of the heated object, while useful heat output scales

with the volume of the heated object. Because of this square-cube relationship—where volume grows as the cube of linear dimensions while surface area grows only as the square—energy losses per unit of output will decline as the amount of output increases with an elasticity equal to $-1/3$. Driven by this force, we estimate a scale elasticity around -0.15 for fired systems like furnaces (where surface losses are most important) and -0.03 for steam systems (where they are much lower).

Mechanical work. Motor efficiency is typically higher for higher horsepower motors because of the lower burden of resistive forces in relative terms. Friction losses (in bearings, seals, and other moving parts) and electrical resistance losses (in windings) represent roughly fixed energy drains that do not scale proportionally with motor size. As motor power increases, useful mechanical work grows faster than these parasitic losses, improving the ratio of output to input energy. We model a combined motor-pump system, and find an elasticity of -0.06.

Non-process. A large share of non-process energy (energy not directly used for producing a unit of output) is used for facilities heating and cooling/refrigeration. These facility-level energy uses exhibit similar scale economies to process heating. In particular for heating, larger facilities have lower surface-area-to-volume ratios, reducing heat loss per unit of conditioned space. The scale elasticity comes from these surface losses, as for heating, at $-1/3$, and we multiply this by the share of surface losses in total energy used, and then by the share of facility heating, ventilation, and air conditioning (HVAC) in total non-process energy use, to arrive at an elasticity of -0.1.

Averaging across end uses, we find that these physical forces can rationalize an average "within-technology" scale elasticity around -0.10. The largest elasticity is for heat generation in fired systems. In Table B.17, we split our sample by the share of energy used for heat generation in fired systems (defined at the industry-level), and we do find more negative scale elasticities for industries most intensive in this energy use.

Technology Investment. Many energy-saving technologies, such as waste heat recovery systems, improved furnace insulation, or combined heat and power units, require substantial upfront investment. Because these costs are largely fixed, they are more easily justified at large scales where savings can be spread over higher output. In addition, smaller firms often face financial constraints and higher borrowing costs, making even profitable efficiency investments difficult to undertake.[8]

In Figure 6, we document this fact using data from the World Bank Enterprise Survey (WBES) module on energy efficiency investments for India in 2014 and 2022. For four major types of energy efficiency investments, we find that the probability to undertake these investments is increasing in firm size. In particular, larger firms (as measured by sales) are much more likely to actively monitor their energy use, and invest in technologies to improve the efficiencies of heating and cooling.

To understand why this pattern arises, in Figure 7 we present evidence that firms in these manufacturing industries face a menu of fixed cost investments, in which energy savings scale

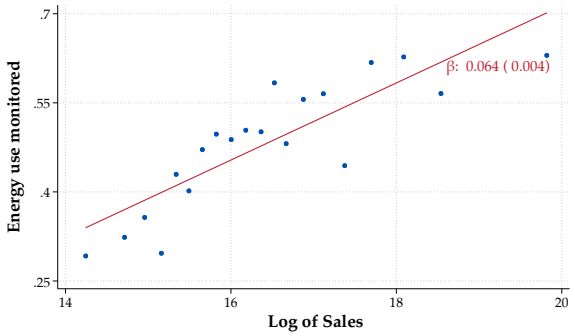
[8]Jordà, Nechio, Phan, and Schwartzman (2026) examines the link between financial conditions and energy-efficient investment, showing that lower discount rates tilt the capital stock towards greater energy efficiency.

in the size of the upfront investment. We gather data from the United Nations Industrial Development Organization (UNIDO) Energy Efficiency and Renewable Energy Technology Compendiums for 2022. These compendiums document, for a set of key industry clusters and locations in India (including foundries, ceramics, brass, dairies and handtools), the availability of technologies to improve energy efficiency at the plant level. Moreover, they are coupled with case studies of these technologies being implemented, and include data on annual energy savings and upfront costs (see details in Appendix A.3).

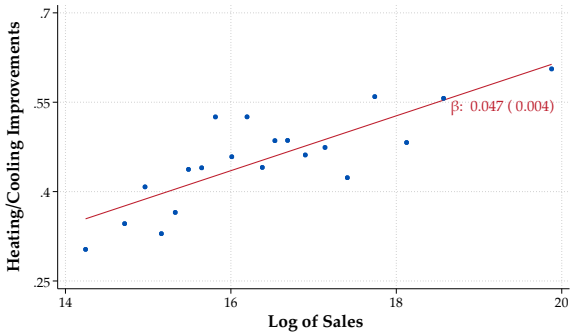
Figure 7 shows how energy savings scale with the size of the investment. Take, for example, the case of foundries. Foundries heat and melt metals to produce castings, for both use in the end-products being manufactured by the firm, and made to order and sold on the market. The majority of energy wastage is in heat loss, and the compendium details several technologies to mitigate heat loss. At the smaller end, one technology is to place a lid mechanism over the mouth of the furnace, which is generally open, thereby substantially reducing surface losses through the mouth. This is costed at INR 0.35 million, and almost halves losses per heating cycle, generating a saving of 293 gigajoules per year for a typical plant. At the larger end, an older cupola furnace can be replaced with an electric induction furnace. This allows to bypass the fuel combustion step where a large fraction of energy is lost, allows for finer temperature control, and reduces the manufacturing rejection rate. This is costed at INR 4 million, and saves 4229 gigajoules annually. A similar range of technologies exists across industries, with lower cost technologies generally giving lower energy savings.

When investment costs rise less than proportionally with scale (i.e., the investment has a fixed cost component), larger firms will optimally choose to adopt higher fixed cost and more energy-efficient technologies, contributing to a negative scale elasticity of energy demand. We formalize this intuition in the model presented in the next section.

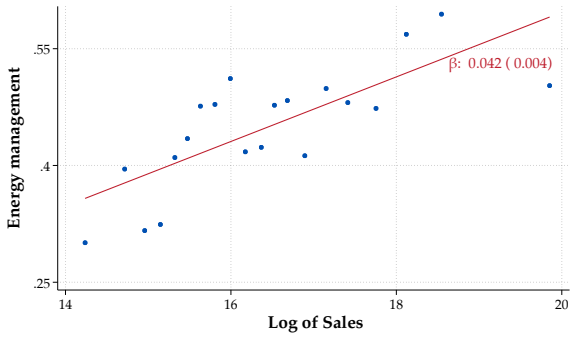
FIGURE 6: Energy Efficiency Investments by Firm Size



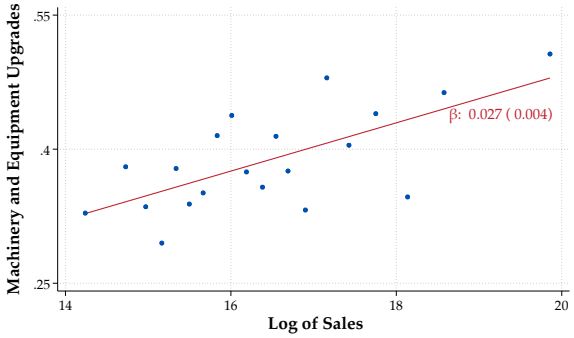
(A) Energy Monitored



(B) Heating/cooling Improvements



(C) Better Energy Management

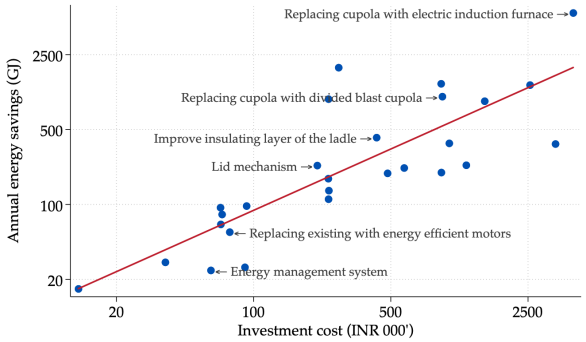


(D) Machinery Improvements

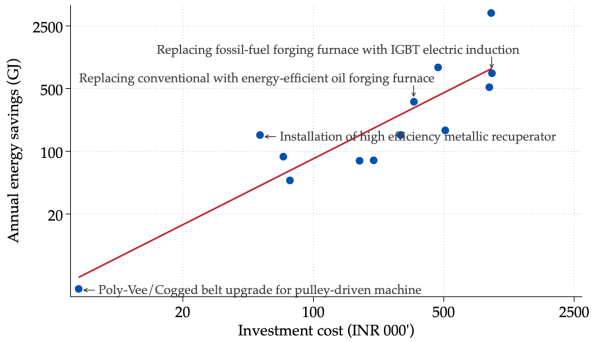
Notes: Data from the WBES energy module. Panel (a): dependent variable indicates whether the establishment monitored its energy consumption over the past three years. Panel (b): dependent variable indicates adoption of heating/cooling improvements. Panel (c): dependent variable indicates whether the establishment is doing active energy management. Panel (d): dependent variable indicates adoption of machinery improvements. Output variable is reported revenue of the firm. Plots show binned means with OLS fits, controlling for 2-digit industry fixed effects and applying analytic weights as specified. Source: World Bank Enterprise Surveys (India 2014, 2022).

FIGURE 7: Energy Technology Menu: Key Industries

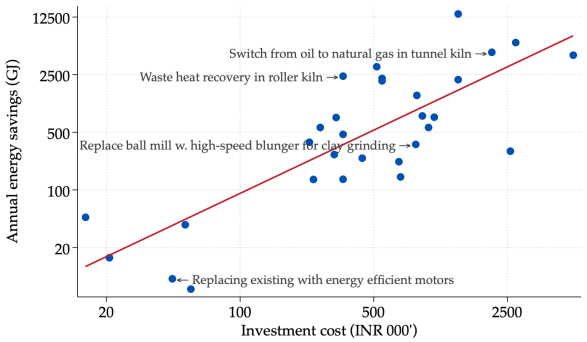
(A) Foundries



(B) Hand-tools



(C) Ceramics



Notes: Data from UNIDO Energy Efficiency and Renewable Energy Technology Compendiums and SAMEEEKSHA. The x-axis is upfront investment cost in INR (000s). The y-axis is annual energy savings in gigajoules. When for a given industry×technology (e.g. "Waste heat recovery in tunnel kiln" in ceramics) we have several estimates (different case studies in the same industrial cluster, or estimates from different clusters), we take the average value. All axes are on a log scale.

4. A Theory of Aggregate Energy Intensity

In this section, we develop a model of manufacturing technologies that turn raw materials into finished products. In this model, at the micro-level, the energy intensity of production for a given technology displays physical returns to scale, in accordance with our discussion above. In addition, firms can invest in a menu of technology choices that improve energy efficiency, by choosing among a menu of fixed costs of different sizes that then lower marginal costs of production.

We first describe the choices of manufacturing firms, and then analyze how these choices affect aggregate energy intensity as economy-wide productivity improves.

We make a number of simple assumptions here for analytical tractability. In particular, we assume fixed proportions between input materials and output, so that they are not at all substitutable with other inputs and do not experience non-homotheticities with scale. We relax these assumptions to bring the model to the data in Section 5.

4.1 Setting

Time is discrete, and indexed by t. There is a continuum of identical households of total mass L_t. These households have static preferences over a final production good, whose output is given by

$$Y_t = \left[\int_0^{N_t} a_i y_{it}^{\frac{\lambda-1}{\lambda}} di \right]^{\frac{\lambda}{\lambda-1}}, \tag{5}$$

where N_t is the total mass of varieties operating in the economy. a_i is a preference shifter that induces firms to produce at different scales. Households own all firms, and supply their labor inelastically to firms for a wage w_t. We take the aggregate price index as our numeraire.

4.2 Firms

Firms produce output y_{it} of a single variety i. For a given technology, the output of the firm is given implicitly by

$$y_{it} = A_{Yt} \min\{m_{it}, \phi \frac{e_{it}}{y_{it}^{\epsilon}}\} \tag{6}$$

where m_{it} is the amount of materials used, and e_{it} is the amount of energy used. A_{Yt} denotes an aggregate productivity shifter for the firm. We let $\epsilon < 0$ govern a physical scale elasticity such that producing larger amounts of output can be done with a lower unit requirement of energy.

The technology of the firm is indexed by ϕ, which denotes an energy efficiency shifter, and a total capacity $\bar{y}$. A higher value for ϕ lowers the unit requirement of energy to produce y_{it}, and higher value of $\bar{y}$ raises the total output that the firm can produce, since we suppose that $y_{it} \leq \bar{y}$.

Before producing, the firm chooses investment in a capital and a labor bundle (e.g. a larger boiler or kiln with attendant staffing levels) that corresponds to a pair $(\phi, \bar{y})$. We suppose that

the cost of this technology bundle is given by

$$\tilde{c}_t(\phi, \bar{y}) = \phi^\gamma \bar{y}^\delta \tilde{C}_t = \frac{1}{\gamma A_{Yt}} \phi^\gamma \bar{y}^\delta w_t^\alpha r_t^{1-\alpha}$$

where w_t is the wage per unit of labor, and r_t is the return on a unit of the capital good. In this way, firms can invest in a larger capital-labor bundle as a fixed cost, in order to both lower their marginal costs of producing their output y_{it}, and to increase the total capacity of their plant.

Firms can be created by hiring ν units of labor. Once a firm is created, it draws its efficiency a_i from a distribution $\Gamma(a)$. We suppose for the baseline that firms exist for a single period.

We suppose that energy can be created at rate Ξ^e using the final good, and similarly for materials and capital at rates Ξ^m and Ξ^k . As such, $p_t^e = 1/\Xi^e$, $p_t^m = 1/\Xi^m$, and $r_t = 1/\Xi^k$.

4.3 Optimality

In what follows, we suppress indexing by i to ease the burden of notation. One can show that the energy efficiency of the technology chosen by the firm is given by

$$\phi = \left(\frac{p_t^e y^{1+\epsilon-\delta}}{w_t^\alpha r_t^{1-\alpha}} \right)^{\frac{1}{1+\gamma}} \tag{7}$$

As such, as long as the physical scale efficiencies are not too strong, and the investment cost doesn't rise too quickly with capacity so that $\epsilon - \delta < 1$, the firm will choose to invest in a more efficient technology as it increases its scale of output.

Firm cost minimization leads to a cost function over output y given by

$$c(A_{Yt}; y) = \frac{1}{A_{Yt}} \left(\frac{\gamma+1}{\gamma} (w_t^\alpha r_t^{1-\alpha})^{\frac{\gamma}{1+\gamma}} (p_t^e)^{\frac{1}{1+\gamma}} y^{\frac{(1+\epsilon)\gamma+\delta}{1+\gamma}} + p_t^m y \right) \tag{8}$$

Importantly, if we define the non-material share of total cost at the firm level as

$$\Omega_t^{NM}(y) \equiv \frac{p_t^e e + w_t l + r_t k}{c(A_Y; y)},$$

then we can derive a useful expression relating this share to firm scale, as in

$$\frac{\Omega_t^{NM}(y)}{1 - \Omega_t^{NM}(y)} = \frac{\gamma+1}{\gamma} (w_t^\alpha r_t^{1-\alpha})^{\frac{\gamma}{1+\gamma}} (p_t^e)^{\frac{1}{1+\gamma}} y^{\frac{\epsilon\gamma+\delta-1}{1+\gamma}} / p_t^m.$$

Note that given $\epsilon < 0$, with $\delta < 1 - \epsilon\gamma$ the non-material share of output is falling. Note also that if we define the energy share of cost as

$$\Omega_t(y) \equiv \frac{p_t^e e}{c(A_Y; y)}$$

we have $\Omega_t(y) = \frac{1}{1+\gamma} \Omega_t^{NM}(y)$, so that the energy cost share inherits movements in the non-

material cost share. A feature of non-homothetic production functions is that returns to scale vary across firms. To see this, define the cost elasticity (the inverse of returns to scale):

$$\mathcal{E} \equiv \frac{d \log C}{d \log y} = \frac{\epsilon\gamma + \delta - 1}{1+\gamma}\Omega^{NM}(y) + 1 \tag{9}$$

When $\frac{\epsilon\gamma+\delta-1}{1+\gamma} < 0$, larger firms, which have lower non-material shares, have lower cost elasticities, i.e., higher returns to scale.

The standard CES demand function gives us an expression for inverse demand as $p(a,y) = ay^{-\frac{1}{\lambda}}Y_t^{\frac{1}{\lambda}}$. We can then write revenue as: $R(y) = ayy^{-\frac{1}{\lambda}}Y_t^{\frac{1}{\lambda}} \equiv ay^{\zeta}Y_t^{\frac{1}{\lambda}}$, where for convenience $\zeta \equiv \frac{\lambda-1}{\lambda}$. Then the problem of the firm is

$$V_t(a) = \max_y \quad ay^{\zeta}Y_t^{\frac{1}{\lambda}} - c(A_{Yt}; y) \tag{10}$$

For a given distribution $\Gamma(a)$, the free entry condition then requires

$$w_t\nu = \int_a V_t(a)d\Gamma(a) \tag{11}$$

Lastly, we define energy intensity of output as energy per dollar of revenue, so that

$$\Theta(a) \equiv \frac{e}{p(y)y} = \frac{1}{aA_{Yt}Y_t^{\frac{1}{\lambda}}}y^{(1+\epsilon)\frac{\gamma+\delta}{1+\gamma}-\zeta}\left(\frac{w_t^{\alpha}r_t^{1-\alpha}}{p_t^e}\right)^{\frac{1}{1+\gamma}}.$$

Accounting for the estimated causal elasticity. In the model, the scale elasticity in energy demand arises from two channels: within technology, energy intensity may decline with scale if $\epsilon < 0$, and there may be a scale elasticity of endogenous technological improvements if $\epsilon + 1 - \delta < 0$.

$$\frac{d \log\left(\frac{e}{y}\right)}{d \log y} = \underbrace{\tilde{\epsilon}}_{\text{Overall}} = \underbrace{\epsilon}_{\text{Within-tech.}} - \underbrace{\frac{\epsilon + 1 - \delta}{1+\gamma}}_{\text{Tech. improvement}}$$

In the data, this elasticity is equal to approximately -0.4 (Table 2, column 6). Can the two mechanisms quantitatively account for this empirical estimate? In the previous section, we show that a reasonable estimate of ϵ is -0.10 (Table 5). That is, the within technology channel can rationalize approximately 25% of the observed effect. The data on technology menus allows us to estimate the cost elasticity to energy savings γ (see Appendix F). Our baseline estimate of γ is $\gamma \approx 0.10$. Finally, we use data on equipment prices by capacity to estimate $\delta \approx 0.6$. This implies that the technological improvement channel can account for an elasticity of -0.27, approximately 70% of our estimated effect.

This analysis yields two conclusions. First, the two mechanisms together can account for our estimated causal elasticity. Second, most of the effect likely comes from technological switching as firms grow, as opposed to the within-technology returns to scale.

4.4 Equilibrium

We now formally define an equilibrium.

Definition 1. *An equilibrium is a wage w_t, and an allocation $\{Y_t, N_t, E_t, M_t, K_t\}$, such that*

1. Given prices and aggregate output, firms solve (10), with resulting input choices $\{l_t(a), e_t(a), m_t(a), k_t(a)\}$ and output choice $y_t(a)$,

2. The free entry condition (11) holds,

3. The labor market clears,

$$N_t \int_a l_t(a) d\Gamma(a) + \nu N_t = L_t.$$

4. The energy market clears,

$$E_t = N_t \int_a e_t(a) d\Gamma(a)$$

5. The materials market clears,

$$M_t = N_t \int_a m_t(a) d\Gamma(a)$$

6. The capital market clears,

$$K_t = N_t \int_a k_t(a) d\Gamma(a)$$

7. The goods market clears,

$$Y_t = Consumption_t + K_t/\Xi^k + M_t/\Xi^m + E_t/\Xi^e \tag{12}$$

4.5 Characterization

We now discuss how economic growth affects movements in aggregate energy intensity. We will consider general growth in Hicks-neutral TFP, A_Y, as exogenous technical progress. Due to the lack of dynamics, we will suppress indexing on t.

We begin by characterizing the ratio of energy to gross output, and then show how this translates into movements in energy intensity in GDP. The energy-to-gross-output ratio, which we call *gross energy intensity*, is given by

$$\bar{\Theta} = \frac{N \int_0^\infty e(a) d\Gamma(a)}{Y}$$

This can be written as

$$\bar{\Theta} = \int_0^\infty \underbrace{\Theta(a)}_{\text{Micro energy intensities}} \cdot \underbrace{\frac{p(a)y(a)}{Y/N}}_{\text{Revenue weights}} d\Gamma(a) \tag{13}$$

where the micro energy intensity of production is the ratio of energy in joules to dollars of revenue. As such, for any change in the energy expenditure share due to changing fundamentals, we have

$$d\bar{\Theta} = \int_0^\infty \underbrace{d\Theta(a) \cdot \frac{p(a)y(a)}{Y/N} d\Gamma(a)}_{\text{Micro intensity changes}} + \int_0^\infty \underbrace{\Theta(a) \cdot d\frac{p(a)y(a)}{Y/N} d\Gamma(a)}_{\text{Reallocation}} \tag{14}$$

Given the structure of production above, the relationship between micro energy intensity $\Theta(a)$ and output y is log-linear, and we can write :

$$d\log\Theta(a) = \left(\frac{(1+\epsilon)\gamma+\delta}{1+\gamma} - \zeta\right) d\log y(a) - d\log\left(A_Y Y^{\frac{1}{\lambda}}\right) + \frac{1}{1+\gamma} d\log\left(\frac{w^\alpha r^{1-\alpha}}{p^e}\right) \tag{15}$$

As such, there are three forces that change the micro-cost share when Hicks-neutral productivity A_Y increases. First, firm size increasing can reduce energy intensity. This operates via decreasing the energy cost share $\Omega(a)$, through both the physical scaling coefficient $\epsilon < 0$ for a given technology, and because investment in energy-efficient technology increases in firm output y. This is partially offset by the fact that energy intensity is measured relative to revenue, and given the firm-specific demand curve, revenue does not rise one for one with output (which is why ζ appears in the first term).

Second, increasing Hicks-neutral productivity increases output y holding inputs constant, and so causes a fall in energy intensity directly across all firms in the economy. There is similarly an aggregate demand effect through the shifter $Y^{\frac{1}{\lambda}}$ which acts to raise all output prices for given outputs, further depressing energy intensity over revenue.

Lastly, increasing Hicks-neutral productivity at the economy-wide level may change relative factor prices between labor, capital and energy, as in standard models of growth.

We now use these results to understand how the aggregate energy intensity of gross output changes as aggregate productivity A_Y changes. Combining (14) and (15) , we can arrive at the following proposition.

Proposition 1. *The change in the aggregate gross energy intensity $\bar{\Theta}$ with a rise in factor-neutral productivity A_Y is*

$$\begin{aligned} d\log(\bar{\Theta}) = & \underbrace{\left(\frac{(1+\epsilon)\gamma+\delta}{1+\gamma} - \zeta\right) \mathbb{E}_R[d\log y(a)]}_{\textit{Micro scale up}} \\ & \underbrace{- d\log\left(A_Y Y^{\frac{1}{\lambda}}\right) + \frac{1}{1+\gamma} d\log\left(\frac{w^\alpha r^{1-\alpha}}{p^e}\right)}_{\textit{Aggregate Changes}} \\ & \underbrace{+ \left(\mathbb{E}_R[d\log p(y)y(a)] - d\log(Y/N)\right) + \frac{(1+\epsilon)\gamma+\delta}{1+\gamma} Cov_R[\frac{\Theta(a)}{\bar{\Theta}}, d\log y(a)]}_{\textit{Reallocation}} \end{aligned} \tag{16}$$

where expectations $\mathbb{E}_R$ are taken over revenue weights.

The first term is the effect of average energy cost shares decreasing due to increases in output, discussed above. It includes the demand side curvature $\zeta = \frac{\lambda-1}{\lambda} < 1$, since increasing output y will push down the output price, and since we are working with energy intensity of revenue this partially pushes back against the scale up channel. The second line contains the effect of factor price changes working against Hicks-neutral productivity, as highlighted above. The last term is the effect of reallocation. Our production function implies falling returns to scale as output increases. As such, small firms will increase output more in response to a general rise in TFP. Reallocation will be towards smaller firms, that have higher energy intensities, and the final line will be positive. As such, general equilibrium reallocation will act to push up the aggregate energy share.

To understand the operation of the general equilibrium forces in the second line of (16), we can use the free entry condition to derive a simple relationship between the changes in the wage and aggregate TFP movements.

Proposition 2. *The general equilibrium response of wages to a factor neutral productivity shock is*

$$\frac{wL}{Y} d\log(w) + \frac{\Pi}{Y} d\log(A_Y) = d\log\left(A_Y Y^{\frac{1}{\lambda}}\right)$$

The logic of this result is straightforward, and indeed applies in a broader class of models than that presented here. Given that energy, capital and materials are produced linearly with the numeraire final goods, there is truly only one variable factor price operating in the model, that of wages. Given that the entry cost is denominated in labor, when greater productivity causes variable profits to rise for fixed factor prices, in general equilibrium the wage must rise to offset this and restore free entry. The extent to which it does so is reflected in the equilibrium labor share of gross output: a more labor intensive economy will see a smaller rise in the wage from TFP growth, all else equal.

We can use the parameters above and this result to get a rough size of the effect in the second line. Given $\alpha/(1+\gamma)$ is roughly equal to the labor share of gross output, and the rental rate of capital and electricity price do not change in equilibrium with an increase in TFP, we are left with the profit share Π/Y times the increase in TFP. In the calibrated model below, the profit share is approximately equal to 20%. Meanwhile, the average change in output $\mathbb{E}_R[d\log y(a)]$ will be significantly greater than the change in TFP, suggesting the second line is dominated quantitatively by the first line.

This answers the question of what happens to gross energy intensity with a rise in aggregate productivity. As for *net* energy intensity, or energy-to-value-added/GDP, that is given by

$$\frac{Energy}{GDP} = \bar{\Theta}\frac{Y}{Y - p^e E - p^m M} = \bar{\Theta}\frac{1}{1 - p^e\bar{\Theta} - p^m M/Y} \tag{17}$$

This falls even faster than gross energy intensity because the denominator increase with aggregate output. In particular, the ratio M/Y falls with aggregate output. The key intuition is that with a Leontief production function, TFP growth implies that less units of material are required

to produce a unit of output (without any potentially offsetting effect due to substitution).[9]

5. Quantifying the Mechanism

We now add several quantitative extensions that allow the model to be taken to the data. In particular, we allow a more flexible cost function that nests our model above as a special case.

We begin with equation (8), and make four strict generalizations. First, we allow energy and the capital-labor bundle to have a non-unitary elasticity of substitution, denoted by σ_e. Second, we allow materials to have a non-zero elasticity of substitution with productive factors, denoted by σ_m. Third, we allow for general scale elasticities that may be different between energy and the capital-labor bundle, denoted ϵ_e and ϵ_l. These parameters regulate how the shares of energy and capital/labor evolve with firm scale. Finally, we allow for a parameter, ϵ_m, that controls the overall returns to scale of the production function. This leaves us with a generalized non-homothetic CES cost function[10] given by

$$\mathcal{C}(y) = A_{Yt}^{-1}\left((A_{Mt}^{-1}\ p_t^m y^{\epsilon_m})^{1-\sigma_m} + \left((A_{Et}^{-1}\ y^{\epsilon_e}\ p_t^e)^{1-\sigma_e} + (y^{\epsilon_l} w_t^{\alpha} r_t^{1-\alpha})^{1-\sigma_e}\right)^{\frac{1-\sigma_m}{1-\sigma_e}}\right)^{\frac{1}{1-\sigma_m}} \tag{18}$$

In addition, we allow for the time-varying aggregate shifters A_{Yt}, A_{Et} and A_{Mt}, as exogenous forms of technical progress. A_{Yt} is a factor neutral shifter, while A_{Et} and A_{Mt} are factor-augmenting forms of technical change which benefit energy and materials, respectively.

5.1 Estimation Strategy

Our estimation strategy proceeds in three stages. First, we use standard estimates from the firm dynamics literature to discipline the demand elasticity ($\lambda = 4$) (see e.g. Peters and Walsh (2022)). Second, we calibrate all of our production function parameters using the empirical estimates from the micro data. We obtain the parameters ϵ_e and ϵ_l from our scale elasticity estimation in section 3. We are able to identify these parameters in our model because we have closed-form analytical mappings from the estimated elasticities in the micro data to the production function parameters. In addition, we estimate the two elasticities of substitution in production, σ_m and σ_e, which we obtain by estimating the response of relative energy demand to shocks to the relative price of the energy inputs.

Third, given our calibrated production function parameters, we quantitatively calibrate the remaining model parameters to match six targeted moments from the micro data: (1) the aggregate energy share, (2) aggregate materials share, (3) aggregate labor share, (4) average firm size, (5) the average degree of returns to scale in the economy, and (6) a measure of the dispersion in firm size. These six moments identify the productivity shifters in the base year 1990 ($A_{E,1990}$, $A_{M,1990}$), the labor share of the inner-most nest α, the entry cost parameter ν, the parameter ϵ_m, which governs overall returns to scale in the model, and the shape parameter θ

[9] See proof in Appendix C.4. With the CES aggregator with $\lambda < +\infty$ (i.e., when the number of firms affects aggregate output), the proof requires that the number of firms N increase with aggregate output growth. This condition is always satisfied when the aggregate non-material share declines with firm scale, as we document empirically.

