Research Institute for Environmental Economics and Management

Discussion Paper Series No.2603

Energy-Augmenting Productivity and Carbon Pricing: Evidence from Production Microdata

Tatsuya Abe, Arlan Brucal & Yuta Toyama

June 2026

WASEDA

Energy-Augmenting Productivity and Carbon Pricing: Evidence from Production Microdata*

Tatsuya Abe† Arlan Brucal‡ Yuta Toyama§

June 2, 2026

Abstract

How does directed technological change toward energy efficiency shape the effects and design of carbon pricing? We estimate a structural production function that separates energy-augmenting productivity from Hicks-neutral productivity using Indonesian manufacturing microdata. Exploiting energy-price variation from fossil-fuel subsidy reforms, we find that higher energy prices induce energy-augmenting productivity growth. Counterfactual simulations show that heterogeneity in energy-augmenting productivity, rather than Hicks-neutral productivity, drives the welfare advantage of carbon pricing over uniform regulation. When carbon pricing induces energy-augmenting innovation, aggregate emissions fall further, although a rebound effect partially offsets this additional abatement. Effective carbon-pricing design should account for productivity heterogeneity, induced innovation, and rebound effects.

JEL Classification: D24, O33, Q41, Q54

Keywords: Energy-augmenting productivity, production function, rebound effect, induced innovation, structural estimation, carbon pricing, production microdata

*We are grateful to Naoshi Doi, Jackson Dorsey, Yoshi Konishi, Daisuke Nakajima, Hiroaki Sakamoto, Yoichi Sugita, and seminar participants at Waseda University, EARIE 2021, JEMIOW Autumn 2024, Otaru University of Commerce, Kobe University, the J-TREE Workshop, the JEA Spring Meeting 2025, and the International Workshop on Industrial Organization of Energy and Environmental Markets for their comments. Keigo Makino provided superb research assistance. We acknowledge financial support from a Waseda University Grant for Special Research Projects (Project number: 2023C-581) and JSPS KAKENHI Grant Number 25K00643. All remaining errors are our own. The findings, interpretations, and conclusions expressed in this article are entirely those of the authors and do not necessarily represent the views of their affiliated institutions.

†Graduate School of Economics, Hitotsubashi University, 2-1 Naka, Kunitachi, Tokyo, Japan. E-mail: tatsuya.abe149@gmail.com

‡The World Bank Group, 1818 H Street, NW, Washington, DC, USA. E-mail: abrucal@worldbank.org

§School of Political Science and Economics, Waseda University, 1-6-1 Nishi-Waseda, Shinjuku-ku, Tokyo, Japan. E-mail: yuta.toyama@gmail.com

1 Introduction

Climate regulation has become one of the most consequential and contested areas of economic policy. In the United States, recent policy shifts have rolled back emissions-reduction commitments and green industrial subsidies, reflecting the view that decarbonization imposes unsustainable costs on the economy. Proponents counter that well-designed carbon pricing and environmental regulation can improve economic performance over time by inducing firms to innovate and eliminate wasteful energy use (e.g., Porter and Linde, 1995). Developing economies face a related tension, as governments weigh the fiscal cost of energy subsidies against the economic disruption that can follow from removing them (e.g., Rentschler and Kornejew, 2017). Whether climate regulation ultimately proves costly or beneficial depends largely on how firms adjust their technology in response.

Resolving this debate requires distinguishing between two ways in which firms respond to higher energy prices, such as those induced by carbon pricing. Higher prices may induce genuine technological improvements that conserve energy, or they may simply lead firms to contract output while leaving production processes largely unchanged. This distinction is central to evaluating whether climate regulation offers a low-cost path to decarbonization. In the manufacturing sector, where end-of-pipe abatement technologies are often limited or impractical, carbon pricing and environmental regulation operate primarily through changes in production processes (Greenstone et al., 2012; Dechezleprêtre and Sato, 2017; Yamazaki, 2022; Colmer et al., 2025).

Yet the empirical literature has largely characterized productivity in a way that does not isolate its most policy-relevant dimension. Most studies treat technological change as Hicks-neutral, meaning that it raises output proportionally for all input combinations. This approach is useful for measuring overall productivity growth, but it does not capture whether productivity growth actually conserves energy or merely expands production. Theoretical research, by contrast, has long emphasized that innovation can be biased toward particular factors, as in the induced innovation hypothesis (Hicks, 1932) and models of directed technical change (Acemoglu, 2002). This bias is especially consequential for energy and climate policy. Technological change biased toward energy efficiency directly affects the energy required to produce a given level of output, and thus the scope for rebound effects when efficiency gains lower effective input costs (Jevons, 1865; Gillingham et al., 2016). We refer to this dimension of technical change as energy-augmenting productivity.

In this paper, we examine how energy-augmenting productivity shapes firms' responses to carbon pricing and the implications for climate policy design. We develop and estimate a structural production function that separates Hicks-neutral from energy-augmenting productivity using Indonesian manufacturing microdata. To identify how energy-augmenting productivity responds to changes in energy costs, we exploit large variation in energy prices induced by fossil-fuel subsidy reforms. We then use the estimated model to simulate the welfare effects of carbon pricing and quantify how this margin of productivity shapes policy outcomes.

Our analysis begins with a structural production function embedded in a plant's static optimization problem. We define energy-augmenting productivity as an input-specific shifter that increases the effective energy services obtained from a given quantity of energy. Rather than using patent counts as proxies for energy-saving technology, as much of the existing literature does (e.g., Newell et al., 1999; Popp, 2002; Aghion et al., 2016), we identify this technological dimension directly as a primitive of the production technology. This formulation allows us to conduct counterfactual policy analysis and quantify the welfare consequences of alternative climate policies.

Incorporating energy-augmenting productivity into the analysis of carbon pricing has two main advantages. First, it allows us to treat energy-augmenting productivity as a source of plant heterogeneity. Accounting for such heterogeneity is central to evaluating carbon pricing and emissions trading policies, which reallocate production from dirtier to cleaner plants and thereby deliver cost-effective abatement (Coase, 1960; Montgomery, 1972; Toyama, 2024). Second, our framework captures the interaction between different types of productivity and energy demand. Although productivity improvements can reduce emissions, they can also trigger rebound effects by lowering effective energy costs, thereby offsetting some of the gains (Gillingham et al., 2016). In our theoretical framework, improvements in Hicks-neutral productivity increase energy use through output expansion. By contrast, when inputs are complementary, improvements in energy-augmenting productivity can reduce energy use and thus emissions.

We estimate the production function and its associated productivity components by building on recent advances in the literature. First, to distinguish Hicks-neutral from energy-augmenting productivity, we use an inversion method that exploits plants' optimal input choices together with detailed data on input quantities and prices. This approach builds on the idea in Doraszelski and Jaumandreu (2018). Second, our estimation approach avoids imposing a parametric law of motion on unobserved productivity, in the spirit of Grieco et al. (2016) and Caselli et al. (2026).

Avoiding a parametric productivity process is particularly important when the response of productivity to policy is itself the object of interest. A common approach in applied work is to first estimate productivity under a first-order Markov assumption and then regress the recovered productivity measure on policy variables. As De Loecker (2013) emphasizes, however, such two-step analyses can be sensitive to assumptions about the productivity process. This concern is especially relevant in our setting. By estimating the production function without specifying how productivity evolves over time, our approach allows us to recover Hicks-neutral and energy-augmenting productivity and then examine how the latter responds to changes in energy prices.

We conduct our empirical analysis using plant-level production data from Indonesia covering the period 1990–2009. Indonesia provides an ideal setting for examining energy-augmenting productivity and its policy implications for two reasons. First, the microdata contain detailed information on energy use and prices, which is essential for implementing our estimation strategy and identifying unobserved energy-augmenting productivity. Second, energy prices exhibit substantial variation across plants and over time. Cross-sectional price differences arise from logistical constraints and infrastructure gaps in energy distribution (Rentschler and Kornejew, 2017), whereas time-series variation is largely driven by Indonesia's fossil-fuel subsidy reforms, which induced large and abrupt increases in plant-level energy prices that were plausibly exogenous to plant-level production decisions (Chelminski, 2018; Brucal and Dechezleprêtre, 2021). We exploit this variation to estimate our structural production function and quantify induced innovation.

We estimate the production function separately for each two-digit industry. The estimated elasticities of substitution range from 0.45 to 0.85 across industries, indicating that production inputs are complementary. We then recover plant-level Hicks-neutral and energy-augmenting productivity and construct aggregate productivity measures. We find that energy-augmenting productivity increased by 123% over the period 1990–2009. Applying the decomposition method of Melitz and Polanec (2015), we find that within-plant productivity growth among surviving plants (92.1%) and the entry of energy-efficient plants (90.4%) are the dominant positive contributions, partly offset by negative reallocation and exit terms.

Building on these plant-level estimates, we analyze how changes in energy prices affect productivity. Higher past energy prices are positively associated with energy-augmenting productivity. Specifically, a 1% increase in the lagged energy price is associated with a 0.2–0.75% increase in energy-augmenting productivity. The energy price at entry has an even stronger effect, with elasticities of 0.8–2.5%. This pattern is consistent with the induced innovation hypothesis. Plants that face higher energy prices at entry have stronger incentives to adopt energy-efficient technologies.

Using the estimated model, we conduct a series of counterfactual simulations of climate policies. Although Indonesia did not implement carbon pricing during our sample period, our framework can simulate a range of policies that raise the effective price of energy, including carbon taxes, allowance prices under cap-and-trade programs, and implicit shadow costs arising from emissions constraints. This simulation exercise allows us to quantify their welfare consequences, which can inform the design of future climate policies.

In our first simulation (Section 6.1), we compare the welfare effects of carbon pricing with those of a uniform regulation under which all plants are required to reduce emissions proportionally (e.g., by 10%). The results indicate that economy-wide carbon pricing achieves the same aggregate reduction in emissions as uniform regulation, while reducing welfare loss by about 86%. This gain in cost-effectiveness arises from plant heterogeneity. Under carbon pricing in which all plants face the same shadow price of emissions, plants with higher energy-augmenting productivity reduce emissions more. Consistent with this mechanism, the welfare advantage of carbon pricing across industries rises with plant-level dispersion in energy-augmenting productivity, but it is unrelated to dispersion in Hicks-neutral productivity.

To quantify the role of this heterogeneity (Section 6.2), we conduct a counterfactual analysis that homogenizes energy-augmenting productivity across plants. We find that abstracting from this heterogeneity leads to an underestimation of aggregate carbon abatement under a given carbon tax. Specifically, at a tax rate of 1 million Indonesian rupiah (IDR) per tCO_2, or about 69 EUR at the 2009 exchange rate, CO_2 abatement is 3.5% lower in the counterfactual than in the baseline.

These first two simulations hold productivity fixed at its estimated level. We conduct a third simulation to examine how carbon pricing changes productivity itself (Section 6.3), by allowing energy-augmenting productivity to respond endogenously through induced innovation. In this simulation, we incorporate the estimated effect of energy prices on energy-augmenting productivity. We find that, with endogenous productivity change driven by the higher energy cost induced by the carbon tax, aggregate emissions are 1.4–2.1% lower than in the case with exogenous (fixed) energy-augmenting productivity. Although an increase in energy-augmenting productivity may lead to higher emissions through the rebound effect, we find that most plants reduce emissions through improvements in energy-augmenting productivity. This suggests that carbon pricing may boost emissions abatement through induced innovation in the long run.

The paper proceeds as follows. We first briefly review the related literature. Section 2 introduces a model of production with multidimensional productivity. Section 3 describes the data used in the analysis and presents descriptive evidence that motivates our focus on energy-augmenting productivity. Section 4 introduces the estimation strategy. Section 5 presents the estimation results and analyzes the properties of the estimated energy-augmenting productivity. Section 6 presents a simulation analysis that examines the role of energy-augmenting productivity in carbon pricing. Section 7 concludes.

1.1 Related Literature

Our paper contributes to three strands of literature. First, we contribute to the literature on energy-saving technology and induced innovation. Building on the theoretical literature on induced innovation (e.g., Hicks, 1932; Acemoglu et al., 2012, 2016), a large body of empirical work studies how energy prices affect innovation and technology adoption using patent-based measures (e.g., Newell et al., 1999; Popp, 2002; Aghion et al., 2016; Calel and Dechezleprêtre, 2016) and energy intensity (e.g., Linn, 2008). We differ by measuring energy-saving technology as energy-augmenting productivity recovered from a structural production function estimated with plant-level microdata.[1] Although some recent studies also estimate energy-augmenting productivity using microdata (e.g., Ryan, 2018; Hawkins-Pierot and Wagner, 2025; Murray-Leclair, 2024; Butters et al., 2025), we quantify how energy prices affect both energy-augmenting and Hicks-neutral productivity and consider the implications for climate policy.

Second, our paper relates to the literature on the effects and design of environmental policies, especially carbon pricing. A broad strand of empirical literature evaluates environmental regulation and its effects on productivity and firm performance (e.g., Greenstone et al., 2012; Dechezleprêtre and Sato, 2017; Yamazaki, 2022; Colmer et al., 2025). Our contribution is to examine the role of energy-augmenting productivity in determining the efficiency and welfare consequences of carbon pricing. In particular, we quantify how heterogeneity in energy-augmenting productivity across plants affects the cost-effectiveness of carbon pricing through the reallocation of abatement toward low-cost plants. Moreover, by linking carbon prices to induced innovation in energy-augmenting productivity, we clarify the long-run consequences of carbon pricing for both productivity and emissions.

Lastly, our paper builds on the industrial organization literature on estimating production functions using production microdata.[2] In particular, we draw on two strands of work: studies that allow for multidimensional (factor-augmenting) productivity in production functions (e.g., Doraszelski and Jaumandreu, 2018; Zhang, 2019; Demirer, 2025) and those that identify production functions without imposing a Markovian law of motion (e.g., Grieco et al., 2016; Caselli et al., 2026). By incorporating insights from these studies, our empirical framework flexibly estimates productivity dynamics and quantifies induced innovation.

[1]Some papers (e.g., Hassler et al., 2021; Fried, 2018; Inoue et al., 2022; Casey, 2024) use macro-level data to measure energy-augmenting productivity at the national level and examine its macroeconomic implications.

[2]See Ackerberg et al. (2007) and De Loecker and Syverson (2021) for surveys.

2 Model

We introduce a model of production technology with factor-augmenting productivity. Section 2.1 presents the production function specification, and Section 2.2 describes the optimization problem. Section 2.3 presents comparative statics showing how Hicks-neutral and energy-augmenting productivity affect energy demand. Lastly, Section 2.4 describes how our structural model can be used to simulate the effects of environmental regulation.

2.1 Production Function

We consider a plant-level production decision. The production technology of a plant j is described by

$$Q_{jt} = F\left(K_{jt}, L_{jt}, \exp(\omega^E_{jt})E_{jt}, M_{jt}\right)\exp(\omega^H_{jt}),$$

where Q_{jt} is the output of plant j ($j = 1, \ldots, J$) in period (year) t ($t = 1, \ldots, T_j$), K_{jt} is capital, L_{jt} is labor, E_{jt} is energy, and M_{jt} is intermediate input. The production technology has two types of productivity, namely energy-augmenting productivity ω^E_{jt} and Hicks-neutral productivity ω^H_{jt}.

We consider the following constant elasticity of substitution (CES) production function:

$$\begin{aligned} &F\left(K_{jt}, L_{jt}, \exp(\omega^E_{jt})E_{jt}, M_{jt}\right) \\ &= \left\{\beta_K\left(\frac{K_{jt}}{\bar{K}}\right)^{\gamma} + \beta_L\left(\frac{L_{jt}}{\bar{L}}\right)^{\gamma} + \beta_E\left[\frac{\exp(\omega^E_{jt})}{\overline{\exp(\omega^E)}}\left(\frac{E_{jt}}{\bar{E}}\right)\right]^{\gamma} + \beta_M\left(\frac{M_{jt}}{\bar{M}}\right)^{\gamma}\right\}^{\frac{\kappa}{\gamma}} \frac{\bar{Q}}{\overline{\exp(\omega^H)}}. \end{aligned}$$

The parameter $\kappa > 0$ captures the returns to scale. The parameter $\gamma \in (-\infty, 1]$ determines the elasticity of substitution σ, which is given by $\sigma = 1/(1-\gamma)$. The production function encompasses three special cases: a Leontief technology if $\gamma \to -\infty$, a Cobb–Douglas technology if $\gamma = 0$, and a linear technology if $\gamma = 1$.[3] Lastly, the set of parameters $(\beta_K, \beta_L, \beta_E, \beta_M)$ represents distributional parameters.

For the CES production function, a normalization is required to determine the specific member of the CES family (León-Ledesma et al., 2010; Grieco et al., 2016). Following León-Ledesma et al. (2010) and Grieco et al. (2016), we set baseline points for each variable in the production function to their geometric mean $\bar{Y} = \sqrt[N]{\Pi^J_{j=1}\Pi^{T_j}_{t=1}Y_{jt}}$, where N is the sample

[3]Specifically, as $\gamma \to 0$, the CES production function takes the following Cobb–Douglas functional form: $\log(Q_{jt}) = \frac{1}{\beta}\left[\beta_K\log(K_{jt}) + \beta_L\log(L_{jt}) + \beta_E\log(E_{jt}) + \beta_M\log(M_{jt})\right] + \omega_{jt}$, where $\beta = \beta_K + \beta_L + \beta_E + \beta_M$ and $\omega_{jt} = \omega^E_{jt} + \omega^H_{jt}$.

size (i.e., $N = \sum_{j=1}^{J} T_j$), and $Y_{jt} \in \{K_{jt}, L_{jt}, E_{jt}, M_{jt}, \exp(\omega_{jt}^E), Q_{jt}\}$. The term $\exp(\omega^H)$ in the production function represents the level of Hicks-neutral productivity that allows the input and productivity combination $(\bar{K}, \bar{L}, \bar{E}, \bar{M}, \overline{\exp(\omega^E)})$ to produce $\bar{Q}$.[4] Note that, since we estimate the production function separately for each industry, these normalization factors above are specific to a given industry.

The productivity of plant j in period t is given by $(\omega_{jt}^E, \omega_{jt}^H)$. Unlike the existing literature on production function estimation (e.g., Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Ackerberg et al., 2015), we do not assume a transition process for productivity. Later, in the estimation, we propose a method that does not require an assumption about the evolution of productivity.

2.2 Optimization Problem

We consider a plant's within-period decision problem. A capital K_{jt} is determined by a prior investment decision, the structure of which we leave outside the model.[5] Given K_{jt}, the plant chooses labor L_{jt}, energy E_{jt}, and intermediate input M_{jt} to maximize its current-period profit. Plants face monopolistic competition in the output market. The market demand that plant j faces is specified as

$$Q_{jt} = P_{jt}^{\eta},$$

where η is the demand elasticity. We assume $\eta < -1$ to ensure the existence of optimal solution.

