



// NO.25-047 | 09/2025

DISCUSSION PAPER

// KATHRINE VON GRAEVENITZ, JOSCHA KRUG,
AND ELISA ROTTNER

Rising Energy Prices Without Falling Consumption? The Role of Energy Price Dispersion in a Multi-Product World

Rising Energy Prices Without Falling Consumption? The Role of Energy Price Dispersion in a Multi-Product World

Kathrine von Graevenitz^{*1, 2} Joscha Krug^{†1, 2} Elisa Rottner^{‡1}

¹ ZEW – Leibniz Centre for European Economic Research

² University of Mannheim

12 September 2025

Governments around the world are under pressure to reduce industrial energy use and emissions without losing out to international competition. For this reason, climate policies often come with exemptions or additional support for large energy-intensive firms, increasing the heterogeneity in energy prices. We document such a rising dispersion in industrial energy prices in the German manufacturing sector that coincides with rising average energy prices. Surprisingly, we observe an increase in industrial energy intensity, while at the same time, manufacturing firms have shifted toward producing less energy intensive products. We develop a model of multi-product firms with heterogeneous energy prices and heterogeneous products that can partially explain this puzzle via a ‘reshuffling’ among producers: If energy prices rise only for a share of firms, those firms will drop energy-intensive products. But the remaining low energy price firms will increase their market share of these products and produce them in a less energy-efficient way. Empirical analyses based on German administrative firm data suggest that such a ‘reshuffling’ is indeed taking place. We show in a simple quantification that reshuffling can have sizable effects on aggregate energy intensity.

Keywords: Product choice, Energy intensity, Carbon emissions, Manufacturing

JEL classification: Q41, D21, D22

1 Introduction

In a world of accelerating climate change, reducing carbon emissions and (fossil) energy consumption is crucial across all sectors of the economy. This is particularly true for the manufacturing sector, which made up 29% of global greenhouse gas emissions in 2016 (Ritchie & Roser, 2023) but has proven to be highly difficult to decarbonise. One key difficulty for policy makers has been to balance local

^{*}Email: kathrine.vongraevenitz@zew.de

[†]Email: joscha.krug@zew.de (corresponding author)

[‡]Email: elisa.rottner@zew.de

climate policy ambition with international competitive pressure. This has led many governments to provide generous exemptions or subsidies for large firms in particularly energy-intensive industries. While these may help selected firms to compete on international product markets, they also have the potential to undermine industrial climate action and domestic allocative efficiency.

Past research has empirically examined the effects of such exemptions and subsidies on firm-level outcomes such as energy use, revenue or employment (Basaglia et al., 2024; Gerster & Lamp, 2024). However, this existing research does not allow us to fully understand the aggregate implications of the resulting energy price dispersion for two reasons: Conceptually, studies that analyse shifts in industrial production and their implications for emissions and energy use (like Barrows & Ollivier, 2018) usually do not account for heterogeneous energy prices between firms. That means that they cannot capture the effect of policies that affect the energy prices of some firms more than others, or subsidise the energy use of a particular group of firms. Empirically, many of the approaches used to analyse the behaviour of firms in the face of changes to input factor supply or production technology either abstract from product heterogeneity (e.g. studies like Gerster & Lamp, 2024, or Marin & Vona, 2021, that focus on firm-level effects of energy prices only, or Barrows et al., 2024, who analyse the effect of heterogeneous energy prices but implicitly assume that products are homogeneous). Or they can only be applied to the subset of single-product firms (like most of the standard production function estimation literature, e.g. Akerberg et al., 2015).

Disregarding product heterogeneity is problematic because, around the world, a substantial share of actual energy and emissions intensity reductions in the manufacturing sector comes from an aggregate shift towards ‘cleaner’ products (e.g. Barrows & Ollivier, 2018, for India, Pan et al., 2022, for China, or Rottner & von Graevenitz, 2024, for Germany) and it has been shown empirically that firms may engage in product switching in response to energy prices and regulation (Abeberese, 2017). Concentrating on single-product firms is problematic because, even though the majority of firms are relatively small, single-product firms, the vast majority of energy use and emissions comes from large multi-product firms: In our case of Germany, the data reveal that the 36% of firms that produce multiple products account for around 80% of energy use and emissions.

For this reason, we analyse the relationship between energy prices and energy intensity in a way that incorporates both multi-product production and energy price heterogeneity, empirically as well as theoretically. Using German administrative firm data from 2005 to 2017, we show that the country’s manufacturing sector has undergone a composition shift towards less energy-intensive products but that this has not reduced energy intensity. Over the same period, firm energy prices have become substantially more dispersed, with low energy-price firms moving into markets for more energy-intensive products and high energy-price firms moving out of them.

We propose a model of multi-product firms in a world with energy price and product heterogeneity to offer a theoretical explanation for this phenomenon: a reshuffling among the producers of energy-intensive products. When energy prices rise for some firms more than for others – for example if a country introduces climate policies but compensates large energy-intensive firms –, firms with increasing energy prices move out of more energy-intensive products. This reduces competitive pressure in the markets for these products, allowing firms with a lower energy price to increase their market shares. Since firms with a lower energy price produce the same products in a more energy-intensive way, this ‘reshuffling’ among producers increases aggregate energy intensity.