[10] See Lashkari et al. (2024), Eckert, Ganapati, and Walsh (2025) and Trottner (2019) for other recent examples of production functions of this form in aggregate analysis.

from the Pareto distribution of firms' idiosyncratic preference shifters.

Scale Elasticity Parameters (ϵ_l, ϵ_e). The parameters governing the scale elasticities ϵ_e and ϵ_l come from our empirical estimates of (i) how the *relative energy share* changes with firm size ($\eta_{E/ML}$), and (ii) how the *energy share relative to labor* changes with firm size ($\eta_{E/L}$). In the model, the scale elasticities are defined with respect to firm physical output (as opposed to revenues); we estimate these elasticities by using our demand shifter as an instrument for the change in firms' physical output in Table B.9. We obtain:

$$\eta_{E/ML} \equiv \frac{d \log\left(\frac{\Omega^E}{\Omega^M+\Omega^L}\right)}{d \log y} = -0.51$$

$$\eta_{E/L} \equiv \frac{d \log\left(\frac{\Omega^E}{\Omega^L}\right)}{d \log y} = 0$$

We derive analytical formulas for these scale elasticities (see details in Appendix D).

$$\eta_{E/ML} = (1-\sigma_e)(\epsilon_e - \epsilon_l)\left[1 - \frac{\Omega^M \Omega^{E,ELK}}{\Omega^M + \Omega^L}\right] + B\left[\frac{\Omega^M}{\Omega^M + \Omega^L}\right] \tag{19}$$

$$\eta_{E/L} = (1-\sigma_e)(\epsilon_e - \epsilon_l) \tag{20}$$

where $B = (1-\sigma_m)\left[\epsilon_e \Omega^{E,ELK} + \epsilon_l(1-\Omega^{E,ELK}) - \epsilon_m\right]$, and Ω^M is defined as the average cost weighted share of materials expenditures in total costs in the base year of 1990. Ω^E is defined similarly and $\Omega^{E,ELK} \equiv \frac{\Omega^E}{\Omega^E+\Omega^L+\Omega^K}$.

Given these formulas for our scale elasticities, we can solve for ϵ_e and ϵ_l in closed form. For a given value of ϵ_m (which we calibrate to match overall degrees of returns to scale as explained below), the scale elasticities on energy and labor are:

$$\epsilon_l = \epsilon_m - \frac{\eta_{E/L}\Omega^{E,ELK}}{1-\sigma_e} + \frac{\eta_{E/ML} - \eta_{E/L} + \frac{\Omega^M}{\Omega^M+\Omega^L}\Omega^{E,ELK}\eta_{E/L}}{\frac{\Omega^M}{\Omega^M+\Omega^L}(1-\sigma_m)} \tag{21}$$

$$\epsilon_e = \epsilon_l + \frac{\eta_{E/L}}{(1-\sigma_e)} \tag{22}$$

Since $\eta_{E/L} = 0$ we set $\epsilon_l = \epsilon_e$. Then, given our calibrated value of $\epsilon_m = 1.3$, our empirical estimate of $\eta_{E/ML} = -0.51$, and cost shares of $\Omega^{E,ELK}_{1990} = 0.315$, $\Omega^M_{1990} = 0.714, \Omega^L_{1990} = 0.06$ as measured in the micro data, we find that $\epsilon_l = \epsilon_e = 0.27$. The values for these scale parameters are small relative to the overall returns to scale parameter ϵ_m, and imply that energy intensity reduces with firm scale.

Elasticity of Substitution (σ_e, σ_m). We estimate the elasticity of substitution between energy and the capital-labor bundle, σ_e, and between materials and other factors, σ_m. More details of our procedure for this is documented in Appendix D.2.1. We identify the elasticity of relative energy demand in response to changes in the relative price of energy using an instrumental variables approach. We use a shift-share instrument which interacts changes in the aggregate prices of different fuels (coal, electricity, oil) and firms' ex-ante exposure to these different fuel

TABLE 6: Model Calibration

	Model Parameters	Value	Targeted Moments	Data	Model
$A_{E,1990}$	Energy shifter	24.43	Energy share Ω^E_{1990}	0.09	0.09
$A_{M,1990}$	Materials shifter	0.75	Materials share Ω^M_{1990}	0.71	0.71
α	Labor share in inner-most nest	0.28	Aggregate Labor share Ω^L_{1990}	0.06	0.06
ν	Entry cost	50.83	Avg. workers/firm	58	59
ϵ_m	Overall RTS parameter	1.30	Avg. returns to scale	1.00	1.01
ϵ_l	Scale elasticity of labor/capital bundle	0.27	Output elasticity $\eta_{E/L}$	0.0	0.0
ϵ_e	Scale elasticity of energy	0.27	Output elasticity $\eta_{E/ML}$	-0.51	-0.51
θ	Pareto shape - preference shifters	4.10	Size distribution (p75-p25 $\log(\Omega^M)$)	0.17	0.18
σ_e	E.O.S. energy and labor	0.59	Directly estimated		
σ_m	E.O.S. materials and other factors	0.47	Directly estimated		
λ	Demand elasticity	4	Set externally		

types. Given our estimates of regressing changes in the log energy share relative to the labor and capital bundle, $\Delta \log \Omega^{E/L+K}$, on $\Delta \log p_e$ in Table D.1, we find a coefficient of $(1-\sigma_e) = 0.41$. This provides us with an estimate of the elasticity of substitution of approximately $\sigma_e = 0.59$, which suggests that energy and other inputs are complements. In addition we estimate the elasticity of the energy-to-materials expenditure ratio to the price of energy, finding estimates similar in magnitude to those obtained using the labor-capital bundle as the denominator. Analytical mappings from these two elasticities to the structural parameters σ_e, σ_m then imply $\sigma_m = 0.47, \sigma_e = 0.59$. Details are provided in Appendix D.2.1 and D.2.2.

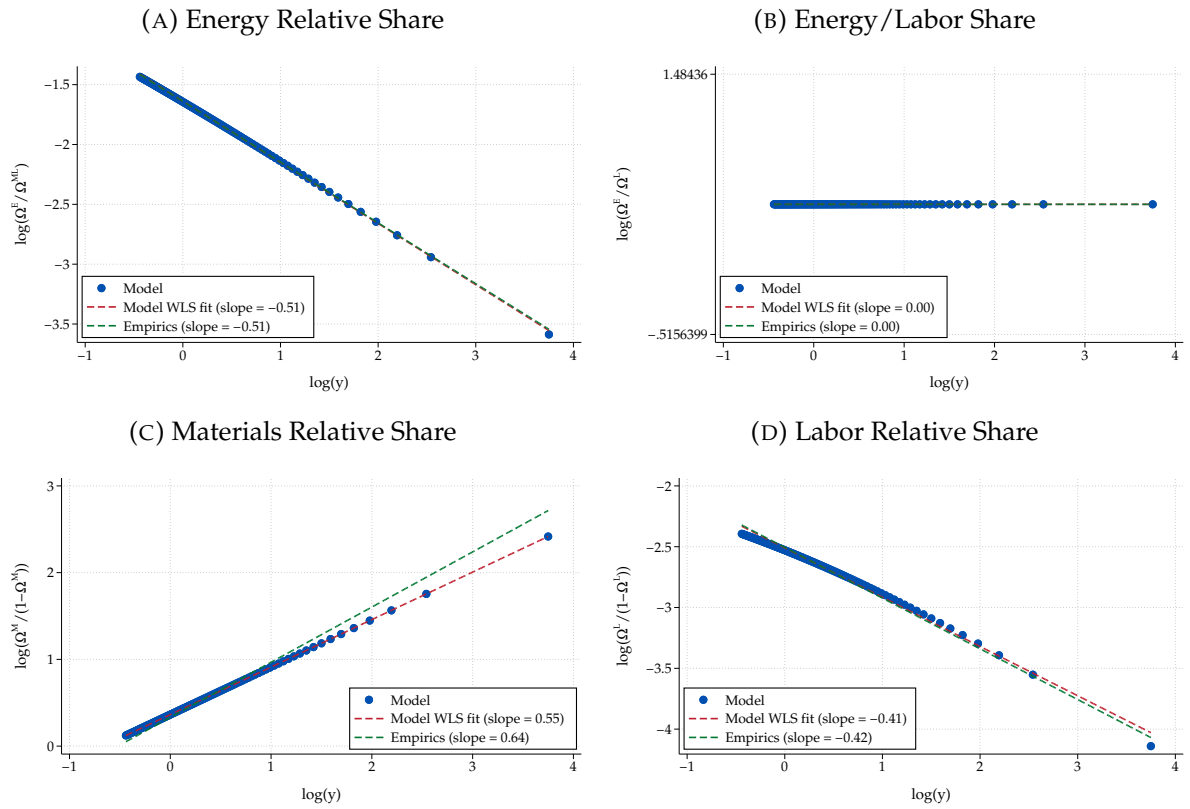
Other Parameters. We calibrate the exogenous shifters A_M, A_E in the base period to match the expenditure share of materials $\Omega^M_{1990} = 0.71$ and energy $\Omega^E_{1990} = 0.09$, respectively, in the year 1990. Similarly, α is calibrated to match the expenditure share of labor in the base year, $\Omega^L_{1990} = 0.06$. We normalize $A_Y = 5$.

To calibrate the entry costs, ν, measured in units of labor, we use the average firm size, defined as total workers over total firms, which is 58 workers in the base year of 1990.

The parameter ϵ_m which scales with the materials share regulates the *overall degree of returns to scale* at the firm-level. In Appendix D.2.3, we estimate average cost elasticities by using our demand shifter D^h_{it} to trace out how cost moves with increases in scale. Across several specifications, we estimate an average return to scale of approximately 1. To match this in the model, we compute for each firm the average cost elasticity $1/RTS_i = dlogC_i/dlogy_i$, and then take a cost-weighted average across firms. We calibrate ϵ_m to match this target of $\overline{RTS} = 1$.

Finally, we assume that the distribution of the idiosyncratic preference shifters, a_i, follows a Pareto distribution with shape parameter θ and scale parameter a_{min}. Firm heterogeneity matters to the extent that it shapes the reallocation term highlighted in equation (16). This reallocation term depends on the dispersion in firm-level growth rates in response to a change in aggregate productivity. This is itself fully determined by the dispersion in returns to scale, with returns to scale being functions of the firm-level input shares. We thus calibrate θ as follows: we project the log material share on firm size (since in our model input shares solely depend on firm size), and then target the interquartile range (p75-p25) of these fitted material shares. Finally, we choose a value of the scale parameter of the Pareto distribution, $a_{min} = 1.78$, such that no firms in our model have negative profits.

FIGURE 8: Firm-level Elasticities: Model vs. Data



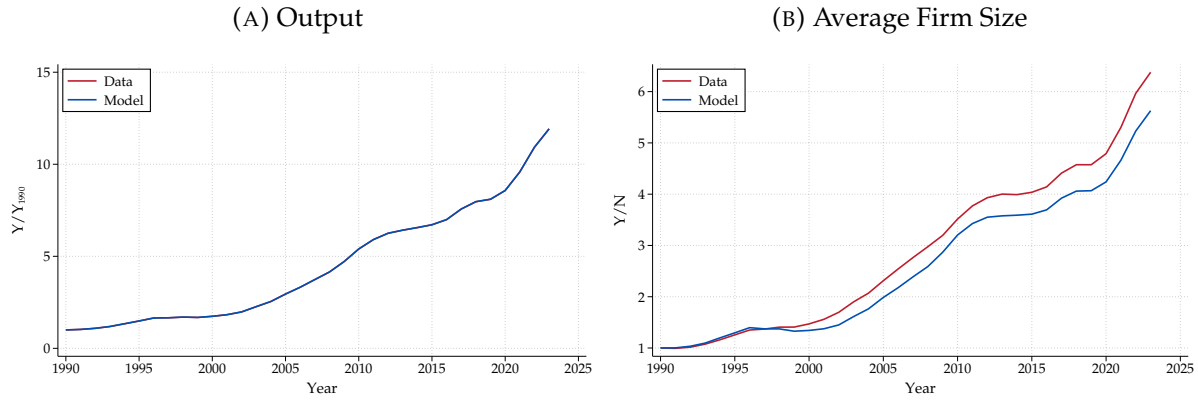
Note: This figure presents the empirical slope estimates (represented by the green dashed line) from quantity regressions in Table B.9 against the predictions of the model. Each blue dot in the graph represents a grid point in the state vector of preference shifters from the quantitative model. The model slope is represented by the red dashed line, and is estimated using weighted least squares with weights corresponding to the total cost share of each node in the model.

The calibrated parameters are presented in Table 6.

Model Fit. At the firm level, the model matches the micro elasticities of input expenditures with respect to firm scale. Since we have calibrated ϵ_e and ϵ_l to match $\eta_{E/ML}$ and $\eta_{E/L}$, we have that, at the firm level, a change in firm log output leads to a decrease in the log energy share with a slope coefficient of -0.51, which matches the estimates from the micro data (Figure 8(A)). The energy share relative to the share of labor is invariant to firm size, which is a targeted moment in our model (panel B). In addition, we can match the elasticity of the materials expenditures relative to non-materials expenditures $\eta_{M/ELK}$ closely in our model, without explicitly targeting the relationship. We see from Figure 8(C) below, that the expenditure share of materials is rising in firm size. Finally, our model's estimates of the relative labor share across the firm distribution declines with firm scale, and closely resembles the relationship documented in the data (Figure 8(D)).

To assess the ability of our model to capture the dispersion in returns to scale—which is crucial for the reallocation effect—we compare the output-weighted and unweighted average returns to scale in the model and in the data. We calibrate the model so that the cost-weighted average returns to scale in 1990 in the model match our estimated weighted returns to scale of about 1 in the data (see appendix D.2.3 for our estimation results). The unweighted returns to scale,

FIGURE 9: Aggregate Output and Average Firm Size: Model vs. Data



Note: This figure plots the model implied trajectory of average firm size in response to an increase in output driven by scaling the aggregate productivity shifter A_Y. The solid blue follows the model calibration assuming a non-homothetic CES production function described in section 5.1. The red line represents the aggregate ASI time series. The time series trends in the model assume that the relative prices p_m, p_e, r are held constant and the labor supply growth matches the growth in the labor force measured in the UN data. The productivity shifter A_Y in the model is chosen each period such that growth in output in the model matches growth in output in the data (as shown in panel A).

which we do not explicitly target, are 1.08 in 1990 in the model compared to 1.13 estimated in the data. These estimates in both the model and data suggest that the returns to scale are higher for small firms compared to large firms. The model matches this pattern both qualitatively and quantitatively: the difference between the unweighted and weighted returns to scale of 0.08 in the model exactly matches the difference of 0.08 in the data.

5.2 Results

With our estimated model, we ask: what is the effect of growth on the aggregate energy share and the energy intensity of production?

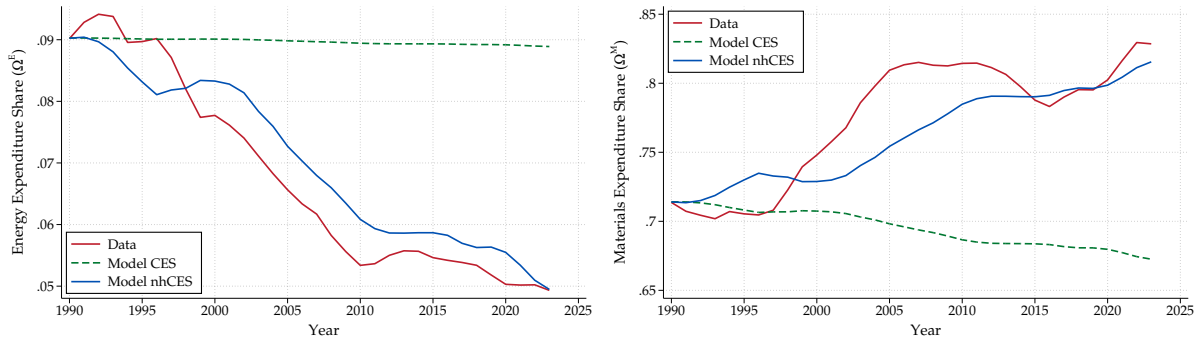
We model growth as driven by growth of the labor force L_t and productivity growth A_{Yt}. Specifically, we feed in the growth rate of the labor force L_t (taken to be the population in India aged 15-64, from UN data) and estimate A_{Yt} to exactly match the aggregate growth rate of output in the ASI data over 1990-2023.[11] Figure 9(A) shows that we (by construction) perfectly match this series. We analyze how the aggregate expenditure share of energy and energy intensity change in our model in response to output growth. In doing so, we solve the full general equilibrium solution, including endogenous wages and firm choices. In our baseline calibration, we hold fixed the technical efficiencies Ξ^k, Ξ^m and Ξ^e, which govern the relative prices of capital, materials, and energy, respectively. We return to these momentarily.

As shown in Figure 9(B), the model closely mirrors the increase in average firm size in the data. In the model, this emerges because entry costs are paid in units of labor, and hence rise with TFP growth, limiting the entry of new firms. This parsimonious modeling, with the entry cost calibrated to 1990 data, fits the Indian time series remarkably well. Fitting this time series is very important for our purpose, since the increase in average firm size is the key forcing variable of

[11] Note that although the sample of the ASI micro data used in our empirical estimates starts in 1998, ASI aggregate data extends back to 1990, which we use as a benchmark for our model.

FIGURE 10: Energy and Materials Cost Shares: Model vs. Data

(A) Energy Share (B) Materials Share



Note: This figure plots the model implied trajectory of energy and materials cost shares in response to an increase in output driven by scaling the aggregate productivity shifter A_Y. The solid blue follows the model calibration assuming a non-homothetic CES production function described in section 5.1. The dashed green line assumes a homothetic (CES) production function. The red line represents the aggregate ASI time series. The time series trends assume that the relative prices p_m, p_e, r are held constant and the labor supply growth matches the growth in the labor force measured in the UN data. The productivity shifter A_Y in the model is chosen each period such that growth in output in the model matches growth in output in the data (as shown in figure 9(A)).

our scale elasticity channel.

Figure 10(A) presents our key result: our estimated model predicts a 45% decline in the aggregate energy expenditure share on energy from 1990 to 2023. This essentially corresponds to all of the decline in the energy expenditure share during this time period, as illustrated by the gap between the blue and red lines in this figure. This decline in the energy expenditure share is entirely attributable to the non-homotheticity in the production function and firms scaling up. To isolate this mechanism, the dotted green reports a counterfactual with a homothetic CES production function (obtained by setting $\epsilon_m = \epsilon_e = \epsilon_l$ in equation (18)). In this case, the energy share barely changes (by less than 2%). Note that the change is not exactly 0 because TFP growth raises the wage relative to other input prices (equal to the final good's price) and substitution is non-unitary.

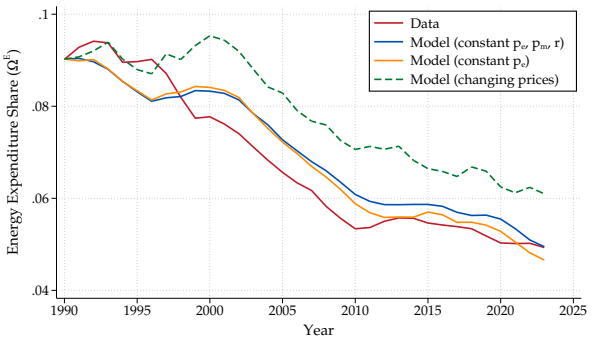
The non-homothetic CES model predicts a macro-elasticity of the relative energy expenditure share with respect to total output equal to -0.41 (Table D.4). We can also estimate the macro-elasticity of the energy share relative to materials and labor, the macro-counterpart to our firm-level elasticity $\eta_{E/ML}$. We find a macro-elasticity equal to -0.45, very close to $\eta_{E/ML} = -0.51$. This implies that general equilibrium forces are quantitatively small compared to the direct micro scaling up channel.

In addition, the model matches well the trend in the share of materials expenditures (Figure 10(B)). As with energy, this pattern is entirely driven by non-homotheticity. Matching materials shares ensures that the model captures both direct and indirect energy use, and therefore provides a consistent account of aggregate energy demand.

We then consider a scenario in which the relative price of energy changes in our model to match the rise in the relative energy price observed in the data from 1990-2023 (Figure B.2). Figure 11 displays three model-predicted time series. In blue, we keep the price of capital, intermediates

FIGURE 11: Impact of Changing Energy Price on Modeled Energy Share

(A) Changing Energy Prices



Note: This figure plots the model implied decline in the aggregate energy share under four scenarios. The blue line simulates the model given that there are no changes in relative prices from 1990-2023 horizon (i.e. holding exogenous technical efficiencies fixed). The orange line allows the price of materials and the price of capital to match the data series. The green line simulates the model also assuming that the relative price of energy to the final good follows the trend seen in the data. The red line represents the aggregate ASI time series for manufacturing.

and energy equal to that of the final good, reproducing Figure 9(C). In orange, we match the time series for the relative price of capital and intermediates, but leave the relative price of energy constant. The model prediction remains very similar. In dotted green, we show the model prediction when we additionally match the rise in the relative energy price. This causes the overall energy share to decline less than in a scenario assuming constant prices. The reason is that energy is estimated to be complementary with other inputs in our model, so a rising price, all else equal, implies an *increasing share* of energy expenditure at the micro level, and this carries over to the aggregate. However, even after including the energy price changes, the model still explains about 70 percent of the overall decline in the energy cost share. This shows that—had the relative price of energy in India stayed constant—the energy expenditure share would have fallen even more.

Finally, we analyze what our model predicts for the path of energy intensity from 1990-2023 compared to the data. In Figure 12(A), energy intensity is defined as the physical quantity of energy consumed in production over aggregate output (E/Y). Allowing for the price of energy to match the rising trend in energy price seen in the data, we find that the model does a good job in matching the aggregate trends. Similarly, when energy intensity is measured as physical energy quantity over value added, our model incorporating both the scale dependence of energy cost shares and the changing prices of energy closely matches the time series (Figure 12(B)).

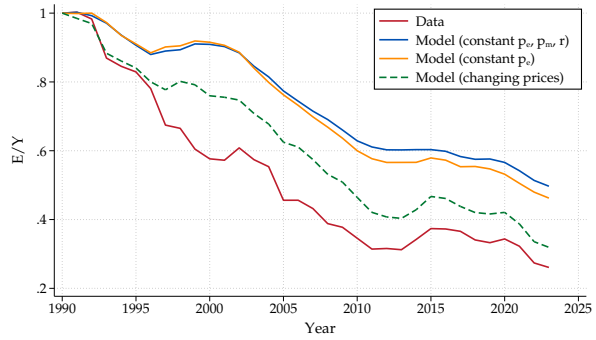
6. ENERGY AND EMISSIONS IN INDIA: 2025-2050

In this final section we consider the implications of our findings for the future trajectory of energy use and emissions in India. For this exercise, we analyze what GDP and labor supply growth projections through 2050 would imply for total energy demand in the medium-term.

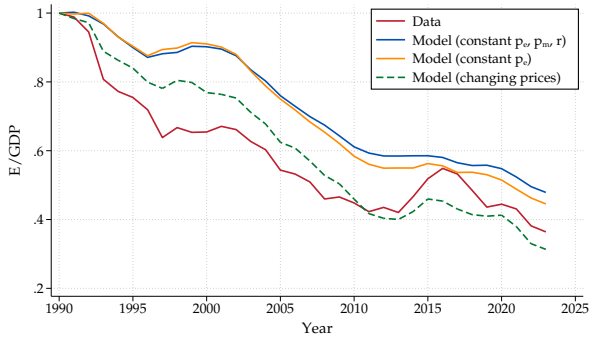
In order to project forward, we assume a continued 2% growth in the labor force, which matches the average growth from 1990-2025. For aggregate output growth, we use the IMF's GDP

FIGURE 12: Energy Intensity: Model vs Data

(A) Energy / Output



(B) Energy / Value Added



Note: This figure plots the model implied paths of energy intensity measured as Energy/Aggregate Output in panel (A), and Energy/Value Added, or GDP, in panel (B). The blue line simulates the model given that there are no changes in relative prices besides wages over the 1990-2023 horizon. The orange line simulates the model assuming that the relative prices of capital and materials to the final good follow the trend seen in the data. The green line additionally assumes that that the relative price of energy to the final good follows the trend seen in the data. The red line represents the aggregate time series for energy over output and energy over value added within manufacturing.

growth projections from 2025-2030, and assume that the annual growth rate in GDP from 2030-2050 matches the compound annual growth rate from 2025-2030. Lastly, we assume for this exercise that our structural analysis of the manufacturing sector is representative of trends in energy expenditure for the economy as a whole.

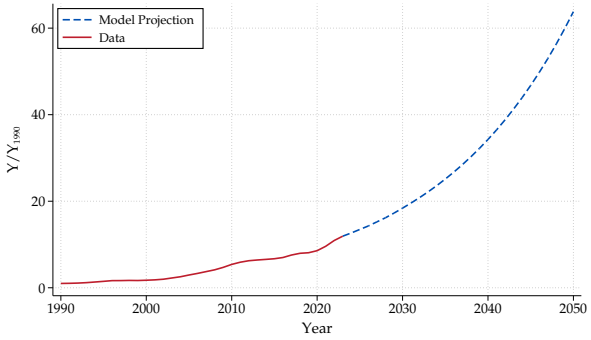
Figure 13 shows the projections of both the energy expenditure share and output to 2050 through the lens of our model. The model predicts an additional two percentage-point decline in the energy share given the growth projections for India through 2050. The figures reveal that catch-up growth in LMICs over the next few decades could be an important driver of the reduction in the energy intensity of production.

In addition, Figure 14 shows the projected path of energy demanded under two scenarios. The first scenario predicts future energy demand according to our mechanism, embedding the scale dependence of energy use. In this scenario, represented by the dashed blue line, we allow the energy cost share in production to decline from 4.9 percent of total costs to 3.2 percent by 2050 as predicted in Figure 13. The second scenario assumes that the energy cost share remains flat at 4.9 percent from 2023-2050, represented by the dashed orange line. We find that embedding our mechanism into the model causes the growth in energy demand from 2024-2050 to be about 40 percent smaller than in a version of the model where the energy share is assumed to stay constant.

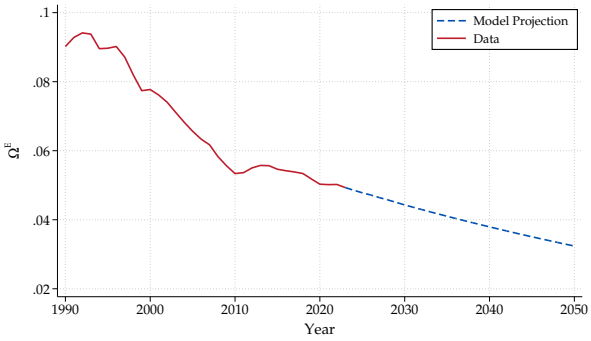
Finally, we analyze what our model predicts for the ratio of CO2/Energy intensity that is required to meet the Paris Climate Agreement goals (Figure 15). The UN publishes the Nationally Determined Contributions (NDCs) for each country which state India's goal to reduce the emissions intensity of its GDP by 45 percent by 2030. Given these goals, we use our model to project what the ratio of CO2/energy must be to achieve this 45 percent decline through the end of the decade. Additionally, we back out what ratio would be required to meet a 60, 70, or 80 percent reduction in CO2/GDP by 2050. Panel (A) shows the reductions required according to

FIGURE 13: Projections to 2050: Output Growth and Energy Share

(A) Output Growth

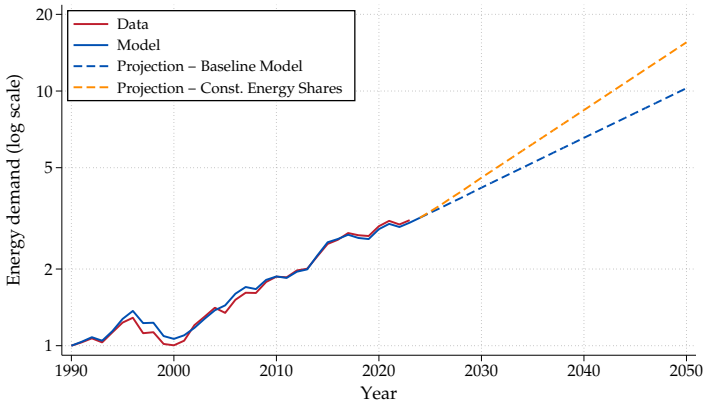


(B) Energy Expenditure Share



Note: Panel (A) plots projections of output through 2050, set equal to the projections from the IMF for GDP from 2025-2030, and applying a 5 year CAGR from that period for the years 2030-2050. Panel (B) plots the model paths of the energy share through the year 2050, where we allow the labor supply to increase by 2 percent per year and the A_{Yt} shifter to change such that output in the model matches the data projections. The red line represents the aggregate ASI time series through 2023. The dashed blue line represents the projections from the model to 2050 when applying our growth assumptions.