Our demand specification treats each plant as facing an independent isoelastic demand curve, $Q_{jt} = P_{jt}^{\eta}$. This abstracts from the cross-price substitution across plants that would arise under a CES aggregator with an endogenous industry price index. We adopt this simpler structure for two reasons. First, our primary object of interest is the plant-level response to carbon pricing. Modeling strategic cross-plant substitution would add an additional channel without altering the margins we seek to identify. Second, our data record plant-level revenue but not the disaggregated price data that would be needed to construct a

[4]Consider the CES production function with four production factor inputs and two types of productivity before normalization: $Q_{jt} = \left\{\beta_K K_{jt}^{\gamma} + \beta_L L_{jt}^{\gamma} + \beta_E \left[\exp(\omega_{jt}^E) E_{jt}\right]^{\gamma} + \beta_M M_{jt}^{\gamma}\right\}^{\frac{\kappa}{\gamma}} \exp(\omega_{jt}^H)$. Then, the normalization factor ω^H is defined as the value that satisfies the following: $\bar{Q} = \left\{\beta_K \bar{K}^{\gamma} + \beta_L \bar{L}^{\gamma} + \beta_E \left[\overline{\exp(\omega^E)} \bar{E}\right]^{\gamma} + \beta_M \bar{M}^{\gamma}\right\}^{\frac{\kappa}{\gamma}} \exp(\omega^H)$. For a detailed discussion of the normalization procedure, see Online Appendix 4 of Grieco et al. (2016). Their normalization method can be straightforwardly extended to our production function that incorporates four production factors and energy-augmenting productivity.

[5]Investment is a dynamic decision with capital accumulating as $K_{j,t+1} = (1-\delta)K_{jt} + I_{jt}$, where δ is the depreciation rate and I_{jt} is investment. We leave the structure of this dynamic problem unspecified, since our analysis operates at the within-period level given the predetermined K_{jt}.

CES industry aggregator. In Section 6, we discuss how this scope affects the interpretation of our counterfactual results, which should be read as a lower bound on the full equilibrium response.

The plant's static optimization problem can be written as follows:

$$\begin{aligned} \max_{L_{jt},E_{jt},M_{jt}} \quad & P_{jt}Q_{jt} - W_{jt}L_{jt} - P^E_{jt}E_{jt} - P^M_{jt}M_{jt} \\ \text{s.t.} \quad & Q_{jt} = F\left(K_{jt}, L_{jt}, \exp(\omega^E_{jt})E_{jt}, M_{jt}\right)\exp(\omega^H_{jt}) \\ & Q_{jt} = P^\eta_{jt} \end{aligned} \tag{2.1}$$

where W_{jt}, P^E_{jt}, and P^M_{jt} are the wage and the prices of energy and intermediate inputs faced by plant j in period t, respectively. The first-order conditions (FOCs) with respect to labor, energy, and intermediate inputs are given by

$$\begin{aligned} \frac{d(P_{jt}Q_{jt})}{dQ_{jt}}\frac{\partial Q_{jt}}{\partial L_{jt}} &= W_{jt}, \\ \frac{d(P_{jt}Q_{jt})}{dQ_{jt}}\frac{\partial Q_{jt}}{\partial E_{jt}} &= P^E_{jt}, \\ \frac{d(P_{jt}Q_{jt})}{dQ_{jt}}\frac{\partial Q_{jt}}{\partial M_{jt}} &= P^M_{jt}. \end{aligned}$$

2.3 Comparative Statics and Rebound Effect

For the model introduced above, we consider comparative statics regarding how productivity affects energy demand. Proposition 1 summarizes the results.[6]

Proposition 1. *Suppose that each plant solves the profit maximization problem described above and determines its input levels of production factors, L_{jt}, E_{jt}, and M_{jt}, in each period. Then, the following relationships hold:*

$$\frac{d\log E_{jt}}{d\omega^H_{jt}} = \frac{1+\eta}{\eta}\frac{A}{B} > 0 \tag{2.2}$$

$$\frac{d\log E_{jt}}{d\omega^E_{jt}} = \frac{1+\eta}{\eta}\frac{A_E}{B} + \frac{\gamma}{1-\gamma}\left(1 - \frac{-\frac{1}{\eta}A_E}{B}\right)\begin{cases} > 0 & \text{if } 0 < \gamma < 1 \\ \gtreqless 0 & \text{if } \gamma < 0. \end{cases} \tag{2.3}$$

where

$$A_K = \beta_K\left(\frac{K_{jt}}{\bar{K}}\right)^\gamma,\ A_L = \beta_L\left(\frac{L_{jt}}{\bar{L}}\right)^\gamma,\ A_E = \beta_E\left[\frac{\exp(\omega^E_{jt})}{\overline{\exp(\omega^E)}}\left(\frac{E_{jt}}{\bar{E}}\right)\right]^\gamma,\ A_M = \beta_M\left(\frac{M_{jt}}{\bar{M}}\right)^\gamma,$$

[6] Appendix B provides the proof.

$A = A_K + A_L + A_E + A_M$, *and*

$$B = -\frac{1}{\eta}(A_L + A_E + A_M) + (1-\gamma)A_K > 0.$$

Equation (2.2) shows that a higher Hicks-neutral productivity ω_{jt}^H unambiguously increases a plant's energy demand through a standard rebound mechanism. Intuitively, an increase in ω_{jt}^H lowers the marginal cost and thus expands the output Q_{jt}. Since energy is a productive input, higher output translates into greater energy use E_{jt}.

By contrast, Equation (2.3) implies that the effect of energy-augmenting productivity ω_{jt}^E on energy demand is ambiguous because it reflects two opposing forces. The first is a production-expansion effect (analogous to the Hicks-neutral productivity case), represented by $\frac{1+\eta}{\eta}\frac{A_E}{B}$, that increases E_{jt}. The second is an input-substitution effect $\frac{\gamma}{1-\gamma}\left(1 - \frac{-\frac{1}{\eta}A_E}{B}\right)$. Its sign and magnitude depend on the CES substitution parameter γ. When inputs are substitutes ($0 < \gamma < 1$), this term is always positive so that a higher ω_{jt}^E increases the overall energy use. However, when inputs are sufficiently complementary ($\gamma < 0$), the substitution channel can dominate, so a higher ω_{jt}^E may reduce energy demand.

These comparative-statics results lead to two key conclusions. First, the direction of technological change has different implications for energy demand and, therefore, carbon emissions. Second, production parameters, in particular, whether inputs are substitutes or complements, affect how energy-augmenting productivity influences energy demand. Assessing the quantitative importance of these channels is ultimately an empirical question, which we address in the following sections.

2.4 Effects of Environmental Regulation

This subsection describes how our structural model can be used to evaluate the effects of environmental regulation. We consider two regulatory regimes. The first is an emissions-cap regulation under which each plant faces a plant-level constraint on emissions. Specifically, we study the following constrained profit-maximization problem:

$$\max_{L_{jt},E_{jt},M_{jt}} \quad P_{jt}Q_{jt} - W_{jt}L_{jt} - P_{jt}^E E_{jt} - P_{jt}^M M_{jt} \quad \text{s.t.} \quad \phi_{jt}E_{jt} \leq \bar{e}, \tag{2.4}$$

where ϕ_{jt} denotes emissions intensity, measured in metric tons of CO_2 per one million British thermal units (MBtu), so that plant-level emissions are given by $e_{jt} = \phi_{jt}E_{jt}$. Let $V_{jt}(\bar{e})$ denote plant j's maximized variable profit in year t when its CO_2 emissions are restricted to be no greater than $\bar{e}$, which is measured in metric tons of CO_2. The value function $V_{jt}(\bar{e})$

therefore summarizes the maximum variable profit attainable under a given emissions cap.

Next, we consider carbon pricing, which encompasses both a carbon tax and a cap-and-trade system. In this case, the plant solves

$$\max_{L_{jt},E_{jt},M_{jt}} \quad P_{jt}Q_{jt} - W_{jt}L_{jt} - P^E_{jt}E_{jt} - P^M_{jt}M_{jt} - \tau^{\text{carbon}}\phi_{jt}E_{jt}, \tag{2.5}$$

where τ^{carbon} denotes the carbon price. Under a carbon tax, τ^{carbon} is exogenously specified by the regulator. Under a cap-and-trade system, by contrast, τ^{carbon} is determined endogenously in equilibrium so that aggregate emissions satisfy a predetermined aggregate cap. We do not explicitly model the initial allocation of emission allowances across plants. Under the maintained assumption of a frictionless and efficient permit market, the initial allocation affects the distribution of rents but not the equilibrium allocation of abatement across plants (e.g., Coase, 1960; Fowlie and Perloff, 2013). As discussed above, we use the term "carbon pricing" to refer broadly to both carbon taxes and cap-and-trade systems.

Two clarifications about the emissions coefficient ϕ_{jt} are in order. First, as we describe in Section 3, our analysis aggregates fuels and electricity into a single energy input E_{jt} measured in British thermal units, but ϕ_{jt} is constructed from plant-level fuel-specific quantities and standard emissions factors. The coefficient, therefore, retains plant-level heterogeneity in fuel composition. Plants that rely more heavily on coal carry a higher ϕ_{jt}, whereas those using more electricity or natural gas have a lower value. Second, in our counterfactual simulations, we hold ϕ_{jt} fixed at its observed value. This is appropriate for the short-run analysis we conduct because fuel mix is largely determined by capital-embodied combustion technology, such as boiler type, that does not change within our counterfactual horizon. An explicit model of fuel substitution, which would allow ϕ_{jt} to respond endogenously to relative fuel prices under carbon pricing, would require tracking fuel-specific capital stocks and is beyond the scope of this paper. We return to the implications of this choice when we discuss the simulation design in Section 6.

3 Data and Descriptive Analysis

3.1 Data Source

The primary data source is Indonesia's Annual Manufacturing Survey (*Industri Besar dan Sedang*).[7] This survey collects detailed plant-level data on the manufacturing sector, covering

[7]For empirical studies using the same data source, see, e.g., Amiti and Konings (2007); Kasahara et al. (2016); Javorcik and Poelhekke (2017); Brucal et al. (2019); Brucal and Dechezleprêtre (2021).

all manufacturing plants with 20 or more employees across 23 industries classified according to the two-digit International Standard Industrial Classification of All Economic Activities (ISIC).[8] Our analysis covers the period 1990–2009 for 11 industries.[9]

The dataset includes standard variables that are commonly available in production microdata, such as output (measured by revenue), capital, investment, employment, wages, and expenditures on intermediate inputs. Importantly, it provides detailed information on energy use, reporting both fuel and electricity in monetary terms and physical units. In the analysis, we aggregate these into a single energy input by converting them into British thermal units. Although the production function treats energy as a single aggregate, we use the underlying fuel-level detail in two ways. First, the plant-level CO_2 emissions intensity ϕ_{jt}, defined in Section 2.4, is constructed from fuel-specific quantities and standard emissions factors. Second, the shift-share instrument used to identify the production-function parameters in Section 4.4 draws on fuel-specific baseline shares and location-specific fuel prices.

By contrast, for intermediate inputs, only expenditure data are available. We address this limitation by following Grieco et al. (2016). Specifically, we exploit the optimality conditions from profit maximization to infer unobserved input prices and derive additional moment restrictions for consistently recovering the production function parameters. We return to this point in Section 4.

Nominal figures in IDR are deflated using the national consumer price index, with 2015 as the base year. We exclude observations that deviate by more than three standard deviations from the industry-year mean for each variable. In addition, we restrict the sample to plants observed for at least three consecutive years. Descriptive statistics are reported in Table 1.

3.2 Descriptive Analysis

Variation in energy prices across plants and over time is crucial for identifying energy-augmenting productivity. In our dataset, energy prices vary substantially across plants, largely reflecting Indonesia's geographic characteristics (Rentschler and Kornejew, 2017). In particular, fuel prices are strongly influenced by transportation costs, which vary markedly across regions because of variations in fuel-transport infrastructure. Electricity prices also vary with the generation capacity of local providers and the development of high-voltage transmission lines in each area.

Energy prices in our dataset also exhibit substantial variation across plants and over time.

[8]Throughout the entire sample period, industrial sectors were classified according to ISIC Revision 3.

[9]For these 11 industries, a panel dataset covering the majority of surveyed plants is constructed over the sample period.

Table 1: Descriptive Statistics

	Observations	Mean	St. Dev.	25p	75p
Revenue (billion IDR)	93,535	39.46	101.97	2.03	26.04
Capital (billion IDR)	93,535	4.84	21.88	0.23	2.64
Labor	93,535	177.59	341.70	32	163
Energy (billion Btu)	93,535	13.79	46.08	0.47	8.08
Wage (million IDR)	93,535	12.88	9.48	7.23	15.77
Energy price (1000 IDR/MBtu)	93,535	118.90	59.07	80.45	141.62
Labor expenditure (billion IDR)	93,535	2.66	6.09	0.29	2.19
Energy expenditure (billion IDR)	93,535	1.39	4.60	0.05	0.84
Material expenditure (billion IDR)	93,535	22.21	60.66	0.90	13.47
CO_2 emissions (tCO_2)	93,535	514.28	2,225.16	14.81	213.23
CO_2 intensity (tCO_2/MBtu)	93,535	0.06	0.13	0.02	0.06

Note: This table covers the variables in our dataset from 1990 to 2009.

Panel (a) of Figure 1 plots average energy prices from 1990 to 2009, and Panel (b) shows the variation in energy prices across plants in each period. Indonesia has provided large fuel subsidies since the 1960s, keeping domestic energy prices below international levels. The average energy price remained around 110,000 IDR per MBtu throughout the 1990s until the Asian financial crisis. Over this period, subsidies kept prices low and encouraged rising domestic energy consumption, contributing to a steady expansion of fiscal subsidy costs. In 2004, Indonesia transitioned from a net oil exporter to a net importer, making the subsidy system increasingly difficult to sustain (Chelminski, 2018). As a result, subsidies were cut sharply in 2005, leading to a pronounced increase in energy prices. We exploit this rapid price increase to examine how energy prices affect plant-level energy-augmenting productivity.

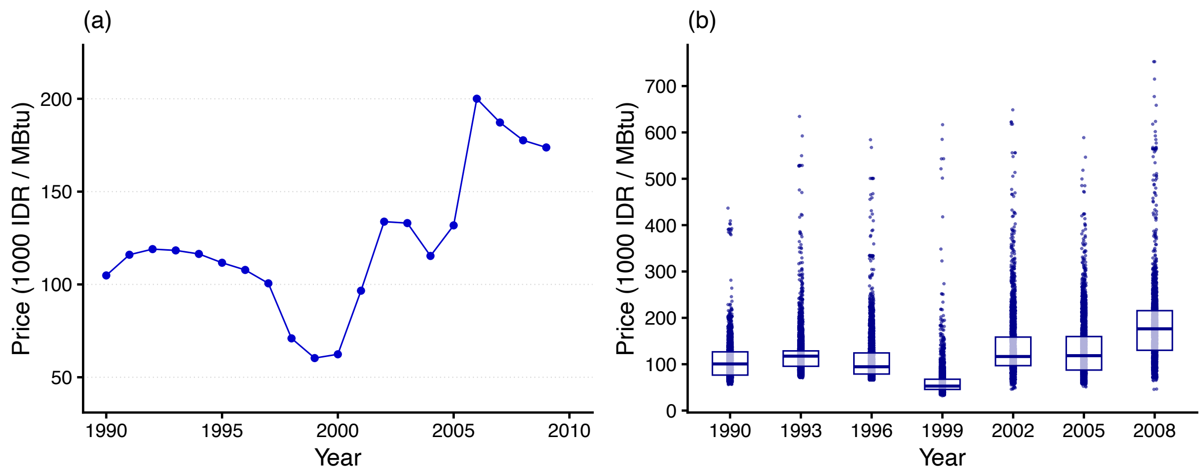

Figure 1: Variation in Energy Prices Across Plants and over Time

Note: Average prices in 1000 IDR per one million British thermal units (MBtu).

Figure 2 illustrates the relationship between energy prices and energy expenditures. In Panel (a), the solid line indicates the average share of energy expenditures in revenue, and the dashed line shows the average energy price. Over the period from 1990 to 2009, the energy expenditure share fluctuated between 4% and 6%. During the 1990s, the energy expenditure share declined alongside a decrease in energy prices. After 2000, the energy expenditure share increased over time alongside energy prices. Overall, during the sample period, the energy price and the energy expenditure share moved closely together. This pattern suggests that, on average, increases in energy prices did not lead to substantial substitution away from energy toward other inputs.

To examine the substitutability between production factors, we consider two production inputs x_1 and x_2 with corresponding prices w_1 and w_2. The ratio of factor expenditures and the ratio of input prices satisfy the following relationship:

$$\frac{d\left(\frac{w_1 x_1}{w_2 x_2}\right)}{d\left(\frac{w_1}{w_2}\right)} = (1-\sigma)\frac{x_1}{x_2}, \tag{3.1}$$

where $\sigma = -\frac{d(x_1/x_2)}{d(w_1/w_2)}\frac{(w_1/w_2)}{x_1/x_2}$.[10] This expression implies that when $0 < \sigma < 1$, the factor expenditure ratio increases with the input price ratio, and when $\sigma > 1$, the factor expenditure ratio decreases as the input price ratio rises. Panel (b) of Figure 2 shows the time evolution of the expenditure ratio between energy and labor (solid line) and their price ratio (dashed line),

[10]Taking the total differential of the factor expenditure ratio yields $d\left(\frac{w_1 x_1}{w_2 x_2}\right) = \frac{w_1}{w_2} d\left(\frac{x_1}{x_2}\right) + \frac{x_1}{x_2} d\left(\frac{w_1}{w_2}\right)$. Applying the definition of the elasticity of substitution σ, we obtain Equation (3.1).

with both series normalized to unity in 1990. Over the sample period, the two series exhibit broadly similar dynamics, suggesting that the elasticity of substitution between energy and labor in Indonesian manufacturing was likely below one during 1990–2009.

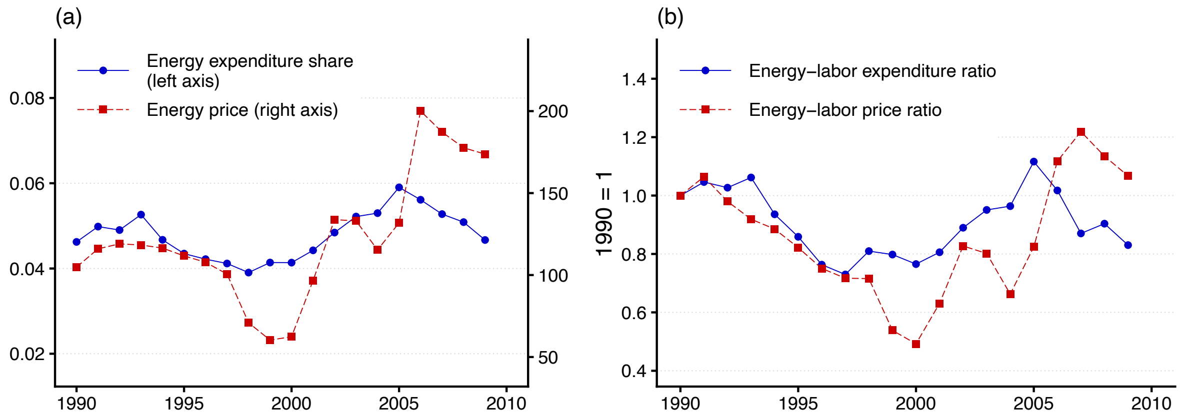

Figure 2: Time-Series Variation in Expenditure Shares and Input Prices

Note: Panel (a) shows the time-series variation in average energy expenditure share $(P^E_{jt}E_{jt})/R_{jt}$ and the average energy price (1000 IDR per MBtu). Panel (b) shows the time-series variation in the average of the ratio of energy expenditure to labor expenditure and the ratio of energy price to labor price, both normalized to one in 1990.