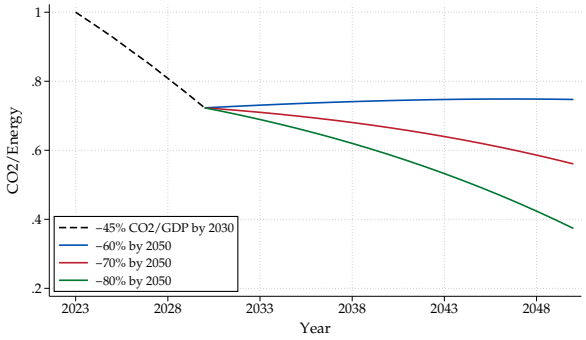
FIGURE 14: Projected Path of Energy Demand



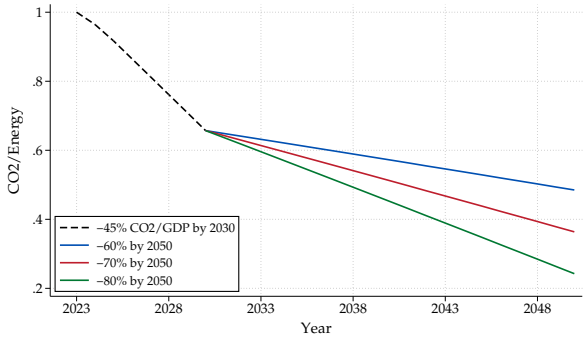
Note: This figure plots energy demand in the data vs. the model predictions through 2023. The dashed blue line is what our model predicts for energy demand according to our baseline scenario described in Figure 13, and incorporating our mechanism which allows the energy share to decline with firm scale. The dashed orange line uses our model to predict the path of energy out to 2050, under the assumption that the energy expenditure share stays constant and equal to 0.049 from 2023 to 2050.

FIGURE 15: CO2/Energy Required to Meet Paris Agreement Goals

(A) Baseline Model



(B) Constant Energy Share



Note: Panel (A) plots the decline in the CO2/Energy ratio required to meet Paris Climate Goals of a 45% reduction in CO2/GDP from 2005 levels by 2030 using our baseline model with a declining energy share. In addition, we plot the reduction in CO2/Energy required to reduce CO2/GDP 60%, 70%, or 80% from 2005 levels by 2050. Panel (B) plots the same graph but under the assumption that the energy expenditure share stays constant and equal to 0.05 from 2023 to 2050.

our baseline model. Panel (B) shows the reductions required in the alternative scenario where we assume that the energy share remains constant. We find that the CO2/Energy ratio would be required to fall significantly more, without accounting for our mechanism.

7. CONCLUSION

This paper uncovers a new mechanism linking economic growth to declining energy intensity: as economies expand, firms grow larger, and larger firms use energy more efficiently. Using microdata from India, together with a causal research design, we show that the energy expenditure share decreases sharply with firm scale. This relationship is not driven by input price differences or compositional shifts, but instead reflects technological forces—physical scaling laws that reduce energy losses at higher capacity, and fixed-cost efficiency investments. Embedding these forces into a structural heterogeneous-firm framework, we find that scale-dependent energy demand can account for almost all of the observed fall in India's energy intensity since 1990, and roughly 70 percent of the decline when incorporating energy price changes. Looking forward, our projected path for India's growth implies substantial additional reductions in energy intensity through 2050—reductions that significantly moderate the rise in total energy demand relative to a benchmark without scale effects.

Our findings have broad implications for growing energy demand in low- and middle-income countries. They suggest that the environmental costs of development may be lower than conventionally believed, as growth itself induces cleaner production through firm-level scaling. They also point toward the value of incorporating scale-dependent production technologies into Integrated Assessment Models and policy frameworks that evaluate the energy transition. Doing so would allow such models to capture an important and previously overlooked channel through which economic development can support, rather than hinder, long-run decarbonization.

REFERENCES

ABEBERESE, A. B. (2017): "Electricity cost and firm performance: Evidence from India," *Review of Economics and Statistics*, 99, 839–852.

ACEMOGLU, D., P. AGHION, L. BARRAGE, AND D. HÉMOUS (2023): "Climate change, directed innovation, and energy transition: The long-run consequences of the shale gas revolution," Tech. rep., National Bureau of Economic Research.

ACEMOGLU, D., P. AGHION, L. BURSZTYN, AND D. HEMOUS (2012): "The environment and directed technical change," *American economic review*, 102, 131–166.

ACEMOGLU, D., P. AGHION, AND D. HÉMOUS (2014): "The environment and directed technical change in a North–South model," *Oxford Review of Economic Policy*, 30, 513–530.

AGHION, P., T. BOPPART, M. PETERS, M. SCHWARTZMAN, AND F. ZILIBOTTI (2024): "A Theory of Endogenous Degrowth and Environmental Sustainability," .

ARKOLAKIS, C. AND C. WALSH (2023): "Clean Growth," .

BARROWS, G. AND H. OLLIVIER (2018): "Cleaner firms or cleaner products? How product mix shapes emission intensity from manufacturing," *Journal of Environmental Economics and Management*, 88, 134–158.

——— (2021): "Foreign demand, developing country exports, and CO2 emissions: Firm-level evidence from India," *Journal of Development Economics*, 149, 102587.

BATRAKOVA, S. AND R. B. DAVIES (2012): "Is there an environmental benefit to being an exporter? Evidence from firm-level data," *Review of World Economics*, 148, 449–474.

BENTO, P. AND D. RESTUCCIA (2017): "Misallocation, Establishment Size, and Productivity," *American Economic Journal: Macroeconomics*, 9, 267–303.

BOEHM, J., S. DHINGRA, AND J. MORROW (2022): "The Comparative Advantage of Firms," *Journal of Political Economy*, 130, 3025–3100.

BOEHM, J. AND E. OBERFIELD (2023): "Growth and the Fragmentation of Production," Tech. rep., Tech. rep., mimeo. 7.

BORUSYAK, K., P. HULL, AND X. JARAVEL (2022): "Quasi-experimental shift-share research designs," *The Review of Economic Studies*, 89, 181–213.

BRUCKNER, B., K. HUBACEK, Y. SHAN, H. ZHONG, AND K. FENG (2022): "Impacts of poverty alleviation on national and global carbon emissions," *Nature Sustainability*, 5, 311–320.

CARON, J. AND T. FALLY (2022): "Per capita income, consumption patterns, and CO2 emissions," *Journal of the Association of Environmental and Resource Economists*, 9, 235–271.

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

COLMER, J. M., D. LAGAKOS, AND M. SHU (2024): "Is the electricity sector a weak link in development?" .

COMIN, D., D. LASHKARI, AND M. MESTIERI (2021): "Structural change with long-run income and price effects," *Econometrica*, 89, 311–374.

ECKERT, F., S. GANAPATI, AND C. WALSH (2025): "Urban-biased growth: a macroeconomic analysis," Tech. rep., National Bureau of Economic Research.

FORSLID, R., T. OKUBO, AND K. H. ULLTVEIT-MOE (2018): "Why are firms that export cleaner? International trade, abatement and environmental emissions," *Journal of Environmental Economics and Management*, 91, 166–183.

GAO, Y. (2025): "Beyond Emissions: The Implications of Capital Misallocation on Optimal Carbon Tax," .

GERTLER, P. J., O. SHELEF, C. D. WOLFRAM, AND A. FUCHS (2016): "The demand for energy-using assets among the world's rising middle classes," *American Economic Review*, 106, 1366–1401.

GHODBANAN, S., R. ALIZADEH, S. SHAFIEI, AND J. RAHBAR SHAHROUZI (2024): "Improving energy

efficiency and heat recovery performance in paper machine drying section," *Journal of Computational & Applied Research in Mechanical Engineering (JCARME)*, 14, 47–62.

HASSLER, J., P. KRUSELL, AND C. OLOVSSON (2021): "Directed technical change as a response to natural resource scarcity," *Journal of Political Economy*, 129, 3039–3072.

HSIEH, C.-T. AND P. J. KLENOW (2014): "The life cycle of plants in India and Mexico," *The Quarterly Journal of Economics*, 129, 1035–1084.

JAMES, R. E., S. LEPTINSKY, M. TURNER, ET AL. (2022): "Quality Guidelines for Energy System Studies: Capital Cost Scaling Methodology: Revision 4a Report," Tech. Rep. DOE/NETL-2022/3340, National Energy Technology Laboratory (NETL), U.S. Department of Energy, technical report; Capital cost scaling guidance for energy system studies (Revision 4a).

JORDÀ, Ò., F. NECHIO, T. PHAN, AND F. SCHWARTZMAN (2026): "Financial Conditions and Capital Investment Choices," Federal Reserve Bank of San Francisco.

KABIR, P. AND E. TAN (2025): "Green Financing, Capital (Mis) allocation, and Unintended Consequences," *Available at SSRN 5849842*.

KIM, S. (2025): "Optimal carbon taxes and misallocation across heterogeneous firms," *Available at SSRN 4844030*.

KLENOW, P. J., E. PASTÉN, AND C. RUANE (2024): "Carbon taxes and misallocation in Chile," Tech. rep., Mimeo, Stanford.

LASHKARI, D., A. BAUER, AND J. BOUSSARD (2024): "Information technology and returns to scale," *American Economic Review*, 114, 1769–1815.

LEE, K., E. MIGUEL, AND C. WOLFRAM (2020): "Does household electrification supercharge economic development?" *Journal of Economic Perspectives*, 34, 122–144.

LEVINSON, A. (2015): "A direct estimate of the technique effect: changes in the pollution intensity of US manufacturing, 1990–2008," *Journal of the Association of Environmental and Resource Economists*, 2, 43–56.

MARTIN, L. A. (2011): "Energy efficiency gains from trade: greenhouse gas emissions and India's manufacturing sector," *Mimeograph, Berkeley ARE*.

MARTIN, L. A., S. NATARAJ, AND A. E. HARRISON (2017): "In with the Big, Out with the Small: Removing Small-Scale Reservations in India," *American Economic Review*, 107, 354–86.

MERTENS, M. AND B. SCHOEFER (2024): "From labor to intermediates: Firm growth, input substitution, and monopsony," Tech. rep., National Bureau of Economic Research.

PETERS, M. AND C. WALSH (2022): "Population growth and firm-product dynamics," *NBER Working Paper*, 29424.

PETROLEUM PLANNING & ANALYSIS CELL (PPAC) (2014): "All India Study on Sectoral Demand of Diesel & Petrol," Tech. rep., Ministry of Petroleum and Natural Gas, Government of India, accessed 2025-09-04.

RIJESH, R. (2022): "Concordance of Manufacturing Industries Using NIC-2008 and NIC-2004: A Disaggregate Level Study," .

SATO, R. (1974): "ON THE CLASS OF SEPARABLE NON-HOMOTHETIC CES FUNCTIONS," *The Economic studies quarterly*, 25, 42–55.

——— (1977): "Homothetic and non-homothetic CES production functions," *The American Economic Review*, 67, 559–569.

SHANKER, A. AND D. I. STERN (2018): "Energy intensity, growth and technical change," .

SHAPIRO, J. S. AND R. WALKER (2018): "Why is pollution from US manufacturing declining? The roles of environmental regulation, productivity, and trade," *American Economic Review*, 108, 3814–3854.

SINGER, G. (2024): "Complementary inputs and industrial development: can lower electricity prices improve energy efficiency?" Tech. rep.

TRIBE, M. AND R. ALPINE (1986): "Scale economies and the "0.6 rule"," *Engineering Costs and Production*

Economics, 10, 271–278.

TROTTNER, F. (2019): "Who gains from scale?" Tech. rep., Working Paper.

TYBOUT, J. R. (2000): "Manufacturing firms in developing countries: How well do they do, and why?" *Journal of Economic literature*, 38, 11–44.

US DEPARTMENT OF ENERGY (1997): "Determining electric motor load and efficiency," *Factsheet*.

U.S. DEPARTMENT OF ENERGY (2022): "2018 Manufacturing Static Energy Sankey Diagrams," Tech. rep., U.S. Department of Energy, page updated June 2022; accessed 2026-02-02.

WOLLBURG, P., S. HALLEGATTE, AND D. G. MAHLER (2023): "Ending extreme poverty has a negligible impact on global greenhouse gas emissions," *Nature*, 623, 982–986.

A. DATA SOURCES AND VARIABLES CONSTRUCTION

A.1 Annual Survey of Industries: description and variable definitions

The ASI is a dataset put together by India's Ministry of Statistics and Programme Implementation (MOSPI). The reference period for each survey is the accounting year, which in India begins on the 1st of April and ends on the 31st of March the following year. Throughout the paper we reference the surveys by the earlier of the two years covered.

Coverage and sampling methodology. The ASI contains information on a representative sample of manufacturing establishments, conditional on them taking part of the organized sector, and either employing more than 20 employees, or employing more than 10 employees and using electricity. We call the subpopulation of firms satisfying this criteria the ASI population. Within the ASI population, ASI defines a Census sector which is sampled exhaustively and a Sample sector for which the microdata contains only a representative sample. Details of how the sampling methodology for the ASI changes over time are shown in Table A.1. ASI provides sampling weights, which we use to weight all data moments.

TABLE A.1: Sampling Methodology for Indian ASI

Period	Census Sector	Sample Sector
1998	Complete enumeration states, plants with $>$ 200 workers, all joint returns	Stratified within state $\times$ 4-digit industry (NIC-98), minimum of 8 plants per stratum
1999-2003	Complete enumeration states, plants with $\geq$ 100 workers, all joint returns	Stratified within state $\times$ 4-digit industry (NIC-98), 12% sampling fraction (20% in 2002), minimum of 8 plants per stratum
2004-2006	6 less industrially developed states, 100 or more workers, all joint returns, all plants within state $\times$ 4-digit industry with $<$ 4 units	Stratified within state $\times$ 4-digit industry, 20% sampling, minimum of 4 plants
2007	5 less industrially developed states, 100 or more workers, all joint returns, all plants within state $\times$ 4-digit industry with $<$ 6 units	Stratified within state $\times$ 4-digit industry, minimum 6 plants, 12% sampling fraction: exceptions
2008-2013	6 less industrially developed states, 100 or more workers, all joint returns, all plants within state $\times$ 4-digit industry with $<$ 4 units	Stratified within district $\times$ 4-digit industry, minimum 4 plants, 20% sampling fraction

Note: Baseline sampling fractions are shown, not accounting for state-specific exceptions.

Sample selection. We start with 1,068,114 plant $\times$ year observations. We subsequently employ multiple sample selection rules. First, we restrict the sample to factory $\times$ year observations with either positive reported gross sales, or positive reported sales at the factory gate. This drops one third of all observations (360,145). Next, we disregard all observations that exactly copied their sales from the previous year, suspecting these plants to be actually closed. This drops 1,179 additional observations. Third, we drop all plant $\times$ years that reported either no days worked,

or no persons employed, dropping a supplementary 206 observations. These cleaning steps follow Martin, Nataraj, and Harrison (2017).

Construction of key variables. We detail below the construction of the main variables for our analysis.

Firm-level variables.

- Output: We construct output as the sum of the gross value of products sold (including distribution expenses, as well as taxes and subsidies). We include the value of electricity sold by the plant (consistent with how we define fuels).
- Capital: Capital is the closing book value of fixed assets (net of depreciation). These include all types of assets deployed for production and transportation, as well as living or recreational facilities (hospitals, schools, etc.) for factory personnel. It excludes intangible assets and current assets.
- Cost of capital: Depreciation rate and average interest rate. We include repair and maintenance costs (plant/machinery, building, etc.), as well as the cost of rented capital.
- Intermediates: We construct intermediates as the sum of the value of non-fuels materials consumed and other intermediate expenses. Other intermediate expenses include operating expenses (freight and transportation charges, taxes paid), non-operating expenses (communication, accounting, advertising), and insurance charges.
- Fuels: Sum of expenses on electricity, oil, coal, gas, and other fuels. We do not include the purchase value of electricity generated within the firm, since we account for the fuels used to generate that electricity.
- Labor: Total days worked.
- Labor Cost: We construct labor costs as total payments to labor over the course of the year. These payments include wages and salaries, bonuses, contributions to old-age pension funds (and other funds), and all welfare expenses. We also include the costs of contract and commission work.

Firm×product-level variables and firm-level price and quantity indices. The key advantage of the ASI is that both for the products that manufacturing plants produce and the inputs they buy, we observe information on sales, quantities, and unit values, with products classified at the 5-digit NPC level (around 1,200 distinct product codes for our sample of manufacturing firms). Some sections of our analysis exploit this data. We detail the construction firm×product- and firm-level price and quantity variables here, and provide more details on the product classification below.

We construct a panel of firm-product prices and quantities. We denote $\Delta \log p_{ijt}$ and $\Delta \log y_{ijt}$ as the changes in the log of price p and log of quantity y of product j sold by firm i at time t, respectively. The key cleaning steps are: (i) harmonizing product codes within firms, as sometimes firm report different codes for the same product in consecutive years; (ii) correcting observations for which the product of unit values and quantity sold differ significantly from the reported sales value; (iii) correcting unit mistakes: the data contains incorrect entries due

to misplaced commas, which we address by rescaling values up or down when the price was multiplied by 10^n and the quantity was multiplied by 10^{-n} with respect to the previous year.

Even working with narrowly-defined product categories, unobserved heterogeneity could prevent a meaningful comparison of prices across firms. We always work with within firm×product changes in the price (or quantity), largely alleviating this concern. The set of products for which we observe a valid price (quantity) change for two consecutive years, which we denote $\mathcal{J}_i$, account for, on average, 75% of firm-level total sales.

We define the firm-level price index as the Törnqvist-weighted change in the observed firm×product-level price changes: $\underline{\Delta} \log p_{it} = \sum_{j \in \mathcal{J}_i} \bar{s}_{ijt} \Delta \log p_{ijt}$. We use the convention of placing a bar on top of the share to denote that these shares are the mid-point of the shares in $t-1$ and t, and the bar under the Δ sign indicates that we take the average price change over the set of *observed* products. We construct the firm-level change in quantities as $\Delta \log y_{it} = \Delta \log \mathcal{R}_{it} - \underline{\Delta} \log p_{it}$.[12]

Similarly, we observe purchase value, unit price, and quantity purchased for materials classified in the same 1,200 products. We denote by $\Delta \log w_{ikt}$ and $\Delta \log x_{ikt}$ the log change in prices and quantities of input k used by firm i. We perform the same cleaning steps as described for prices. The inputs for which we observe a valid price (quantity) change across two consecutive years, which we denote $\mathcal{K}_i$, account for on average 57% of firm-level total input purchases. We define the firm-level input quantity change as the Törnqvist-weighted change in the observed firm×input-level quantity changes: $\underline{\Delta} \log x_{it} = \sum_{k \in \mathcal{K}_i} \bar{s}_{ikt} \Delta \log x_{ijt}$. We construct the firm-level intermediate input price index as $\Delta \log w^x_{it} = \Delta \log \mathcal{C}^x_{it} - \underline{\Delta} \log x_{it}$.[13]

Industry classification. Our data relies on three distinct industry classification systems: NIC-98 (1998–2003), NIC-04 (2004–2007), and NIC-08 (2008 and beyond). We first address issues with the 5-digit industry codes in NIC-98, where codes are sometimes masked with zeroes or absent from the official documentation, by replacing them with the most frequent 5-digit code within each 4-digit grouping following the approach of Martin et al. (2017). Next, we apply concordances from NIC-08 to NIC-04 using the mapping provided by Rijesh (2022), and from NIC-04 to NIC-98 using the mapping provided by Martin et al. (2017). We manually supplement the mappings for industries not covered by these concordances. In case of 1:m mappings, we select the appropriate industry based on transition matrices in the microdata.

Product classification. Our analysis standardizes product classifications across four distinct classifications used in our sample: NPCMS 2015 (2016-2017), NPCMS 2011 (2010-2015), ASICC 2009 (2008-2009), and ASICC 2008 (pre-2008). We harmonize all product codes to NPCMS 2011, as it provides a well-defined five-digit structure that balances granularity and coverage. Given the absence of an official concordance between NPCMS 2015 and NPCMS 2011, we constructed a mapping using fuzzy matching (based on product codes, descriptions, and units) and semantic embeddings (OpenAI's AA2 model). For ASICC 2009, we utilize the official

[12] Because we do not observe the price and quantity change for all total sales, in general $\underline{\Delta} \log p_{it} + \underline{\Delta} \log y_{it} \neq \Delta \log \mathcal{R}_{it}$. We assume that the price change of observed products is on average equal to the price change for all products.

[13] The assumption is that $\underline{\Delta} \log x_{it}$ (the average increase in input quantity for the inputs $\underline{\mathcal{K}_i}$ for which we observe input-level data) is equal to the average change in input quantity for all inputs $\Delta \log x_{it}$. This assumption is the most natural when different material inputs are strong complements (and it is exactly true if production is Leontief).

concordance to NPCMS 2011 but address its limitations (such as missing mappings and invalid classification) by leveraging ASI data and semantic embeddings. The harmonization of ASICC 2008 to ASICC 2009 follows the concordance from Boehm, Dhingra, and Morrow (2022). Table A.2 shows an excerpt of the product classification.

TABLE A.2: Example of NPC-MS 2011 5-digit classification

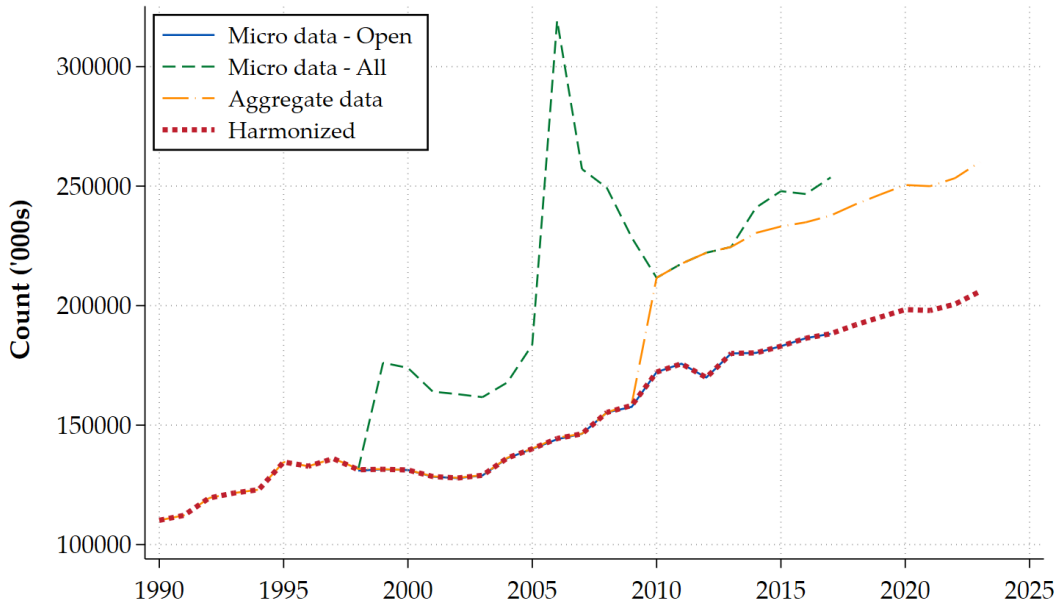
Code	Description
35	Other chemical products; man-made fibres
351	Paints and varnishes and related products; artists' colours; ink
35110	Paints and varnishes and related products
35120	Artists', students' or signboard painters' colours, modifying tints, amusement colours and the like
35130	Printing ink
35140	Writing or drawing ink and other inks
352	Pharmaceutical products
353	Soap, cleaning preparations, perfumes and toilet preparations
354	Chemical products n.e.c.
355	Man-made fibres
36	Rubber and plastics products
361	Rubber tyres and tubes
36111	New pneumatic tyres, of rubber, of a kind used on motor cars
36112	New pneumatic tyres, of rubber, of a kind used on motorcycles or bicycles
36113	Other new pneumatic tyres, of rubber
36114	Inner tubes, solid or cushion tyres, interchangeable tyre treads and tyre flaps, of rubber
36115	Camel back strips for retreading rubber tyres
36120	Retreaded pneumatic tyres, of rubber
362	Other rubber products
36210	Reclaimed rubber
36220	Unvulcanized compounded rubber, in primary forms or in plates, sheets or strip; unvulcanized rubber in forms other than primary forms or plates, sheets or strip
36230	Tubes, pipes and hoses of vulcanized rubber other than hard rubber
36240	Conveyor or transmission belts or belting, of vulcanized rubber
36250	Rubberized textile fabrics, except tyre cord fabric
36260	Articles of apparel and clothing accessories (including gloves) of vulcanized rubber other than hard rubber
36270	Articles of vulcanized rubber n.e.c.; hard rubber; articles of hard rubber
363	Semi-manufactures of plastics
364	Packaging products of plastics
369	Other plastics products

Note: For codes other than 351, 361 & 362, the 5-digit classifications are not shown.

Number of factories in ASI. We construct the number of factories as follows. From 1990 to 2009, we use the number of factories reported in the ASI aggregated annual series. This perfectly corresponds to the number of factories declared as Open in the micro data. In 2010, the aggregate series switches to counting all factories, including those declared as closed. To obtain a consistent time series, for years after 2009, we use the number of factories declared as Open in the micro data. This series is available until 2017. For 2018-2023, we extend this series by using the growth rate in the number of factories found in the aggregate series. Our harmonized series is plotted in dotted red on Figure A.1.

A.2 Annual Survey of Industries: data validation

FIGURE A.1: Number of factories in ASI



Note: The figure shows the number of factories in ASI from 1990 to 2023 using three underlying series. "Micro data – Open" and "Micro data – All" plot, respectively, the number of factories declared open and the total number of factories in the ASI micro data. "Aggregate data" plots the number of factories reported in the ASI aggregated annual series, which counts open factories up to 2009 and all factories thereafter. The dotted red "Harmonized" series is our preferred measure: it uses the aggregate series (which coincides with open factories) up to 2009, the number of open factories in the micro data from 2010–2017, and then extrapolates to 2023 using the growth rate of the aggregate series.

Manufacturing in ASI and in the national accounts. ASI only covers the organized sector. Compared to the national accounts (which include the informal sector), ASI covers 64% of manufacturing value added and 83% of manufacturing output on average across years.

Figure A.2 plots gross output and value added growth in ASI, in KLEMS, and in the national accounts, indexed to 1 in 1990. The figure reveals that manufacturing growth in ASI closely tracks the national accounts aggregates.

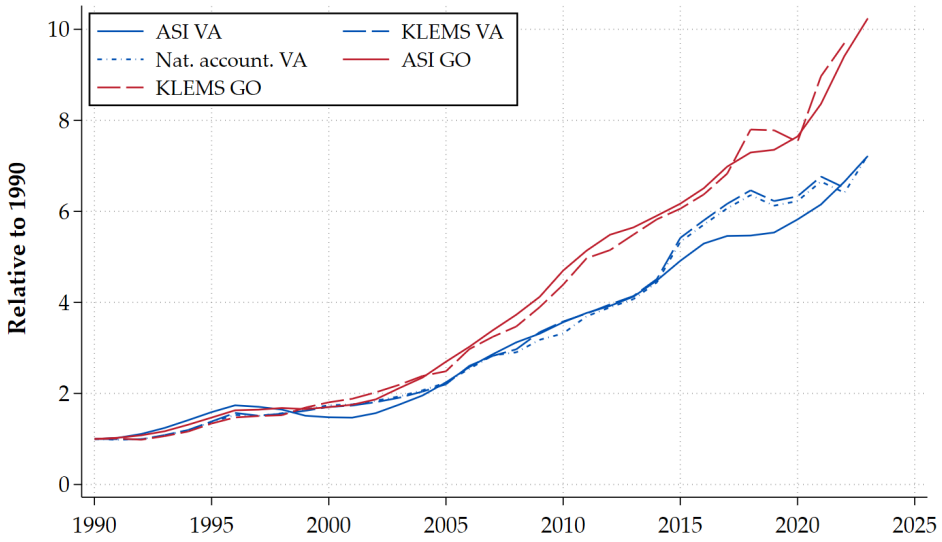
A.3 Other data sources

International Energy Agency. We exploit the World Energy Balance dataset produced by the International Energy Agency (link) to obtain energy use by broad sectors for all countries.

Indian national accounts and KLEMS. For macroeconomic series, we use the Indian national accounts. For some manufacturing specific series, we use the KLEMS database produced by the Reserve Bank of India (link).

World Bank Enterprise Survey. We exploit the "Green Economy" module of the World Bank Enterprise Survey administered in India in 2022 (link). The World Bank Enterprise Survey is a firm-level survey of a representative sample of an economy's private sector. The Enterprise Surveys are conducted across all geographic regions and cover small, medium, and large companies, and are representative of firms in the non-agricultural formal private economy.

FIGURE A.2: Value added and output growth in ASI



Note: The figure plots gross output (GO) and value added (VA) indices for Indian manufacturing constructed from ASI micro data, India KLEMS, and the national accounts. All series are expressed in 2011 constant prices and normalized to 1 in 1990. ASI series are obtained by aggregating plant-level outcomes from the annual ASI files; KLEMS and national-account series use official manufacturing aggregates.

UNIDO Technology compendia. The United Nations Industrial Development Organization (UNIDO), in collaboration with the Indian Bureau of Energy Efficiency (BEE), undertook a project titled "Promoting energy efficiency and renewable energy in selected MSME clusters in India" over 2020-2022. The project aimed to promote energy efficiency and renewable energy technologies in process applications in energy-intensive industrial clusters, comprising micro, small and medium enterprises. In this context, technology compendia were produced to provide a list of available energy-efficient and renewable energy technologies. There is a compendia for each industry×cluster where the project was implemented (e.g., ceramics in Morbi). Each compendia details a menu of technologies available to firms in this industry×cluster. For each technology, the document details the baseline scenario, the energy efficient alternative, and provides a cost-benefit analysis. Some cost-benefit analysis are estimates for the typical firm in the industry×cluster, while others rely on data from actual implementations.