3.3 Suggestive Evidence for Energy-Augmenting Technological Change

This subsection provides suggestive evidence of factor-augmenting technological change in Indonesian manufacturing. Specifically, we consider an implication of a model that abstracts from factor-augmenting productivity and examine whether it holds in our dataset.

Unlike the theoretical model introduced in Section 2, we consider a CES production function with Hicks-neutral productivity only:

$$Q_{jt} = \exp(\omega^H_{jt}) \left(\beta_K K^\gamma_{jt} + \beta_L L^\gamma_{jt} + \beta_E E^\gamma_{jt} + \beta_M M^\gamma_{jt}\right)^{\frac{1}{\gamma}}, \qquad \gamma \in (-\infty, 1].$$

Combining the FOCs for labor and energy from Section 2.2, we have

$$\frac{P^E_{jt}}{W_{jt}} = \frac{\beta_E}{\beta_L}\left(\frac{E_{jt}}{L_{jt}}\right)^{\gamma-1}, \tag{3.2}$$

where P^E_{jt} and W_{jt} denote the energy price and the wage rate, respectively. This expression implies that, absent factor-biased technological change, the energy–labor input ratio is determined entirely by the relative input price P^E_{jt}/W_{jt} and the constants β_E/β_L and γ. In particular, conditional on the input price ratio, no further plant-level heterogeneity should appear in the input ratio.

We examine this relationship empirically. Motivated by Equation (3.2), we estimate the following regression for plant j in year t:

$$\log\left(\frac{E_{jt}}{L_{jt}}\right) = -\sigma \log\left(\frac{P_{jt}^E}{W_{jt}}\right) + \nu_j + \nu_t + \varepsilon_{jt},$$

where $\sigma = 1/(1-\gamma)$ is the elasticity of substitution, ν_j and ν_t denote plant and year fixed effects, respectively, and ε_{jt} is an error term.[11]

Table 2 reports the regression results. The estimated elasticity of substitution σ deviates from one in the specification with plant and year fixed effects (Specification 3), suggesting that the production technology is not well approximated by a Cobb–Douglas production function (which would imply $\sigma = 1$). In addition, the R^2 value is low, especially in specifications without fixed effects (e.g., Specification 1). Under the Hicks-neutral-only specification considered here, the input price ratio is the sole determinant of the input ratio, so the unexplained variation cannot be attributed to omitted observables. We instead read it as evidence of plant-level factor-augmenting heterogeneity, which our structural estimation in Section 4 explicitly recovers as ω_{jt}^E.

Table 2: Preliminary Analysis of Factor-Augmenting Technological Change

	(1)	(2)	(3)
$-\log(P_{jt}^E/W_{jt})$	1.011	1.051	0.527
	(0.011)	(0.012)	(0.010)
Plant FE	No	No	Yes
Year FE	No	Yes	Yes
R^2	0.267	0.276	0.811
Observations	93,535	93,535	93,535

Note: Standard errors clustered at the plant level are shown in parentheses.

4 Estimation

4.1 Overview of Estimation

We proceed to estimate the parameters γ, β_K, β_L, β_E, β_M, κ, and η of the production function introduced in Section 2 and recover the productivity components ω_{jt}^E and ω_{jt}^H. As

[11]The plant fixed effect ν_j absorbs the constant term $\log(\beta_E/\beta_L)$ in Equation (3.2).

discussed in the industrial organization literature, estimating production functions poses well-known challenges. A primary challenge is the endogeneity of inputs (Olley and Pakes, 1996; Ackerberg et al., 2015). Plants decide their input usage to maximize their profits, and these input decisions are naturally influenced by the plant's productivity components ω_{jt}^E and ω_{jt}^H, which cannot be observed directly by econometricians.

A further concern is the presence of multiple unobservables in our model. Specifically, it incorporates both energy-augmenting productivity ω_{jt}^E and Hicks-neutral productivity ω_{jt}^H (Doraszelski and Jaumandreu, 2018). In addition, the dataset does not contain information on the price of intermediate inputs P_{jt}^M, making it difficult to separately obtain intermediate-input quantities M_{jt} and prices P_{jt}^M from the observed expenditure $P_{jt}^M M_{jt}$ (Grieco et al., 2016).

To overcome these challenges, we follow the approach in the recent industrial organization literature to estimate production functions with multiple productivity components. The basic idea is to use the optimality conditions with respect to input choices, along with detailed input price data, to invert unobserved productivity. We also exploit the optimality conditions to recover the quantity of intermediate inputs following Grieco et al. (2016). To see this, recall that the FOCs with respect to L_{jt}, E_{jt}, and M_{jt} are given by

$$\frac{1+\eta}{\eta} Q_{jt}^{\frac{1}{\eta}} \cdot F_L\left(K_{jt}, L_{jt}, \exp(\omega_{jt}^E) E_{jt}, M_{jt}\right) \exp(\omega_{jt}^H) = W_{jt}, \tag{4.1}$$

$$\frac{1+\eta}{\eta} Q_{jt}^{\frac{1}{\eta}} \cdot F_E\left(K_{jt}, L_{jt}, \exp(\omega_{jt}^E) E_{jt}, M_{jt}\right) \exp(\omega_{jt}^H) = P_{jt}^E, \tag{4.2}$$

$$\frac{1+\eta}{\eta} Q_{jt}^{\frac{1}{\eta}} \cdot F_M\left(K_{jt}, L_{jt}, \exp(\omega_{jt}^E) E_{jt}, M_{jt}\right) \exp(\omega_{jt}^H) = P_{jt}^M, \tag{4.3}$$

where F_L, F_E, and F_M represent the partial derivatives of the production function F with respect to inputs L_{jt}, E_{jt}, and M_{jt}, respectively. These three FOCs allow us to derive ω_{jt}^E, ω_{jt}^H, and M_{jt} up to the parameters. Details of the inversion process are provided in Section 4.2. Subsequently, Section 4.3 describes the derivation of the estimating equations used to construct moment conditions based on unobserved productivity. Finally, a generalized method of moments (GMM) estimator is introduced in Section 4.4 on the basis of these equations.

4.1.1 Normalization, Identification, and Caveat

Before we proceed with the details of the estimation, we discuss two parameter restrictions. The first restriction is a normalization on the distributional parameters: $\beta_K + \beta_L + \beta_E + \beta_M = 1$. This restriction is innocuous because, given the geometric-mean baseline introduced in

Section 2.1, any common rescaling of $(\beta_K, \beta_L, \beta_E, \beta_M)$ can be absorbed into $\bar{Q}/\exp(\omega^H)$, leaving the production function and the implied input choices unchanged. The second restriction concerns identification. Given that only revenue data R_{jt} are available in our dataset, the returns-to-scale parameter κ and the demand elasticity η are not separately identified (Klette and Griliches, 1996; De Loecker, 2011; Chow et al., 2026). We therefore set $\kappa = 1$, imposing constant returns to scale. Unlike the sum-to-one restriction, this is a substantive identification assumption rather than a free normalization.

One caveat of working with revenue data is that the Hicks-neutral component ω_{jt}^H recovered in our framework should be interpreted as a revenue-based productivity measure rather than a pure physical efficiency term. In particular, ω_{jt}^H conflates physical Hicks-neutral efficiency with plant-level heterogeneity in output price stemming from quality, unmodeled markup variation, or idiosyncratic demand shifters (e.g., Foster et al., 2008). We therefore interpret levels and cross-plant comparisons of ω_{jt}^H with appropriate caution. In our welfare analysis in Section 6, the relative welfare differences between policy scenarios are unaffected by this measurement issue, since the counterfactual comparisons hold the cross-plant distribution of ω_{jt}^H fixed across the scenarios.[12]

4.2 Inversion

We first explain the inversion of ω_{jt}^E. Taking the ratio of the expressions for the FOCs in Equations (4.1) and (4.2), we have

$$\frac{E_{jt}}{\bar{E}} = \left(\frac{\beta_L}{\beta_E}\right)^{\frac{1}{\gamma}} \left(\frac{\exp(\omega_{jt}^E)}{\overline{\exp(\omega^E)}}\right)^{-1} \left(\frac{P_{jt}^E E_{jt}}{W_{jt} L_{jt}}\right)^{\frac{1}{\gamma}} \frac{L_{jt}}{\bar{L}}. \tag{4.4}$$

This expression allows us to obtain ω_{jt}^E up to parameters. Taking the logarithm of the above equation and rearranging it, we have

$$\omega_{jt}^E - \log\left(\overline{\exp(\omega^E)}\right) = \frac{1}{\gamma}\log\left(\frac{\beta_L}{\beta_E}\right) + \log\left(\frac{\bar{E}}{\bar{L}}\right) + \frac{1-\gamma}{\gamma}\log\left(\frac{E_{jt}}{L_{jt}}\right) + \frac{1}{\gamma}\log\left(\frac{P_{jt}^E}{W_{jt}}\right).$$

We denote this inversion by $\omega_{jt}^E \equiv h^E\left(L_{jt}, E_{jt}, W_{jt}, P_{jt}^E\right)$.

Note that the inversion of ω_{jt}^E depends only on the relative input price ratio P_{jt}^E/W_{jt} and

[12] An alternative approach would be to deflate plant-level revenue by an industry-year output price index to absorb between-industry price heterogeneity before inversion (e.g., Amiti and Konings, 2007, who use the wholesale price indices at the 5-digit ISIC level for Indonesian manufacturing). We do not pursue this route here because deflation does not remove within-industry price dispersion, the primary concern for ω_{jt}^H. Extending the framework with a structural demand model that separately identifies output-price heterogeneity (e.g., De Loecker, 2011) is left for future work.

the input quantity ratio E_{jt}/L_{jt}, both of which are directly observed at the plant level. As a result, the identification of ω_{jt}^{E} does not rely on the level of output prices and is robust to the within-industry output-price heterogeneity that contaminates ω_{jt}^{H}.

We now consider the inversion of the intermediate input $M_{jt}/\bar{M}$. Using the ratio of the FOCs, as expressed by Equations (4.1) and (4.3), we can recover the intermediate input quantity as follows:

$$\frac{M_{jt}}{\bar{M}} = \left(\frac{\beta_L}{\beta_M}\right)^{\frac{1}{\gamma}} \left(\frac{P_{jt}^{M} M_{jt}}{W_{jt} L_{jt}}\right)^{\frac{1}{\gamma}} \frac{L_{jt}}{\bar{L}}. \tag{4.5}$$

The intermediate input M_{jt} (normalized by its geometric mean $\bar{M}$) in Equation (4.5) is expressed by observables up to parameters. We denote this inversion by $\frac{M_{jt}}{\bar{M}} = h^{M}\left(L_{jt}, W_{jt}, P_{jt}^{M} M_{jt}\right)$.

Lastly, we consider the inversion of ω_{jt}^{H}. Plugging the inversion for ω_{jt}^{E} and $M_{jt}/\bar{M}$ into the FOC for L_{jt} in Equation (4.1), we obtain the inversion for ω_{jt}^{H}. The details of the derivation are provided in Appendix C.1. We denote this by

$$\omega_{jt}^{H} = h^{H}\left(K_{jt}, L_{jt}, E_{jt}, W_{jt}, P_{jt}^{E}, P_{jt}^{M} M_{jt}\right). \tag{4.6}$$

4.3 Estimating Equations

Given the inversion results obtained in Section 4.2, we now derive estimating equations for the underlying parameters. These equations are used to define the system GMM estimator in Section 4.4.

We first use the ratio of the FOCs for E_{jt} and L_{jt} given in Equation (4.4), which can be rewritten as

$$\begin{aligned}\log\left(\frac{E_{jt}}{L_{jt}}\right) =& \frac{1}{1-\gamma}\log\left(\frac{\beta_E}{\beta_L}\right) - \frac{\gamma}{1-\gamma}\log\left(\frac{\bar{E}}{\bar{L}}\right) - \frac{1}{1-\gamma}\log\left(\frac{P_{jt}^{E}}{W_{jt}}\right)\\ &+ \frac{\gamma}{1-\gamma}\left(\omega_{jt}^{E} - \log\left(\overline{\exp(\omega^{E})}\right)\right).\end{aligned}$$

Note here that since the logarithm of the geometric mean of $\exp(\omega_{jt}^{E})$, that is, $\log\left(\overline{\exp(\omega^{E})}\right)$, is a simple average of ω_{jt}^{E}, the error term in the above equation, $\omega_{jt}^{E} - \log\left(\overline{\exp(\omega^{E})}\right)$, satisfies the zero-expectation condition.

For estimation, we transform this equation to incorporate fixed effects:

$$\log\left(\frac{E_{jt}}{L_{jt}}\right) = -\frac{1}{1-\gamma}\log\left(\frac{P_{jt}^{E}}{W_{jt}}\right) + \nu_j + \nu_{pt} + \tilde{\omega}_{jt}^{E}, \tag{4.7}$$

where ν_j and ν_{pt} denote plant and province-by-year fixed effects, respectively, and $\tilde{\omega}_{jt}^{E}$ captures the remaining variation in unobserved energy-augmenting productivity. We treat $\tilde{\omega}_{jt}^{E}$ as the econometric error term for identification and estimation, as discussed in Section 4.4.

Next, we use the ratios of the FOCs given in Equations (4.4) and (4.5). Following Grieco et al. (2016), we can directly obtain the ratios β_E/β_L and β_M/β_L. Taking the geometric mean of both sides of Equations (4.4) and (4.5), respectively, we obtain

$$\frac{\beta_E}{\beta_L} = \frac{\overline{P^E E}}{\overline{WL}}, \tag{4.8}$$

$$\frac{\beta_M}{\beta_L} = \frac{\overline{P^M M}}{\overline{WL}}, \tag{4.9}$$

where $\overline{P^E E}$, $\overline{P^M M}$, and $\overline{WL}$ denote the geometric means of the energy expenditure $P_{jt}^{E} E_{jt}$, intermediate input expenditure $P_{jt}^{M} M_{jt}$ (which is observable in the data), and labor expenditure $W_{jt} L_{jt}$, respectively. Appendix C.2 provides further details of this derivation.

We now consider an estimating equation based on the revenue equation. We begin by assuming that the revenue R_{jt} is observed with error and is given by

$$R_{jt} = P_{jt} Q_{jt} \exp(\varepsilon_{jt}), \tag{4.10}$$

where ε_{jt} is a measurement error or an ex-post unexpected shock that does not affect input choices. Allowing for such ex-post shocks is common in production function estimation (e.g., Ackerberg et al., 2015).

Given the demand specification, Equation (4.10) for the revenue can be rewritten as

$$R_{jt} = P_{jt} Q_{jt} \exp(\varepsilon_{jt}) = Q_{jt}^{\frac{1+\eta}{\eta}} \exp(\varepsilon_{jt}).$$

Taking logarithms yields

$$\log R_{jt} = \frac{1+\eta}{\eta} \log F\big(K_{jt},\, L_{jt},\, \exp\big(\omega_{jt}^{E}\big)\, E_{jt},\, M_{jt}\big) + \frac{1+\eta}{\eta} \omega_{jt}^{H} + \varepsilon_{jt}. \tag{4.11}$$

By substituting in the expressions for ω_{jt}^{E} from Equation (4.4), $M_{jt}/\bar{M}$ from Equation (4.5), and ω_{jt}^{H} from Equation (4.6) into Equation (4.11), we obtain the following estimating equation:[13]

$$\log R_{jt} = -\log\left(\kappa\frac{1+\eta}{\eta}\right) + \log\left[W_{jt}L_{jt}\cdot\frac{\beta_K}{\beta_L}\left(\frac{K_{jt}/\bar{K}}{L_{jt}/\bar{L}}\right)^{\gamma} + W_{jt}L_{jt} + P^E_{jt}E_{jt} + P^M_{jt}M_{jt}\right] + \varepsilon_{jt}. \tag{4.12}$$

Note that the error term ε_{jt} is an unexpected shock that is uncorrelated with contemporaneous variables. Thus, this equation can be estimated without additional instruments.

4.4 System GMM

We estimate the model parameters using the GMM and the estimating equations and conditions in Equations (4.7), (4.8), (4.9), and (4.12). A key feature of our moment conditions is that they do not rely on assumptions about the evolution of unobserved productivity, in contrast to prior work (e.g., Olley and Pakes, 1996; Ackerberg et al., 2015). This approach enables us to estimate the production function without imposing a specific structure on the productivity dynamics. Once the production function and the corresponding productivity measures are recovered, we can flexibly examine how productivity responds to external factors such as energy prices. We conduct this analysis in Section 5.3.

First, recall Equation (4.7):

$$\log\left(\frac{E_{jt}}{L_{jt}}\right) = -\frac{1}{1-\gamma}\log\left(\frac{P^E_{jt}}{W_{jt}}\right) + \nu_j + \nu_{pt} + \tilde{\omega}^E_{jt},$$

where ν_j and ν_{pt} denote plant and province-by-year fixed effects, respectively. The moment condition is based on $\tilde{\omega}^E_{jt}$, which captures the remaining variation in unobserved energy-augmenting productivity:

$$E\left[Z_{jt}\tilde{\omega}^E_{jt}\right] = 0,$$

where Z_{jt} is an instrumental variable. The moment condition requires the instrument to be orthogonal to the component of energy-augmenting productivity that remains after partialling out plant-level and province-by-year heterogeneity, rather than to the level of ω^E_{jt} itself.

To address the potential correlation between plant-level energy prices and unobserved productivity, we construct a fixed-weight energy price index as an instrument, in the spirit

[13] This estimating equation can also be obtained using the FOC for labor input. Specifically, substituting the revenue from Equation (4.10) and Equation (C.1) from the appendix into the FOC for labor, as expressed by Equation (4.1), immediately yields Equation (4.12).

of Bartik (1991):

$$Z_{jt} = \sum_{k} s^{k}_{j,0} p^{k}_{l,t}, \tag{4.13}$$

where $s^{k}_{j,0}$ denotes the share of fuel type k in the total energy use in the baseline year, and $p^{k}_{l,t}$ is the average price of fuel type k in location l of plant j.[14]

Our identification relies primarily on the exogeneity of fuel-price shifts rather than that of baseline fuel shares. In the terminology of Borusyak et al. (2025), our design is most naturally interpreted through the exogenous-shifts approach. While the baseline fuel shares serve as predetermined exposure weights, identification results from plausibly exogenous variation in location-specific fuel prices. These prices vary over time because of fuel-subsidy reforms, as illustrated in Figure 1, and across locations because of transportation costs, logistical constraints, and infrastructure disparities in energy distribution (see, e.g., Rentschler and Kornejew, 2017). Thus, the instrument aggregates these fuel-price shocks across fuel types using predetermined baseline shares.

One potential concern is that local energy prices are highly persistent, as documented in Figure 1. The induced innovation literature suggests that plants exposed to persistently high energy prices over the long run may accumulate energy-saving technological adjustments. Thus, the level of ω^{E}_{jt} could, in principle, reflect such long-run local price environments. If these accumulated responses were driving the identifying variation, the Bartik instrument Z_{jt} could be correlated with the level of ω^{E}_{jt}.

Two features of our specification mitigate this concern. First, plant fixed effects ν_j absorb time-invariant differences in productivity levels generated by such long-run accumulation. Identification therefore relies on within-plant, within-province-year variation in exposure. Second, the shift-share structure isolates variation driven by plant-level differences in baseline fuel mix, a plant-specific characteristic that is orthogonal to the common regional trends captured by ν_{pt}. The remaining identifying variation is, therefore, the interaction between plant-level baseline exposure and province-by-year fuel-price shocks, which is plausibly separable from the accumulated induced innovation component absorbed by ν_j.

Having established the first moment condition based on Equation (4.7), we now turn to the second moment condition, which exploits the expression for revenue in Equation (4.12). Since the measurement error ε_{jt} is uncorrelated with current inputs and prices, we can estimate the parameters using nonlinear least squares with the following loss function:

$$\mathcal{L}(\theta) = \frac{1}{N} \sum_{j=1}^{J} \sum_{t=1}^{T_j} \left\{ \log R_{jt} + \log\left(\frac{1+\eta}{\eta}\right) - \log\left[W_{jt} L_{jt} \cdot \frac{\beta_K}{\beta_L} \left(\frac{K_{jt}/\bar{K}}{L_{jt}/\bar{L}}\right)^{\gamma} + VC_{jt} \right] \right\}^2$$

[14] Similar instruments are used in previous studies (e.g., Linn, 2008; Marin and Vona, 2021).

where $\theta = \left(\gamma, \frac{1+\eta}{\eta}, \beta_K\right)$ and $VC_{jt} = W_{jt}L_{jt} + P^E_{jt}E_{jt} + P^M_{jt}M_{jt}$.

The FOCs for the nonlinear least-squares problem above are given by

$$\mathbf{E}\left[\frac{\partial \mathcal{L}(\theta)}{\partial \theta}\right] = 0.$$

We incorporate this optimality condition into the system GMM estimator. The sample analogues of the above moment condition and the FOCs can be summarized as

$$\begin{bmatrix} \frac{1}{N}\sum_j \sum_t Z_{jt}\tilde{\omega}^E_{jt}(\theta) \\ \frac{\partial \mathcal{L}(\theta)}{\partial \theta} \end{bmatrix} = \frac{1}{N}\sum_j \sum_t \begin{bmatrix} Z_{jt}\tilde{\omega}^E_{jt}(\theta) \\ \frac{\partial A_{jt}(\theta)^2}{\partial \theta} \end{bmatrix} = \frac{1}{N}\sum_j \sum_t g_{jt}(\theta) = 0,$$

where $A_{jt}(\theta) = \log R_{jt} + \log\left(\frac{1+\eta}{\eta}\right) - \log\left[VC_{jt} + W_{jt}L_{jt} \cdot \frac{\beta_K}{\beta_L}\left(\frac{K_{jt}/\bar{K}}{L_{jt}/\bar{L}}\right)^{\gamma}\right]$.

We then define the following GMM objective function:

$$N\,\bar{g}(\theta)' W \bar{g}(\theta)$$

where $\bar{g}(\theta) = \frac{1}{N}\sum_j \sum_t g_{jt}(\theta)$ and W is a weighting matrix. Note that we impose the parameter restrictions implied by Equations (4.8) and (4.9), that is, $\frac{\beta_E}{\beta_L} = \frac{\overline{P^E E}}{\overline{WL}}$ and $\frac{\beta_M}{\beta_L} = \frac{\overline{P^M M}}{\overline{WL}}$, as well as the adding-up constraint $\beta_K + \beta_L + \beta_E + \beta_M = 1$.

We apply the two-step GMM estimator as follows. First, we obtain an initial estimate $\tilde{\theta}$ by setting $W = I$, where I is the identity matrix. Then, we construct an efficient weighting matrix $\hat{W}$ as

$$\hat{W} = \left(\frac{1}{N}\sum_j \sum_t g_{jt}(\tilde{\theta})g_{jt}(\tilde{\theta})' - \bar{g}(\tilde{\theta})\bar{g}(\tilde{\theta})'\right)^{-1}$$

and reestimate the parameters using this updated weight matrix.

5 Estimation Results

This section presents the estimation results and discusses their implications. First, Section 5.1 reports the GMM estimates of the production-function parameters. Next, Section 5.2 examines the estimated energy-augmenting productivity and the evolution of aggregate productivity. We then test the induced innovation hypothesis in Section 5.3. Finally, using the estimated production function, Section 5.4 quantifies plant-level CO_2 abatement costs.

5.1 Parameter Estimates

Table 3 presents the estimates of the elasticity of substitution σ, the distributional parameters $(\beta_K, \beta_L, \beta_E, \beta_M)$, and the price elasticity of demand η by industry. A comparison of the three distributional parameters for capital (β_K), labor (β_L), and intermediate inputs (β_M) reveals that intermediate input intensity is the highest across all industries, ranging from 0.55 to 0.88, followed by labor intensity. Furthermore, the estimated price elasticity of demand ranges from -3.96 to -2.14, which is consistent with estimates reported in previous studies.[15]

Figure 3 displays the estimated elasticity of substitution by industry. The estimates range from 0.45 to 0.85, which are consistent with those obtained in previous studies using micro-level data (Doraszelski and Jaumandreu, 2018; Raval, 2019).[16] The estimated elasticity of substitution is statistically significantly greater than zero but less than one for all industries, rejecting both the Leontief ($\sigma \to 0$) and Cobb–Douglas ($\sigma = 1$) production functions. Note that our estimation exploits variations in labor and energy prices to identify the elasticity of substitution. Our estimates are consistent with those from prior studies that instead use variations in different factor prices.[17]

The above estimates of the elasticity of substitution suggest that, in light of Proposition 1 in Section 2.3, improvements in energy-augmenting productivity can reduce energy demand. To evaluate this comparative-statics result using our estimates, we calculate the elasticity given in Equation (2.3), $d \log E_{jt}/d\omega_{jt}^E$. Figure 4 plots the distribution of this elasticity. Overall, the elasticity is heterogeneous across plants. A large share of the sample exhibits a negative elasticity, implying that improvements in energy-augmenting productivity reduce energy demand. However, a small fraction of the sample has a positive elasticity, meaning that higher energy-augmenting productivity increases energy demand, a phenomenon known as the rebound effect. Therefore, the overall impact of improvements in energy-augmenting productivity is theoretically ambiguous. In the counterfactual simulation presented in Section 6.3, we revisit this point and quantify the role of the rebound effect in carbon pricing.

[15] For example, Doraszelski and Jaumandreu (2018) report an average price elasticity of -3.20 in the Spanish manufacturing sector during the period 1990–2006, Oberfield and Raval (2021) report a range between -5.22 and -2.91 in the U.S. manufacturing sector, and Zhang (2019) reports a price elasticity of -3.59 in the Chinese steel industry between 2000 and 2007.

[16] Our estimates are also consistent with those reported in previous studies that use macro-level data. For example, Inoue et al. (2022) use country-industry-level data in OECD countries to estimate the elasticity of substitution among capital, labor, and energy inputs, finding that it ranges from 0.43 to 0.68.

[17] For example, Doraszelski and Jaumandreu (2018) use variations in labor and intermediate input prices to identify the elasticity of substitution among capital, labor, and intermediate inputs.

Table 3: Results of GMM Estimation

ISIC	Industry	Observations	Plants	σ	η	β_K	β_L	β_E	β_M
15	Food products and beverages	28065	4022	0.539	−3.658	0.003	0.145	0.039	0.814
				(0.007092)	(0.018529)	(0.000127)	(1.8e-05)	(5e-06)	(0.000104)
17	Textiles	10653	1746	0.542	−3.805	0.002	0.162	0.05	0.786
				(0.016235)	(0.027909)	(0.00021)	(3.4e-05)	(1.1e-05)	(0.000166)
18	Wearing apparel	8014	1472	0.504	−3.316	0.003	0.242	0.024	0.731
				(0.017889)	(0.025597)	(0.000449)	(0.000109)	(1.1e-05)	(0.000329)
19	Leather products	3115	567	0.848	−3.767	0.049	0.197	0.035	0.718
				(0.047712)	(0.114138)	(0.009383)	(0.001946)	(0.000348)	(0.00709)
20	Wood products	9508	1546	0.545	−3.158	0.003	0.173	0.056	0.769
				(0.024324)	(0.022705)	(0.000446)	(7.7e-05)	(2.5e-05)	(0.000343)
21	Paper and paper products	2583	433	0.562	−3.419	0.003	0.139	0.043	0.815
				(0.035294)	(0.053407)	(0.000827)	(0.000115)	(3.5e-05)	(0.000676)
22	Publishing and printing	3490	558	0.497	−2.72	0.007	0.238	0.045	0.709
				(0.024844)	(0.034833)	(0.002166)	(0.00052)	(9.9e-05)	(0.001547)
24	Chemicals and chemical products	6591	948	0.445	−2.352	0.0005	0.181	0.052	0.767
				(0.028257)	(0.01684)	(0.000165)	(3e-05)	(9e-06)	(0.000127)
25	Rubber and plastics products	8726	1326	0.544	−3.87	0.004	0.14	0.064	0.791
				(0.01428)	(0.036192)	(0.000435)	(6.1e-05)	(2.8e-05)	(0.000346)
26	Other non-metallic mineral products	8652	1362	0.505	−2.491	0.006	0.389	0.081	0.525
				(0.010069)	(0.013825)	(0.000451)	(0.000176)	(3.7e-05)	(0.000238)
28	Fabricated metal products	4138	677	0.543	−3.225	0.001	0.205	0.05	0.744
				(0.03839)	(0.036304)	(0.000192)	(3.9e-05)	(1e-05)	(0.000143)

Note: Standard errors are shown in parentheses.

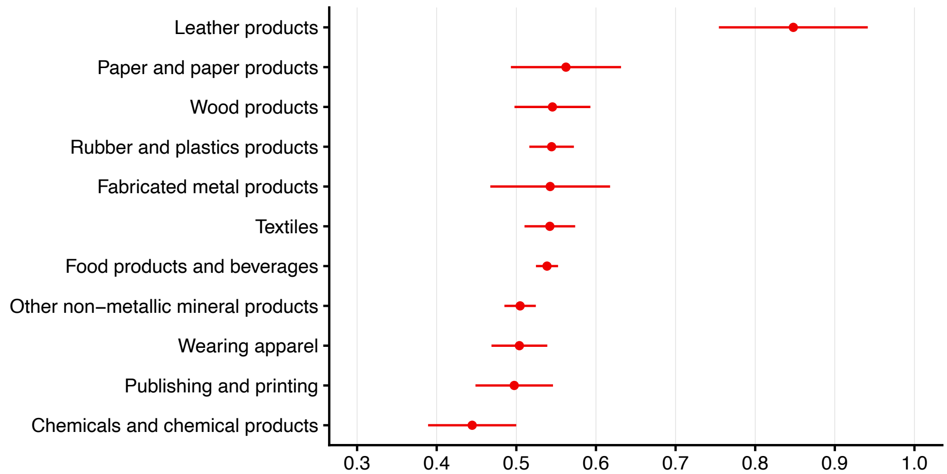

Figure 3: Estimated Elasticity of Substitution by Industry

Note: This figure shows the estimated elasticity of substitution σ obtained via the GMM estimation for each industry, along with a 95% confidence interval.

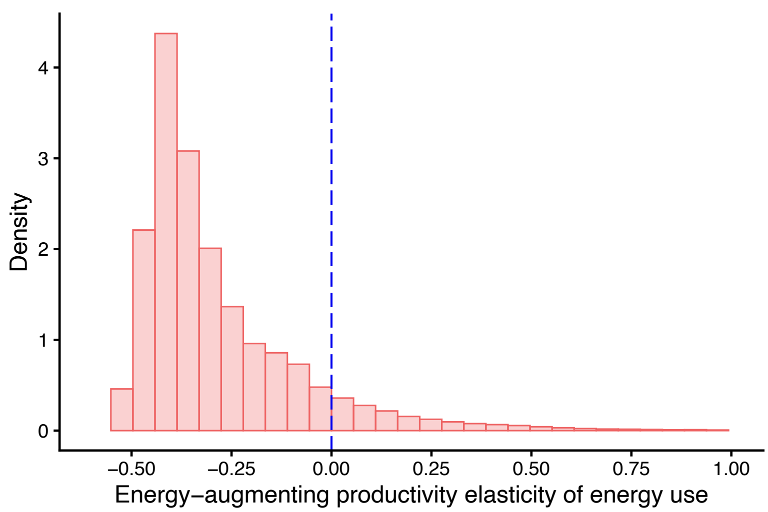

Figure 4: Histogram of Energy-Augmenting Productivity Elasticity of Energy Use

5.2 Productivity Estimates

We now provide a descriptive analysis of energy-augmenting and Hicks-neutral productivity. First, using the estimated production function, we recover the energy-augmenting productivity ω_{jt}^E and Hicks-neutral productivity ω_{jt}^H for each observation. Then, we examine the relationship between these quantities at the plant level. Figure 5 shows a binned scatter plot of these two productivity measures for each industry, which indicates an overall negative correlation between the two.[18] One might be concerned that this negative correlation arises mechanically from the inversion algebra, since ω_{jt}^H is recovered from the FOC for labor after substituting in ω_{jt}^E. As discussed in Section 4.2, however, the two productivity components are identified from distinct sources of variation. Specifically, ω_{jt}^E is identified from the observed labor–energy price ratio and the corresponding input ratio in Equation (4.4), whereas ω_{jt}^H is identified from the revenue-based inversion in Equation (4.6), conditional on ω_{jt}^E. The negative correlation, therefore, reflects a cross-plant pattern in the data rather than a purely algebraic consequence of the inversion.

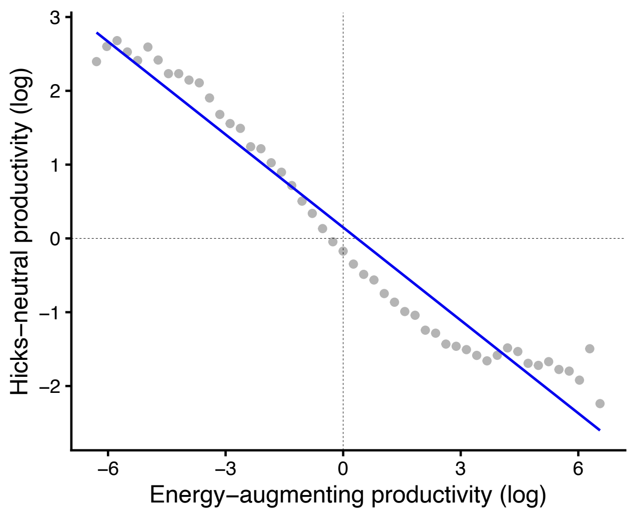

Figure 5: Estimated Energy-Augmenting Productivity and Hicks-Neutral Productivity

Note: The figure shows a binned scatter plot of $(\hat{\omega}_{jt}^E, \hat{\omega}_{jt}^H)$. To compare estimated productivity levels across industries, we regress $\hat{\omega}_{jt}^E$ and $\hat{\omega}_{jt}^H$ on industry, province, and year fixed effects and plot the residuals. The top and bottom 1% of each distribution are excluded.

We then examine the evolution of productivity at the aggregate level. For $k \in \{E, H\}$, we define aggregate productivity as

$$\Omega_t^k = \sum_{j=1}^{J_t} s_{jt} \, \zeta_{jt}^k,$$

where J_t represents the number of plants in year t, the weight s_{jt} is a market share defined

[18] Figure A.1 in the appendix presents an analogous plot using growth rates, which also shows a negative relationship.

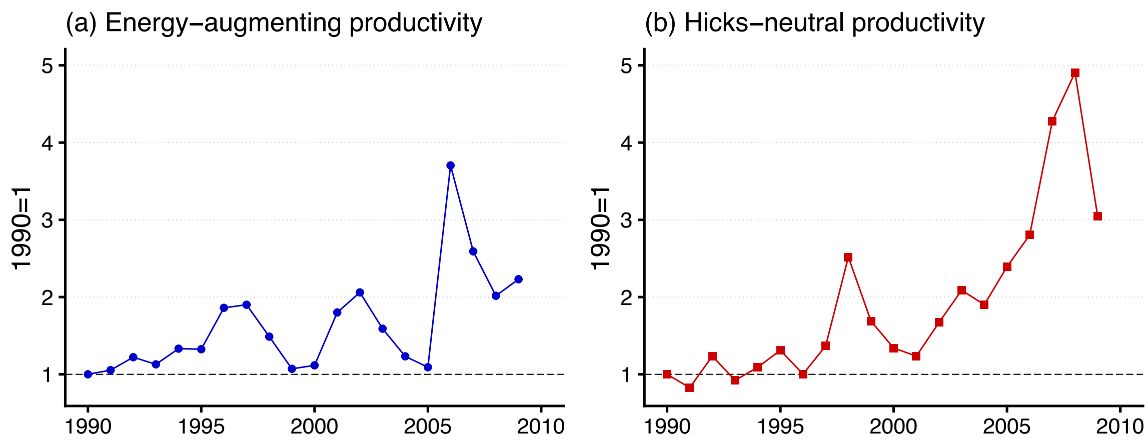

Figure 6: Changes in Estimated Productivity over Time

Note: This figure shows the changes in market-share-weighted averages of (a) $\exp(\omega^E_{jt})/(\overline{\exp(\omega^E)}\bar{E})$ and (b) $\exp(\omega^H_{jt})/(\exp(\omega^H)\bar{Q})$ relative to 1990.

by $s_{jt} = R_{jt}/\sum_{j=1}^{J_t} R_{jt}$, and ζ^k_{jt} is the normalized plant-level productivity component:

$$\zeta^E_{jt} = \frac{\exp\left(\omega^E_{jt}\right)}{\overline{\exp(\omega^E)}\ \bar{E}}, \qquad \zeta^H_{jt} = \frac{\exp\left(\omega^H_{jt}\right)}{\exp(\omega^H)\ \bar{Q}}.$$

The normalization factors $\overline{\exp(\omega^E)}\bar{E}$ and $\exp\left(\omega^H\right)\bar{Q}$ are specific to the industry of plant j. Figure 6 plots time series of two measures of aggregate productivity: energy-augmenting productivity Ω^E_t (blue circles) and Hicks-neutral productivity Ω^H_t (red squares). These are both normalized to unity in 1990 (i.e., $\Omega^E_{1990} = \Omega^H_{1990} = 1$). Over time, both indices generally increase, indicating sustained productivity growth relative to the 1990 baseline. However, they exhibit distinct dynamics. The energy-augmenting productivity Ω^E_t rises gradually through the mid-1990s, dips around the early 2000s, and then displays a sharp spike in 2006, which coincides with an energy price surge due to a substantial reduction in fuel subsidies. In contrast, Ω^H_t shows larger medium-run swings and a pronounced surge in the late 2000s.