We systematically collect all the data on technology menus for a number of industries for which sufficiently many clusters are available:

- Foundries: Faridabad, Belgaum, Howrah, Ahmedabad, Agra, Ahmedabad, Rajkot
- Hand tools: Jalandhar, Nagaur, Ludhiana
- Ceramics: Khurja, Morbi, Thangadh

After dropping duplicates (for an industry, the exact same cost-benefit analysis is sometimes used in several clusters), we obtain 57 distinct technology data points for foundries, 31 for hand-tools, and 54 for ceramics.

For this data, we obtain the required investment $\Delta k_{\tau is}$ in INR, and associated annual energy

savings $\Delta e_{\tau,is}$ in gigajoules, for a number of alternative technologies τ available to firms i in a given industry s.

Output, materials, and capital deflators. We construct a materials deflator by combining: (i) WPI product-level price indices at the level of the WPI sub-headings (30 indices), and (ii) the share of individual products within total material inputs, obtained from the ASI data over 1998-2017 and extrapolated for other years.

We construct a capital deflator by combining: (i) Penn Word Tables capital deflators for India for four capital types (structures, machinery, transport equipment, and a residual other asset group), and (ii) the share of each type of capital within the total capital stock, obtained from the ASI data over 1998-2017 and extrapolated for other years.

Time series for the price of energy inputs. We construct price series for electricity, oil, coal, and gas for 1990-2023. We have several candidate sources for the price of energy inputs: (i) the average price paid by firms in ASI (available for coal, electricity for 1998-2017, for gas for 2008-2017); (ii) the item-level Wholesale Price Index (available for oil, electricity, coal for 1993-2013, and for gas for 2011-2023); the average price paid by firms in Prowess (available for coal, electricity, oil, and gas for 1990-2020). We find that for electricity, coal, and gas, the average price paid by firm in ASI is better approximated by the Prowess series than the WPI series (on their respective common samples). We therefore use the Prowess series as a baseline. We extend these series for 2021-2023 by using price growth from the WPI series for these three years. All series are deflated by the final goods deflator.

We construct the energy price index as the weighted average of the price index of each energy inputs, with expenditure weights as shares. The expenditure weights are constructed using ASI and are assumed to be constant over time. Until 2008, the "Other" category groups gas and other fuels, while after 2008 we separately observe "Gas" and "Other". We assume that the weight of gas in the combined gas and other fuels category has remained constant over time. Firms have a 9% expenditure share on other fuels for which we do not have a price index. We assume that the relative price of other fuels has remained constant over time.

A.4 Measurement of firm-level productivity growth

We measure the change in physical productivity as follows. Consider firm i at time t producing a single physical output:

$$y_{it} = z_{it}\, F(\phi_{1,it} x_{1,it},\ \ldots,\ \phi_{J,it} x_{J,it})\,, \tag{A.1}$$

where z_{it} is Hicks-neutral physical productivity, $x_{j,it}$ is the physical quantity of input j, $\phi_{j,it}$ is input-specific (input-augmenting) productivity, and $F(\cdot)$ is a differentiable production function. F can have arbitrary curvature (Cobb–Douglas, CES, nested CES, translog, etc.), arbitrary returns to scale, and arbitrary (non-)homotheticity. Let p_{it} denote the output price faced by the firm and $w_{j,it}$ the firm-specific price of input j. Firms may face arbitrary firm-specific input prices.

Define the output elasticity of physical input $x_{j,it}$:

$$\theta_{j,it} \equiv \frac{\partial \log y_{it}}{\partial \log x_{j,it}} \tag{A.2}$$

Taking log differences of (A.1):

$$\Delta \log y_{it} = \Delta \log z_{it} + \sum_{j=1}^{J} \theta_{j,it} \Delta \log x_{j,it} + \sum_{j=1}^{J} \theta_{j,it} \Delta \log \phi_{j,it} \tag{A.3}$$

Define measured physical TFPQ growth as:

$$\Delta \widehat{\log TFPQ}_{it} \equiv \Delta \log y_{it} - \sum_{j=1}^{J} \widehat{\theta}_{j,it} \Delta \log x_{j,it}. \tag{A.4}$$

with $\widehat{\theta}_{j,it}$ an estimate of the output elasiticity. If $\widehat{\theta}_{j,it} = \theta_{j,it}$, then:

$$\Delta \widehat{\log TFPQ}_{it} = \Delta \log z_{it} + \sum_{j} \theta_{j,it} \Delta \log \phi_{j,it},$$

so measured TFPQ growth captures both Hicks-neutral and input-augmenting productivity growth.

Note that physical productivity measurement is made possible by the availability of price and quantity data separately, for both outputs and inputs.

Measurement of output elasticities. Under the assumption that firms are price-takers on the input and output markets, and for any flexible input j chosen at an interior optimum of the static cost minimization problem, we can show that:

$$\theta_{j,it} = \frac{w_{j,it} x_{j,it}}{p_{it} y_{it}} \equiv s_{j,it}, \tag{A.5}$$

That is, output elasticities are equal to output cost shares. In the presence of markups μ_{it} and non-priced input wedges $\tau_{j,it}$, we instead have

$$\theta_{j,it} = \mu_{it}(1 + \tau_{j,it})\, s_{j,it}. \tag{A.6}$$

We measure output elasticities in three different ways. Option 1 uses firm-level revenue shares, and option 2 uses average revenue shares by industry×year×size decile. Firm-level shares are conceptually preferable under heterogeneity (e.g., if the production function is CES or non-homotheticity) but may reflect firm-specific wedges (markups, adjustment costs). In addition, firm-level shares are unreliable for quasi-fixed inputs such as capital, so we always use industry-level shares for capital. On the other hand, industry×year×size decile-level shares average out input wedges and statistical noise, but may not capture the true output elasticity for any individual firm.

For both firm-level and industry-level shares, we implement the following correction for the

markup-wedge multiplicative factor. We exploit the fact that $\sum_j \theta_{j,it} = \text{RTS}_{it}$, the returns to scale for firm i. In section 5, where we calibrate our quantitative model, we estimate average returns to scale $\widehat{\text{RTS}}$ by estimating the elasticity of the cost function with respect to output. Assuming a constant combined distortion $\mu(1+\tau)$, we can back out $\mu(1+\tau)$ from the relationship $\mu(1+\tau) = \widehat{\text{RTS}}/\mathbb{E}[\sum_j s_{j,it}]$.

Finally, in option 3, we use the estimated production function from the quantitative model.

We consider four inputs: raw materials, energy, capital, and labor.

Validation of measured productivity growth. In Table A.3, we show that, as predicted by theory, our measure of physical productivity growth predicts a decline in firm-level prices, an increase in firm-level quantities, and an increase in firm-level revenues.

TABLE A.3: Validation of measured productivity growth

Panel A: Effect of the change in TFPQ growth on the change in price

	$\Delta^h \log(p)$								
	$h=1$			$h=3$			$h=5$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Δ^h log(tfpq)	-0.645***	-0.607***	-0.629***	-0.869***	-0.821***	-0.852***	-0.864***	-0.817***	-0.849***
	(0.008)	(0.008)	(0.008)	(0.015)	(0.015)	(0.015)	(0.014)	(0.014)	(0.014)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓	✓	✓	✓
Def. TFPQ	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)
R-squared	0.565	0.542	0.555	0.635	0.615	0.627	0.702	0.678	0.695
Observations	266,142	266,142	266,142	166,176	166,176	166,176	114,191	114,191	114,191

Panel B: Effect of the change in TFPQ growth on the change in quantity

	$\Delta^h \log(y)$								
	$h=1$			$h=3$			$h=5$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Δ^h log(tfpq)	0.793***	0.733***	0.754***	0.936***	0.875***	0.907***	0.904***	0.850***	0.882***
	(0.008)	(0.008)	(0.008)	(0.014)	(0.014)	(0.014)	(0.013)	(0.013)	(0.013)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓	✓	✓	✓
Def. TFPQ	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)
R-squared	0.568	0.533	0.540	0.644	0.616	0.627	0.710	0.680	0.696
Observations	266,142	266,142	266,142	166,176	166,176	166,176	114,191	114,191	114,191

Panel C: Effect of the change in TFPQ growth on the change in revenues

	$\Delta^h \log(py)$								
	$h=1$			$h=3$			$h=5$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Δ^h log(tfpq)	0.122***	0.100***	0.097***	0.072***	0.059***	0.061***	0.052***	0.044***	0.045***
	(0.005)	(0.005)	(0.005)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓	✓	✓	✓
Def. TFPQ	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)
R-squared	0.220	0.214	0.213	0.264	0.261	0.261	0.295	0.292	0.292
Observations	266,142	266,142	266,142	166,176	166,176	166,176	114,191	114,191	114,191

Note: This table presents the results of regressions of firm-level variables on firm-level physical productivity growth, defined in (A.4). The outcome variable is the change in the firm-level price index (panel A), the change in the firm-level quantity index (panel B), and the change in firm-level revenues (panel C). "Def. TFPQ" indicates the chosen measure for output elasticities, as detailed in the text. Regressions are weighted by firm-level total costs. Standard errors are clustered at the firm and main product×year level. ***, ** and * indicate significance at the 1%, 5% and 10% levels, respectively.

B. ADDITIONAL EMPIRICAL RESULTS

B.1 Descriptive Statistics: Aggregate Data

B.1.1 Energy intensity

To analyze the sources of change in aggregate energy intensity, we apply a standard shift-share (or structural decomposition) approach. Denote E_{jt} energy use in sector j at time t, Y_{jt} sectoral output, and $s_{jt} = Y_{jt}/Y_t$ is the sectoral output share. The change in aggregate energy intensity between two periods can be decomposed as follows:

$$\Delta \frac{E_t}{Y_t} = \underbrace{\sum_{j \in J} s_{j,t-1} \, \Delta \left(\frac{E_{jt}}{Y_{jt}} \right) + \Delta s_{jt} \, \Delta \left(\frac{E_{jt}}{Y_{jt}} \right)}_{\text{Within + Cross}} + \underbrace{\sum_{j \in J} \Delta s_{jt} \, \frac{E_{j,t-1}}{Y_{j,t-1}}}_{\text{Reallocation}}$$

The reallocation term reflects changes in aggregate energy intensity attributable to shifts in the composition of output across sectors, keeping sectoral energy intensities fixed. This term is orthogonal to our theory. The within-sector term captures changes in sectoral energy intensity, keeping sectoral shares constant. The cross term captures the fact that faster-growing industries may experience a larger decline in sectoral energy intensity. Our theory is consistent with both terms.

Estimating this decomposition requires to match sectoral energy use and sectoral output. We proceed as follows. Sectoral energy use comes from the IEA. Sectoral output shares come from the UN. We match the sectors in these two datasets as described in Table B.1.

Transport energy use in the IEA data corresponds to both the household sector (energy use that cannot be linked to value added in the transportation sector, like personal car use) and the production sector. We estimate energy used by the transportation sector as follows. For road transport, we exclude fuel use by passenger cars, two-wheelers and three-wheelers. We obtain their share of total road fuel use in India in Petroleum Planning & Analysis Cell (PPAC) (2014). We use the same share for all countries. All non-road uses (rail, pipeline, domestic aviation, world aviation bunkers, world marine bunkers, domestic navigation, n.e.c.) are treated as production of goods and services.[14]

India. The results for India are presented in main text Figure 2.

Global. We document the same decomposition at the global level in Figure B.1. The within cell j is defined as a country$\times$sector pair so that the red line excludes reallocation across both sectors and countries. In terms of sectoral shifts, the small decline in highly energy-intensive industry and large decline in moderately energy-intensive agriculture are offset by a rise in both low-energy services and highly energy-intensive transport. In addition, reallocation of economic activity across countries tends to push up the aggregate energy intensity.

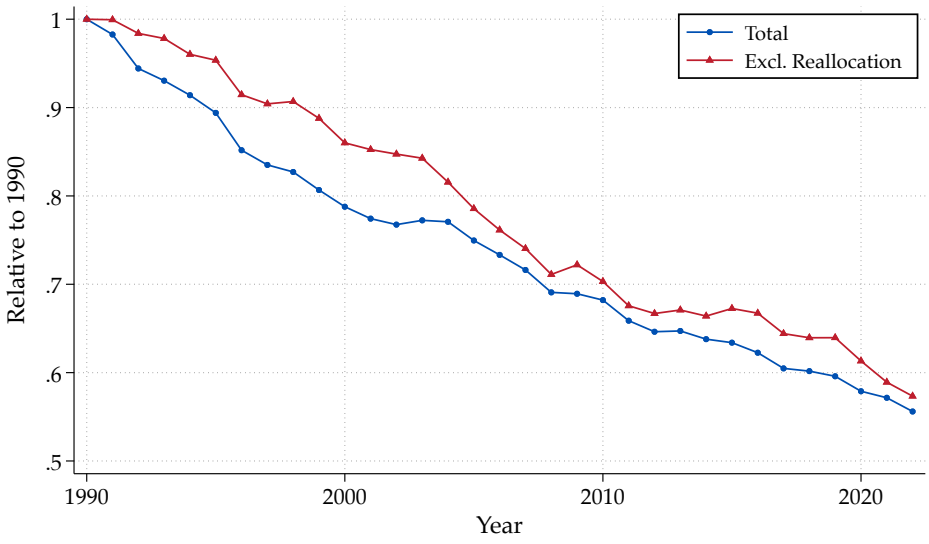
[14] Apart from road, the largest end-uses are domestic navigation, aviation and rail. For these three categories, the share of the fleet that is individually-owned and hence does not contribute to the transport sector GDP is negligible.

TABLE B.1: Sectors for the Within-Across Decomposition

Sector	IEA	UN
Agriculture, forestry and fishing	Agriculture, forestry (01–02) + Fishing (03)	Agriculture, forestry and fishing (01–03)
Industry (mining, manufacturing, construction)	Mining (07–08–099), Manufacturing (10–32, except 19), Construction (41–43)	Mining (05–09), Manufacturing (10–33), Construction (41–43), Utilities (35)
Transport and communications	Transport* (49–51)	Transport, storage and communications (49–53)
Services	Commercial and Public Services (33, 36–39, 45–47, 52–53, 55–56, 58–66, 68–75, 77–82, 84–88, 90–96, 99)	Wholesale and Retail Trade, Restaurants and Hotels (45–47), Other Activities (58–66, 68–75, 84–88, 90–96, 97–98).

Note: This table shows the sectors used in the within-across sectors decomposition of energy intensity. Sector codes correspond to ISIC Rev. 4. For transport in the IEA data, please refer to explanations in the main text. Utilities energy use is excluded because the utility sector is not an end-user of energy.

FIGURE B.1: The global decline in energy intensity

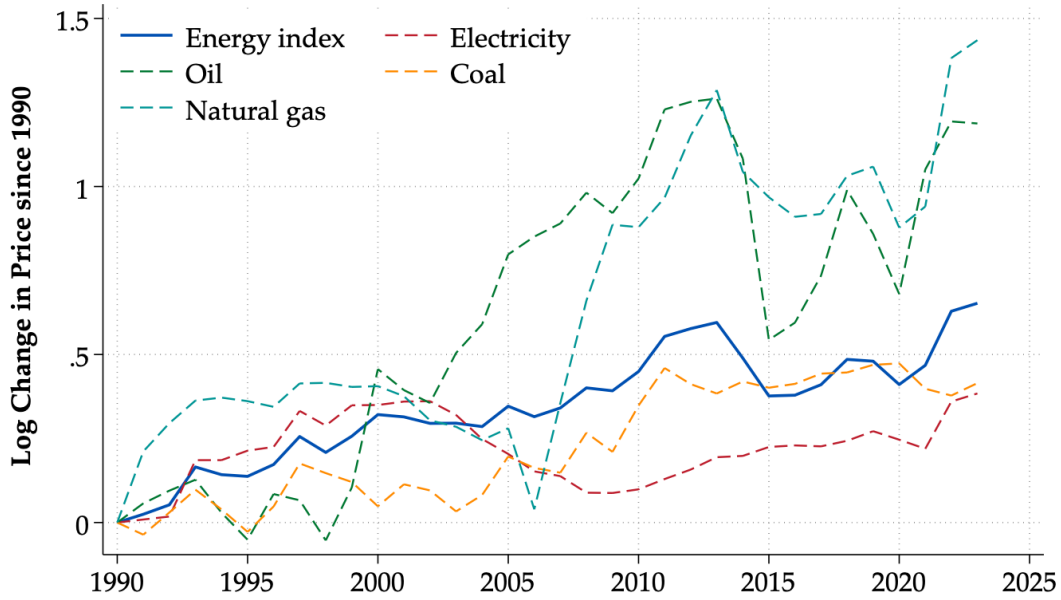


Note: This figure displays the change in energy intensity at the global level. It is defined as energy used in the production of goods and services divided by GDP (measured in constant 2021 USD PPP), and is normalized to 1 in 1990. The red line excludes the reallocation component (see equation (B.1.1)). The sectors are: Agriculture and Fishing, Industry (mining, manufacturing, construction), Transport and Communications, and Other Services.

B.1.2 Other descriptive statistics on aggregate data

Relative Price of Energy Inputs. Figure B.2 plots the price of energy inputs, deflated by the GDP deflator, indexed to 1 in 1990. Energy index refers to the average price of energy inputs, weighted by their cost share in the ASI data.

FIGURE B.2: Relative price of energy inputs



Note: The figure plots real price indices for electricity, oil, coal, and natural gas used by Indian manufacturing plants, together with a composite energy price index constructed as a fixed-weight average using ASI expenditure shares on each fuel. Fuel-specific indices are built by splicing Prowess firm-level price indices with WPI energy price series after 2020, deflating by the GDP deflator, and normalizing all series to 1 in 1990. The sample covers fiscal years 1990–2025.

Energy by End Use. Table B.2 uses data from US 2018 Manufacturing Energy Consumption Survey energy flowcharts to calculate the fraction of energy by end use across the manufacturing industries present in the ASI.

TABLE B.2: Energy by End Use

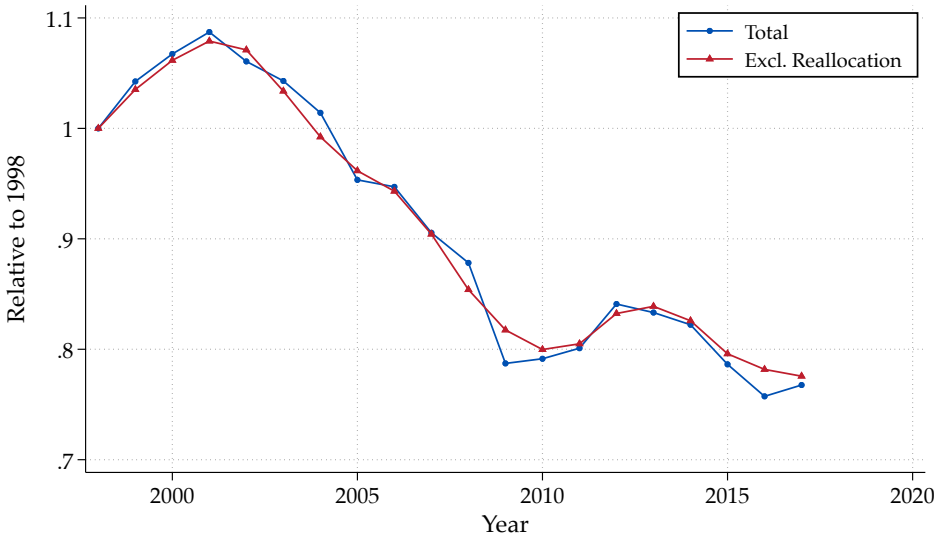
	Process				Non-Process
	Heat: fired	Heat: steam	Mechanical	Other	
All Manufacturing	0.51	0.09	0.17	0.12	0.11
Alumina and Aluminum	0.46	0.01	0.11	0.31	0.10
Automobile and Light Duty Motor Vehicle	0.36	0.03	0.21	0.10	0.30
Cement	0.75	0.00	0.21	0.03	0.01
Chemicals	0.31	0.23	0.19	0.17	0.09
Computers, Electronics, and Electrical Equipment	0.20	0.01	0.17	0.17	0.45
Fabricated Metals	0.38	0.01	0.22	0.07	0.32
Food and Beverage	0.19	0.24	0.18	0.17	0.21
Forest Products	0.13	0.44	0.24	0.06	0.13
Foundries	0.62	0.00	0.15	0.05	0.18
Glass and Glass Products	0.79	0.00	0.11	0.03	0.08
Iron and Steel	0.69	0.05	0.09	0.09	0.08
Machinery	0.18	0.01	0.34	0.06	0.42
Petrochemicals	0.46	0.23	0.12	0.13	0.05
Petroleum Refining	0.68	0.14	0.11	0.05	0.02
Plastic Material and Resins	0.19	0.29	0.20	0.24	0.08
Plastics and Rubber Products	0.21	0.05	0.36	0.10	0.27
Semiconductor and Related Devices	0.12	0.02	0.16	0.29	0.41
Textiles	0.23	0.11	0.38	0.05	0.22
Transportation Equipment	0.24	0.02	0.24	0.09	0.41

Note: The table reports end-use energy shares by industry, obtained from the US 2018 Manufacturing Energy Consumption Survey energy flowcharts. The "All Manufacturing" row aggregates industries using Indian ASI total energy as weights. Industries without an ASI match are excluded. Row shares sum to one by construction.

B.2 Descriptive Statistics: Micro-data

Manufacturing energy share in ASI: Excluding reallocation across sectors. Figure B.3 shows the change in the energy expenditure share in ASI excluding the reallocation component. Reallocation is defined as shifting sectoral shares across 2-digit industries. This decomposition requires the micro-data, hence we report results for 1998-2017.

FIGURE B.3: Manufacturing energy share in ASI



Note: This figure displays the change in the energy expenditure share in ASI. It is normalized to 1 in 1998. The red line excludes the reallocation component (see equation (B.1.1)). The sectors are 2-digit industries.

Analysis of Continuing Firms. One potential concern with our estimation of η is that our identification strategy exploiting within-firm changes restricts the sample to firms present in two consecutive periods (i.e., continuing firms). Here, we show that changes in the energy share of continuing firms are representative of the dynamics of the energy share in the whole sample.

Let $\mathcal{N}_t$ be the set of firms active at time t. Consider the energy expenditure share with a given denominator X. Here we use total costs as a denominator. The aggregate energy share Ω_t^E can be written as:

$$\Omega_t^E = \sum_{i \in \mathcal{N}_t} \omega_{it} \Omega_{it}^E \tag{B.1}$$

where $\omega_{it} = \frac{X_{it}}{X_t}$. Let $\mathcal{C}_t = \mathcal{N}_t \cap \mathcal{N}_{t-1}$ be the set of continuing firms between $t-1$ and t. Let $\mathcal{E}_t = \mathcal{N}_t \backslash \mathcal{N}_{t-1}$ be the set of entering firms. Let $\mathcal{X}_t = \mathcal{N}_{t-1} \backslash \mathcal{N}_t$ be the set of exiting firms.

First, we show that in levels continuing firms are representative of the broader sample. Write:

$$\Omega_t^E = \omega_{\mathcal{C}_t,t} \Omega_{\mathcal{C}_t,t}^E + \omega_{\mathcal{E}_t,t} \Omega_{\mathcal{E}_t,t}^E \tag{B.2}$$

where $\Omega_{\mathcal{C}_t,t}^E = \sum_{i \in \mathcal{C}_t,t} \frac{\omega_{it}}{\omega_{\mathcal{C}_t,t}} \Omega_{it}^E$ the energy share within group $\mathcal{C}_t$, and we use similar notations for other groups of firms.

The left panel of Figure B.4 plots the share of continuing firms $\omega_{\mathcal{C}_t,t}$ in each time period. Continuing firms account for around 50% of the sample in 1999 and reach approximately 75% by the end of the sample. This ratio is far below 1 mostly because of entry *in the sample* as opposed to new plant creations. To show this, we exploit the fact that the ASI reports the year of creation of each establishment and reports the share of establishments created before year $t-1$. This share is significantly higher, averaging 96%. This also suggests that the increase in $\omega_{\mathcal{C}_t,t}$ over time is mostly attributable to an increase in the sampling rate, as opposed to an increase in the share of incumbents in total costs.

The right panel of Figure B.4 plots Ω_t^E (blue line) and $\Omega_{\mathcal{C}_t,t}^E$ (pink line) over time. The figure shows that energy intensity among continuing firms closely tracks aggregate energy intensity. Note that here the definition of incumbents changes in each period, hence two consecutive points do not include the same set of firms.

Second, we show that changes in energy intensity among continuing firms account for most of the decline in the aggregate energy intensity. Note that,

$$\begin{aligned}
\Omega_t^E &= \omega_{\mathcal{C}_t,t} \Omega_{\mathcal{C}_t,t}^E + \omega_{\mathcal{E}_t,t} \Omega_{\mathcal{E}_t,t}^E = \Omega_{\mathcal{C}_t,t}^E + \omega_{\mathcal{E}_t,t} (\Omega_{\mathcal{E}_t,t}^E - \Omega_{\mathcal{C}_t,t}^E) \\
\Omega_{t-1}^E &= \omega_{\mathcal{C}_t,t-1} \Omega_{\mathcal{C}_t,t-1}^E + \omega_{\mathcal{X}_t,t-1} \Omega_{\mathcal{X}_t,t-1}^E = \Omega_{\mathcal{C}_t,t-1}^E + \omega_{\mathcal{X}_t,t-1} (\Omega_{\mathcal{X}_t,t-1}^E - \Omega_{\mathcal{C}_t,t-1}^E)
\end{aligned}$$

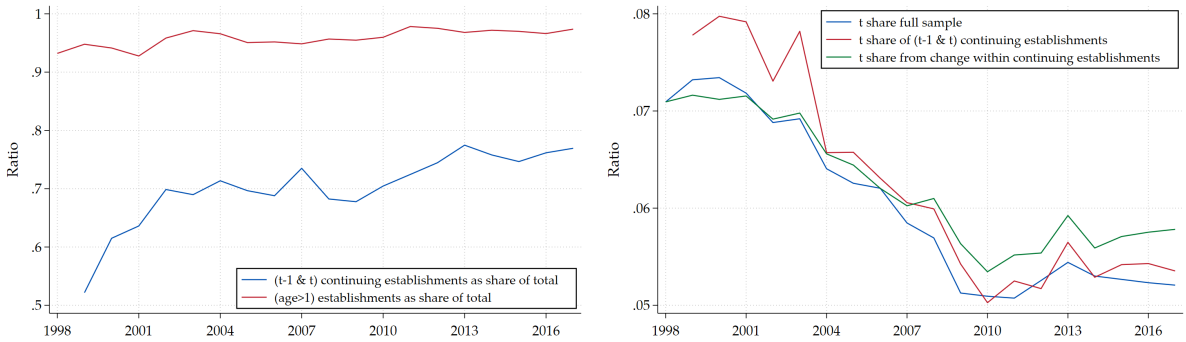
Therefore,

$$\Omega_t^E - \Omega_{t-1}^E = \underbrace{\Omega_{\mathcal{C}_t,t}^E - \Omega_{\mathcal{C}_t,t-1}^E}_{\text{Change among incumbents}} + \underbrace{\omega_{\mathcal{E}_t,t} (\Omega_{\mathcal{E}_t,t}^E - \Omega_{\mathcal{C}_t,t}^E) - \omega_{\mathcal{X}_t,t-1} (\Omega_{\mathcal{X}_t,t-1}^E - \Omega_{\mathcal{C}_t,t-1}^E)}_{\text{Contribution of entry and exit}} \tag{B.3}$$

FIGURE B.4: Energy Share Decomposition

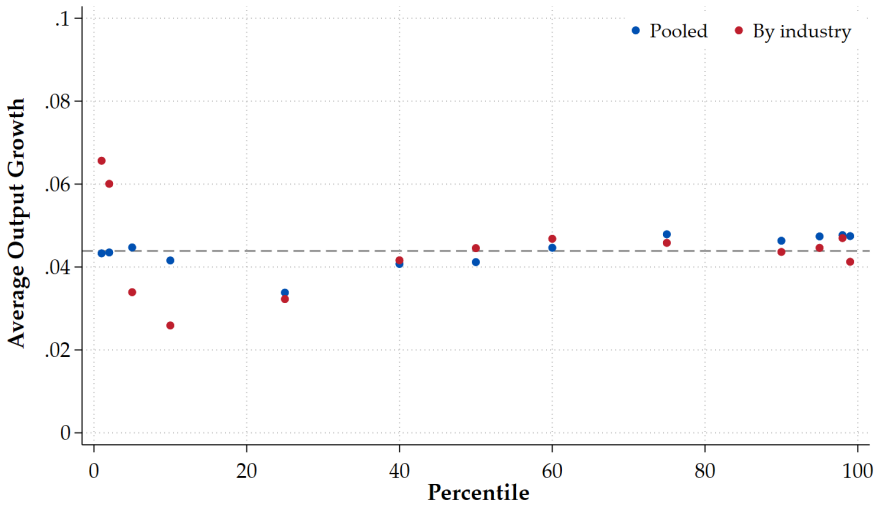
(A) Share of continuing establishments

(B) Energy share and role of continuing establishments



Note: Panel (a) reports the ratio of total costs by continuing establishments (defined as establishments present in the sample in year $t-1$ and t) over total costs for all establishments (blue line) and the ratio of total costs by establishments created before year $t-1$ over total costs for all establishments. Panel (b) reports three versions of the energy expenditure share. The blue line is Ω_t^E. The pink line is $\Omega_{\mathcal{C}_t,t}^E$. The green line is $\tilde{\Omega}_t^E$.