Following Melitz and Polanec (2015), we decompose the change in aggregate productivity from year t_1 to year t_2, $\Delta\Omega^k = \Omega^k_{t_2} - \Omega^k_{t_1}$, for $k \in \{E, H\}$. The decomposition separates the change into four contributions: within-plant productivity growth among surviving plants, reallocation of market shares among survivors, entry of new plants, and exit of incumbents. We report each contribution as a share of the baseline aggregate productivity $\Omega^k_{t_1}$. The Online Appendix D provides the formal expression.

Table 4 presents the results of the decomposition analysis. Panel A focuses on energy-augmenting productivity. Over the full period 1990–2009, aggregate energy-augmenting productivity increased by 123%, with the dominant contributions coming from within-plant productivity growth among surviving plants (92.1%) and new entry of energy-efficient plants (90.4%). In contrast, the reallocation term is negative (-0.382), indicating that market shares shifted away from more energy-efficient plants toward less efficient ones. Exiting plants also contributed negatively, although the effects are quantitatively modest relative to within-plant growth. The subperiod analysis reveals substantial heterogeneity. In particular, the sharp increase in aggregate energy-augmenting productivity during 2005–2009 was driven primarily by strong within-plant improvements among surviving plants.

Panel B presents the corresponding decomposition for Hicks-neutral productivity. In contrast to Panel A, aggregate Hicks-neutral productivity growth over 1990–2009 (231%) was largely driven by reallocation among surviving plants (191%), with a comparatively smaller role for within-plant productivity growth (55.3%). Entry and exit effects were, again, small on average. This pattern was especially pronounced in the early 1990s and after 2000, when reallocation accounted for most of the observed aggregate growth. Thus, although energy-augmenting productivity growth is primarily driven by within-plant improvements, Hicks-neutral productivity growth is driven mainly by reallocation across surviving plants.

5.3 Productivity and Energy Prices

In this subsection, we examine how past energy prices affect productivity. As emphasized earlier, our production function estimation is agnostic about the evolution of productivity over time. This gives us flexibility to study how the recovered productivity measures respond to external variables, such as energy prices, within a regression model.

To examine how energy prices affect plant-level productivity, we estimate the following regression:

$$\omega_{jt} = \beta_1 \log(P^E_{j,t-1}) + \beta_2 \log(P^E_{j,t_{j0}}) + \nu_{i(j)} + \nu_{p(j)} + \nu_t + \nu_{t_0} + \varepsilon_{jt}.$$

Here, ω_{jt} denotes either energy-augmenting productivity $\hat{\omega}^E_{jt}$ or Hicks-neutral productivity $\hat{\omega}^H_{jt}$, as estimated in Section 5.2. The terms $\nu_{i(j)}$, $\nu_{p(j)}$, ν_t, and ν_{t_0} are industry, province, year, and entry year fixed effects, respectively. The regressors of interest are the lagged energy price $P^E_{j,t-1}$ and the energy price at plant j's entry year t_{j0}, denoted by $P^E_{j,t_{j0}}$. The coefficients β_1 and β_2 capture the elasticities of productivity with respect to these past energy prices. To address potential endogeneity of energy prices, we use the instruments $Z_{j,t-1}$ and $Z_{j,t_{j0}}$ defined in Equation (4.13) for $\log(P^E_{j,t-1})$ and $\log(P^E_{j,t_{j0}})$, respectively.

Table 5 reports the estimation results for energy-augmenting productivity. The lagged energy price is positively associated with current energy-augmenting productivity: depending

Table 4: Productivity Growth Decomposition

		Survivors		Entrants		Exiters	
	Aggregate productivity	Within growth	Reallocation		Unweighted		Unweighted
Period: t_1–t_2	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A. Energy-augmenting productivity							
1990–2009	1.225	0.921	−0.382	0.904	(1.527)	−0.218	(−0.594)
1990–1995	0.317	0.892	−0.600	0.093	(0.336)	−0.067	(−0.996)
1995–2000	−0.157	−0.800	0.675	0.023	(0.183)	−0.055	(−0.581)
2000–2005	−0.021	−0.371	0.377	0.044	(0.575)	−0.071	(−0.650)
2005–2009	1.047	1.381	−0.646	0.310	(2.133)	0.003	(0.031)
Panel B. Hicks-neutral productivity							
1990–2009	2.313	0.553	1.910	−0.103	(−0.174)	−0.047	(−0.128)
1990–1995	0.425	0.087	0.440	−0.102	(−0.368)	−0.001	(−0.016)
1995–2000	0.022	−0.018	−0.028	0.042	(0.340)	0.026	(0.275)
2000–2005	0.745	0.054	0.500	0.136	(1.788)	0.055	(0.497)
2005–2009	0.303	0.065	0.258	0.038	(0.259)	−0.057	(−0.596)

Note: This table reports a productivity growth decomposition for the overall manufacturing sample, following Melitz and Polanec (2015). Each entry is expressed as a share of baseline aggregate productivity Ω_{t_1}. Column (1) reports the change in aggregate productivity between t_1 and t_2, which equals the sum of Columns (2), (3), (4), and (6). Columns (2) and (3) report the contributions of surviving plants, namely within-plant productivity growth and the reallocation of market shares among survivors. Columns (4) and (6) report the contributions of entry and exit, each weighted by the market share of entering and exiting plants. Columns (5) and (7) report the corresponding unweighted productivity gaps, that is, the average productivity of entrants relative to survivors and of survivors relative to exiters, respectively. The Online Appendix D provides the formal derivation.

on the specification, a 1% increase in the lagged energy price is associated with a 0.2–0.75% increase. We also find that the energy price at entry has the largest estimated effect on energy-augmenting productivity. The elasticity of ω^E_{jt} with respect to the energy price in the entry year ranges from 0.8–2.5%. This pattern is consistent with prior studies such as Linn (2008) and Hawkins-Pierot and Wagner (2025). In particular, when plants face higher energy prices at entry, they are incentivized to adopt energy-efficient technologies, which increase their energy-augmenting productivity.[19]

[19] Table A.2 reports analogous regressions for the Hicks-neutral productivity $\hat{\omega}^H_{jt}$, which shows a negative association with past energy prices, mirroring the negative cross-plant correlation between $\hat{\omega}^E_{jt}$ and $\hat{\omega}^H_{jt}$ in Figure 5. Because ω^H_{jt} is a revenue-based measure that conflates physical efficiency with output-price heterogeneity, we treat this association with caution. Section 6.3 discusses why we hold ω^H_{jt} fixed in our long-run carbon pricing simulation.

Table 5: Effects of Energy Prices on Energy-Augmenting Productivity

	log of energy-augmenting productivity $\hat{\omega}^E_{jt}$			
	(1)	(2)	(3)	(4)
$\log(P^E_{j,t-1})$	0.538***	0.209***	0.754***	0.677***
	(0.050)	(0.051)	(0.253)	(0.252)
$\log(P^E_{j,t_{j0}})$		0.839***		2.46***
		(0.075)		(0.216)
Method	OLS	OLS	IV	IV
Industry FE	Yes	Yes	Yes	Yes
Province FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Entry year FE	No	Yes	No	Yes
Observations	74,474	74,091	74,474	74,091
R^2	0.03741	0.05786	0.03676	0.01274

Note: Standard errors clustered at the plant level are shown in parentheses.

5.4 Abatement Cost of CO_2 Emissions

We now use the estimated production function to quantify the abatement cost of CO_2 emissions. We first describe how we measure abatement costs and then summarize the implied shadow costs.

Recall the value function $V_{jt}(\bar{e})$ defined by Equation (2.4), which gives the maximum variable profit attainable under a plant-level emissions cap $\bar{e}$. We define the abatement cost as the loss in profit associated with reducing emissions by one metric ton relative to the observed business-as-usual (BAU) emissions level e^{BAU}_{jt}.[20] Using the value function, we find that the one-metric-ton abatement cost at the BAU level is

$$V_{jt}\left(e^{\text{BAU}}_{jt}\right) - V_{jt}\left(e^{\text{BAU}}_{jt} - 1\right).$$

[20]Note that the marginal abatement cost can, in principle, be measured directly as the Lagrange multiplier associated with the constrained optimization problem in Equation (2.4). Specifically, define the Lagrangian as

$$\mathcal{L} = P_{jt}Q_{jt} - W_{jt}L_{jt} - P^E_{jt}E_{jt} - P^M_{jt}M_{jt} + \lambda_{jt}\left(\bar{e} - \phi_{jt}E_{jt}\right).$$

By the envelope theorem, we have $\frac{dV_{jt}}{d\bar{e}} = \lambda_{jt}$. However, because there is no regulation on carbon emissions in the BAU scenario observed in our sample, firms are not subject to any emissions constraints. In other words, the observed emissions level e^{BAU}_{jt} is realized in the absence of a binding constraint, implying $\lambda_{jt} = 0$. Therefore, rather than calculating the *marginal* BAU abatement cost, we calculate the cost of reducing CO_2 emissions by one metric ton from the BAU level and report that measure instead.

Figure 7 presents the distribution of the estimated abatement costs. The mean and standard deviation are approximately 147,000 IDR and 501,000 IDR, respectively, indicating substantial heterogeneity across plants. This dispersion implies that uniform regulations are likely to be inefficient, whereas carbon pricing can exploit cross-plant differences in abatement costs to reduce aggregate emissions. We examine this point through counterfactual simulations in Section 6.1.

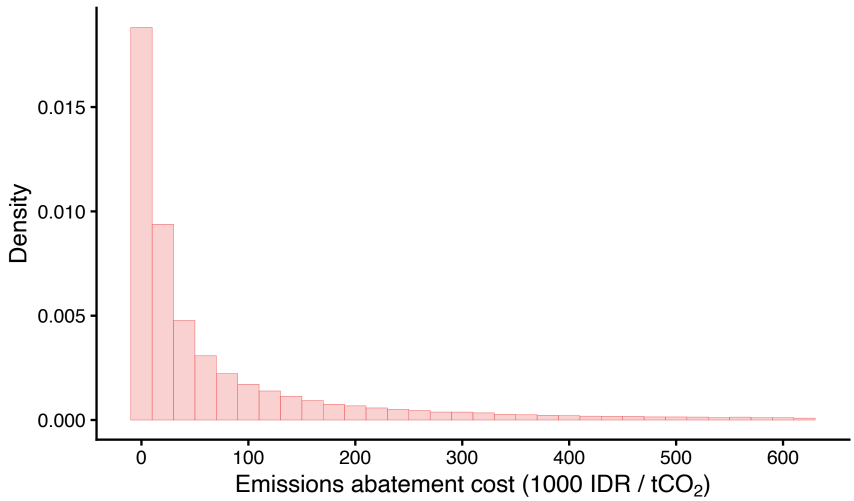

Figure 7: Estimated Costs of Emissions Abatement

Note: The figure plots the histogram of estimated shadow costs of CO_2 abatement at the plant level. The simulation uses plants emitting more than 10 metric tons of CO_2. The top 5% of the distribution is excluded.

6 Policy Simulations

Using the estimated model, we conduct a series of simulation exercises to investigate the policy implications of energy-augmenting productivity, using the 2009 sample. We first simulate the effects of carbon pricing and compare them with those of uniform regulation in Section 6.1. We then analyze the role of energy-augmenting productivity in Section 6.2. Finally, we incorporate induced innovation into the analysis of carbon pricing in Section 6.3. Throughout this section, we use the term "carbon pricing" to encompass both cap-and-trade systems (allowance trading) and carbon taxes. Although these are distinct policy instruments in principle, they can generate equivalent outcomes under certain conditions.[21]

[21]Specifically, in the absence of transaction costs or other frictions, a carbon tax and a cap-and-trade system with an appropriately chosen emissions cap both induce firms to equate their marginal abatement costs, thereby achieving a cost-effective allocation of emissions reductions. See, for example, Toyama (2024) for an empirical analysis that explicitly considers transaction costs.

6.1 Effects of Carbon Pricing

We begin with simulations of carbon-pricing policies and consider three scenarios: (1) uniform regulation, (2) an industry-wise cap-and-trade program, and (3) an economy-wide cap-and-trade program. In the first scenario, all plants are required to reduce their CO_2 emissions by 10% relative to the BAU level. In the second scenario, we impose industry-specific permit prices so that each industry's CO_2 emissions fall by 10%. In the third scenario, we implement an economy-wide cap-and-trade system with an aggregate emissions target of a 10% reduction relative to BAU.

To simulate the effects of these counterfactual policies, we solve the problems introduced in Section 2.4. Specifically, for the uniform regulation, we solve the problem in Equation (2.4) with $\bar{e}_{jt} = 0.9\, e_{jt}^{\text{BAU}}$. For the carbon-pricing simulations, we consider the plant's optimization problem given in Equation (2.5). We then solve for the equilibrium carbon price τ^{carbon} in each case such that CO_2 emissions are reduced to 90% of their BAU levels at either the industry or aggregate level.

Several modeling choices shape the interpretation of the results below. The short-run simulation holds capital K_{jt}, energy-augmenting productivity ω_{jt}^{E}, and Hicks-neutral productivity ω_{jt}^{H} fixed at their 2009 values. It also uses independent demand curves rather than a CES aggregator and holds the emissions coefficient ϕ_{jt} fixed at its observed value. Section 6.3 extends the analysis by allowing ω_{jt}^{E} to respond endogenously to the carbon price. We discuss these modeling choices in Section 6.4, where we argue that the results below should be read as a conservative lower bound on the cost-effectiveness advantage of carbon pricing.

Table 6 presents the simulation results. Relative to uniform regulation, an economy-wide cap-and-trade system achieves the same 10% emissions reduction with substantially smaller welfare losses, implying greater cost-effectiveness.[22] An industry-wise cap-and-trade system delivers an intermediate outcome between uniform regulation and economy-wide cap-and-trade. It attains the same aggregate reduction in CO_2 externality while mitigating welfare losses: total surplus falls by 0.36%, compared with 0.88% under uniform regulation. However, it is less cost-effective than an economy-wide cap-and-trade system, under which total surplus falls by only 0.12%.

This ranking is consistent with the extent to which each policy equalizes marginal abatement costs. Uniform regulation constrains all plants to achieve the same proportional reduc-

[22] Consumer surplus and producer surplus are given by $\text{CS} = \sum_j \left[\int_0^{Q_{jt}} P_{jt}(z)\, dz - P_{jt}Q_{jt}\right]$ and $\text{PS} = \sum_j \pi_{jt}$, respectively. The CO_2 externality is given by $\sum_j \text{SCC} \cdot e_{jt}$, where SCC denotes the social cost of carbon. In the simulations, we set SCC = 80 EUR, a value close to the EU ETS allowance price in November 2025.

tion, whereas industry-wise cap-and-trade reallocates abatement toward lower-cost plants within an industry via permit trading, improving efficiency relative to uniform regulation. However, because the cap is enforced separately by industry, marginal abatement costs need not be equalized across industries, leaving additional gains from trade unrealized. An economy-wide cap-and-trade program relaxes this segmentation and further improves cost-effectiveness. Consistent with this interpretation, the changes in consumer surplus are also monotone across regimes, from −1.60% under uniform regulation to −0.78% under industry-wise cap-and-trade and −0.39% under economy-wide cap-and-trade, whereas those in producer surplus are small in all cases.

Table 6: Welfare Effects

	Consumer surplus	Producer surplus	CO_2 externality	Total surplus
Business-as-usual (billion IDR)	36627.63	25168.33	712.53	61083.43
Uniform regulation (%)	−1.60	−0.08	−10.00	−0.88
Industry-wise cap-and-trade (%)	−0.78	−0.04 (−0.77)	−10.00	−0.36
Economy-wide cap-and-trade (%)	−0.39	−0.02 (−0.40)	−10.00	−0.12

Note: SCC = 80 EUR per metric ton of CO_2. 1 EUR = 14,451.82 IDR (annual average exchange rate in 2009). The numbers in parentheses represent the rate of change in producer surplus under an equivalently priced carbon tax, excluding the rent from freely allocated permits. Note that total surplus is unaffected by this distinction, since carbon tax payments constitute government revenue and therefore represent a pure transfer within the economy.

Figure 8 shows the distribution of realized marginal abatement costs after the implementation of the uniform regulation. It reveals substantial dispersion. The blue dashed vertical line indicates the equilibrium permit price under an economy-wide cap-and-trade system, which is approximately 173,000 IDR per metric ton. Under carbon pricing, this is equal to the marginal abatement cost. The observed dispersion implies that uniform regulation fails to equalize marginal abatement costs across plants, generating avoidable efficiency losses relative to carbon pricing. By contrast, carbon pricing equalizes marginal incentives across plants, reallocates abatement toward lower-cost plants, and achieves the same aggregate abatement at a lower total cost.

Figure 9 relates the relative welfare gain from a cap-and-trade system to the degree of productivity dispersion across plants. The horizontal axis measures cross-plant dispersion (the standard deviation) of productivity, and the vertical axis reports, for each industry, the ratio of total surplus under economy-wide cap-and-trade to that under uniform regulation. Panel (a) shows a clear positive relationship for energy-augmenting productivity. Industries with greater dispersion in ω_{jt}^E tend to have a higher ratio of total surplus under economy-wide cap-and-trade relative to uniform regulation, indicating that this policy delivers greater efficiency gains when plants differ more in their ability to use energy productively. A common permit price equalizes marginal abatement costs across plants by reallocating abatement toward those with lower costs, and this margin is stronger when energy-related technologies are more heterogeneous. Panel (b) shows little relationship for Hicks-neutral productivity, suggesting that dispersion in overall efficiency is less informative about the welfare advantage of a cap-and-trade program over uniform regulation. The implication is that the key source of heterogeneity governing the gains from cap-and-trade is energy-specific technology rather than general productivity. Therefore, evaluating cap-and-trade systems requires measuring heterogeneity in energy-augmenting productivity.