FIGURE B.5: Output Growth Along the Firm Size Distribution



To assess the contribution of incumbents to the overall dynamics of the energy share, we define:

$$\tilde{\Omega}_{1998}^E = \Omega_{1998}^E$$
$$\tilde{\Omega}_t^E = \tilde{\Omega}_{t-1} + \Omega_{\mathcal{C}_t,t}^E - \Omega_{\mathcal{C}_t,t-1}^E \text{ for } t > 1998$$

$\tilde{\Omega}_t^E$ is plotted in the green line in panel (b) of Figure B.4. $\tilde{\Omega}_t^E$ closely tracks Ω_t^E. That is the dynamics of the labor share among continuing firms for each pair of consecutive dates $(t-1, t)$ closely tracks the change in the aggregate labor share. This suggests that the contribution of entry and exit (the residual) is minimal.

Output Growth Along the Firm Size Distribution. In this section we show how output growth varies along the firm size distribution in the period of study.

In each year y and for each percentile p, we compute $Sales(p)_t$ the p-th percentile of the

distribution of real output (sales value) in year t. We define percentile p growth in year t as $\log(Sales(p)_t) - \log(Sales(p)_{t-1})$, and average these values over time for each percentile p. In the second version, we perform this procedure by 3-digit industries and aggregate using fixed industry sales weight to remove any effect driven by changes in the composition of industries.

Growth in the percentiles of the distribution is roughly even, except for some slight deviation at the bottom of the distribution. In the upper half, annual sales growth for the very largest firms is approximately equal to growth at the median from 1990 to 2020.

B.3 Firm Scale and Energy Demand

B.3.1 Instrument Validity: First Stage

Table B.3 reports the first stage results. The estimating equation is:

$$\Delta^h \log(Output)_{it} = \alpha_{st} + \beta^h \, \Delta^h \mathcal{D}_{it} + \varepsilon_{it}$$

The F-stat varies between 297 and 502 across horizons and specifications.

TABLE B.3: First Stage

	$h=1$		$h=3$		$h=5$	
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta^h \mathcal{D}_{it}$	0.285***	0.273***	0.288***	0.276***	0.272***	0.254***
	(0.013)	(0.012)	(0.014)	(0.014)	(0.016)	(0.015)
Year × Ind. FE	✓	✓	✓	✓	✓	✓
Year × Ind. × Cohort FE		✓		✓		✓
Observations	372,244	368,027	263,884	260,197	188,447	185,081
R-squared	0.066	0.19	0.092	0.22	0.11	0.25
F-stat	502.3	498.4	416.6	414.2	296.9	296.6

Note: This table presents the results of estimating the first stage of the IV regression in equation (3), applying analytic weights. Standard errors are clustered at the firm level. ***, ** and * indicate significance at the 1%, 5% and 10% levels, respectively.

B.3.2 Instrument Validity: Exogeneity

Orthogonality condition for the identification of the scale elasticity of energy demand. To clarify the discussion of identification, we relate our estimating equation to a simplified version of the production function used in the quantitative analysis of section 5, where we collapse the inner nest to include only energy. This simpler version suffices to highlight the concerns for the identification of scale elasticity of energy demand and discuss the validity of our instrument. We go back to identification of the full production function in section 5.

Consider a non-homothetic CES production function, implicitly defined through the constraint:

$$\left(\frac{e^{z_i} m_i}{y_i^{\gamma}}\right)^{\frac{\sigma-1}{\sigma}} + \left(\frac{e^{z_i+\phi_i} e_i}{y_i^{\gamma+\epsilon}}\right)^{\frac{\sigma-1}{\sigma}} = 1$$

e_i is energy, m_i is materials, and y_i is output. z_i and ϕ_i are factor-symmetric and energy-biased

(log) productivity states, potentially heterogeneous across firms.

Solving cost minimization leads to the following equation for the relative energy cost share:

$$\log\left(\frac{\Omega_{it}^{E}}{1-\Omega_{it}^{E}}\right) = \underbrace{(1-\sigma)\epsilon\,\log y_{it}}_{\text{Scale}} + \underbrace{(1-\sigma)\,\log\left(\frac{w_{it}^{M}}{w_{it}^{L}}\right)}_{\text{Substitution}} + \underbrace{(\sigma-1)\,\phi_{it}}_{\substack{\text{Energy-specific}\\ \text{productivity shock}}}$$

The first term captures the scale elasticity of relative energy demand: as long as $(1-\sigma)\epsilon \neq 0$, the energy expenditure share depends on scale. The goal of our empirical exercise is to identify $(1-\sigma)\epsilon$. The equation shows that the relative energy cost share depends on two additional forces: a substitution effect dependent on relative prices, and energy-biased productivity. This highlights that this requires exploiting variation in firm scale that is orthogonal to: (i) the relative price of energy, (ii) energy-specific productivity shocks.

Empirical specification and instrument. We estimate:

$$\Delta^{h}\log\left(\frac{\Omega_{it}^{E}}{1-\Omega_{it}^{E}}\right) = \alpha_{st} + \eta^{h}\,\Delta^{h}\log Output_{it} + \varepsilon_{it} \tag{B.4}$$

where $\Delta^{h}\log Output_{it}$ is instrumented with:

$$\mathcal{D}_{it}^{h} = \sum_{j\in\mathcal{J}} \omega_{ijt}\,\Delta^{h}\log Output_{jt} \tag{B.5}$$

The coefficient η^{h} provides an estimate of the structural scale elasticity of energy demand $(1-\sigma)\epsilon$ as long as the following orthogonality condition is satisfied:

$$\mathcal{D}_{it}^{h} \perp\!\!\!\perp \varepsilon_{it} \mid \alpha_{st} \tag{B.6}$$

The error term of our empirical specification ε_{it} captures the firm-specific determinants of the energy share, in particular the relative price of energy inputs and energy-specific productivity shocks. In what follows, we refer to both the relative price of energy inputs and energy-specific productivity shocks as "energy supply shocks".

Instrument Validity. It is useful to distinguish two types of potential issues. A first concern is that we do not observe the true product-level demand shocks, but instead estimate them as $\Delta^{h}\log Output_{jt}$, which may induce a mechanical correlation between $\mathcal{D}_{it}^{h}$ and ε_{it}. Second, even if we observed true product-level demand shocks, assumption (B.6) requires that firms exposed to products with large demand shocks are not subject to systematically different energy price or productivity shocks. This is the identifying assumption of the shift-share design. We discuss both concerns in turn.

Measurement of the shifters. The shifters $\Delta^{h}\log Output_{jt}$ are not exogenous shocks but product-level sales growth, an equilibrium outcome. In particular, product-level sales growth aggregates firm-specific shocks that may directly enter the residual ε_{it}. As shown in Borusyak et al. (2022), this mechanical bias tends to be problematic in settings where there are only few

firms contributing to each shifter estimate. We alleviate this concern in two ways. First, we construct $\Delta^h \log Output_{jt}$ excluding any firms with market share above 20% in the production of product j. Second, we construct a leave-one-out version of the instrument $\mathcal{D}_{it}^h$ where the shifters $\Delta^h \log Output_{jt}$ exclude observation i. Both variants yield highly similar results, suggesting that our strategy is relatively immune to this concern.

Shift-share identification. Assumption (B.6) requires that firm exposure to product-level growth is orthogonal to the unobserved determinants of energy shares ε_{it}, in particular any firm-specific energy supply shock. That is, firms must not sort across products such that firms with high (low) energy supply shocks systematically have high shares in high (low) growth products (Borusyak et al. 2022).

In our context, the main threat to identification is that product-level sales growth is correlated with product-level energy-specific supply shocks. For instance, it may be that product j is growing fast because of an energy-specific productivity improvement in the production process for this product. Then, firms specialized in j would benefit from this energy-specific productivity shock, and $\mathcal{D}_{it}^h$ would be correlated with firm-level energy-specific productivity shocks.

The most direct test supporting assumption (B.6) is firm-level balance on observables. Figure B.6(A) shows that firms with high and low $\mathcal{D}_{it}^h$ are similar on variables that are likely correlates of energy supply or productivity, e.g., factor shares, factor prices, having an electricity generator, or facing electricity shortages. Figure B.6(A) also reports correlations between the instrument and the lagged changes in the energy share, showing the absence of pre-trends. Balanced firm-level characteristics make it less likely that high $\mathcal{D}_{it}^h$ firms are systematically subject to different energy-specific supply shocks.

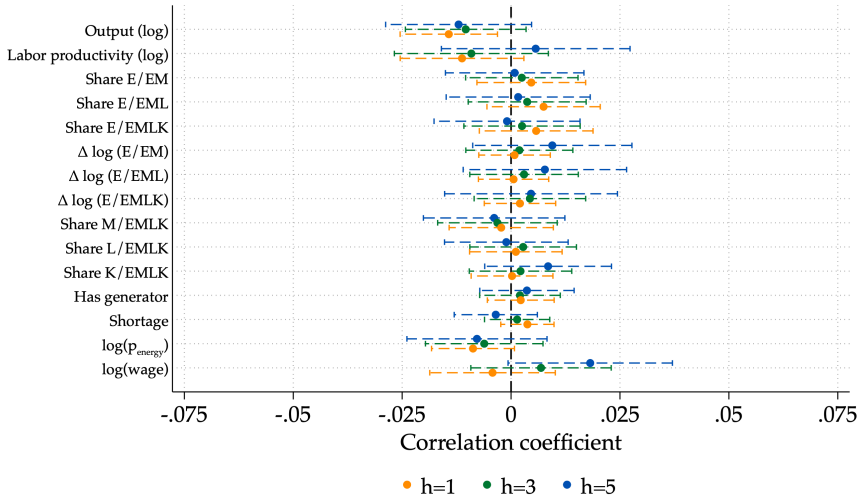
Following the argument in Borusyak et al. (2022), the orthogonality condition (B.6) can be expressed at the shifter-level: the necessary requirement for identification is that product-level shifters are uncorrelated with the average firm-level determinants of the energy share for the firms most exposed to each product. A sufficient condition for this to be true is that the product-level shifters are "as good as random". Figure B.7 shows product-level orthogonality tests: the change in product-level sales is not systematically correlated with lagged product-level factor shares, or the lagged change in those shares (i.e., product-level pre-trends). We define product-level variables as the sales-weighted average of firm-level variables. Since all our variables are defined at the firm-level, we find it more intuitive to focus on the firm-level balance tests described above.

Demand vs. productivity shifters. We further alleviate identification concerns by proposing tests that our demand instrument is indeed a demand shift as opposed to a productivity shock. Remember that it is *not* a problem for identification if scale is shifted by a factor-neutral productivity shock. However, it is plausible that if our instrument were correlated to energy-specific productivity, then it would be correlated to overall productivity too. It is therefore useful to validate that our instrument affects scale via a demand shift, as opposed to a factor-neutral productivity changes.

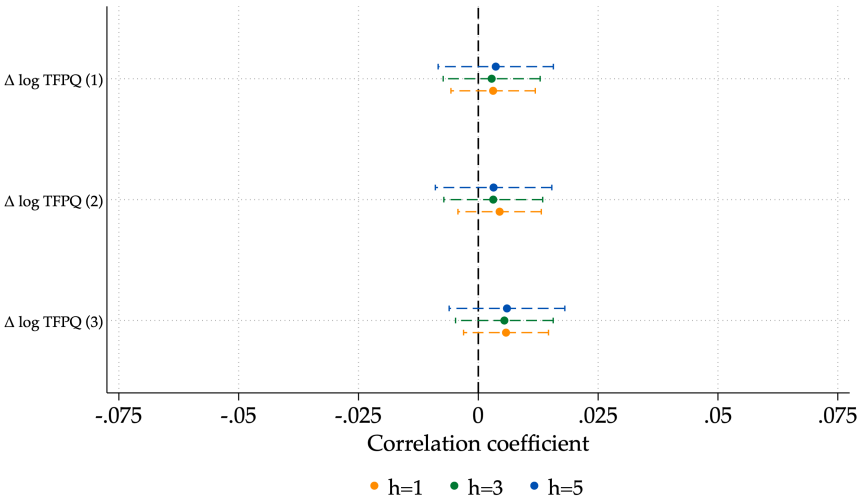
We produce estimates of the change in physical productivity and show that it is uncorrelated

FIGURE B.6: Correlation of Demand Shifter with Firm Characteristics

(A) Correlation with lagged firm characteristics



(B) Correlation with contemporaneous productivity shocks

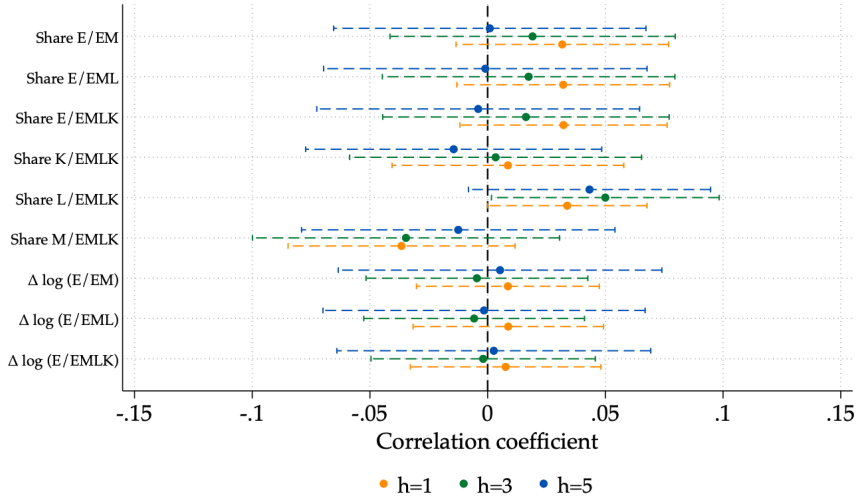


Note: This figure shows the coefficients of firm-level regressions of the firm-level demand shifter (defined in (4)) on firm characteristics. All variables are standardized so that the coefficients can be interpreted as correlation coefficients. Panel (A) reports correlations with lagged characteristics. Panel (B) reports correlations with the contemporaneous change in productivity (TFPQ), with the three alternative versions defined in section A.4. Regressions include industry×year×cohort fixed effects. Observations are weighted by firm-level total costs. Standard errors are clustered at the firm and main product×year level. The dot is the point estimate and the bar is the 95% confidence interval.

with our demand instrument. We construct the change in physical productivity as detailed in section A.4. Note that we observe the change in firm-level quantities sold so that we measure the change in physical as opposed to revenue productivity (that is, TFPQ as opposed to TFPR). These orthogonality tests are shown in Figure B.6(B).

Consistency. Exposure to common product-level shocks induce dependencies across firms with similar exposure shares, so that the setting is not *iid*. Borusyak et al. (2022) show that

FIGURE B.7: Shifter-level Orthogonality Tests



Note: This figure shows the coefficients of product-level regressions of the product-level shifter $\Delta^h \log Output_{jt}$ on lagged product-level characteristics. All variables are standardized so that the coefficients can be interpreted as correlation coefficients. Regressions include year fixed effects. Observations are weighted by product-level total sales. Standard errors are clustered at the product level. The dot is the point estimate and the bar is the 95% confidence interval.

the conditions for consistency are that (i) there is a sufficiently large number of shocks with sufficient shock-level variation, and (ii) that shocks exposure is not too concentrated. Panel A of Table B.4 documents a large dispersion in $\Delta^h \log Output_{jt}$, which persists when residualizing on fixed effects. Besides, exposure shares are not too concentrated. Define product-level weights as $s_{jt} = \sum_i Sales_{it}\ \omega_{ijt}$. Panel B shows that the largest weight is small (0.43%) and the inverse Herfindahl index is large (1,742). We report the same statistics when exposure weights are aggregated at the product-level, and there is sufficient product-level dispersion even when shocks are allowed to be serially correlated.[15]

[15] A a benchmark, Borusyak et al. (2022) show that their methodology is relevant in the canonical "China shock" setting where the inverse Herfindahl is 58.4 and the largest share is 6.5%.

TABLE B.4: Shock-level Summary Statistics

Panel A: Summary statistics on municipality-level shocks

	count	mean	sd	p25	p75
Product-level Δ^5 log(sales)	10,302	-0.023	0.671	-0.282	0.340
Residualized on time FE	10,302	0.000	0.652	-0.262	0.369
Residualized on industry $\times$ time FE	10,230	0.000	0.563	-0.183	0.268
Residualized on product FE	10,240	-0.000	0.560	-0.236	0.263

Panel B: Summary statistics on exposure shares

	Across products and dates	Across products
Inverse HHI	1742.451	132.657
Largest weight	.43%	3.06%

Note: This table presents descriptive statistics relevant for the shift-share design. Panel A presents summary statistics of the product-level shocks. Panel B presents summary statistics of product-level weights $s_{jt} = \sum_i Sales_{it}\ \omega_{ijt}$. Weights are normalized to sum to 1 for the whole sample. We compute the inverse Herfindahl index and the largest weight, and then the same quantities when weights are aggregated across time for a given product.

B.3.3 Main regressions

TABLE B.5: Relative energy expenditure share and firm size

Panel A: $\Delta^h \log\left(\Omega^E/(1-\Omega^E)\right)$; Ω^E = E/(E+M)

	OLS			IV		
	$h=1$	$h=3$	$h=5$	$h=1$	$h=3$	$h=5$
	(1)	(2)	(3)	(4)	(5)	(6)
Δ^h log(output)	-0.386***	-0.395***	-0.385***	-0.527***	-0.453***	-0.486***
	(0.008)	(0.008)	(0.008)	(0.042)	(0.037)	(0.041)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓
F-Stat				507.2	426.3	297.5
R-squared	0.225	0.275	0.301	0.078	0.123	0.124
Observations	367,500	260,301	185,305	366,569	259,125	184,285

Panel B: $\Delta^h \log\left(\Omega^E/(1-\Omega^E)\right)$; Ω^E = E/(E+M+L)

	OLS			IV		
	$h=1$	$h=3$	$h=5$	$h=1$	$h=3$	$h=5$
	(1)	(2)	(3)	(4)	(5)	(6)
Δ^h log(output)	-0.315***	-0.319***	-0.313***	-0.481***	-0.409***	-0.451***
	(0.008)	(0.008)	(0.008)	(0.033)	(0.029)	(0.034)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓
F-Stat				382.8	393.8	295.0
R-squared	0.074	0.103	0.114	0.055	0.093	0.088
Observations	392,889	282,570	205,479	370,868	262,870	187,690

Panel C: $\Delta^h \log\left(\Omega^E/(1-\Omega^E)\right)$; Ω^E = E/(E+M+L+K)

	OLS			IV		
	$h=1$	$h=3$	$h=5$	$h=1$	$h=3$	$h=5$
	(1)	(2)	(3)	(4)	(5)	(6)
Δ^h log(output)	-0.219***	-0.222***	-0.218***	-0.387***	-0.323***	-0.356***
	(0.008)	(0.008)	(0.009)	(0.032)	(0.029)	(0.034)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓
F-Stat				382.9	393.7	295.2
R-squared	0.039	0.054	0.060	0.018	0.043	0.034
Observations	392,571	282,229	205,163	370,582	262,553	187,411

Note: This table reports estimates of equation (3). Columns (1)–(3) report OLS estimates, while columns (4)–(6) report IV estimates where output growth $\Delta^h \log(\text{output})$ is instrumented with the firm-level demand shock defined in equation (4). The dependent variable is indicated in the panel headings. E, M, L, and K denote expenditures on energy, materials, labor, and capital, respectively. Regressions are weighted by firm-level total cost. Standard errors are clustered at the firm and main product×year level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

TABLE B.6: Energy expenditure share and firm size

Panel A: $\Delta^h \log \Omega^E$; $\Omega^E = \frac{w^E e}{py}$ (energy expenditures over gross output)

	OLS			IV		
	$h=1$	$h=3$	$h=5$	$h=1$	$h=3$	$h=5$
	(1)	(2)	(3)	(4)	(5)	(6)
Δ^h log(output)	-0.581***	-0.515***	-0.491***	-0.637***	-0.501***	-0.530***
	(0.007)	(0.008)	(0.008)	(0.038)	(0.034)	(0.039)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓
F-Stat				507.1	426.1	297.6
R-squared	0.330	0.358	0.379	0.214	0.229	0.230
Observations	367,613	260,383	185,352	366,682	259,207	184,332

Panel B: $\Delta^h \log \Omega^E$; $\Omega^E = \frac{w^E e}{py - w^M m}$ (energy expenditures over value added)

	OLS			IV		
	$h=1$	$h=3$	$h=5$	$h=1$	$h=3$	$h=5$
	(1)	(2)	(3)	(4)	(5)	(6)
Δ^h log(output)	-0.561***	-0.461***	-0.433***	-0.549***	-0.477***	-0.485***
	(0.014)	(0.012)	(0.012)	(0.062)	(0.044)	(0.050)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓
F-Stat				338.9	359.7	257.7
R-squared	0.073	0.084	0.089	0.071	0.082	0.083
Observations	342,026	245,025	177,950	322,646	227,760	162,418

Note: This table reports estimates of equation (3). Columns (1)–(3) report OLS estimates, while columns (4)–(6) report IV estimates where output growth $\Delta^h \log(\text{output})$ is instrumented with the firm-level demand shock defined in equation (4). The dependent variable is indicated in the panel headings. E, M, L, and K denote expenditures on energy, materials, labor, and capital, respectively. Regressions are weighted by firm-level total cost. Standard errors are clustered at the firm and main product×year level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

TABLE B.7: Energy expenditure share and firm size

Panel A: $\Delta^h \log \Omega^E$; Ω^E = E/M

	OLS			IV		
	$h=1$ (1)	$h=3$ (2)	$h=5$ (3)	$h=1$ (4)	$h=3$ (5)	$h=5$ (6)
Δ^h log(output)	-0.386***	-0.395***	-0.385***	-0.527***	-0.453***	-0.486***
	(0.008)	(0.008)	(0.008)	(0.042)	(0.037)	(0.041)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓
F-Stat				507.2	426.3	297.5
R-squared	0.225	0.275	0.301	0.078	0.123	0.124
Observations	367,500	260,301	185,305	366,569	259,125	184,285

Panel B: $\Delta^h \log \Omega^E$; Ω^E = E/L

	OLS			IV		
	$h=1$ (1)	$h=3$ (2)	$h=5$ (3)	$h=1$ (4)	$h=3$ (5)	$h=5$ (6)
Δ^h log(output)	0.044***	0.034***	0.026***	-0.025	0.012	-0.052
	(0.006)	(0.007)	(0.008)	(0.040)	(0.035)	(0.040)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓
F-Stat				506.6	425.9	297.5
R-squared	0.142	0.165	0.186	-0.002	0.001	-0.006
Observations	367,351	260,188	185,213	366,422	259,013	184,193

Panel C: $\Delta^h \log \Omega^E$; Ω^E = E/K

	OLS			IV		
	$h=1$ (1)	$h=3$ (2)	$h=5$ (3)	$h=1$ (4)	$h=3$ (5)	$h=5$ (6)
Δ^h log(output)	0.219***	0.212***	0.176***	0.161***	0.150***	0.122**
	(0.008)	(0.008)	(0.009)	(0.042)	(0.042)	(0.049)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓
F-Stat				507.3	426.7	297.1
R-squared	0.188	0.220	0.244	0.024	0.028	0.021
Observations	367,162	259,956	184,989	366,233	258,780	183,970

Panel D: $\Delta^h \log \Omega^E$; Ω^E = E/(L + K)

	OLS			IV		
	$h=1$ (1)	$h=3$ (2)	$h=5$ (3)	$h=1$ (4)	$h=3$ (5)	$h=5$ (6)
Δ^h log(output)	0.145***	0.142***	0.124***	0.088**	0.090**	0.057
	(0.007)	(0.007)	(0.008)	(0.038)	(0.036)	(0.041)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓
F-Stat				507.3	426.6	297.3
R-squared	0.164	0.198	0.222	0.013	0.017	0.012
Observations	367,304	260,052	185,067	366,374	258,876	184,047

Note: This table reports estimates of equation (3). Columns (1)–(3) report OLS estimates, while columns (4)–(6) report IV estimates where output growth $\Delta^h \log(\text{output})$ is instrumented with the firm-level demand shock defined in equation (4). The dependent variable is indicated in the panel headings. E, M, L, and K denote expenditures on energy, materials, labor, and capital, respectively. Regressions are weighted by firm-level total cost. Standard errors are clustered at the firm and main product×year level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

B.3.4 Quantity Regressions

TABLE B.8: Unweighted

	Targeted		Untargeted							
	$\tilde{\Omega}^{E}_{EML}$ (1)	Ω^{E}_{L} (2)	$\tilde{\Omega}^{E}_{EM}$ (3)	$\tilde{\Omega}^{E}_{EMLK}$ (4)	Ω^{E}_{LK} (5)	Ω^{E}_{VA} (6)	Ω^{E}_{Y} (7)	$\tilde{\Omega}^{M}_{EMLK}$ (8)	$\tilde{\Omega}^{L}_{EMLK}$ (9)	$\tilde{\Omega}^{K}_{EMLK}$ (10)
Δ^5 log(quantity)	-0.393***	-0.079**	-0.470***	-0.336***	0.010	-0.531***	-0.502***	0.474***	-0.263***	-0.473***
	(0.052)	(0.033)	(0.060)	(0.047)	(0.030)	(0.097)	(0.061)	(0.057)	(0.038)	(0.061)
FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Weighted										
F-Stat	73.0	73.1	72.9	72.2	72.2	44.0	73.2	71.0	71.0	71.3
Observations	138,804	138,704	138,778	138,011	138,006	119,905	138,815	138,341	138,267	138,332

Note: This table reports IV estimates of equation (3), where the main regressor is the five-year change in firm size, measured as $\Delta^5 \log(\text{quantity})$. The endogenous change in firm size is instrumented with the firm-level demand shock defined in equation (4). The dependent variable in each column is the five-year change in the log-odds of an expenditure share, with targeted cost shares in columns (1)–(5) and untargeted cost shares in columns (6)–(10); numerators and denominators are indicated in the column headings. All variables are winsorized at the 1% level. Regressions use ASI sampling weights only and include year × industry × cohort fixed effects. Standard errors are clustered at the firm and product-by-cohort levels. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

TABLE B.9: Weighted

	Targeted		Untargeted							
	$\tilde{\Omega}^{E}_{EML}$ (1)	Ω^{E}_{L} (2)	$\tilde{\Omega}^{E}_{EM}$ (3)	$\tilde{\Omega}^{E}_{EMLK}$ (4)	Ω^{E}_{LK} (5)	Ω^{E}_{VA} (6)	Ω^{E}_{Y} (7)	$\tilde{\Omega}^{M}_{EMLK}$ (8)	$\tilde{\Omega}^{L}_{EMLK}$ (9)	$\tilde{\Omega}^{K}_{EMLK}$ (10)
Δ^5 log(quantity)	-0.508***	-0.041	-0.595***	-0.427***	0.067	-0.801***	-0.654***	0.636***	-0.417***	-0.613***
	(0.088)	(0.046)	(0.098)	(0.079)	(0.049)	(0.200)	(0.103)	(0.097)	(0.071)	(0.102)
FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
F-Stat	42.1	42.1	42.1	42.0	42.1	19.9	42.1	41.7	41.7	41.7
Observations	138,245	138,148	138,219	138,011	138,006	119,473	138,256	138,341	138,267	138,332

Note: This table reports IV estimates of equation (3), where the main regressor is the five-year change in firm size, measured as $\Delta^5 \log(\text{quantity})$. The endogenous change in firm size is instrumented with the firm-level demand shock defined in equation (4). The dependent variable in each column is the five-year change in the log-odds of an expenditure share, with targeted cost shares in columns (1)–(5) and untargeted cost shares in columns (6)–(10); numerators and denominators are indicated in the column headings. All variables are winsorized at the 1% level. Regressions apply analytic weights based on input expenditures and include year × industry × cohort fixed effects. Standard errors are clustered at the firm and product-by-cohort levels. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

B.3.5 Robustness Checks

TABLE B.10: Heterogeneity: By single product and single plant

	Single Product?			Single Plant?		
	EM	EML	EMLK	EM	EML	EMLK
	(1)	(2)	(3)	(4)	(5)	(6)
Δ^5 log(output) × No	-0.385***	-0.305***	-0.216***	-0.329***	-0.253***	-0.188***
	(0.011)	(0.011)	(0.012)	(0.018)	(0.017)	(0.017)
Δ^5 log(output) × Yes	-0.395***	-0.313***	-0.222***	-0.402***	-0.321***	-0.226***
	(0.009)	(0.009)	(0.009)	(0.008)	(0.008)	(0.009)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓
R-squared	0.374	0.343	0.308	0.344	0.312	0.275
Observations	199,903	199,946	199,636	199,046	199,087	198,780

Note: This table reports estimates of equation (3). The dependent variable is the relative energy expenditure share, $\Delta^h \log(\Omega^E/(1-\Omega^E))$. Since the variable used to define subsamples is a t to $t+h$ change, we present only OLS results. The denominator of the energy expenditure share Ω^E varies by specification and is indicated in each column heading. E, M, L, and K denote expenditures on energy, materials, labor, and capital, respectively. Regressions are weighted by firm-level total cost. Standard errors are clustered at the firm and main product×year level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

TABLE B.11: Heterogeneity: By change in input adjustments

	By input adjustments		
	EM	EML	EMLK
	(1)	(2)	(3)
Δ^5 log(output) × None	-0.399***	-0.322***	-0.244***
	(0.010)	(0.010)	(0.010)
Δ^5 log(output) × Add only	-0.354***	-0.286***	-0.201***
	(0.013)	(0.013)	(0.013)
Δ^5 log(output) × Del only	-0.402***	-0.302***	-0.202***
	(0.014)	(0.013)	(0.013)
Δ^5 log(output) × Both	-0.398***	-0.326***	-0.221***
	(0.034)	(0.032)	(0.031)
Year × Ind. × Cohort FE	✓	✓	✓
Weighted	✓	✓	✓
R-squared	0.452	0.427	0.398
Observations	193,732	193,769	193,453

Note: This table reports estimates of equation (3). The dependent variable is the relative energy expenditure share, $\Delta^h \log(\Omega^E/(1-\Omega^E))$. Since the variable used to define subsamples is a t to $t+h$ change, we present only OLS results. The denominator of the energy expenditure share Ω^E varies by specification and is indicated in each column heading. E, M, L, and K denote expenditures on energy, materials, labor, and capital, respectively. Regressions are weighted by firm-level total cost. Standard errors are clustered at the firm and main product×year level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

TABLE B.12: Heterogeneity: By $h = 0$ output

	By $h = 0$ output					
	EM		EML		EMLK	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ^5 log(output) × Bin 1	-0.406***	-0.549***	-0.310***	-0.434**	-0.281***	-0.438**
	(0.111)	(0.185)	(0.102)	(0.172)	(0.106)	(0.184)
Δ^5 log(output) × Bin 2	-0.563***	-0.566***	-0.479***	-0.456***	-0.424***	-0.389***
	(0.064)	(0.111)	(0.061)	(0.105)	(0.060)	(0.104)
Δ^5 log(output) × Bin 3	-0.459***	-0.475***	-0.406***	-0.425***	-0.353***	-0.352***
	(0.036)	(0.053)	(0.034)	(0.051)	(0.033)	(0.051)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓
Weighted		✓		✓		✓
R-squared	0.182	0.116	0.135	0.079	0.079	0.029
Observations	180,180	179,515	180,214	179,549	179,270	179,270

Note: This table reports estimates of equation (3). The dependent variable is the relative energy expenditure share, $\Delta^h \log(\Omega^E/(1-\Omega^E))$. Since the variable used to define subsamples is a t to $t+h$ change, we present only IV results. The denominator of the energy expenditure share Ω^E varies by specification and is indicated in each column heading. E, M, L, and K denote expenditures on energy, materials, labor, and capital, respectively. Regressions are weighted by firm-level total cost. Standard errors are clustered at the firm and main product×year level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

TABLE B.13: Heterogeneity: By change in days worked and output sign

	By % of days worked			By output growth sign		
	EM	EML	EMLK	EM	EML	EMLK
	(1)	(2)	(3)	(4)	(5)	(6)
Δ^5 log(output) × Negative	-0.371***	-0.287***	-0.181***			
	(0.013)	(0.013)	(0.013)			
Δ^5 log(output) × Close-zero	-0.412***	-0.340***	-0.275***			
	(0.010)	(0.010)	(0.009)			
Δ^5 log(output) × Positive	-0.369***	-0.292***	-0.185***			
	(0.012)	(0.013)	(0.013)			
Δ^5 log(output) × $\Delta \ln Y < 0$				-0.416***	-0.299***	-0.179***
				(0.013)	(0.013)	(0.013)
Δ^5 log(output) × $\Delta \ln Y \geq 0$				-0.331***	-0.276***	-0.190***
				(0.014)	(0.014)	(0.014)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓
R-squared	0.393	0.365	0.336	0.380	0.351	0.318
Observations	197,313	197,353	197,038	200,101	200,143	199,831

Note: This table reports estimates of equation (3). The dependent variable is the relative energy expenditure share, $\Delta^h \log(\Omega^E/(1-\Omega^E))$. Since the variable used to define subsamples is a t to $t+h$ change, we present only OLS results. The denominator of the energy expenditure share Ω^E varies by specification and is indicated in each column heading. E, M, L, and K denote expenditures on energy, materials, labor, and capital, respectively. Regressions are weighted by firm-level total cost. Standard errors are clustered at the firm and main product×year level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

B.3.6 Firm Scope and Indirect Energy Consumption.

We define a firms indirect energy purchase as the energy embodied in its purchases of material inputs. The total energy share is the sum of the direct and indirect use of energy.

Expenditures on material inputs $w_{it}^M m_{it}$ can be decomposed into distinct varieties indexed by k: $w_{it}^M m_{it} = \sum_k w_{kit}^M m_{kit}$. We can then write the direct and indirect energy expenditure share of firm i as:

$$\begin{aligned}
\underbrace{\Omega_{it}^{E,Total}}_{\substack{\text{Total energy}\\ \text{share for firm } i}} &= \underbrace{\frac{w_{it}^E e_{it}}{\sum_{X\in\mathcal{X}} w_{it}^X x_{it}}}_{\substack{\text{Direct energy}\\ \text{share for firm } i}} + \sum_{k=1}^{\mathcal{K}} \underbrace{\frac{w_{kit}^M m_{kit}}{\sum_{X\in\mathcal{X}} w_{it}^X x_{it}}}_{\substack{\text{Material variety } k\\ \text{share for firm } i}} \underbrace{\Omega_{kt}^{E,Total}}_{\substack{\text{Total energy}\\ \text{share for variety } k}} \\
&= \underbrace{\Omega_{it}^E}_{\substack{\text{Direct energy}\\ \text{share for firm } i}} + \underbrace{\Omega_{it}^M}_{\substack{\text{Direct material}\\ \text{share for firm } i}} \times \underbrace{\sum_{k=1}^{\mathcal{K}} \omega_{ikt} \times \Omega_{kt}^{E,Total}}_{\substack{\text{Average energy share}\\ \text{of materials}\\ \Psi_{it}^E}}
\end{aligned}$$

where $\omega_{kit} = \frac{w_{ikt}^M m_{ikt}}{w_{it}^M M_{it}}$ is the share of variety k in total material inputs of firm i.

We solve for $\Omega_{kt}^{E,Total}$ by aggregating the preceding equation at the product-code level, and solving the matrix equation:

$$\begin{aligned}
\mathbf{\Omega^{E,Total}} &= \mathbf{\Omega^E} + \text{diag}(\mathbf{\Omega^M})\, \mathbf{\Omega^{M,M}}\, \mathbf{\Omega^{E,Total}} \\
&= \left(\mathbf{I} - \text{diag}(\mathbf{\Omega^M})\, \mathbf{\Omega^{M,M}}\right)^{-1}\, \mathbf{\Omega^E}
\end{aligned}$$

where $\mathbf{\Omega^E}$ and $\mathbf{\Omega^M}$ are $\mathcal{K} \times 1$ vectors of product-level direct energy and material shares, and $\mathbf{\Omega^{M,M}}$ is the $\mathcal{K} \times \mathcal{K}$ input-output matrix at the product level.

The effects for the total energy share are lower than the direct because the decline in Ω_{it}^E is partially offset by the increase in Ω_{it}^M (even if Ψ_{it}^E does not react).

TABLE B.14: Firm Growth and Direct and Indirect Energy Share

	$\bar{\Psi}^E_{it}$	Energy share					
		E/EM		E/EML		E/EMLK	
		Direct	Total	Direct	Total	Direct	Total
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Δ^5 log(output)	0.001	-0.028***	-0.022***	-0.016***	-0.011***	-0.012***	-0.006**
	(0.001)	(0.004)	(0.004)	(0.003)	(0.003)	(0.002)	(0.003)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓	✓
F-Stat	289.989	334.837	289.771	334.843	289.771	326.124	284.009
Observations	176,161	185,500	176,106	185,502	176,106	184,527	175,284

Note: Column (1) regresses the average energy product energy share of the firm's intermediate inputs on the change in log output at a 5-year horizon, instrumented according to equation (4). Columns (2)-(7) regress the direct and total energy shares on 5-year changes in log output, instrumented according to equation (4).

B.3.7 Heterogeneity

TABLE B.15: Heterogeneity: By change in capital

	By change in capital							
	EM				EML			
	K^{full}	K^{equip}	K^{full}/Y	K^{equip}/Y	K^{full}	K^{equip}	K^{full}/Y	K^{equip}/Y
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ^5 log(output) × ↓↓	-0.329***	-0.212***	-0.320***	-0.316***	-0.210**	-0.116**	-0.227***	-0.225***
	(0.100)	(0.055)	(0.035)	(0.031)	(0.093)	(0.051)	(0.032)	(0.031)
Δ^5 log(output) × ↓	-0.307***	-0.537***	-0.203***	-0.278***	-0.215***	-0.393***	-0.139***	-0.221***
	(0.080)	(0.039)	(0.046)	(0.046)	(0.079)	(0.038)	(0.044)	(0.044)
Δ^5 log(output) × ≈	-0.433***	-0.438***	-0.294***	-0.299***	-0.343***	-0.349***	-0.253***	-0.255***
	(0.009)	(0.009)	(0.011)	(0.011)	(0.009)	(0.010)	(0.010)	(0.010)
Δ^5 log(output) × ↑	-0.436***	-0.415***	-0.282***	-0.244***	-0.379***	-0.363***	-0.215***	-0.187***
	(0.021)	(0.028)	(0.019)	(0.018)	(0.021)	(0.026)	(0.017)	(0.017)
Δ^5 log(output) × ↑↑	-0.333***	-0.388***	-0.360***	-0.371***	-0.270***	-0.317***	-0.218***	-0.238***
	(0.024)	(0.023)	(0.018)	(0.017)	(0.023)	(0.022)	(0.017)	(0.016)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓	✓	✓
R-squared	0.407	0.416	0.432	0.437	0.379	0.388	0.407	0.412
Observations	191,481	189,579	189,573	188,301	191,522	189,623	189,616	188,343

Note: This table reports estimates of equation (3). The dependent variable is the relative energy expenditure share, $\Delta^h \log(\Omega^E/(1-\Omega^E))$. Since the variable used to define subsamples is a t to $t+h$ change, we present only OLS results. The denominator of the energy expenditure share Ω^E varies by specification and is indicated in each column heading. E, M, L, and K denote expenditures on energy, materials, labor, and capital, respectively. Regressions are weighted by firm-level total cost. Standard errors are clustered at the firm and main product×year level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

TABLE B.16: Firm growth and energy sources

	Energy sources				On-site elec. generation
	Coal	Oil	Elec.	Gas	
	(1)	(2)	(3)	(4)	(5)
Δ^5 log(output)	0.006	-0.001	-0.008	0.001	0.017
	(0.007)	(0.014)	(0.015)	(0.009)	(0.015)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓
F-Stat	297.6	297.6	297.6	135.6	297.2
Observations	184,218	184,218	184,218	75,819	184,967

Note: This table reports estimates of equation (3) where output growth Δ^h log(output) is instrumented with the firm-level demand shock defined in equation (4). In columns (1)-(4), the dependent variable is the change in the cost share of energy source f in total energy expenditures, where the energy source f is indicated in the column heading. In column (5), the dependent variable is the change in the share of electricity generated on site. We use changes in shares (as opposed to log shares) to accommodate the many zeros in energy source-specific shares. Data on expenditures on natural gas is available only since 2008. Regressions are weighted by firm-level total cost. Standard errors are clustered at the firm and main product×year level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

TABLE B.17: By fraction of energy used in fired systems

	Δ^5 log($\Omega^E/(1-\Omega^E)$)						Δ^5 log(e/y)	
	EM		EML		EMLK			
	Low	High	Low	High	Low	High	Low	High
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ^5 log(py)	-0.436***	-0.663***	-0.373***	-0.612***	-0.305***	-0.529***		
	(0.042)	(0.132)	(0.040)	(0.126)	(0.041)	(0.128)		
Δ^5 log(y)							-0.364***	-0.498***
							(0.113)	(0.182)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓	✓	✓
F-Stat	253.5	43.1	253.5	43.2	253.2	43.2	27.7	10.0
Observations	133,790	38,331	133,800	38,349	133,581	38,310	99,666	28,283

Note: This table reports estimates of equation (3) where output growth is instrumented with the firm-level demand shock defined in equation (4). Results are reported separately for firms below (above) the median of the fraction of total energy used in fired systems (defined at the industry level, see Table B.2). Regressions are weighted by firm-level total cost. Standard errors are clustered at the firm and main product×year level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

B.4 Results in the Combined ASI-NSS Sample

We construct two versions of the energy expenditure share that can be defined in the same way in ASI and in the NSS survey. First, we consider energy expenditures over output. Second, we consider energy expenditures over cost of goods sold (roughly equivalent to using materials plus energy as a denominator in our baseline version).

TABLE B.18: Energy expenditure shares and firm size: ASI vs. NSS

	Log of Energy/Expenses			Log of Energy/Output		
	(1) ASI	(2) NSS	(3) Pooled	(4) ASI	(5) NSS	(6) Pooled
log(output)	-0.291***	-0.371***	-0.342***	-0.091***	-0.224***	-0.053***
	(0.002)	(0.002)	(0.001)	(0.002)	(0.001)	(0.001)
Ind. × Cohort FE	✓	✓	✓	✓	✓	✓
R-squared	0.398	0.375	0.370	0.417	0.340	0.376
Observations	158,080	308,761	466,842	156,888	308,684	465,573

Note: This table reports OLS regressions of energy use on firm size in the combined ASI–NSS sample. In all specifications the regressor is log(output), defined as the log of real sales. The first three columns use as dependent variable the log share of energy expenditures in cost of goods sold, $\log(E/X)$, where E denotes energy expenditures and X denotes cost of goods sold. The last three columns use the log of energy intensity, $\log(E/Y)$, where Y denotes output. Columns labeled "ASI" and "NSS" estimate separate regressions for organized (ASI) and unorganized (NSS) manufacturing plants, while the "Pooled" columns stack both datasets. All regressions include industry-by-cohort fixed effects, and standard errors are clustered at the firm level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

B.5 Is the mechanism specific to LMICs? Suggestive evidence from the United States

In this section we provide suggestive evidence that our results generalize beyond LMICs like India to high income countries.

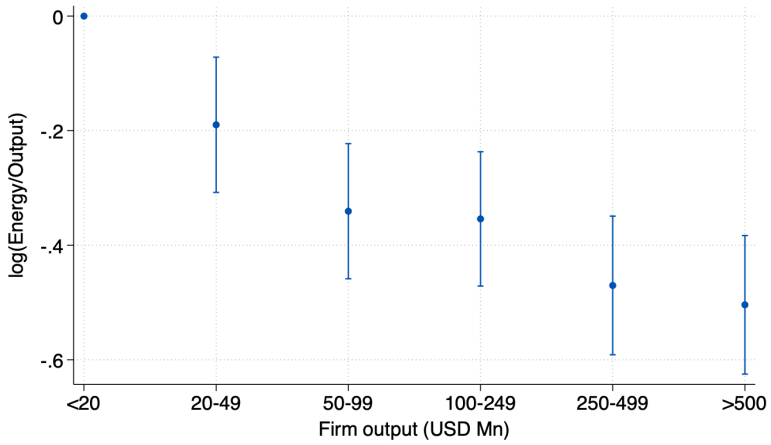
The Manufacturing Energy Consumption Survey for the United States reports energy intensity (energy in physical unit per value of shipments in USD millions) by bins of firm size (6 bins defined by value of shipments in USD millions). The data is available at intervals of four years. We use the waves 2002 through 2018. The data is provided for a selected number of industries: in the modal wave we have forty-five 6-digit industries, one 5-digit industry, seven 4-digit industries, and eleven 3-digit industries. We estimate the following specification:

$$\log(E/Y)_{qst} = \alpha_{st} + \sum_{x=1}^{6} \beta_x [Bin = x]_{qst} + \varepsilon_{qst} \tag{B.7}$$

where q indexes the six size bins, and α_{st} are industry×time fixed effects.

The results are presented in Figure B.8. The figure shows that energy intensity declines in firm size in the United States too. While suggestive, one should note several caveats to this analysis. First, we can only estimate this relationship in the cross-section of firms, which is sensitive to the endogeneity concerns mentioned above. Second, we do not observe the energy expenditure

FIGURE B.8: Energy intensity by firm size in the United States



Note: The figure shows the results of estimating equation (B.7). The dot is the point estimate and the bar is the 95% confidence interval.

share but only the energy-to-output ratio, which may be affected by differential energy prices or differential output markups across the firm size distribution.

C. Proofs and Derivations

C.1 Derivation of Equation (7) for Optimal Technology Choice

The problem of the firm is to choose technology to minimize the cost of output y, subject to producing at capacity. In other words, they solve

$$\min_{\phi} p^e e + \frac{1}{\gamma A_{Yt}} \phi^{\gamma} \bar{y}^{\delta} w^{\alpha} r^{1-\alpha}$$

subject to

$$y = A_{Yt} \phi \frac{e}{y^{\epsilon}}$$

and $y = \bar{y}$ for a given $\bar{y}$. This has the associated first order condition

$$w^{\alpha} r^{1-\alpha} \phi^{\gamma-1} y^{\delta} = \frac{p^e y^{1+\epsilon}}{\phi^2},$$

which we can solve for optimal ϕ as

$$\phi = \left(\frac{p^e y^{1+\epsilon-\delta}}{w^{\alpha} r^{1-\alpha}} \right)^{\frac{1}{1+\gamma}}.$$

Energy demand for a given level of output y is then

$$e = \frac{y^{1+\epsilon}}{\phi A_Y} = \frac{1}{A_Y} y^{\frac{(1+\epsilon)\gamma+\delta}{1+\gamma}} \left(\frac{w^{\alpha} r^{1-\alpha}}{p^e} \right)^{\frac{1}{1+\gamma}}$$

Then energy intensity measured in joules per dollar of output, is

$$\Theta(a) = \frac{e}{p(y)y} = \frac{1}{a A_Y Y^{\frac{1}{\lambda}}} y^{\frac{(1+\epsilon)\gamma+\delta}{1+\gamma} - \zeta} \left(\frac{w^{\alpha} r^{1-\alpha}}{p^e} \right)^{\frac{1}{1+\gamma}}$$

C.2 Proof of Proposition 1

We begin from the expression for the change in energy intensity at the aggregate level:

$$d\bar{\Theta} = \int_0^{\infty} \underbrace{d\Theta(a) \cdot \frac{p(y)y(a)}{Y/N} d\Gamma(a)}_{\text{Micro intensity changes } \equiv d\bar{\Theta}_1} + \int_0^{\infty} \underbrace{\Theta(a) \cdot d\frac{p(y)y(a)}{Y/N} d\Gamma(a)}_{\text{Reallocation } \equiv d\bar{\Theta}_2}.$$

We note that output of the firm is given from the production function by

$$y = A_Y \phi \frac{e}{y^{\epsilon}}.$$

Using the optimal solution to firm-level efficiency ϕ from equation (7), as well as the expression

for revenue $p(a)y(a) = ay^{\zeta}Y^{\frac{1}{\lambda}}$, we have

$$\Theta(a) = \frac{1}{aA_Y Y^{\frac{1}{\lambda}}} y(a)^{\frac{(1+\epsilon)\gamma+\delta}{1+\gamma}-\zeta} \left(\frac{w^{\alpha} r^{1-\alpha}}{p^e}\right)^{\frac{1}{1+\gamma}},$$

and so, given a general equilibrium movement in aggregate productivity A_Y, we have

$$d\log\Theta(a) = \left(\frac{(1+\epsilon)\gamma+\delta}{1+\gamma} - \zeta\right) d\log y(a) - d\log\left(A_Y Y^{\frac{1}{\lambda}}\right) + \frac{1}{1+\gamma} d\log\left(\frac{w^{\alpha} r^{1-\alpha}}{p^e}\right).$$

Then we can write

$$\begin{aligned} d\bar{\Theta}_1 &= \left(\frac{(1+\epsilon)\gamma+\delta}{1+\gamma} - \zeta\right) \mathbb{E}_R[\Theta(a) d\log y(a)] - \mathbb{E}_R[\Theta(a)] d\log\left(A_Y Y^{\frac{1}{\lambda}}\right) \\ &+ \frac{1}{1+\gamma}\mathbb{E}_R[\Theta] d\log\left(\frac{w^{\alpha} r^{1-\alpha}}{p^e}\right) \\ &= \left(\frac{(1+\epsilon)\gamma+\delta}{1+\gamma} - \zeta\right)\left(\mathbb{E}_R[\Theta(a)]\mathbb{E}_R[d\log y(a)] + Cov_R[\Theta(a) d\log y(a)]\right) \\ &- \mathbb{E}_R[\Theta] d\log\left(A_Y Y^{\frac{1}{\lambda}}\right) + \frac{1}{1+\gamma}\mathbb{E}_R[\Theta(a)] d\log\left(\frac{w^{\alpha} r^{1-\alpha}}{p^e}\right). \end{aligned} \tag{C.1}$$

For the second term, we have

$$\begin{aligned} d\bar{\Theta}_2 &= \int_0^{\infty} \Theta(a) d\log(p(a)y(a)) \frac{p(a)y(a)}{Y/N} d\Gamma(a) - \int_0^{\infty} \Theta(a) d\log(Y_t/N_t) \frac{p(a)y(a)}{Y/N} d\Gamma(a) \\ &= \mathbb{E}_R[\Theta(a)]\mathbb{E}_R[d\log p(a)y(a)] + Cov_R[\Theta(a) d\log p(a)y(a)] - \mathbb{E}_R[\Theta(a)] d\log(Y/N) \\ &= \mathbb{E}_R[\Theta(a)]\mathbb{E}_R[d\log p(a)y(a)] + \zeta Cov_R[\Theta(a) d\log y(a)] - \mathbb{E}_R[\Theta(a)] d\log(Y/N). \end{aligned} \tag{C.2}$$

Combining (C.1) and (C.2) we have

$$\begin{aligned} d\log(\bar{\Theta}) &= \left(\frac{(1+\epsilon)\gamma+\delta}{1+\gamma} - \zeta\right) \mathbb{E}_R[d\log y(a)] \\ &- d\log\left(A_Y Y^{\frac{1}{\lambda}}\right) + \frac{1}{1+\gamma} d\log\left(\frac{w^{\alpha} r^{1-\alpha}}{p^e}\right) \\ &+ \left(\mathbb{E}_R[d\log p(a)y(a)] - d\log(Y_t/N_t)\right) + \frac{(1+\epsilon)\gamma+\delta}{1+\gamma} Cov_R[\frac{\theta(a)}{\bar{\Theta}}, d\log y(a)]. \end{aligned}$$

C.3 Proof of Proposition 2

Write the free entry condition as

$$\begin{aligned} wv &= \int_a V(a) d\Gamma(a) \\ &= \int_a \left(\max_y \quad ay^{\zeta}Y^{\frac{1}{\lambda}} - c(A_Y; y)\right), \end{aligned}$$

where the cost function is given by equation (8).

By the envelope theorem

$$\frac{\partial V(a)}{\partial Y^{\frac{1}{\lambda}}} = ay^{\zeta} \quad \frac{\partial V(a)}{\partial w} = -\frac{\partial c(a, A_Y; y)}{\partial w} \quad \frac{\partial V(a)}{\partial p^e} = -\frac{\partial c(a, A_Y; y)}{\partial p}$$
$$\frac{\partial V(a)}{\partial p^m} = -\frac{\partial c(a, A_Y; y)}{\partial p^m} \quad \frac{\partial V(a)}{\partial r} = -\frac{\partial c(a, A_Y; y)}{\partial r}.$$

Totally differentiating the free entry condition and using these expressions gives

$$\begin{aligned} vdw = \int_a \Bigg(& aY^{\frac{1}{\lambda}} y^{\zeta} d\log\left(Y^{\frac{1}{\lambda}}\right) - \frac{\partial c(a, A_Y; y)}{\partial w} dw - \frac{\partial c(a, A_Y; y)}{\partial r} dr \\ & - \frac{\partial c(a, A_Y; y)}{\partial p^e} dp^e - \frac{\partial c(a, A_Y; y)}{\partial p^m} dp^m - \frac{\partial c(a, A_Y; y)}{\partial A_Y} dA_Y \Bigg) d\Gamma(a). \end{aligned} \tag{C.3}$$

Note also by Shepard's Lemma we have

$$\begin{aligned} vw = \int_a \Bigg(& aY^{\frac{1}{\lambda}} y^{\zeta} - \frac{\partial c(a, A_Y; y)}{\partial w} w - \frac{\partial c(a, A_Y; y)}{\partial r} \\ & - \frac{\partial c(a, A_Y; y)}{\partial p^e} p^e - \frac{\partial c(a, A_Y; y)}{\partial p^m} p^m \Bigg) d\Gamma(a). \end{aligned} \tag{C.4}$$

Combine (C.4) and (C.3) to get

$$\begin{aligned} & vwd\log(w) + \int_a \Bigg(\frac{\partial c(a, A_Y; y)}{\partial w} wd\log(w) + \frac{\partial c(a, A_Y; y)}{\partial p^e} p^e d\log(p^r) \\ & + \frac{\partial c(a, A_Y; y)}{\partial r} rd\log(r) + \frac{\partial c(a, A_Y; y)}{\partial p^m} p^m d\log(p^m) \Bigg) d\Gamma(a) \\ & = \left(vw + \int_a \left(\frac{\partial c(a, A_Y; y)}{\partial w} w + \frac{\partial c(a, A_Y; y)}{\partial r} r + \frac{\partial c(a, A_Y; y)}{\partial p^e} p^e + \frac{\partial c(a, A_Y; y)}{\partial p^m} p^m \right) d\Gamma(a) \right) d\log\left(Y^{\frac{1}{\lambda}}\right) \\ & - \int_a \left(\frac{\partial c(a, A_Y; y)}{\partial A} \right) d\Gamma(a) dA_Y. \end{aligned}$$

Multiply by the number of firms, use the fact $d\log(p^e) = d\log(r) = d\log(p^m) = 0$ in general equilibrium, and again use Shephard's Lemma, and this becomes

$$(\Phi_V^w + \Phi_E^w) d\log(w) = (\Phi_E^w + \Phi_V^w + \Phi_V^E + \Phi_V^r + \Phi_V^m) d\log\left(Y^{\frac{1}{\lambda}}\right) + \left(\Phi_V^w + \Phi_V^E + \Phi_V^r + \Phi_V^m\right) d\log(A). \tag{C.5}$$

where Φ_V^w is total payments to labor in economy-wide variable costs, with analogous notation for the other terms, and Φ_E^w is total payments to labor in economy-wide entry costs. This, in turn, is equal to total profits Π via free entry. Hence, (C.5) can be rearranged to yield

$$\frac{wL}{Y} d\log(w) + \frac{\Pi}{Y} d\log(A_Y) = d\log\left(A_Y Y^{\frac{1}{\lambda}}\right).$$

C.4 Derivation of Materials Share

At the micro level, we have materials per dollar of revenue as

$$\frac{m(a)}{p(a)y(a)} = \frac{y(a)^{1-\zeta}}{aA_Y Y^{\frac{1}{\lambda}}}.$$

In the aggregate we have

$$d\frac{M}{Y} = \int_0^\infty d\frac{m(a)}{p(a)y(a)} \cdot \frac{p(y)y(a)}{Y/N} d\Gamma(a) + \int_0^\infty \frac{m(a)}{p(a)y(a)} \cdot d\frac{p(y)y(a)}{Y/N} d\Gamma(a).$$

Repeating the derivations above for the energy intensity, we have for the first term

$$\begin{aligned}
\int_0^\infty d\frac{m(a)}{p(a)y(a)} \cdot \frac{p(y)y(a)}{Y/N} d\Gamma(a) &= (1-\zeta)\mathbb{E}_R[\frac{m(a)}{p(a)y(a)} d\log y(a)] - \mathbb{E}_R[\frac{m(a)}{p(a)y(a)}] d\log\Big(A_Y A_Y Y^{\frac{1}{\lambda}}\Big) \\
&= (1-\zeta)\mathbb{E}_R[\frac{m(a)}{p(a)y(a)}]\mathbb{E}_R[d\log y(a)] + (1-\zeta)Cov_R(\frac{m(a)}{p(a)y(a)} d\log y(a)) \\
&- \mathbb{E}_R[\frac{m(a)}{p(a)y(a)}] d\log\Big(A_Y Y^{\frac{1}{\lambda}}\Big),
\end{aligned}$$

and for the second term

$$\begin{aligned}
\int_0^\infty \frac{m(a)}{p(a)y(a)} \cdot d\frac{p(y)y(a)}{Y/N} d\Gamma(a) &= \int_0^\infty \frac{m(a)}{p(a)y(a)} \frac{p(y)y(a)}{Y/N} \cdot d\log\frac{p(y)y(a)}{Y/N} d\Gamma(a) \\
&= \mathbb{E}_R[\frac{m(a)}{p(a)y(a)} d\log p(y)y(a)] - \mathbb{E}_R[\frac{m(a)}{p(a)y(a)}] d\log(Y/N) \\
&= \mathbb{E}_R[\frac{m(a)}{p(a)y(a)}]\mathbb{E}_R[d\log p(y)y(a)] + \zeta Cov_R(\frac{m(a)}{p(a)y(a)} d\log y(a)) \\
&- \mathbb{E}_R[\frac{m(a)}{p(a)y(a)}] d\log(Y/N).
\end{aligned}$$

Combining these, we find

$$\begin{aligned}
d\log\Big(\frac{M}{Y}\Big) &= (1-\zeta)\mathbb{E}_R[d\log y(a)] - d\log\Big(A_Y Y^{\frac{1}{\lambda}}\Big) \\
&+ Cov_R(\frac{m(a)}{p(a)y(a)} / (\frac{M}{Y}), d\log y(a)) + \Big(\mathbb{E}_R[d\log p(y)y(a)] - d\log(Y/N)\Big).
\end{aligned}$$

To see that this is negative if the number of firms increases, first note that we can write aggregate output as

$$Y = (N \int_0^\infty a y(a)^{\frac{\lambda-1}{\lambda}} d\Gamma(a))^{\frac{\lambda}{\lambda-1}},$$

so that

$$\begin{aligned} d\log(Y) &= \frac{\lambda}{\lambda-1} d\log(N) + \left(\int_0^\infty a \frac{y(a)^{\frac{\lambda-1}{\lambda}}}{Y/N} d\log(y) d\Gamma(a) \right) \\ &= \frac{\lambda}{\lambda-1} d\log(N) + E_R[d\log(y)], \end{aligned}$$

which implies

$$E_R[d\log(y)] = d\log(Y/N) - \frac{1}{\lambda-1} d\log(N).$$

So, recalling that $(1-\zeta) = \frac{1}{\lambda}$ we can write,

$$d\log\left(\frac{M}{Y}\right) = -d\log(A_Y) - \frac{1}{\lambda-1} d\log(N_t) + Cov_R(\frac{m(a)}{p(a)y(a)}/(\frac{M}{Y}), d\log y(a))$$

C.5 Determination of the Number of Firms N_t

First note that the labor market clearing condition can be written

$$\begin{aligned} w_t L_t &= w_t \nu N_t + w_t L_t^P \\ &= \Pi_t + w_t L_t^P \\ &= \mathbb{E}_R[S_\pi(y)]Y + \mathbb{E}_R[S_l(y)]Y. \end{aligned}$$

where $S_\pi(y) \equiv \frac{\pi(y)}{py}$ and $S_l(y) \equiv \frac{w_t l(y)}{py}$ are the profit and labor shares in firm revenue respectively.