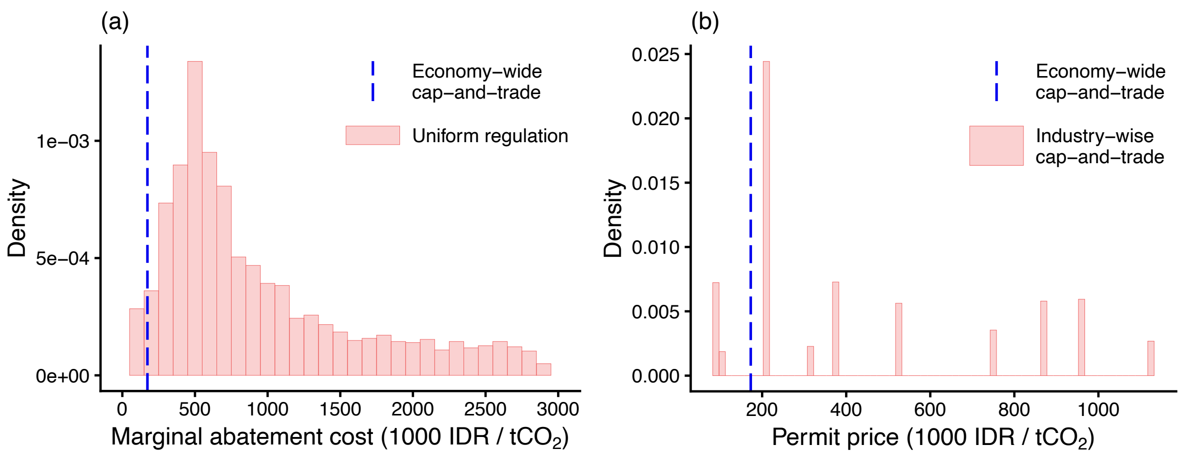

Figure 8: Marginal Abatement Costs and Permit Prices

Note: Panel (a) shows a histogram of the marginal abatement costs after abating CO_2 emissions by 10% under uniform regulation and an economy-wide cap-and-trade program. Panel (b) shows a histogram of permit prices under industry-wise cap-and-trade programs.

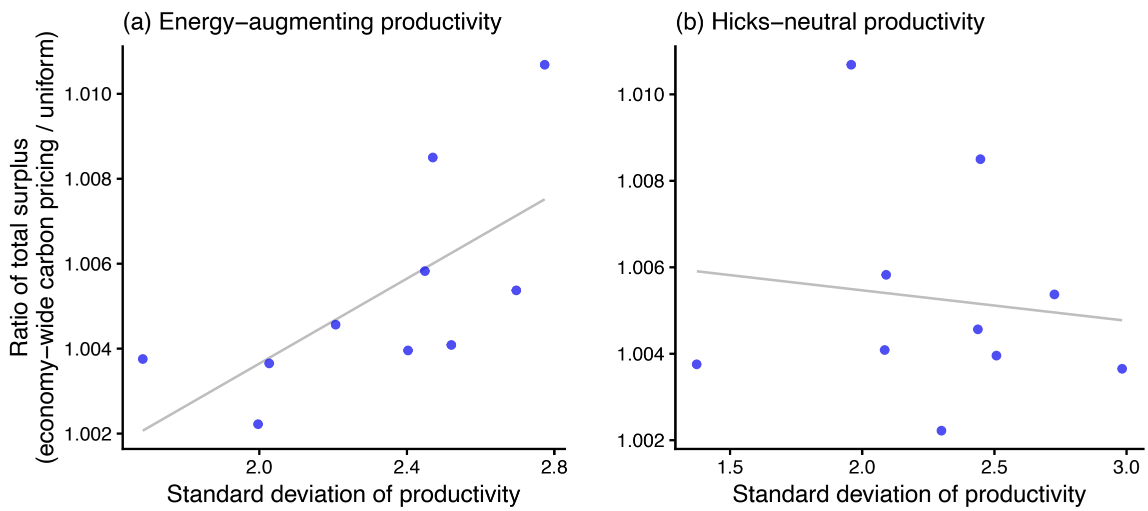

Figure 9: The Effectiveness of Carbon Pricing and Productivity Heterogeneity

Note: The figure shows, for each industry, the relationship between productivity heterogeneity and the ratio of total surplus under economy-wide cap-and-trade to that under uniform regulation. The horizontal axes in panels (a) and (b) represent the standard deviation of energy-augmenting and Hicks-neutral productivity, respectively.

6.2 Role of Heterogeneity in Energy-Augmenting Productivity

To investigate the role of energy-augmenting productivity in carbon pricing, we construct aggregate emissions as a function of the carbon tax rate. We consider two cases: (i) a baseline with heterogeneous energy-augmenting productivity across plants and (ii) a counterfactual with homogeneous energy-augmenting productivity. In the counterfactual, we equalize energy-augmenting productivity across plants (setting it to a common value) and simulate the resulting emissions–tax curve.[23]

Figure 10 plots aggregate emissions against the carbon tax rate. The solid line shows the baseline case with heterogeneous energy-augmenting productivity, and the dashed line shows the counterfactual case in which energy-augmenting productivity is equalized across plants. Overall, the baseline curve lies below the homogeneous-productivity curve, implying that accounting for heterogeneity in energy-augmenting productivity yields greater emissions abatement at a given tax rate. For example, at a tax rate of 1,000,000 IDR per tCO_2, the

[23]Specifically, we replace the term

$$\frac{\exp(\omega^E_{jt})}{\overline{\exp(\omega^E)}} \cdot \frac{1}{\bar{E}}$$

with the reciprocal of the geometric mean of E_{jt} across plants, thereby removing heterogeneity in energy-augmenting productivity. Note that $\bar{E}$ and $\overline{\exp(\omega^E)}$ are industry-specific geometric means.

degree of CO_2 abatement is 3.5% greater in the baseline than in the counterfactual.

These results have an important policy implication. The effectiveness of carbon pricing depends not only on the average level of energy-saving technology but also on its dispersion across plants. When energy-augmenting productivity is heterogeneous, a given tax rate induces greater aggregate abatement because adjustments concentrate in plants with greater scope for energy-saving responses. Consequently, analyses that ignore heterogeneity can understate expected abatement and miscalibrate the carbon tax needed to meet a given emissions target.

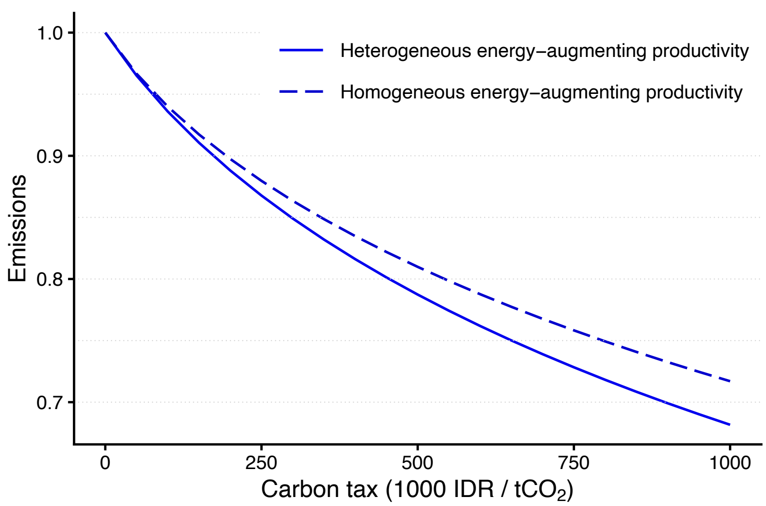

Figure 10: Aggregate Emissions and Carbon Tax: Role of Heterogeneous Energy-Augmenting Productivity

Note: This figure shows the relationship between aggregate emissions and the carbon tax. We simulate plant-level emissions under each carbon tax rate and then aggregate them. The solid curve corresponds to the baseline case with heterogeneous energy-augmenting productivity. The dashed curve shows the counterfactual case in which plant-level energy-augmenting productivity is equalized.

6.3 Carbon Pricing with Induced Innovation

We extend the short-run counterfactuals of Section 6.1 to incorporate induced innovation in energy-augmenting productivity. We treat this as an illustrative long-run projection that focuses only on changes in ω_{jt}^E induced by the carbon price. The energy-augmenting productivity ω_{jt}^E is allowed to respond endogenously to the carbon price, whereas the capital K_{jt}, the set of active plants, and the Hicks-neutral productivity ω_{jt}^H remain fixed at their 2009 values.

Note that the projected response of ω_{jt}^E reflects a mixture of capital-embodied technology, process improvements, and operational and organizational adjustments. Holding K_{jt} fixed, therefore, restricts the projection to the components of energy-augmenting productivity that respond within plants without new capital installation, rather than to the full long-run response, which would also involve capital turnover and the adoption of new vintages.

To implement this projection, we use the lagged-price coefficient from Column (4) of Table 5 and map a 1% increase in the effective energy price into a 0.677% increase in energy-augmenting productivity. We use only the lagged-price coefficient and omit the entry-year coefficient, which would require an explicit model of plant entry and technology adoption at birth. A full long-run counterfactual would also require modeling capital accumulation and firm dynamics jointly, which is beyond the scope of this paper.

One natural question is why we hold ω_{jt}^H fixed in this exercise. Although our estimates in Table A.2 indicate a negative association between past energy prices and ω_{jt}^H, we do not interpret this association as a structural response of physical Hicks-neutral efficiency to energy prices. As discussed in Section 4.1, ω_{jt}^H is a revenue-based measure that conflates physical efficiency with output-price heterogeneity. The estimated negative association may, therefore, reflect price-side adjustments, such as changes in markups or quality, rather than changes in physical efficiency.

A key objective of this counterfactual analysis is to quantify the tension between induced innovation and the rebound effect. As shown in Table 5, a higher carbon price improves energy-augmenting productivity. However, to the extent that higher energy-augmenting productivity increases energy use and emissions, these improvements may partially offset the abatement induced by the carbon price. Our simulations demonstrate this tension and clarify its implications for policy design.

Figure 11 shows aggregate emissions as a function of the carbon tax rate, with and without endogenous productivity change. The solid curve corresponds to the case with estimated (and fixed) productivity, whereas the dashed curve corresponds to the case in which energy-augmenting productivity is allowed to respond to the imposed carbon tax. Overall, with endogenous productivity change, aggregate emissions are lower, and abatement is therefore greater, at a given tax rate. When the carbon tax is set at 200,000 IDR per metric ton, aggregate emissions are 88.8% and 87.4% of BAU emissions under exogenous and endogenous productivity, respectively. The figure also suggests that the discrepancy between the two curves widens as the tax rate increases. When the carbon tax is set at 1,000,000 IDR per metric ton, aggregate emissions under endogenous productivity are 2.1% lower than those under exogenous productivity.

To examine the relative size of the rebound effect more directly, we decompose the carbon-tax effect into a direct and an indirect channel. Specifically, with endogenous productivity change, the effect of the carbon tax τ on CO_2 emissions e_{jt} is given by

$$\frac{de_{jt}\left(\tau,\omega_{jt}^{E}\right)}{d\tau} = \underbrace{\frac{\partial e_{jt}}{\partial \tau}}_{(1)\ \text{direct effect}} + \underbrace{\frac{\partial e_{jt}}{\partial \omega_{jt}^{E}}\frac{d\omega_{jt}^{E}}{d\tau}}_{(2)\ \text{indirect effect}} . \tag{6.1}$$

The first term, the direct effect, arises from the higher effective cost of energy and is captured in the baseline simulation with exogenous productivity. The second term captures the effect of the carbon tax through induced innovation.

We now quantify this second term. We simulate emissions under cases with exogenous and endogenous productivity at a carbon tax rate of 173,000 IDR, which is required to achieve a 10% reduction in emissions in the baseline case (i.e., with exogenous productivity). We then compute the difference in abatement at the plant level. Figure 12 shows the distribution of this indirect effect. It is positive for a large majority of plants (91.4%), suggesting that induced innovation further increases emissions abatement, whereas a smaller share of plants (8.6%) exhibits a negative indirect effect, indicating a rebound effect. The rebound effect is, therefore, present at the plant level but quantitatively limited in the aggregate.

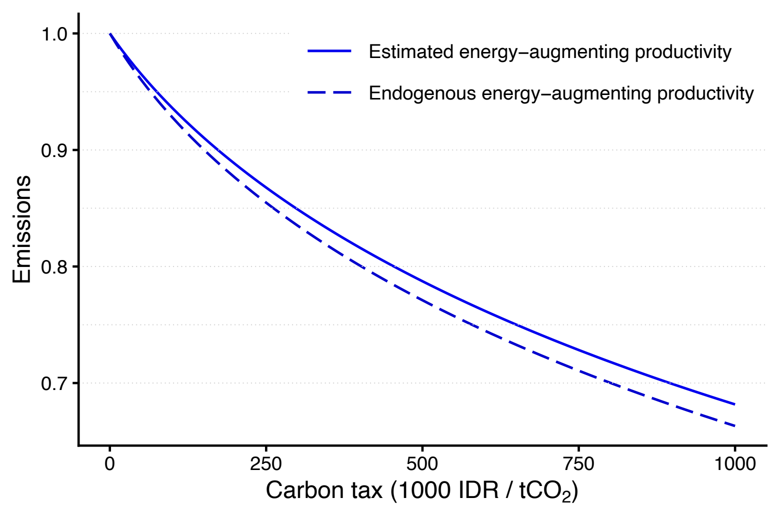

Figure 11: Aggregate Emissions and Carbon Tax with Induced Innovation

Note: This figure shows the relationship between aggregate emissions and the carbon tax. We simulate plant-level emissions under each carbon tax rate and then aggregate them. The solid curve corresponds to the baseline case with exogenous energy-augmenting productivity. The dashed curve shows the counterfactual case in which plant-level energy-augmenting productivity is allowed to respond to increases in the carbon tax.

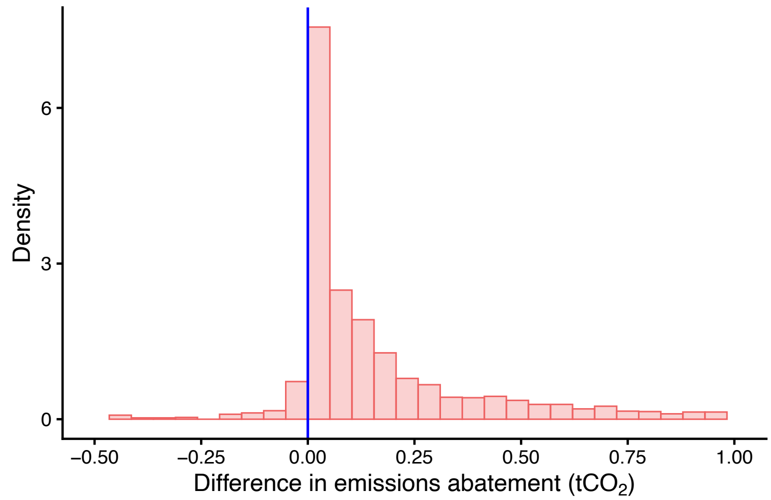

Figure 12: Indirect Effect of Energy-Augmenting Productivity Increase on Emissions Abatement

Note: This figure shows the distribution of the plant-level indirect effect of carbon pricing through induced innovation in energy-augmenting productivity. The indirect effect is defined as the difference in emissions abatement between the cases with endogenous and exogenous productivity, evaluated at a carbon tax rate of 173,000 IDR. Positive values indicate that induced innovation increases emissions abatement, whereas negative values indicate a rebound effect.

6.4 Discussion of Modeling Choices

This subsection discusses the rationale for the modeling choices introduced in Section 6.1 and explains why the cost-effectiveness gains we report should be interpreted as a conservative lower bound.

The first set of choices concerns variables held fixed at their 2009 values. Regarding capital K_{jt}, adjustment is costly and occurs at a lower frequency than the within-period choices of variable inputs. To assess the empirical plausibility of treating capital as fixed, Appendix A examines whether plants systematically adjust observed capital in response to energy-price variation. Specifically, we estimate

$$\log K_{jt} = \beta_0 + \beta_1 \log P^E_{jt} + \nu_j + \nu_t + \varepsilon_{jt},$$

where ν_j and ν_t denote plant and year fixed effects, respectively. Table A.1 shows that the estimated response of capital to energy prices is small and statistically indistinguishable from zero, suggesting that capital is approximately predetermined with respect to short-run energy-price movements in our sample. Regarding the productivity components ω^E_{jt} and ω^H_{jt}, they capture cumulative technology and efficiency adjustments that materialize over longer horizons, similar to capital. Holding all three fixed together preserves internal consistency across margins in the short-run analysis.

Our simulation also uses independent demand curves rather than a demand system with CES aggregation (see Section 2.2). We emphasize that carbon pricing still generates aggregate reallocation toward cleaner plants through differential cost pass-through, since plants with higher ω_{jt}^E face smaller marginal cost increases and contract their output less than more polluting competitors. A demand system with CES aggregation would amplify this channel through consumer-side substitution. The reallocation effects we report should be interpreted as a conservative lower bound on the full equilibrium response. Importantly, the cost-effectiveness comparison we conduct in Section 6.1 does not depend on the magnitude of reallocation. Its driver is the equalization of marginal abatement costs across plants under pricing, which operates irrespective of the demand structure.

Finally, our simulation holds the emissions coefficient ϕ_{jt} fixed at its observed value and therefore abstracts from fuel substitution. In reality, a carbon price would raise the relative price of emissions-intensive fuels and induce plants to shift toward cleaner fuels or electricity over the medium to long run, an additional abatement margin documented in the fuel-switching literature (Murray-Leclair, 2024). Because our simulation abstracts from this margin, carbon pricing operates only through reductions in total energy use and, in Section 6.3, through induced changes in ω_{jt}^E. Carbon pricing can exploit cross-plant heterogeneity in fuel-switching costs by reallocating abatement toward low-cost plants, whereas uniform regulation does not. Our results should, therefore, be interpreted as a conservative lower bound on the cost-effectiveness advantage of carbon pricing.

7 Conclusion

This paper develops and estimates a structural production function that separates Hicks-neutral and energy-augmenting productivity, using plant-level manufacturing data from Indonesia for 1990–2009. Exploiting plausibly exogenous variation in energy prices driven by fossil-fuel subsidy reforms, we find that higher energy prices raise energy-augmenting productivity, consistent with induced innovation along the energy dimension. Policy simulations show that economy-wide carbon pricing achieves the same aggregate emissions abatement as uniform regulation with substantially lower welfare losses. This cost-effectiveness advantage is driven primarily by heterogeneity in energy-augmenting productivity across plants. In an endogenous-productivity simulation, induced productivity improvements reduce aggregate emissions relative to the fixed-productivity benchmark and increase plant-level abatement for 91.4% of plants, although 8.6% exhibit a negative indirect effect consistent with rebound.

These findings have implications for the evaluation and design of climate policies. Standard analyses that summarize technological change as a single Hicks-neutral component miss both the heterogeneity that makes carbon pricing cost-effective and the innovation channel through which it operates. Accounting for energy-augmenting productivity is, therefore, essential for calibrating carbon taxes, predicting the distribution of abatement, and assessing the net effect of induced innovation after accounting for rebound.