Combining, we have

$$Y = \frac{wL_t}{\mathbb{E}_R[S_\pi(y)] + \mathbb{E}_R[S_l(y)]}.$$

Note also that via free entry we have

$$N = \frac{\Pi}{w\nu} = \frac{\mathbb{E}_R[S_\pi(y)]Y}{w\nu}.$$

Rearranging, we can write

$$N = \nu^{-1} \frac{\mathbb{E}_R[S_\pi(y)]}{\mathbb{E}_R[S_\pi(y)] + \mathbb{E}_R[S_L(y)]} L. \tag{C.6}$$

From this we recover the usual logic found in Melitz and Krugman models that with CES and constant returns to scale (so that profit and expenditure shares are constant), the number of firms scales only with population, and is invariant to productivity.

As for the micro-level profit shares of revenue, we can write these as

$$\frac{\pi(y)}{p(y)y} = 1 - \frac{\lambda - 1}{\lambda}\frac{1}{\mathcal{E}_i}$$

where $\mathcal{E}_i$ is the firm-level cost elasticity, defined as

$$\mathcal{E}_i \equiv \frac{d\log C}{d\log y} = \frac{AC(y)}{MC(y)}$$

and by the second equality is equal to the ratio of average and marginal cost. We can also show that for our firms, the cost elasticity is

$$\mathcal{E}_i = (\chi - 1)\Omega^{NM}(y) + 1 < 1 \tag{C.7}$$

so that cost rises less than one-for-one with scale y, but approaches 1 as firms grow large, where $\chi \equiv \frac{(1+\epsilon)\gamma+\delta}{1+\gamma} < 1$.

Now we present a key aggregation result, following the logic and notation in Lashkari et al. (2024). First, given that markups are constant across firms, the allocations of factors to firms is efficient. As such, we can define an aggregate cost function for output Y, and for a given set of firms i (i.e. holding fixed the mass of firms N) in terms of factor prices $\mathbf{W}$, which satisfies

$$\bar{C}(Y, \mathbf{W}) \equiv \min_{y_i} \int C_i(y_i, \mathbf{W}), \quad \text{such that} \quad Y = \left(\int a_i y_i^{\frac{\lambda-1}{\lambda}} di \right)^{\frac{\lambda}{\lambda-1}}$$

The allocation of output Y_i across firms coincides with the market equilibrium with factor prices $\mathbf{W}$ and output Y. Similarly we can define an aggregate cost elasticity

$$\bar{\mathcal{E}} \equiv \frac{d\log \bar{C}(Y, \mathbf{W})}{d\log Y}$$

Given efficiency, we must have the same ratio of revenue to costs $1 + \Pi = PY/\bar{C}$ in this equilibrium and the market equilibrium, we can split this ratio into the markup $\mu = \frac{\lambda}{\lambda-1}$ and some number which must be equal to the aggregate cost elasticity, so that

$$PY/\bar{C} = \mu\bar{\mathcal{E}}$$

Now defining the cost weight as

$$\Lambda_i^c \equiv \frac{C_i}{\bar{C}} = \frac{p_i y_i/(\mu\mathcal{E}_i)}{PY/(\mu\bar{\mathcal{E}})} = \frac{\bar{\mathcal{E}}}{\mathcal{E}_i}\frac{p_i y_i}{PY}$$

We can integrate both sides of these to find that

$$\bar{\mathcal{E}} = \int_i \mathcal{E}_i \Lambda_i^c di$$

So that the economy-wide cost elasticity is the cost-weighted average of individual firm cost elasticities. Then, using (C.7), we have

$$\bar{\mathcal{E}} = (\chi - 1)\mathbb{E}_C[\Omega^{NM}(y)] + 1$$

So under the condition that the cost-weighted average non-material cost share decreases with growth, then the economy-wide profit share increases. As such, the number of firms must increase.

D. Quantitative Model Expressions

Consider the cost function as described in the quantitative model:

$$\text{(D.1)}\quad C(A_Y, y) = A_{Yt}^{-1}\left((A_{Mt}^{-1}\ p_{mt}y^{\epsilon_m})^{1-\sigma_m} + \left((A_{Et}^{-1}\ y^{\epsilon_e}\ p_{et})^{1-\sigma_e} + (y^{\epsilon_l} w_t^{\alpha} r_t^{1-\alpha})^{1-\sigma_e}\right)^{\frac{1-\sigma_m}{1-\sigma_e}}\right)^{\frac{1}{1-\sigma_m}}$$

To pin down the scale elasticities ϵ_e and ϵ_l we target two moments: (1) the elasticity of the energy share relative to the material and labor bundle with respect to firm size ($\eta_{E/ML} \equiv \frac{d\log\left(\Omega^E/\Omega^{M+L}\right)}{d\log y}$) and (2) the elasticity of energy over labor cost share with respect to firm size ($\eta_{E/L} \equiv \frac{d\log\left(\Omega^E/\Omega^{L}\right)}{d\log y}$). We can derive analytical expressions for these scale parameters, which can be written in terms of the estimated scale elasticities and cost shares only. We solve for them in closed form below.

D.1 Closed Form Expression for ϵ_l, ϵ_e

To derive closed form solutions, we first analytically compute the two required scale elasticities.

Energy to Materials and Labor: $\eta_{E/ML}$.

$$\frac{d\log\left(\frac{\Omega^E}{\Omega^L+\Omega^M}\right)}{d\log y} = \frac{d\log\Omega^E}{d\log y} - \frac{\Omega^M}{\Omega^M+\Omega^L}\frac{d\log\Omega^M}{d\log y} - \frac{\Omega^L}{\Omega^M+\Omega^L}\frac{d\log\Omega^L}{d\log y}.$$

First, we solve for $\frac{d\log\Omega^E}{d\log y}$ as:

$$\eta_E \equiv \frac{d\log\Omega^E}{d\log y} = (1-\sigma_e)(1-\Omega^{E,ELK})(\epsilon_e-\epsilon_l) + \Omega^M(1-\sigma_m)\left[\epsilon_e\Omega^{E,ELK} + \epsilon_l(1-\Omega^{E,ELK}) - \epsilon_m\right].$$

Second, we solve for $\frac{d\log\Omega^M}{d\log y}$ as:

$$\frac{d\log\Omega^M}{d\log y} = -\Omega^{ELK}(1-\sigma_m)\left[\epsilon_e\Omega^{E,ELK} + \epsilon_l(1-\Omega^{E,ELK}) - \epsilon_m\right].$$

Third, we solve for $\frac{d\log\Omega^L}{d\log y}$ as:

$$\frac{d\log\Omega^L}{d\log y} = \Omega^{M,MELK}(1-\sigma_m)\left[\epsilon_e\Omega^{E,ELK} + \epsilon_l(1-\Omega^{E,ELK}) - \epsilon_m\right] - \Omega^{E,ELK}(1-\sigma_e)(\epsilon_e-\epsilon_l).$$

Substituting in the previous expressions:

$$\begin{aligned}\frac{d\log\left(\frac{\Omega^E}{\Omega^L+\Omega^M}\right)}{d\log y} &= (1-\sigma_e)(1-\Omega^{E,ELK})(\epsilon_e-\epsilon_l) + \Omega^{M,MELK}(1-\sigma_m)\left[\epsilon_e\Omega^{E,ELK} + \epsilon_l(1-\Omega^{E,ELK}) - \epsilon_m\right] \\ &+ \frac{\Omega^M}{\Omega^M+\Omega^L}\left[\Omega^{ELK}(1-\sigma_m)\left[\epsilon_e\Omega^{E,ELK} + \epsilon_l(1-\Omega^{E,ELK}) - \epsilon_m\right]\right] \\ &- \frac{\Omega^L}{\Omega^M+\Omega^L}\left[\Omega^{M,MELK}(1-\sigma_m)\left[\epsilon_e\Omega^{E,ELK} + \epsilon_l(1-\Omega^{E,ELK}) - \epsilon_m\right] - \Omega^{E,ELK}(1-\sigma_e)(\epsilon_e-\epsilon_l)\right].\end{aligned}$$

Simplifying, we obtain:

$$\eta_{E,ML} \equiv \frac{d \log\left(\frac{\Omega^E}{\Omega^L+\Omega^M}\right)}{d \log y} = A\left[1 - \frac{\Omega^M \Omega^{E,ELK}}{\Omega^M + \Omega^L}\right] + B\left[\frac{\Omega^M}{\Omega^M + \Omega^L}\right],$$

where

$$A = (1-\sigma_e)(\epsilon_e - \epsilon_l), \qquad B = (1-\sigma_m)\left[\epsilon_e \Omega^{E,ELK} + \epsilon_l(1-\Omega^{E,ELK}) - \epsilon_m\right].$$

Energy to Labor and Capital: $\eta_{E/L}$. Define

$$\eta_{E,L} \equiv \frac{d \log\left(\frac{\Omega^E}{\Omega^L}\right)}{d \log y} = \frac{d \log(\Omega^E)}{d \log y} - \frac{d \log(\Omega^L)}{d \log y}.$$

Plugging in our solved expressions, we have

$$\eta_{E,L} \;=\; (1-\sigma_e)\,(\epsilon_e - \epsilon_l)$$

Note, that the analytical expression for $\eta_{E/LK}$ is identical to that of $\eta_{E/L}$. In our empirical estimates we find that both scale elasticities are insignificant and thus we set equal our estimate of $\eta_{E/L} = 0$ in our calibration which implies that $\epsilon_e = \epsilon_l$.

Analytical Calibration of ϵ_e, ϵ_l Given these formulas for our scale elasticities, we can solve for ϵ_e, and ϵ_l in closed form. For a given value of ϵ_m (which we calibrate to match overall degrees of returns to scale), the scale elasticities on energy and labor are:

$$\epsilon_l = \epsilon_m - \frac{\eta_{E/L}\Omega^{E,ELK}}{1-\sigma_e} + \frac{\eta_{E/ML} - \eta_{E/L} + s\,\Omega^{E,ELK}\eta_{E/L}}{s(1-\sigma_m)} \tag{D.2}$$

$$\epsilon_e = \epsilon_l + \frac{\eta_{E/L}}{(1-\sigma_e)} \tag{D.3}$$

where $s \equiv \frac{\Omega^M}{\Omega^M+\Omega^L}$

Since $\eta_{E/L} = 0$ we set $\epsilon_l = \epsilon_e$. Given our calibrated value of $\epsilon_m = 1.3$, our empirical estimate of $\eta_{E/ML} = -0.51$, and cost shares of $\Omega^{E,ELK}_{1990} = 0.315,\ \Omega^{M}_{1990} = 0.714, \Omega^{L}_{1990} = 0.06$ as measured in the micro data, we find that $\epsilon_l = \epsilon_e = 0.27$.

D.2 Estimation of Additional Production Function Parameters

D.2.1 Elasticity of Substitution Between Energy and Other Inputs

We estimate:

$$\Delta^h \log \widetilde{\Omega^E_{it}} = \alpha_{st} + \beta\, \Delta^h \log w^E_{it} + \varepsilon_{it}, \tag{D.4}$$

where $\Delta^h \log w^E_{it}$ is the log change in the price of the firm-specific energy bundle, and $\widetilde{\Omega^E_{it}}$ is a measure of the relative energy demand. For this analysis we take the energy expenditure relative to labor expenditure $\widetilde{\Omega^E_{it}} = \frac{p_e E}{wL}$. As a reminder, $\Delta^h \log w^E_{it}$ is constructed as the Törnqvist-

weighted change in the observed firm×fuel-level price changes.

The goal is to identify the elasticity of substitution, i.e., the elasticity of relative energy demand in response to changes in the relative price of energy.

The key threat to identification is that firm-specific energy prices themselves respond to firm-level energy demand. To identify σ_e, we construct a supply shifter affecting $\Delta^h \log w_{it}^E$, but plausibly orthogonal to other shifts in firm-specific energy demand.

We construct a shift-share instrument exploiting changes in the aggregate prices of different fuels $f \in \mathcal{E} = \{\text{coal}, \text{electricity}, \text{oil}\}$, as well as firm-level exposure to different fuels:

$$\mathcal{S}_{it}^{E,h} = \sum_{f \in \mathcal{E}} w_{fit} \, \Delta^h \log p_{ft}, \tag{D.5}$$

where $\Delta^h \log p_{ft} = \log p_{ft+h} - \log p_{ft}$. For oil, we use the price of oil from the Petroleum Planning & Analysis Cell. For electricity, given that electricity prices are set by state-level utilities, we use the change in the average state-level electricity price. For coal, we account for the fact that coal markets are highly local and use the district-level change in the price of coal.

The identifying assumption is that firms most exposed to changes in aggregate energy prices due to their fuel mix are not systematically subject to other energy-specific demand or supply shocks. We check that this instrument is uncorrelated with the change in other input prices (price of the material bundle, wage) and with the demand shifter defined in (4). The results are summarized in Table D.1.

TABLE D.1: Estimation of the Elasticity of Substitution

	E / L			E / M		
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta^5 \log(\text{price}^E)$	0.411***	0.449***	0.410***	0.449***	0.501***	0.469***
	(0.064)	(0.066)	(0.074)	(0.072)	(0.088)	(0.077)
$\Delta^5 \log(w)$ controls		✓			✓	
$\Delta^5 \log(y)$ controls			✓			✓
Year × Ind FE	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓
F-stat	534.180	195.235	269.042	534.961	57.398	269.268
Observations	204,508	204,164	186,767	204,587	158,130	186,834

Note: This table presents the results of estimating the elasticity of substitution using different specifications of the relative energy share: (i) energy costs relative to labor costs or (ii) energy costs relative to materials costs. $\Delta^h \log(p_e)$ is instrumented using a shift-share instrument which exploits changes in the aggregate prices of different fuels. $\Delta^5 \log(w)$ controls indicates that the regression controls for the firm-level changes in input prices (wage or raw materials price index). $\Delta^5 \log(y)$ controls indicates that the regression controls for the firm-level changes in output, instrumented by our demand shifter. Regressions are weighted by firm-level total costs. Standard errors are clustered at the firm and state level. ***, ** and * indicate significance at the 1%, 5% and 10% levels, respectively.

D.2.2 Elasticity of Substitution Between Materials and Other Inputs

Using analytical mappings from our key elasticities to model parameters, we can cleanly back out σ_m, the elasticity of substitution between materials and other inputs, from (i) our estimates of σ_e, and (ii) estimates of the elasticity of the energy-to-materials expenditure ratio to the price of energy. To see this, recall that the energy expenditure share relative to labor is:

$$\Omega^{E/L} \equiv \frac{p_e E}{wL} = \frac{A_e^{\sigma_e - 1} p_e^{1-\sigma_e} y^{(1-\sigma_e)\epsilon_e}}{\alpha w^{\alpha(1-\sigma_e)} r^{(1-\alpha)(1-\sigma_e)} y^{(1-\sigma_e)\epsilon_l}}$$

which implies

$$d \log \Omega^{E/L} = (1 - \sigma_e) d \log p_e \tag{D.6}$$

holding fixed any changes in firm scale or general equilibrium effects through the prices of other inputs. As robustness, in our empirical results we also control for changes in firm scale instrumented by our demand shifter and changes in the prices of other inputs, and we find similar results. For the energy-to-materials expenditure ratio, we use:

$$\Omega^{E/M} \equiv \frac{p_e E}{p_m M} = \frac{A_e^{\sigma_e - 1} x^{\frac{\sigma_e - \sigma_m}{1-\sigma_e}} p_e^{1-\sigma_e} y^{(1-\sigma_e)\epsilon_e}}{A_m^{\sigma_m - 1} p_m^{1-\sigma_m} y^{(1-\sigma_m)\epsilon_m}}$$

where $x \equiv (A_e^{-1} p_e y^{\epsilon_e})^{1-\sigma_e} + (w^{\alpha} r^{1-\alpha} y^{\epsilon_l})^{1-\sigma_e}$. This implies

$$d \log \Omega^{E/M} = [(1 - \sigma_e) + (\sigma_e - \sigma_m)\Omega^{E,ELK}] d \log p_e \tag{D.7}$$

Using our empirical estimates from table D.1 and equations D.6 and D.7, we find $\sigma_e = 0.59, \sigma_m = 0.47$.

D.2.3 Elasticity of the Cost Function With Respect to Output

By definition,

$$\text{RTS}_i^{-1} = \mathcal{E}_i = \frac{\partial \log C(y_i, z_i, \mathbf{w}_i)}{\partial \log y_i} \tag{D.8}$$

We obtain the average cost elasticity in our sample by estimating the following regression model:

$$\Delta^h \log C_{it} = \alpha_{st} + \beta \Delta^h \log y_{it} + \varepsilon_{it} \tag{D.9}$$

where $\Delta^h \log y_{it}$ is instrumented by the demand shifter $\mathcal{D}_{it}^h$. The identifying assumption is that $\mathcal{D}_{it}^h$ is orthogonal to unobserved determinants of the change in total costs. In particular, $\mathcal{D}_{it}^h$ must be orthogonal to: (i) the change in firm-level physical productivity, and (ii) the change in firm-level input prices. Note that this is a more stringent identifying assumption that that underlying the identification of the scale elasticity of energy demand, which only required that $\mathcal{D}_{it}^h$ be

orthogonal to energy-specific productivity shocks. To alleviate these identification concerns, we include the change in firm-specific input prices and the change in physical productivity as controls.

TABLE D.2: Estimation of the Cost Elasticity

	$\Delta^5 \log(C)$					
	(1)	(2)	(3)	(4)	(5)	(6)
Δ^5 log(y)	0.948***	0.929***	1.011***	0.957***	1.003***	0.949***
	(0.134)	(0.152)	(0.103)	(0.051)	(0.101)	(0.053)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓
$\Delta^5 \log(w)$ controls		✓		✓		✓
$\Delta^5 \log(\text{tfpq})$ controls			✓ (1)	✓ (1)	✓ (2)	✓ (2)
F-Stat	43.213	34.545	79.060	163.036	81.630	159.488
Observations	138,382	114,295	114,189	114,183	114,189	114,183
Returns to scale	1.055	1.076	0.989	1.045	0.997	1.054

Note: This table presents the results of estimating equation (D.9). $\Delta^h \log(y)$ is instrumented by the firm-level demand shock defined in (4). Returns to scale are defined as the inverse of the estimated coefficient. $\Delta^5 \log(w)$ controls indicates that the regression controls for the firm-level changes in input prices (wage, energy price index, raw materials price index). $\Delta^5 \log(\text{tfpq})$ controls indicates that the regression controls for the firm-level changes in physical productivity. Regressions are weighted by firm-level total costs. Standard errors are clustered at the firm and main product×year level. ***, ** and * indicate significance at the 1%, 5% and 10% levels, respectively.

TABLE D.3: Cost Elasticity: Heterogeneity Analysis

Panel A: Unweighted vs. Weighted Estimation

	$\Delta^5 \log(C)$					
	(1)	(2)	(3)	(4)	(5)	(6)
Δ^5 log(y)	0.880***	0.892***	0.957***	0.949***	0.964***	0.981***
	(0.041)	(0.046)	(0.051)	(0.053)	(0.052)	(0.056)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓
Weight			Cost	Cost	Output	Output
$\Delta^5 \log(w)$ controls	✓	✓	✓	✓	✓	✓
$\Delta^5 \log(\text{tfpq})$ controls	✓ (1)	✓ (2)	✓ (1)	✓ (2)	✓ (1)	✓ (2)
F-Stat	198.491	178.965	163.036	159.488	154.323	151.970
Observations	114,183	114,183	114,183	114,183	114,183	114,183
Returns to scale	1.137	1.122	1.045	1.054	1.038	1.019

Panel B: Heterogeneity Across Samples

	$\Delta^5 \log(C)$							
	By material share				By output			
	Low		High		Low		High	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ^5 log(y)	0.831***	0.860***	1.010***	0.979***	0.830***	0.772***	0.957***	0.948***
	(0.089)	(0.097)	(0.068)	(0.065)	(0.096)	(0.088)	(0.058)	(0.058)
Year × Ind. × Cohort FE	✓	✓	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓	✓	✓
$\Delta^5 \log(w)$ controls	✓	✓	✓	✓	✓	✓	✓	✓
$\Delta^5 \log(\text{tfpq})$ controls	✓ (1)	✓ (2)	✓ (1)	✓ (2)	✓ (1)	✓ (2)	✓ (1)	✓ (2)
F-Stat	42.720	42.066	103.893	106.883	37.540	41.170	132.250	131.253
Observations	55,485	55,485	56,066	56,066	53,728	53,728	57,789	57,789
Returns to scale	1.204	1.163	0.990	1.021	1.205	1.296	1.045	1.055

Note: This table presents the results of estimating equation (D.9). $\Delta^h \log(y)$ is instrumented by the firm-level demand shock defined in (4). Returns to scale are defined as the inverse of the estimated coefficient. In Panel A, we estimate this specification unweighted (only sampling weights), weighted by firm-level total costs, and weighted by firm-level sales. In Panel B, we estimate this specification separately for firms below and above the median material share (sales, respectively). $\Delta^5 \log(w)$ controls indicates that the regression controls for the firm-level changes in input prices (wage, energy price index, raw materials price index). $\Delta^5 \log(\text{tfpq})$ controls indicates that the regression controls for the firm-level changes in physical productivity. Standard errors are clustered at the firm and main product×year level. ***, ** and * indicate significance at the 1%, 5% and 10% levels, respectively.

D.3 Macro (Time-Series) Elasticities

We estimate the macro-elasticity of the energy share using the aggregate time series computed from our model, described in section 5.2. To do this we regress log changes in cost shares on log changes in output, running the following regression in the time series from 1990-2023:

$$\Delta \log(\Omega_t) = \alpha + \eta^{macro} \Delta \log(\text{output}_t) + \varepsilon_t \tag{D.10}$$

where Ω_t are the various cost shares and measures of energy intensity displayed in Table D.4. We find that the macro-elasticity for our main energy share measure in the model ($\eta^{macro}_{E/ML} = -0.45$) is very similar to our empirical estimate of $\eta_{E/ML} = -0.51$.

TABLE D.4: Model Fit: Macro (time-series) Elasticities

	(1) $\Omega_{\frac{E}{ML}}$	(2) $\Omega_{\frac{E}{L}}$	(3) $\Omega_{\frac{E}{M}}$	(4) $\Omega_{\frac{E}{MLK}}$	(5) $\Omega_{\frac{E}{LK}}$	(6) $\Omega_{\frac{E}{VA}}$	(7) $\Omega_{\frac{E}{Y}}$	(8) $\Omega_{\frac{M}{ELK}}$	(9) $\Omega_{\frac{L}{EMK}}$	(10) $\Omega_{\frac{K}{EML}}$
Δ log(quantity)	-0.447	-0.131	-0.467	-0.406	-0.131	-0.468	-0.444	0.375	-0.259	-0.280
	(0.011)	(0.002)	(0.012)	(0.011)	(0.002)	(0.011)	(0.012)	(0.010)	(0.010)	(0.010)
Observations	33	33	33	33	33	33	33	33	33	33

Note: This table displays the model-implied macro elasticities estimated from time-series regressions of changes in log shares on changes in log output, using equation D.10. Standard errors are in parentheses.

E. PHYSICAL MECHANISMS

In this Appendix, we derive the scale elasticities for typical manufacturing technologies corresponding the main energy end uses.

E.1 Process Heating (Fired Systems)

This model estimates the scale elasticity of energy intensity for a crucible furnace used for melting and holding molten metal (see picture in Figure E.1(A)). For the numerical application, we consider the case of aluminum melting. This is a simple system for which we can derive a straightforward characterization. Below, we also provide estimates of the scale elasticity for a more complex system, the billet reheating furnace in Figure E.1(B).

FIGURE E.1: Furnaces

(A) Crucible furnace



(B) Billet reheating furnace



Step 1: Combustion. We estimate the energy input Q_{total} required to deliver Q_{furnace} to the furnace. At the combustion stage, the main source of energy losses is flue gas losses. The derivations below show that these losses are proportional to the total energy input, so that we can write :

$$Q_{\text{total}} = \frac{Q_{\text{furnace}}}{\eta_{\text{comb}}} \tag{E.1}$$

where η_{comb} denotes the efficiency of the combustion process.

To obtain this relationship, start from:

$$Q_{\text{total}} = Q_{\text{furnace}} + L_{\text{flue}} \tag{E.2}$$

where L_{flue} are the flue gas losses. The energy lost in flue gas per unit of time L_{flue}/t depends on the temperature difference between the gas and the environment $\left(T_{\text{flue}} - T_{\text{ref}}\right)$ (K) as the flue gas mass flow rate $\dot{m}_{\text{flue}}$ (kg s^{-1}). Hence,

$$L_{\text{flue}} \;=\; \dot{m}_{\text{flue}}\, c_p \left(T_{\text{flue}} - T_{\text{ref}}\right) t \tag{E.3}$$

where c_p is the mean specific heat of flue gas, with units J kg^{-1} K^{-1}.
The flue gas mass flow rate $\dot{m}_{\text{flue}}$ is proportional to the fuel mass flow rate $\dot{m}_{\text{fuel}}$ (kg s^{-1}):

$$\dot{m}_{\text{flue}} \approx (1+\lambda)\, A_{\text{st}}\, \dot{m}_{\text{fuel}} \tag{E.4}$$

where λ is the excess air ratio (dimensionless), A_{st} is the stoichiometric air-to-fuel mass ratio (kg air / kg fuel). The fuel energy input is then given by:

$$Q_{\text{total}} = Q_{\text{fuel}} = \dot{m}_{\text{fuel}}\, \text{LHV}\, t \tag{E.5}$$

where LHV is the lower heating value of the fuel in J kg^{-1}. Therefore, the flue gas losses as a fraction of fuel energy input is:

$$\frac{L_{\text{flue}}}{Q_{\text{fuel}}} = \frac{(1+\lambda)\, A_{\text{st}}\, c_p \left(T_{\text{flue}} - T_{\text{ref}}\right)}{\text{LHV}} = 1 - \eta_{\text{comb}} \tag{E.6}$$

Define η_{comb} such that $1 - \eta_{\text{comb}} = \frac{(1+\lambda)\, A_{\text{st}}\, c_p \left(T_{\text{flue}} - T_{\text{ref}}\right)}{\text{LHV}}$, a constant for this system. η_{comb} is the combustion efficiency and does not depend on scale. Then,

$$Q_{\text{total}} = \frac{Q_{\text{furnace}}}{\eta_{\text{comb}}} \tag{E.7}$$

This analysis only considers heat losses in dry flue gas, which is a simplification. In practice, there are also (smaller) heat losses due to the evaporation of water formed due to hydrogen in the fuel, as well as due to moisture in the fuel. These losses are also proportional to m_{fuel}. Hence, they further reduce η_{comb}, but preserve the proportional relationship.

Step 2: Energy required by the furnace. We now estimate the energy required by the furnace.

Energy to heat and melt the metal: The sensible heat (heating from initial to melting temperature) is given by:

$$Q_{\text{sensible}} = Mc(\overline{T} - T_{\text{init}}) \tag{E.8}$$

and the latent heat of fusion (melting the metal) is given by:

$$Q_{\text{latent}} = ML_f \tag{E.9}$$

where M is the mass of the metal (kg). c is the specific heat capacity of aluminum, L_f is the latent heat of fusion of aluminum, $\overline{T}$ (operating temperature for molten aluminum) and T_{init} (ambient air temperature). Therefore, using $M = \rho V$ with ρ the density of aluminum, we obtain:

$$Q_{\text{useful}} = Q_{\text{sensible}} + Q_{\text{latent}} = (c(\overline{T} - T_{\text{init}}) + L_f)M \tag{E.10}$$

We can write:

$$Q_{\text{useful}} = a_{\text{useful}} M \tag{E.11}$$

and see that the useful energy scales proportionally with the mass.

Energy required to maintain molten state (holding losses): We now quantify the surface losses. They are of three types:

Radiative loss (electromagnetic heat loss from radiation):

$$\frac{dQ_{\text{rad}}}{dt} = \varepsilon\sigma A_s \left(\overline{T}^4 - T_{\text{init}}^4\right) \tag{E.12}$$

where ε is the emissivity for oxidized molten aluminum, σ is the Stefan-Boltzmann constant ($5.67 \times 10^{-8}\,\text{W}\,\text{m}^{-2}\,\text{K}^{-4}$), and A_s is the surface area of molten metal in m^2.

Convective loss (heat transfer through metal-air contact):

$$\frac{dQ_{\text{conv}}}{dt} = hA_s(\overline{T} - T_{\text{init}}) \tag{E.13}$$

where h is the natural convective heat transfer coefficient.

Conductive loss through furnace walls (heat transfer through direct contact between materials):

$$\frac{dQ_{\text{cond}}}{dt} = \frac{k \cdot A_c}{d} \cdot (\overline{T} - T_{\text{init}}) \tag{E.14}$$

where k is the effective thermal conductivity of the furnace insulation, and d is the wall thickness of the refractory or insulation layer.

Hence, the energy required to maintain the molten state, which equals the holding losses, is given by:

$$L_{\text{surface}} = Q_{\text{holding}} = \left(\frac{dQ_{\text{rad}}}{dt} + \frac{dQ_{\text{conv}}}{dt} + \frac{dQ_{\text{cond}}}{dt}\right) \cdot t \tag{E.15}$$

We assume a cubic furnace for simplicity (results would be similar for a cylinder). The edge length ℓ for a given M is given by $M = \rho V = \rho\ell^3$. Then, the open surface area (top) is $A_s = \ell^2$

and the conductive surface area (sides + bottom) is $A_c = 5\ell^2$. Therefore,

$$L_{\text{surface}} = \left(\left(\varepsilon \cdot \sigma \cdot \left(\overline{T}^4 - T_{\text{init}}^4 \right) + h_c \left(\overline{T} - T_{\text{init}} \right) \right) \cdot t + \frac{k \cdot 5}{d} \cdot \left(\overline{T} - T_{\text{init}} \right) \cdot t \right) \left(\frac{M}{\rho} \right)^{\frac{2}{3}} = a_{\text{surface}} M^{\frac{2}{3}} \tag{E.16}$$

The surface losses scale less than proportionally with the mass M.

Finally, the energy requirement for the furnace (after combusion) is:

$$Q_{\text{furnace}} = Q_{\text{useful}} + L_{\text{surface}} \tag{E.17}$$

Step 3: Scale elasticity of energy intensity The total energy required is:

$$Q_{\text{total}} = \frac{1}{\eta_{\text{comb}}} (Q_{\text{useful}} + L_{\text{surface}}) \tag{E.18}$$

Total energy per unit mass:

$$\frac{Q_{\text{total}}}{M} = \frac{1}{\eta_{\text{comb}}} \left(\frac{Q_{\text{useful}}}{M} + \frac{L_{\text{surface}}}{M} \right) = \frac{1}{\eta_{\text{comb}}} \left(a_{\text{useful}} + a_{\text{surface}} \, M^{-\frac{1}{3}} \right) \tag{E.19}$$

Then,

$$\frac{d \log \frac{Q_{\text{total}}}{M}}{d \log M} = -\frac{L_{\text{surface}}}{Q_{\text{furnace}}} \frac{1}{3} \tag{E.20}$$

Numerical implementation for the aluminum melting furnace. We use the following parameter values for the calculations. These parameter are either physical constants or typical values from the engineering literature.

- Specific heat capacity of aluminum: $c = 900\,\text{J}\,\text{kg}^{-1}\,\text{K}^{-1}$
- Latent heat of fusion of aluminum: $L_f = 397{,}000\,\text{J}\,\text{kg}^{-1}$
- Operating temperature of molten aluminum: $\overline{T} = 973\,\text{K}$
- Initial temperature: $T_{\text{init}} = 293\,\text{K}$
- Density of aluminum: $\rho = 2700\,\text{kg}\,\text{m}^{-3}$
- Emissivity of oxidized molten aluminum: $\varepsilon = 0.2$
- Convective heat transfer coefficient between molten metal and air: $h = 20\,\text{W}\,\text{m}^{-2}\,\text{K}^{-1}$
- Thermal conductivity of insulation: $k = 2\,\text{W}\,\text{m}^{-1}\,\text{K}^{-1}$
- Wall thickness of insulation: $d = 0.25\,\text{m}$
- The holding time is set to $t = 10\,\text{hours} = 36{,}000\,\text{s}$, a typical industrial work shift during which the metal is held at temperature before being cast or transferred being from 4 hours (more advanced technologies like the CRIMSON method) to 16 hours (traditional casting processes).

- We assume we melt a mass of aluminum $M = 250\,\text{kg}$.

We find that $\frac{L_{\text{surface}}}{Q_{\text{furnace}}} = 0.6$. Foundries have a large fraction of surface losses because the furnace needs to hold metal molten for extended period of time, and the outer-surface temperature of crucibles is very high.

Therefore,

$$\frac{d \log \frac{Q_{\text{total}}}{M}}{d \log M} = -0.2 \tag{E.21}$$

Alternative calibrations. *Billet furnace.* The Indian Bureau of Energy Efficiency (link) proposes a calculation for the case of an oil-fired billet reheating furnace for rolling mills, processing 6000 kg / hour. From this analysis, we obtain $\frac{L_{\text{surface}}}{Q_{\text{furnace}}} \approx 0.26$. The billet furnace has much lower surface losses: it does not hold metal molten for hours as a crucible does; in addition, the Bureau of Energy Efficiency gives measured outer-surface temperatures around 100 C°, much cooler than the exposed shell of an aluminum crucible.

This implies that:

$$\frac{d \log \frac{Q_{\text{total}}}{M}}{d \log M} \approx -0.09 \tag{E.22}$$

This magnitude, that we derive from first principles, is very much in line with the one that can be computed from the furnace efficiency by capacity tables provided by the Indian Bureau of Energy Efficiency (link) for a generic pusher-type billet reheating furnace, equal to -0.16.

Forging furnace. The same document provides furnace efficiency by capacity tables for a generic pusher forging furnace, from which we can estimate a scale elasticity equal to -0.23.

Taking stock. From this analysis, we take -0.15 to be a reasonable scale elasticity for fired systems, and use this value in our summary table.

E.2 Process Heating (Steam System)

Step 1. Steam production in a boiler. We estimate the energy input Q_{total} required to deliver Q_{steam} to the process. Input energy must be equal to the energy delivered to the system plus losses occurring in the boiler.

$$Q_{\text{total}} = Q_{\text{steam}} + L_{\text{boiler}} \tag{E.23}$$

In a boiler, most of the energy loss comes from the combustion step (due, as explained above, to heat losses in dry flue gas, and evaporation of the water due to the hydrogen or moisture in the fuel). Surface heat losses occur at the shell of the boiler, but they are quantitatively negligible (of the order of magnitude of 0.25% of the energy input). Hence, we ignore them.

Therefore, as in step 1 in the section above, we can write:

$$\frac{L_{\text{boiler}}}{Q_{\text{fuel}}} = 1 - \eta_{\text{boiler}} \tag{E.24}$$

where η_{boiler} is the boiler efficiency, which does not depend on scale. Then,

$$Q_{\text{total}} = \frac{Q_{\text{steam}}}{\eta_{\text{boiler}}} \tag{E.25}$$

The energy out of the boiler and entering the process can be written as:

$$Q_{\text{steam}} = h\, m_s^{input} \tag{E.26}$$

where m_s^{input} is the mass of steam entering the process and h is total steam enthalpy.

Step 2. Steam-based process heating. We consider a drying task, akin to the use of steam in the paper products industry. The task is to remove r kg water per kg product (drying duty) for a mass of dry product M (kg).

Useful heat: Consider the energy required to vaporize a mass m_{water} of water:

$Q_{vap} = m_{water}\, h_{\text{vap}}$ where h_{vap} is latent heat of vaporization of the water in the product

$m_{water} = r\, M$

$Q_{vap} = r\, M\, h_{\text{vap}}$

Energy entering the dryer heat-exchange interface: let η_{ht} be the fraction of the latent heat that actually reaches the product, capturing the heat-transfer effectiveness of the steam side (incl. exchanger/drum effectiveness). Then, the steam input energy required is:

$$Q_{in} = \frac{1}{\eta_{ht}} Q_{\text{vap}} \tag{E.27}$$

Fractional enthalpy loss: This energy is supplied by a mass m_s of steam. Let h_{fg} (kJ/kg steam) the steam latent enthalpy (energy available from each kg of steam). We denote ϕ the fractional enthalpy loss per kg steam across the loop (dimensionless). This captures condensate losses, flash steam venting, steam traps and valve leakages, all of which are primarily proportional to the mass of steam.

$$L_{\text{flow}} = \phi\, m_s\, h_{fg} \tag{E.28}$$

Surface losses: surface losses are proportional to the area surface of the dryer. Considering that the volume of the dryer scales proportionally with the mass to dry, we obtain:

$$L_{\text{surf}} = c_{\text{surf}}\, M^{2/3} \tag{E.29}$$

The required steam from the boiler is:

$$Q_{\text{steam}} = \frac{h}{h_{fg}} (Q_{\text{in}} + L_{\text{flow}} + L_{\text{surf}}) \tag{E.30}$$

$\frac{h}{h_{fg}}$ converts the latent energy basis to the total enthalpy basis (accounting for the sensible

portion in total steam enthalpy).

Step 3. Total energy per kg product and scale elasticity.

$$\frac{Q_{total}}{M} = \frac{1}{\eta_{\text{boiler}}} \frac{h}{h_{fg}} \left((1+\phi) \frac{r\, h_{\text{vap}}}{\eta_{ht}} + c_{\text{surf}}\, M^{-\frac{1}{3}} \right) \tag{E.31}$$

$$\frac{d \log \frac{Q_{\text{total}}}{M}}{d \log M} = -\frac{L_{\text{surf}}}{Q_{\text{steam}}} \frac{1}{3} \tag{E.32}$$

The fraction of losses in the dryer section due to the surface is about 9% (Ghodbanan, Alizadeh, Shafiei, and Rahbar Shahrouzi 2024). This gives us a scale elasticity around -0.03 for the combined boiler-dryer system.

E.3 Mechanical Work

Motor. We consider a motor performing mechanical work. For instance, a motor is used to lift a mass M at a height h in t seconds. We can write the energy requirement as:

$$Q_{total} = \frac{g\, h\, M}{\eta(\ell, P_r)} \tag{E.33}$$

g is the gravitational constant. $\eta(\ell, P_r)$ is motor efficiency defined as the ratio of the mechanical energy delivered at the rotating shaft to the electrical energy input at its terminals. P_r is the nameplate-rated power of the motor (often expressed in horsepowers, or watts). ℓ is the load, defined by dividing the actual power by P_r.

Therefore, we can write energy intensity as a function of size as:

$$\frac{d \log \frac{Q_{\text{total}}}{M}}{d \log M} = -\frac{d \log \eta(\ell, P_r)}{d \log M} \tag{E.34}$$

Motor efficiency depends on the magnitude of two types of losses. First, magnetic core losses and friction and windage losses are essentially fixed for a given motor design (P_r), and do not scale with load ℓ. They also scale less than linearly with P_r. Second, resistance losses are proportional to the square of the current, and hence vary with the load.

Motor efficiencies by nameplate horsepower and load are given by reference tables such as US Department of Energy (1997).

The scale elasticity of energy intensity depends on whether the increase in mass is accommodated by an increase in P_r at constant load, or an increase in load ℓ at constant nameplate horsepower.

$$\frac{d \log \frac{Q_{\text{total}}}{M}}{d \log M} = -\frac{d \log \eta(\ell, P_r)}{d \log \ell} \frac{d \log \ell}{d \log M} - \frac{d \log \eta(\ell, P_r)}{d \log P_r} \frac{d \log P_r}{d \log M} \tag{E.35}$$

with the constraint that $\frac{d \log \ell}{d \log M} + \frac{d \log P_r}{d \log M} = 1$.

We estimate the elasticities of efficiency with respect to load and nameplate horsepower using the reference tables produced by the U.S. Department of Energy. We obtain:

$$\frac{d\log\eta(\ell,P_r)}{d\log P_r} = 0.03 \tag{E.36}$$

$$\frac{d\log\eta(\ell,P_r)}{d\log\ell} = \begin{cases} 0.07, & \ell\in[0.25,\,0.5], \\ 0.01, & \ell\in(0.5,\,0.75], \end{cases} \tag{E.37}$$

Pumps. Pumps are the largest final users of mechanical work. From simulated data, we obtain a scale elasticity of pump energy intensity, where scale is defined as impeller size, approximately equal to -0.03.

Combined motor-pump system. We assume that in the long-run scale increases have to be accommodated by higher nameplate horsepower motors. This gives us a combined scale elasticity of the motor-pump system equal to $-(0.03+0.03) = -0.06$.

E.4 Other Process

Other process groups include electro-chemical processes. We assume a scale elasticity of 0.

E.5 Non-process

Non-process covers a range of diverse energy uses, summarized in Table E.1. Among these end-uses, only facility heating, ventilation, and air-conditioning (HVAC) exhibits a significant scale elasticity. Indeed, in facility heating, the same negative scale elasticity due to surface losses as the one derived in (E.20) applies, and the share of holding losses in total energy is close to 1 since heating is often continuous (i.e. t tends to infinity). Using a share of holding losses equal to 0.95, we get a negative scale elasticity for facility heating equal to -0.32. This effect is less important for ventilation and air-conditioning, which have a lower fraction of energy lost via the surface (compared to other sources of loss). We assume a scale elasticity of 0 for ventilation and air-conditioning. The Manufacturing Energy Consumption Survey does not allow to separate heating from the rest of HVAC, so we use the fraction of heating in HVAC from the 2018 Commercial Buildings Energy Consumption Survey, which is equal to 65%. This yields the elasticity in the first line of the table.

F. Identification of the Technology Cost Function

We suppose that the cost of a given technology bundle is given by:

$$\tilde{c}_t(\phi,\bar{y}) = \phi^\gamma \bar{y}^\delta \tilde{C}_t = \frac{1}{\gamma A_{\Upsilon t}} \phi^\gamma \bar{y}^\delta w_t^\alpha r_t^{1-\alpha}$$

Note that we also have $\tilde{c}_{it} = w_t l_{it} + r_t k_{it}$ by assumption. At the optimum, $\phi = \frac{y}{e}$ and $\bar{y} = y$. Therefore, we can write:

$$\log\tilde{c}_t = -\gamma\log\frac{e}{y} + \delta\log y + \log\tilde{C}_t$$

TABLE E.1: Non-process energy use by end-use

Category	Share of Total (%)	Scale elasticity
Facility HVAC	49.5	$\approx -0.65 \times 0.95 \times \frac{1}{3}$
Facility Lighting	34.7	≈ 0
Other Facility Support	10.6	≈ 0
Onsite Transportation	2.9	≈ 0
Other Nonprocess Use	2.5	≈ 0
Total Nonprocess	100	≈ -0.10

Note: U.S. Manufacturing Energy Consumption Survey (2018)

Identification of γ. γ is the elasticity with respect to the energy intensity of production $\frac{e}{y}$, keeping production y constant.

This elasticity can be measured using the data on required investment $\Delta k_{\tau is}$ and associated annual energy savings $\Delta e_{\tau,is}$ obtained from the UNIDO technology compendia data. Note that the technology compendia report these amounts for a given firm assumed to produce the same quantity before and after the technology investment, hence they indeed keep production y constant.

The coefficient γ cannot be directly read from Figure 7, because the required investment $\Delta k_{\tau is}$ in INR and associated annual energy savings $\Delta e_{\tau,is}$ in gigajoules are expressed in levels as opposed to log changes. To estimate γ, we now restrict our attention of the observations for which we observe pre-existing energy consumption, allowing us to construct $\Delta \log(e_{\tau is})$, the change in energy consumption upon the adoption of technology τ.

To construct the log change in fixed cost, we proceed as follows. We construct the cost of the pre-existing capital-labor bundle $\tilde{c}^0_{\tau is} = (rk^0_{\tau is} + wl^0_{\tau is})$ by multiplying the production level $y_{\tau is}$ reported in the UNIDO data by the median $(rk + wl)/y$ ratio for this industry found in the ASI data.[16] Then, $\Delta \log(\tilde{c}_{\tau is})$, the change in the fixed cost of the technology, is defined as $\Delta \log(\tilde{c}_{\tau is}) = r\Delta k_{\tau is}/(rk^0_{\tau is} + wl^0_{\tau is})$. This is the user cost of increasing the capital stock by $\Delta k_{\tau is}$, relative to the pre-existing total fixed cost. We use the same r as the one used to construct the cost of capital. This measure implicitly assumes that investing in the more energy-efficient capital does not require new staff or training cost. The fixed cost increase is only the capital amount. We also propose an alternative measure $\Delta \log(\tilde{c}_{\tau is}) = \Delta k_{\tau is}/k^0_{\tau is}$. Here, by contrast, we assume that investing in the more energy-efficient capital comes with an unobserved labor cost proportional to the increase in the capital stock.

We then estimate:

$$\Delta \log \tilde{c}_{\tau is} = \alpha + \beta \Delta \log e_{\tau is} + \varepsilon_{kis}$$

separately for each sector s.

The results are presented in Table F.1. Panel A reports results for the baseline measure of $\Delta \log(\tilde{c}_{\tau is})$, while panel B presents results for the alternative version. Odd columns use only

[16] For instance, if fo industry s, UNIDO reports production in tons, we compute fixed costs in INR per tons in the ASI data, and multiply this value by production in the UNIDO data.

observations for which we observe pre-existing energy consumption and production quantity. In even columns, we expand the sample by imputing these variables with the industry-specific average when they are missing.

Looking at panel A, we obtain $\gamma \approx 0.1$ with coefficients very stable across industries and sample choice. The R-squared are high, in particular in odd columns where we use only the highest quality data points, suggesting that the functional form is appropriate. The results in panel B yield higher slopes by construction. Averaging across the three industries, this set of results implies $\gamma \approx 0.25$.

TABLE F.1: Estimation of γ

Panel A: Baseline $\Delta \log(\tilde{c}_{\tau is})$

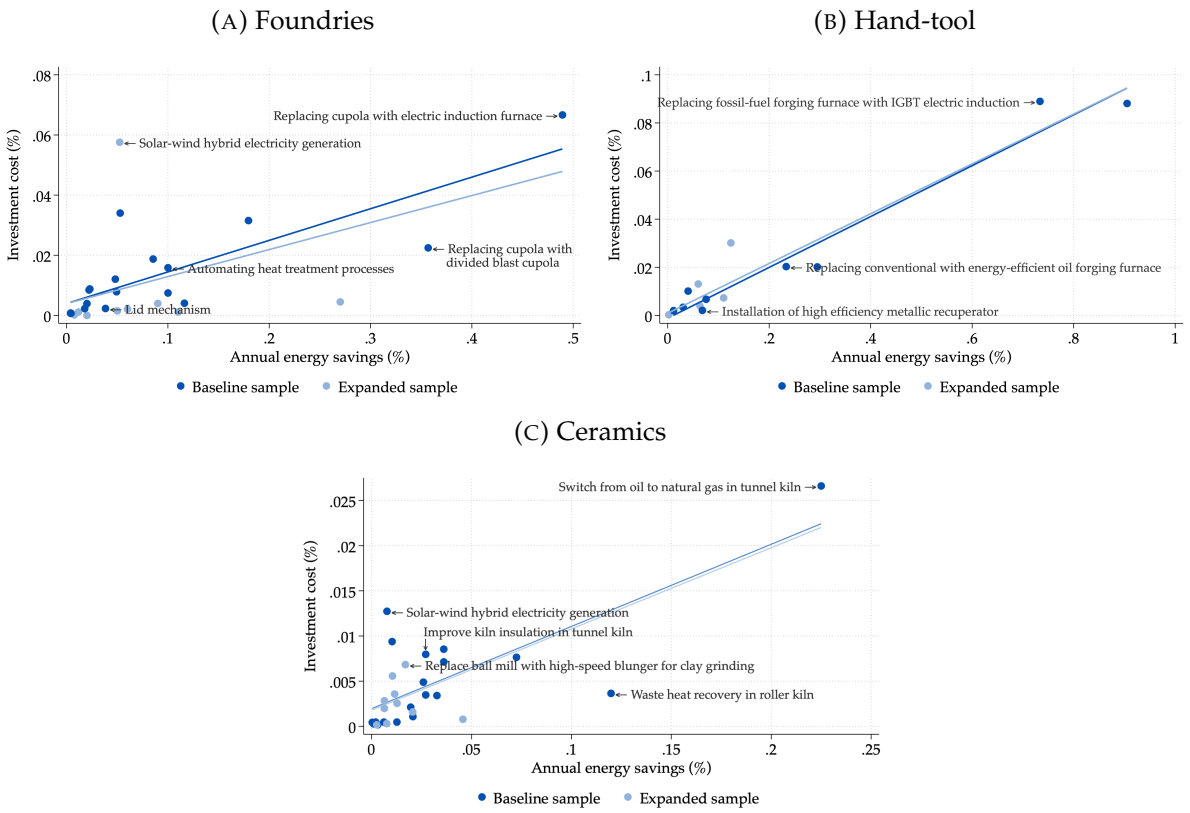
	Investment cost (%)					
	(1)	(2)	(3)	(4)	(5)	(6)
Energy savings (%)	0.105***	0.090***	0.106***	0.104***	0.091***	0.090***
	(0.018)	(0.023)	(0.008)	(0.008)	(0.018)	(0.015)
Industry	Foundry	Foundry	Hand-tool	Hand-tool	Ceramic	Ceramic
Sample	Baseline	Expanded	Baseline	Expanded	Baseline	Expanded
Observations	18	28	9	14	20	30
R-squared	0.68	0.36	0.96	0.94	0.59	0.55

Panel B: Alternative $\Delta \log(\tilde{c}_{\tau is})$

	Investment cost (%)					
	(1)	(2)	(3)	(4)	(5)	(6)
Energy savings (%)	0.245***	0.210***	0.460***	0.451***	0.159***	0.157***
	(0.042)	(0.055)	(0.034)	(0.034)	(0.032)	(0.027)
Industry	Foundry	Foundry	Hand-tool	Hand-tool	Ceramic	Ceramic
Sample	Baseline	Expanded	Baseline	Expanded	Baseline	Expanded
Observations	18	28	9	14	20	30
R-squared	0.68	0.36	0.96	0.94	0.59	0.55

Note: This table reports the results of estimating equation F separately for each industry s. In panel A, $\Delta \log(\tilde{c}_{\tau is}) = r\Delta k_{\tau is}/(rk^0_{\tau is} + wl^0_{\tau is})$. In panel B, $\Delta \log(\tilde{c}_{\tau is}) = \Delta k_{\tau is}/k^0_{\tau is}$. "Sample" refers to the construction of the sample, as detailed in the text. ***, ** and * indicate significance at the 1%, 5% and 10% levels, respectively.

FIGURE F.1: Estimation of γ



Note: This figure shows the scatter plots corresponding to equation F, separately for each industry s. We use the baseline version of the outcome variable $\Delta \log(\tilde{c}_{\tau is}) = r\Delta k_{\tau is}/(rk^0_{\tau is} + wl^0_{\tau is})$. The dark (light) blue dots correspond to the baseline (expanded) sample, as described in the text.

Identification of δ. δ is the elasticity of the fixed cost with respect to capacity, for a given technology quality ϕ. A lower value of δ implies that the capital-labor bundle behaves more like a fixed cost, providing more room for a scale elasticity of technological upgrading.

We estimate δ as follows. Our main data source is data on equipment costs by capacity compiled by the U.S. Department of Energy.[17] The goal of this dataset is to allow engineers to rapidly perform cost estimates when evaluating a new system. It covers a relatively comprehensive set of 20 conventional process equipment types. While the report has a chemical engineering focus, it covers most of the equipment found in the UNIDO compendia for foundries, hand-tools, ceramics, and dairies. An observation is characterized by an equipment type e (e.g., dryer), a model j (e.g., direct contact rotary), and a capacity i (e.g., 100 sq. ft.). The definition of capacity is specific to each type of equipment: for instance, for a cooling tower it is expressed in gallons of water per minute; while for a dryer it is expressed in square feet. For each equipment type e, we then estimate the following model:

$$\log\left(\text{Price}_{eij}\right) = \alpha_{ej} + \beta \log\Big(\text{Capacity}_{eij}\Big) + \varepsilon_{eij} \tag{F.1}$$

α_{ej} are equipment$\times$model fixed effects, so that β is estimated using variation in prices across capacities for a given equipment$\times$model pair (hence keeping model quality ϕ constant). The results are reported in Table F.2. Across all equipment types, we find an average scale elasticity equal to 0.52. Across equipments, all estimates are strictly below 1, with most estimates in the 0.4-0.6 range. The R-squared values are high, providing support for the log-log functional form.

We complement this analysis in several ways. First, we collect data on price by capacity for additional equipments mentioned in the UNIDO compendia but that do not appear in the U.S. Department of Energy data (chiller, bulk milk cooler, pasteurizer, variable frequency drive, motor, and biomass boiler). Across these equipment categories, we find an average elasticity equal to 0.54. This is very close to our baseline estimate, suggesting that our results are robust beyond chemical engineering equipments. Second, we note that these elasticities all relate to process equipments. Scale elasticities for items related to instrumentation and control, site improvements, and facilities are generally much lower—estimated around 0.15 in James, Leptinsky, Turner et al. (2022)—so that our estimate of δ is likely to be on the conservative side. Finally, our baseline estimate is very close to 0.6, a benchmark in the scale elasticity literature (see, e.g., Tribe and Alpine 1986).

[17] DOE/NETL-2002/1169 Process Equipment Cost Estimation Final Report, January 2002

TABLE F.2: Estimation of δ

Equipment type	Coefficient	Std. Error	R-squared	Observations
Vessel	0.314***	(0.023)	0.75	44
Storage Tank	0.544***	(0.039)	0.95	18
Column	0.330***	(0.014)	0.99	230
Heat Exchanger	0.521***	(0.035)	0.94	40
Air Cooler	0.404***	(0.045)	0.91	15
Furnace	0.713***	(0.018)	0.99	9
Cooling Tower	0.345***	(0.042)	0.91	9
Package Steam Boiler	0.434***	(0.033)	0.97	8
Evaporator	0.501***	(0.005)	1.00	20
Crusher	0.951***	(0.116)	0.97	19
Mill	0.574***	(0.021)	0.99	16
Dryer	0.664***	(0.080)	0.98	19
Centrifuge	0.792***	(0.065)	0.99	24
Filter	0.379***	(0.044)	0.99	26
Agitator	0.495***	(0.057)	0.93	6
Pump	0.360***	(0.025)	0.95	81
Compressor	0.509***	(0.031)	0.92	33
Centrifugal Fan	0.568***	(0.064)	0.92	9
Rotary Blower	0.415***	(0.034)	0.97	6
Turbine	0.650***	(0.055)	0.98	24
Average	.523			

Note: This table provides the results of estimating equation (F.1) by equipment type, as well as the average coefficient across equipments.

Taking stock: implied energy intensity scale elasticity. Putting together our estimates of ϵ, γ, and δ, we can estimate the scale elasticity of energy intensity implied by our model:

$$\frac{d\log\left(\frac{e}{y}\right)}{d\log y} = \underbrace{\tilde{\epsilon}}_{\text{Overall}} = \underbrace{\epsilon}_{\text{Within-tech.}} - \underbrace{\frac{\epsilon+1-\delta}{1+\gamma}}_{\text{Tech. improvement}}$$

Table F.3 shows that we obtain scale elasticities highly similar to the one we measure in the data (-0.4).

TABLE F.3: Scale elasticity of energy intensity

ϵ	δ	γ	Within tech.	Tech. improvement	Overall
-0.1	0.5	0.1	-0.1	-0.36	-0.46
-0.1	0.5	0.25	-0.1	-0.32	-0.42
-0.1	0.6	0.1	-0.1	-0.27	-0.37
-0.1	0.6	0.25	-0.1	-0.24	-0.34