References

Acemoglu, Daron, "Directed technical change," *The Review of Economic Studies*, 2002, *69* (4), 781–809.

_ **, Philippe Aghion, Leonardo Bursztyn, and David Hemous**, "The environment and directed technical change," *American economic review*, 2012, *102* (1), 131–66.

_ **, Ufuk Akcigit, Douglas Hanley, and William Kerr**, "Transition to clean technology," *Journal of Political Economy*, 2016, *124* (1), 52–104.

Ackerberg, Daniel A, Kevin Caves, and Garth Frazer, "Identification properties of recent production function estimators," *Econometrica*, 2015, *83* (6), 2411–2451.

Ackerberg, Daniel, C Lanier Benkard, Steven Berry, and Ariel Pakes, "Econometric tools for analyzing market outcomes," *Handbook of econometrics*, 2007, *6*, 4171–4276.

Aghion, Philippe, Antoine Dechezleprêtre, David Hemous, Ralf Martin, and John Van Reenen, "Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry," *Journal of Political Economy*, 2016, *124* (1), 1–51.

Amiti, Mary and Jozef Konings, "Trade liberalization, intermediate inputs, and productivity: Evidence from Indonesia," *American economic review*, 2007, *97* (5), 1611–1638.

Bartik, Timothy J, "Who benefits from state and local economic development policies?," 1991.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "A practical guide to shift-share instruments," *Journal of Economic Perspectives*, 2025, *39* (1), 181–204.

Brucal, Arlan and Antoine Dechezleprêtre, "Assessing the impact of energy prices on plantlevel environmental and economic performance: Evidence from Indonesian manufacturers," *OECD Environment Working Papers*, 2021, (170), 0_1–47.

_ **, Beata Javorcik, and Inessa Love**, "Good for the environment, good for business: Foreign acquisitions and energy intensity," *Journal of International Economics*, 2019, *121*, 103247.

Butters, R. Andrew, Jackson Dorsey, and Ivan Rudik, "How Firms Cut Carbon: Evidence from the European Emissions Market," 2025. Working Paper.

Calel, Raphael and Antoine Dechezleprêtre, "Environmental policy and directed technological change: evidence from the European carbon market," *Review of Economics and Statistics*, 2016, *98* (1), 173–191.

Caselli, Mauro, Arpita Chatterjee, and Shengyu Li, "Productivity and quality of multi-product firms," 2026. Working Paper.

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

Chelminski, Kathryn, "Fossil fuel subsidy reform in Indonesia," in "The Politics of Fossil Fuel Subsidies and Their Reform," Cambridge University Press Cambridge, 2018, pp. 193–211.

Chow, Chun Pang, Hiroyuki Kasahara, and Yoichi Sugita, "Identification and Estimation of Production Function and Consumer Demand Function under Monopolistic Competition from Revenue Data," March 2026. Working Paper, arXiv:2603.01492.

Coase, Ronald H, *The problem of social cost*, Springer, 1960.

Colmer, Jonathan, Ralf Martin, Mirabelle Muûls, and Ulrich J Wagner, "Does pricing carbon mitigate climate change? firm-level evidence from the european union emissions trading system," *Review of Economic Studies*, 2025, *92* (3), 1625–1660.

Dechezleprêtre, Antoine and Misato Sato, "The impacts of environmental regulations on competitiveness," *Review of Environmental Economics and Policy*, 2017.

Demirer, Mert, "Production Function Estimation with Factor-Augmenting Technology: An Application to Markups," *forthcoming at Econometrica*, 2025.

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

Foster, Lucia, John Haltiwanger, and Chad Syverson, "Reallocation, firm turnover, and efficiency: Selection on productivity or profitability?," *American Economic Review*, March 2008, *98* (1), 394–425.

Fowlie, Meredith and Jeffrey M Perloff, "Distributing pollution rights in cap-and-trade programs: are outcomes independent of allocation?," *Review of Economics and Statistics*, 2013, *95* (5), 1640–1652.

Fried, Stephie, "Climate policy and innovation: A quantitative macroeconomic analysis," *American Economic Journal: Macroeconomics*, 2018, *10* (1), 90–118.

Gillingham, Kenneth, David Rapson, and Gernot Wagner, "The rebound effect and energy efficiency policy," *Review of environmental economics and policy*, 2016.

Greenstone, Michael, John A List, and Chad Syverson, "The effects of environmental regulation on the competitiveness of US manufacturing," Technical Report, National Bureau of Economic Research 2012.

Grieco, Paul LE, Shengyu Li, and Hongsong Zhang, "Production function estimation with unobserved input price dispersion," *International Economic Review*, 2016, *57* (2), 665–690.

Hassler, John, Per Krusell, and Conny Olovsson, "Directed technical change as a response to natural resource scarcity," *Journal of Political Economy*, 2021, *129* (11), 3039–3072.

Hawkins-Pierot, Jonathan T and Katherine RH Wagner, "Technology lock-in and costs of delayed climate policy," *forthcoming at Journal of Political Economy Microeconomics*, 2025.

Hicks, John, *The theory of wages*, Springer, 1932.

Inoue, Emiko, Hiroya Taniguchi, and Ken Yamada, "Measuring energy-saving technological change: International trends and differences," *Journal of Environmental Economics and Management*, 2022, *115* (102709).

Javorcik, Beata and Steven Poelhekke, "Former foreign affiliates: Cast out and outperformed?," *Journal of the European Economic Association*, 2017, *15* (3), 501–539.

Jevons, William Stanley, *The Coal Question: An Inquiry Concerning the Progress of the Nation, and the Probable Exhaustion of Our Coal-Mines*, London: Macmillan, 1865.

Kasahara, Hiroyuki, Yawen Liang, and Joel Rodrigue, "Does importing intermediates increase the demand for skilled workers? Plant-level evidence from Indonesia," *Journal of International Economics*, 2016, *102*, 242–261.

Klette, Tor Jakob and Zvi Griliches, "The inconsistency of common scale estimators when output prices are unobserved and endogenous," *Journal of applied econometrics*, 1996, *11* (4), 343–361.

León-Ledesma, Miguel A, Peter McAdam, and Alpo Willman, "Identifying the elasticity of substitution with biased technical change," *American Economic Review*, 2010, *100* (4), 1330–1357.

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

Linn, Joshua, "Energy prices and the adoption of energy-saving technology," *The Economic Journal*, 2008, *118* (533), 1986–2012.

Loecker, Jan De, "Product differentiation, multiproduct firms, and estimating the impact of trade liberalization on productivity," *Econometrica*, 2011, *79* (5), 1407–1451.

_ , "Detecting Learning by Exporting," *American Economic Journal: Microeconomics*, 2013, *5* (3), 1–21.

_ **and Chad Syverson**, "An industrial organization perspective on productivity," in "Handbook of industrial organization," Vol. 4, Elsevier, 2021, pp. 141–223.

Marin, Giovanni and Francesco Vona, "The impact of energy prices on socioeconomic and environmental performance: Evidence from French manufacturing establishments, 1997–2015," *European Economic Review*, 2021, *135*, 103739.

Melitz, Marc J and Sašo Polanec, "Dynamic Olley-Pakes productivity decomposition with entry and exit," *The Rand Journal of Economics*, 2015, *46* (2), 362–375.

Montgomery, W David, "Markets in licenses and efficient pollution control programs," *Journal of economic theory*, 1972, *5* (3), 395–418.

Murray-Leclair, Emmanuel, "Balancing Production and Carbon Emissions with Fuel Substitution," *Available at SSRN 4855214*, 2024.

Newell, Richard G, Adam B Jaffe, and Robert N Stavins, "The induced innovation hypothesis and energy-saving technological change," *The Quarterly Journal of Economics*, 1999, *114* (3), 941–975.

Oberfield, Ezra and Devesh Raval, "Micro data and macro technology," *Econometrica*, 2021, *89* (2), 703–732.

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

Popp, David, "Induced innovation and energy prices," *American economic review*, 2002, *92* (1), 160–180.

Porter, Michael E and Claas van der Linde, "Toward a new conception of the environment-competitiveness relationship," *Journal of Economic Perspectives*, 1995, *9* (4), 97–118.

Raval, Devesh R, "The micro elasticity of substitution and non-neutral technology," *The RAND Journal of Economics*, 2019, *50* (1), 147–167.

Rentschler, Jun and Martin Kornejew, "Energy price variation and competitiveness: Firm level evidence from Indonesia," *Energy Economics*, 2017, *67*, 242–254.

Ryan, Nicholas, "Energy productivity and energy demand: Experimental evidence from Indian manufacturing plants," Technical Report, National Bureau of Economic Research 2018.

Toyama, Yuta, "Dynamic Incentives and Permit Market Equilibrium in Cap-and-Trade Regulation," *American Economic Journal: Microeconomics*, August 2024, *16* (3), 374–423.

Yamazaki, Akio, "Environmental taxes and productivity: lessons from Canadian manufacturing," *Journal of Public Economics*, 2022, *205*, 104560.

Zhang, Hongsong, "Non-neutral technology, firm heterogeneity, and labor demand," *Journal of Development Economics*, 2019, *140*, 145–168.

Appendix for Online Publication

A Additional Table and Figure

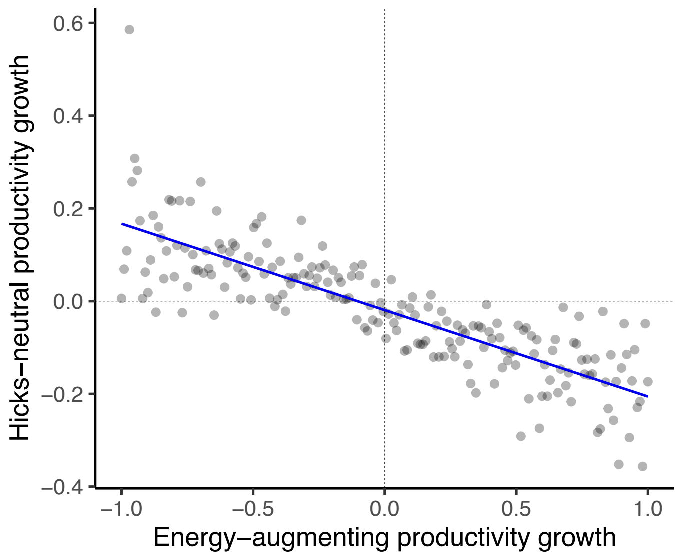

Figure A.1: Growth Rates of Two Types of Productivity

Note: Binned scatter plot. Horizontal axis: $\Delta\hat{\omega}^E_{jt} = \hat{\omega}^E_{jt} - \hat{\omega}^E_{j,t-1}$. Vertical axis: $\Delta\hat{\omega}^H_{jt} = \hat{\omega}^H_{jt} - \hat{\omega}^H_{j,t-1}$.

Table A.1: Effects of Energy Prices on Capital Stock

	Dependent variable: $\log(K_{jt})$		
	(1)	(2)	(3)
$\log(P^E_{jt})$	-0.013		0.004
	(0.013)		(0.014)
$\log(P^E_{j,t-1})$		0.007	0.006
		(0.015)	(0.014)
Plant FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Observations	93,535	74,474	74,474
R^2	0.86207	0.87281	0.87281

Note: Standard errors clustered at the plant level are shown in parentheses.

Table A.2: Effects of Energy Prices on Hicks-Neutral Productivity

	log of Hicks-neutral productivity, $\hat{\omega}^H_{jt}$			
	(1)	(2)	(3)	(4)
$\log(P^E_{j,t-1})$	-0.302^{***}	-0.036	-0.526^{**}	-0.354
	(0.050)	(0.050)	(0.252)	(0.246)
$\log(P^E_{j,t_{j0}})$		-0.701^{***}		-2.07^{***}
		(0.078)		(0.211)
Method	OLS	OLS	IV	IV
Industry FE	Yes	Yes	Yes	Yes
Province FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Entry year FE	No	Yes	No	Yes
Observations	74,474	74,091	74,474	74,091
R^2	0.02886	0.04115	0.02809	0.00784

Note: Standard errors clustered at the plant level are shown in parentheses. The dependent variable $\hat{\omega}^H_{jt}$ is a revenue-based productivity measure recovered from the FOC for labor. See Section 4.1 for a caveat regarding the interpretation of revenue-based productivity.

B Proof

This section provides the proof of Proposition 1, most of which follows Appendix E of Zhang (2019). First, recall that the FOCs for the profit maximization are written as follows:

$$\frac{1+\eta}{\eta} A^{\frac{1}{\gamma}\frac{1+\eta}{\eta}-1} \left[\frac{\exp(\omega^H_{jt})}{\exp(\omega^H)}\right]^{\frac{1+\eta}{\eta}} \bar{Q}^{\frac{1+\eta}{\eta}} \beta_L \left(\frac{1}{\bar{L}}\right)^{\gamma} L_{jt}^{\gamma-1} = W_{jt},$$

$$\frac{1+\eta}{\eta} A^{\frac{1}{\gamma}\frac{1+\eta}{\eta}-1} \left[\frac{\exp(\omega^H_{jt})}{\exp(\omega^H)}\right]^{\frac{1+\eta}{\eta}} \bar{Q}^{\frac{1+\eta}{\eta}} \beta_E \left[\frac{\exp(\omega^E_{jt})}{\exp(\omega^E)}\right]^{\gamma} \left(\frac{1}{\bar{E}}\right)^{\gamma} E_{jt}^{\gamma-1} = P^E_{jt},$$

$$\frac{1+\eta}{\eta} A^{\frac{1}{\gamma}\frac{1+\eta}{\eta}-1} \left[\frac{\exp(\omega^H_{jt})}{\exp(\omega^H)}\right]^{\frac{1+\eta}{\eta}} \bar{Q}^{\frac{1+\eta}{\eta}} \beta_M \left(\frac{1}{\bar{M}}\right)^{\gamma} M_{jt}^{\gamma-1} = P^M_{jt},$$

where

$$A = \beta_K \left(\frac{K_{jt}}{\bar{K}}\right)^{\gamma} + \beta_L \left(\frac{L_{jt}}{\bar{L}}\right)^{\gamma} + \beta_E \left[\frac{\exp(\omega^E_{jt})}{\exp(\omega^E)} \left(\frac{E_{jt}}{\bar{E}}\right)\right]^{\gamma} + \beta_M \left(\frac{M_{jt}}{\bar{M}}\right)^{\gamma}.$$

Taking the logarithm of each FOC yields

$$
\begin{aligned}
&C_L + \left(\frac{1}{\gamma}\frac{1+\eta}{\eta} - 1\right)\log A + \frac{1+\eta}{\eta}\omega_{jt}^H + (\gamma-1)\log L_{jt} = \log W_{jt},\\
&C_E + \left(\frac{1}{\gamma}\frac{1+\eta}{\eta} - 1\right)\log A + \frac{1+\eta}{\eta}\omega_{jt}^H + \gamma\omega_{jt}^E + (\gamma-1)\log E_{jt} = \log P_{jt}^E,\\
&C_M + \left(\frac{1}{\gamma}\frac{1+\eta}{\eta} - 1\right)\log A + \frac{1+\eta}{\eta}\omega_{jt}^H + (\gamma-1)\log M_{jt} = \log P_{jt}^M,
\end{aligned}
$$

where C_L, C_E, and C_M are constants. In addition, taking the total differentiation for the three expressions above, we have

$$
\left(\frac{1}{\gamma}\frac{1+\eta}{\eta} - 1\right)\frac{1}{A}\,dA + \frac{1+\eta}{\eta}\,d\omega_{jt}^H + (\gamma-1)\,d\log L_{jt} = d\log W_{jt}, \tag{B.1}
$$

$$
\left(\frac{1}{\gamma}\frac{1+\eta}{\eta} - 1\right)\frac{1}{A}\,dA + \frac{1+\eta}{\eta}\,d\omega_{jt}^H + \gamma\,d\omega_{jt}^E + (\gamma-1)\,d\log E_{jt} = d\log P_{jt}^E, \tag{B.2}
$$

$$
\left(\frac{1}{\gamma}\frac{1+\eta}{\eta} - 1\right)\frac{1}{A}\,dA + \frac{1+\eta}{\eta}\,d\omega_{jt}^H + (\gamma-1)\,d\log M_{jt} = d\log P_{jt}^M. \tag{B.3}
$$

We take differences between Equations (B.1) and (B.2) and between Equations (B.3) and (B.2), respectively, and obtain the following:

$$
d\log L_{jt} = d\log E_{jt} - \frac{\gamma}{1-\gamma}\,d\omega_{jt}^E - \frac{1}{1-\gamma}\left(d\log W_{jt} - d\log P_{jt}^E\right), \tag{B.4}
$$

$$
d\log M_{jt} = d\log E_{jt} - \frac{\gamma}{1-\gamma}\,d\omega_{jt}^E - \frac{1}{1-\gamma}\left(d\log P_{jt}^M - d\log P_{jt}^E\right). \tag{B.5}
$$

Next, by taking the total differentiation of A, we find that dA can be expressed as follows:

$$
\begin{aligned}
dA =& \gamma\beta_K\left(\frac{K_{jt}}{\bar{K}}\right)^\gamma d\log K_{jt} + \gamma\beta_L\left(\frac{L_{jt}}{\bar{L}}\right)^\gamma d\log L_{jt} + \gamma\beta_E\left[\frac{\exp(\omega_{jt}^E)}{\overline{\exp(\omega^E)}}\left(\frac{E_{jt}}{\bar{E}}\right)\right]^\gamma d\omega_{jt}^E\\
&+ \gamma\beta_E\left[\frac{\exp(\omega_{jt}^E)}{\overline{\exp(\omega^E)}}\left(\frac{E_{jt}}{\bar{E}}\right)\right]^\gamma d\log E_{jt} + \gamma\beta_M\left(\frac{M_{jt}}{\bar{M}}\right)^\gamma d\log M_{jt}\\
=& \gamma A_K\,d\log K_{jt} + \gamma A_L\,d\log L_{jt} + \gamma A_E\,d\omega_{jt}^E + \gamma A_E\,d\log E_{jt} + \gamma A_M\,d\log M_{jt}.
\end{aligned} \tag{B.6}
$$

Then, substituting Equation (B.6) into Equation (B.2) yields

$$
\begin{aligned}
&\left(\frac{1}{\gamma}\frac{1+\eta}{\eta} - 1\right)\frac{1}{A}\left(\gamma A_K\,d\log K_{jt} + \gamma A_L\,d\log L_{jt} + \gamma A_E\,d\omega_{jt}^E + \gamma A_E\,d\log E_{jt} + \gamma A_M\,d\log M_{jt}\right)\\
&+ \frac{1+\eta}{\eta}\,d\omega_{jt}^H + \gamma\,d\omega_{jt}^E + (\gamma-1)\,d\log E_{jt} = d\log P_{jt}^E.
\end{aligned}
$$

Moreover, substituting Equations (B.4) and (B.5) into this expression leads to

$$\left(\frac{1}{\gamma}\frac{1+\eta}{\eta}-1\right)\frac{1}{A}\left[\gamma A_K\, d\log K_{jt}+\gamma A_L\left\{d\log E_{jt}-\frac{\gamma}{1-\gamma}\,d\omega_{jt}^E-\frac{1}{1-\gamma}\left(d\log W_{jt}-d\log P_{jt}^E\right)\right\}\right.$$
$$\left.+\gamma A_E\, d\omega_{jt}^E+\gamma A_E\, d\log E_{jt}+\gamma A_M\left\{d\log E_{jt}-\frac{\gamma}{1-\gamma}\,d\omega_{jt}^E-\frac{1}{1-\gamma}\left(d\log P_{jt}^M-d\log P_{jt}^E\right)\right\}\right]$$
$$+\frac{1+\eta}{\eta}\,d\omega_{jt}^H+\gamma\, d\omega_{jt}^E+(\gamma-1)\,d\log E_{jt}=d\log P_{jt}^E.$$

Solving the equation for $d\log E_{jt}$, we obtain the following:

$$\left\{\left(\frac{1}{\gamma}\frac{1+\eta}{\eta}-1\right)\frac{1}{A}\gamma(A_L+A_E+A_M)+(\gamma-1)\right\}d\log E_{jt}$$
$$=-\left(\frac{1}{\gamma}\frac{1+\eta}{\eta}-1\right)\frac{1}{A}\left[\gamma A_K\, d\log K_{jt}+\gamma\left(\frac{\gamma}{\gamma-1}A_L+A_E+\frac{\gamma}{\gamma-1}A_M\right)d\omega_{jt}^E\right.$$
$$\left.+\frac{\gamma}{\gamma-1}\left(A_L\, d\log W_{jt}+A_M\, d\log P_{jt}^M\right)-\frac{\gamma}{\gamma-1}(A_L+A_M)\,d\log P_{jt}^E\right]-\frac{1+\eta}{\eta}\,d\omega_{jt}^H$$
$$-\gamma\, d\omega_{jt}^E+d\log P_{jt}^E.$$

Finally, using the above equation, we obtain the following partial derivatives:

$$\begin{aligned}\frac{d\log E_{jt}}{d\omega_{jt}^E}&=\frac{-\left(\frac{1}{\gamma}\frac{1+\eta}{\eta}-1\right)\frac{1}{A}\gamma\left(\frac{\gamma}{\gamma-1}A_L+A_E+\frac{\gamma}{\gamma-1}A_M\right)-\gamma}{\left(\frac{1}{\gamma}\frac{1+\eta}{\eta}-1\right)\frac{1}{A}\gamma(A_L+A_E+A_M)+(\gamma-1)}\\
&=\frac{\gamma\left\{(1-\gamma)A-\left(\frac{1+\eta}{\eta}-\gamma\right)(A_L+A_E+A_M)\right\}+\left(\frac{1+\eta}{\eta}-\gamma\right)A_E}{(1-\gamma)\left\{(1-\gamma)A-\left(\frac{1+\eta}{\eta}-\gamma\right)(A_L+A_E+A_M)\right\}}\\
&=\frac{\gamma\left\{-\frac{1}{\eta}(A_L+A_E+A_M)+(1-\gamma)A_K\right\}+\left(\frac{1+\eta}{\eta}-\gamma\right)A_E}{(1-\gamma)\left\{-\frac{1}{\eta}(A_L+A_E+A_M)+(1-\gamma)A_K\right\}}\\
&=\frac{\gamma}{1-\gamma}+\frac{\left(\frac{1+\eta}{\eta}-\gamma\right)A_E}{(1-\gamma)\left\{-\frac{1}{\eta}(A_L+A_E+A_M)+(1-\gamma)A_K\right\}}\end{aligned}$$

and

$$\frac{d\log E_{jt}}{d\omega_{jt}^H} = \frac{-\frac{1+\eta}{\eta}}{\left(\frac{1}{\gamma}\frac{1+\eta}{\eta}-1\right)\frac{1}{A}\gamma(A_L+A_E+A_M)+(\gamma-1)}$$
$$= \frac{\frac{1+\eta}{\eta}A}{(1-\gamma)A-\left(\frac{1+\eta}{\eta}-\gamma\right)(A_L+A_E+A_M)}$$
$$= \frac{\frac{1+\eta}{\eta}A}{-\frac{1}{\eta}(A_L+A_E+A_M)+(1-\gamma)A_K}.$$

Let us define B_1 and B_2 as

$$B_1 = -\frac{1}{\eta}(A_L+A_M)+(1-\gamma)(A_K+A_E),$$
$$B_2 = -\frac{1}{\eta}(A_L+A_E+A_M)+(1-\gamma)A_K.$$

Then, we have

$$\frac{d\log E_{jt}}{d\omega_{jt}^E} = \frac{\gamma}{1-\gamma}+\frac{\left(\frac{1+\eta}{\eta}-\gamma\right)A_E}{(1-\gamma)B_2}$$
$$= \frac{\gamma}{1-\gamma}\left(1-\frac{A_E}{B_2}\right)+\frac{1}{1-\gamma}\frac{1+\eta}{\eta}\frac{A_E}{B_2}$$
$$= \frac{\gamma}{1-\gamma}\frac{1}{B_2}\left\{-\frac{1}{\eta}(A_L+A_M)+(1-\gamma)A_K-\frac{1+\eta}{\eta}A_E\right\}+\frac{1}{1-\gamma}\frac{1+\eta}{\eta}\frac{A_E}{B_2}$$
$$= \frac{\gamma}{1-\gamma}\frac{1}{B_2}\left\{-\frac{1}{\eta}(A_L+A_M)+(1-\gamma)A_K\right\}+\frac{1+\eta}{\eta}\frac{A_E}{B_2}$$
$$= \frac{\gamma}{1-\gamma}\left(1-\frac{-\frac{1}{\eta}A_E}{B_2}\right)+\frac{1+\eta}{\eta}\frac{A_E}{B_2}$$

and

$$\frac{d\log E_{jt}}{d\omega_{jt}^H} = \frac{1+\eta}{\eta}\frac{A}{B_2}.$$

C Estimation Details

C.1 Inversion of ω_{jt}^H

This subsection explains the details of the derivation for the inversion of ω_{jt}^H. First, recall that the production function $F\left(K_{jt}, L_{jt}, \exp(\omega_{jt}^E)E_{jt}, M_{jt}\right)$ can be written as

$$F\left(K_{jt}, L_{jt}, \exp(\omega_{jt}^E)E_{jt}, M_{jt}\right) = (X_{jt})^{\frac{\kappa}{\gamma}} \frac{\bar{Q}}{\exp(\omega^H)},$$

where $X_{jt} = \beta_K \left(\frac{K_{jt}}{\bar{K}}\right)^\gamma + \beta_L \left(\frac{L_{jt}}{\bar{L}}\right)^\gamma + \beta_E \left[\frac{\exp(\omega_{jt}^E)}{\exp(\omega^E)}\left(\frac{E_{jt}}{\bar{E}}\right)\right]^\gamma + \beta_M \left(\frac{M_{jt}}{\bar{M}}\right)^\gamma$.

We first express X_{jt} in terms of observable variables using the inversion. Substituting Equations (4.4) and (4.5) into the term X_{jt}, we have

$$X_{jt} = \beta_K \left(\frac{K_{jt}}{\bar{K}}\right)^\gamma + \beta_L \left(\frac{L_{jt}}{\bar{L}}\right)^\gamma \cdot \frac{W_{jt}L_{jt} + P_{jt}^E E_{jt} + P_{jt}^M M_{jt}}{W_{jt}L_{jt}}, \tag{C.1}$$

which only depends on observed variables and unknown parameters. Using the expression of X_{jt} above and the FOC for L_{jt}, we derive the inversion of ω_{jt}^H. Taking the logarithm of both sides of Equation (4.1), we have

$$\begin{aligned}
&\log\left(\kappa\frac{1+\eta}{\eta}\right) + \left(-1 + \frac{1}{\gamma}\cdot\kappa\frac{1+\eta}{\eta}\right)\log X_{jt} + \log\left[\beta_L\left(\frac{L_{jt}}{\bar{L}}\right)^\gamma \frac{1}{L_{jt}}\right] \\
&+ \frac{1+\eta}{\eta}\left(\omega_{jt}^H - \omega^H\right) + \frac{1+\eta}{\eta}\log\bar{Q} = \log W_{jt}.
\end{aligned}$$

This can be written as

$$\begin{aligned}
\frac{1+\eta}{\eta}\omega_{jt}^H = &-\log\left(\kappa\frac{1+\eta}{\eta}\right) + \left(1 - \frac{1}{\gamma}\cdot\kappa\frac{1+\eta}{\eta}\right)\log X_{jt} - \log\left[\beta_L\left(\frac{L_{jt}}{\bar{L}}\right)^\gamma \frac{1}{L_{jt}}\right] \\
&+ \frac{1+\eta}{\eta}\omega^H - \frac{1+\eta}{\eta}\log\bar{Q} + \log W_{jt} \\
= &-\log\left(\kappa\frac{1+\eta}{\eta}\right) - \frac{1}{\gamma}\cdot\kappa\frac{1+\eta}{\eta}\log X_{jt} + \log\left[\frac{X_{jt}W_{jt}}{\beta_L\left(\frac{L_{jt}}{\bar{L}}\right)^\gamma \frac{1}{L_{jt}}}\right] \\
&+ \frac{1+\eta}{\eta}\omega^H - \frac{1+\eta}{\eta}\log\bar{Q}.
\end{aligned}$$

We now consider the term $\frac{X_{jt}W_{jt}}{\beta_L\left(\frac{L_{jt}}{\bar{L}}\right)^{\gamma}\frac{1}{L_{jt}}}$. Using the inversion results, we have

$$\begin{aligned}
\frac{X_{jt}W_{jt}}{\beta_L\left(\frac{L_{jt}}{\bar{L}}\right)^{\gamma}\frac{1}{L_{jt}}} &= \frac{W_{jt}}{\beta_L\left(\frac{L_{jt}}{\bar{L}}\right)^{\gamma}\frac{1}{L_{jt}}}\left[\beta_K\left(\frac{K_{jt}}{\bar{K}}\right)^{\gamma}+\beta_L\left(\frac{L_{jt}}{\bar{L}}\right)^{\gamma}\cdot\frac{W_{jt}L_{jt}+P^E_{jt}E_{jt}+P^M_{jt}M_{jt}}{W_{jt}L_{jt}}\right] \\
&= W_{jt}L_{jt}\left[\frac{\beta_K}{\beta_L}\left(\frac{K_{jt}/\bar{K}}{L_{jt}/\bar{L}}\right)^{\gamma}+\frac{W_{jt}L_{jt}+P^E_{jt}E_{jt}+P^M_{jt}M_{jt}}{W_{jt}L_{jt}}\right] \\
&= W_{jt}L_{jt}\cdot\frac{\beta_K}{\beta_L}\left(\frac{K_{jt}/\bar{K}}{L_{jt}/\bar{L}}\right)^{\gamma}+W_{jt}L_{jt}+P^E_{jt}E_{jt}+P^M_{jt}M_{jt}.
\end{aligned}$$

Therefore, we have

$$\begin{aligned}
\frac{1+\eta}{\eta}\omega^H_{jt} =& \frac{1+\eta}{\eta}\omega^H - \frac{1+\eta}{\eta}\log\bar{Q} - \log\left(\kappa\frac{1+\eta}{\eta}\right) - \frac{1}{\gamma}\cdot\kappa\frac{1+\eta}{\eta}\log X_{jt} \\
&+ \log\left[W_{jt}L_{jt}\cdot\frac{\beta_K}{\beta_L}\left(\frac{K_{jt}/\bar{K}}{L_{jt}/\bar{L}}\right)^{\gamma}+W_{jt}L_{jt}+P^E_{jt}E_{jt}+P^M_{jt}M_{jt}\right]. \qquad \text{(C.2)}
\end{aligned}$$

C.2 Derivation of β_M/β_L

This subsection explains the derivation of $\frac{\beta_M}{\beta_L} = \frac{\overline{P^M M}}{\overline{WL}}$ following Grieco et al. (2016). We first take the geometric mean of both sides of Equation (4.5), yielding

$$\prod_{j,t}\frac{M_{jt}}{\bar{M}} = \prod_{j,t}\left(\frac{\beta_L}{\beta_M}\right)^{\frac{1}{\gamma}}\left(\frac{P^M_{jt}M_{jt}}{W_{jt}L_{jt}}\right)^{\frac{1}{\gamma}}\frac{L_{jt}}{\bar{L}}.$$

Note that N denotes the sample size of the dataset (i.e., $N \equiv \sum_{j=1}^{J} T_j$). Then,

$$\begin{aligned}
\left(\prod_{j,t}\frac{M_{jt}}{\bar{M}}\right)^{1/N} &= \left[\prod_{j,t}\left(\frac{\beta_L}{\beta_M}\right)^{\frac{1}{\gamma}}\left(\frac{P^M_{jt}M_{jt}}{W_{jt}L_{jt}}\right)^{\frac{1}{\gamma}}\frac{L_{jt}}{\bar{L}}\right]^{1/N} \\
\frac{\left(\prod_{j,t}M_{jt}\right)^{1/N}}{\bar{M}} &= \left(\frac{\beta_L}{\beta_M}\right)^{\frac{1}{\gamma}}\left[\prod_{j,t}\left(\frac{P^M_{jt}M_{jt}}{W_{jt}L_{jt}}\right)^{\frac{1}{\gamma}}\frac{L_{jt}}{\bar{L}}\right]^{1/N} \\
&= \left(\frac{\beta_L}{\beta_M}\right)^{\frac{1}{\gamma}}\left[\prod_{j,t}\left(\frac{P^M_{jt}M_{jt}}{W_{jt}L_{jt}}\right)^{\frac{1}{\gamma}}\right]^{1/N}\left[\prod_{j,t}\frac{L_{jt}}{\bar{L}}\right]^{1/N}.
\end{aligned}$$

Noting that $\bar{M} = \left(\prod_{j,t} M_{jt}\right)^{1/N}$ and $\bar{L} = \left(\prod_{j,t} L_{jt}\right)^{1/N}$, we have

$$\begin{aligned}
1 &= \left(\frac{\beta_L}{\beta_M}\right)^{\frac{1}{\gamma}} \left[\prod_{j,t} \left(\frac{P^M_{jt} M_{jt}}{W_{jt} L_{jt}}\right)^{\frac{1}{\gamma}}\right]^{1/N} \\
&= \left(\frac{\beta_L}{\beta_M}\right)^{\frac{1}{\gamma}} \left[\left(\frac{\prod_{j,t} P^M_{jt} M_{jt}}{\prod_{j,t} W_{jt} L_{jt}}\right)^{1/N}\right]^{1/\gamma} \\
&= \left(\frac{\beta_L}{\beta_M}\right)^{\frac{1}{\gamma}} \left(\frac{\overline{P^M M}}{\overline{WL}}\right)^{\frac{1}{\gamma}}.
\end{aligned}$$

Therefore,

$$\frac{\beta_M}{\beta_L} = \frac{\overline{P^M M}}{\overline{WL}}.$$

D Productivity Growth Decomposition

This section provides the formal derivation of the decomposition introduced in Section 5.2. We focus on the change in aggregate productivity from year t_1 to year t_2, $\Delta\Omega^k = \Omega^k_{t_2} - \Omega^k_{t_1}$, for $k \in \{E, H\}$. To avoid notational clutter, the superscript k is omitted in the derivation below.

To facilitate the analysis, we classify plants into the following three groups:

1. Surviving plants ($\mathcal{A}$): Plants that remain in the industry from t_1 to t_2.
2. Entrant plants ($\mathcal{B}$): Plants that enter the industry after t_1 and before t_2.
3. Exiting plants ($\mathcal{C}$): Plants that exit the industry between t_1 and t_2.

Then, applying the decomposition method of Melitz and Polanec (2015) to $\Delta\Omega$, we have the following expression:

$$\Delta\Omega = \Delta\bar{\zeta}(\mathcal{A}) + \Delta\mathrm{Cov}(\mathcal{A}) + s_{t_2}(\mathcal{B})\left\{\Omega_{t_2}(\mathcal{B}) - \Omega_{t_2}(\mathcal{A})\right\} + s_{t_1}(\mathcal{C})\left\{\Omega_{t_1}(\mathcal{A}) - \Omega_{t_1}(\mathcal{C})\right\}. \quad \text{(D.1)}$$

Here, $s_t(g) = \sum_{j\in g} s_{jt}$ represents the market share of group $g \in \{\mathcal{A}, \mathcal{B}, \mathcal{C}\}$, $s_{jt}(g) = s_{jt}/s_t(g)$ is the market share of a plant within group g, and $\Omega_t(g) = \sum_{j\in g} s_{jt}(g)\zeta_{jt}$ is the weighted average of ζ_{jt} in group g. The first term is given by $\Delta\bar{\zeta}(\mathcal{A}) = \bar{\zeta}_{t_2}(\mathcal{A}) - \bar{\zeta}_{t_1}(\mathcal{A})$ in Equation (D.1), where $\bar{\zeta}_t(\mathcal{A}) = \frac{1}{N_{\mathcal{A}}}\sum_{j\in\mathcal{A}} \zeta_{jt}$ and $N_{\mathcal{A}}$ is the number of plants in group $\mathcal{A}$. This term reflects the improvement in energy-augmenting productivity among surviving plants.

The second term represents the change in $\mathrm{Cov}_t(\mathcal{A})$ from t_1 to t_2, where

$$\mathrm{Cov}_t(\mathcal{A}) = \sum_{j \in \mathcal{A}} (\zeta_{jt} - \bar{\zeta}_t(\mathcal{A}))(s_{jt}(\mathcal{A}) - \bar{s}_t(\mathcal{A}))$$

and $\bar{s}_t(\mathcal{A}) = 1/N_{\mathcal{A}}$. This is referred to as the reallocation term, as it captures the increase in aggregate energy-augmenting productivity resulting from a shift in market share toward plants with higher productivity. The third and fourth terms represent the contributions of entrant and exiting plants, respectively, to the aggregate energy-augmenting productivity. Lastly, rather than reporting the absolute change in aggregate productivity (i.e., $\Delta\Omega$), we report the change relative to the baseline period, defined as

$$\frac{\Delta\Omega}{\Omega_{t_1}} = \frac{\Delta\bar{\zeta}(\mathcal{A})}{\Omega_{t_1}} + \frac{\Delta\mathrm{Cov}(\mathcal{A})}{\Omega_{t_1}} + \frac{s_{t_2}(\mathcal{B})\left\{\Omega_{t_2}(\mathcal{B}) - \Omega_{t_2}(\mathcal{A})\right\}}{\Omega_{t_1}} + \frac{s_{t_1}(\mathcal{C})\left\{\Omega_{t_1}(\mathcal{A}) - \Omega_{t_1}(\mathcal{C})\right\}}{\Omega_{t_1}}. \quad \text{(D.2)}$$