This mechanism can explain how increasing energy price dispersion can lead to an increase in energy intensity that is decoupled from the energy intensity of the product mix. In line with the German data, it implies that energy-intensive products are both added and dropped more frequently than others, and that they are added in particular by firms with a low energy price. A simple preliminary quantification based on a log-normal firm energy price distribution confirms the role of reshuffling but suggests that other factors play a role as well. In such a setting, the demand-reducing increase in

average energy prices that accompanies the rising energy price dispersion dominates and aggregate energy intensity declines.

Our paper connects to several strands of the literature. First, we contribute to the extensive literature on multi-product production (Bernard et al., 2011; Eckel & Neary, 2010; Liu, 2010; Ma et al., 2014; Mayer et al., 2014; Nocke & Yeaple, 2014; Qiu & Zhou, 2013). This literature focuses mostly on the role of trade and investigates how export and import competition drive firms to adapt their product portfolio, typically using a one-factor, homogeneous-product Melitz-type workhorse model. To our knowledge, only two papers from this literature have considered the interaction between multi-product production and energy use or emissions: Barrows and Ollivier (2018) have included energy as a second factor of production into the workhorse model of product choice to explain the emissions-reducing effect of competitiveness shocks in India around the turn of the millennium. We expand on their model by introducing heterogeneous energy prices and heterogeneous products, which allows us to analyse the role of energy price dispersion and product portfolio shifts. Close to our work, Mayr-Dorn (2024) develops a model of multi-product firms to show that the emissions of such firms may increase rather than decrease when an emissions pricing scheme is introduced: If emissions pricing induces firms to focus on their core products, and if these core products are more emissions-intensive than their ‘fringe products’, the shift towards more emissions-intensive core products may outweigh the emissions-reducing effect of a carbon price. Mayr-Dorn studies symmetric firms in an oligopolistic setting, where everyone faces emissions pricing or no-one does. Our model is set up to gauge the consequences of a more realistic policy framework that provides heterogeneous incentives across firms even within industries. We also allow for firm heterogeneity in productivity which is an important feature of the manufacturing sector, and we assess our model predictions empirically using German administrative data.

Second, we speak to the large literature on industrial decarbonisation by focusing on product mix adjustments as a channel through which firms might react to climate policy. Many papers find that climate policies lead to reductions in firm-level carbon emissions, without being able to attribute these emission reductions to either output reductions or to increases in investments (Flues & Lutz, 2015; Gerster & Lamp, 2024; von Graevenitz & Rottner, 2024). Other studies identify output effects that are too small to explain observed emissions reductions in full (Aldy & Pizer, 2015; Wolvertson et al., 2022). Product mix adjustments might play a role in this, complementing technique effects like efficiency improvements or substitution of energy with other inputs. The literature has already established a link between climate policies and product innovation, i.e., R&D aimed at developing new (clean) products (Calel & Dechezleprêtre, 2016). Our paper provides a clear theoretical mechanism of how both the ‘technique’ effect and the ‘product composition’ effect of increasing energy prices come about at the firm level, and provides preliminary evidence that both effects matter for firm energy use and, hence, emissions.

We also connect to a strand of the empirical energy and climate literature that documents how firms change their product portfolios in response to energy prices and regulation (Abeberese, 2017; Elliott et al., 2019; Elrod & Malik, 2017; 2019; Zhang et al., 2023). We contribute to this strand of the literature by producing empirical evidence on product switching in response to energy prices in Germany. But more importantly, we complement this literature by providing a theoretical underpinning for the reduced-form evidence: We show in particular that it is not sufficient for reducing aggregate energy use or emissions if firms with a higher energy price move out of energy-intensive products. In equilibrium, the remaining firms with a lower energy price may ramp up their output of these products and produce them in a more energy-intensive way than the firms which left, potentially undoing the effect of the original energy price increase.

The remainder of this paper is structured as follows: In Section 2 we present our data as well as key descriptive evidence about product portfolios and energy and emissions intensity. In Section 3 we develop our theoretical model that can explain the stylized facts documented in Section 2. In Section 4,

we provide empirical evidence of the mechanisms posited by our model, before we proceed with a simple quantification in Section 5. Finally, we discuss the policy implications of our results in Section 6.

2 Motivating results

The focus of our empirical analysis is on the German manufacturing sector. As one of the world’s largest manufacturing producers and a leading industrial economy, the case of Germany is highly instructive for much of the industrialised world. Moreover, the German statistical offices provide highly granular output and energy use data for almost the entire manufacturing sector, which allow us to look in depth at firm-product-level shifts in production and energy use.

2.1 Data

Our empirical results are chiefly based on firm-product-level administrative microdata from Germany. The so-called *Amtliche Firmendaten für Deutschland* (AFiD) panel is provided by the German Federal Statistical Office and the Statistical Offices of the Federal States. It contains detailed information on each manufacturing firm with more than 20 employees in Germany, including output quantity and value by highly granular 9-digit product codes (see Table 1 for an excerpt from the product classification)¹, and energy consumption by source.

The AFiD panel contains around 30,000 firms each year. The data are available annually between 2003 and 2017. However, we focus on three focal years 2005, 2011 and 2017 and the changes occurring in the six-year intervals between them. This is in line with previous empirical studies on product choice (Bernard et al., 2010; Goldberg et al., 2010; Navarro, 2012) and allows us to concentrate on longer-term changes in firms’ product portfolios.² We define products in terms of 6-digit product codes according to the GP 2009 product classification. By this definition, we distinguish between around 1,400 distinct manufacturing products, which gives us a similar granularity as in most existing studies of multi-product production.³

Table 1: Excerpt from GP 2009 product classification

GP 2009 code	Description
10	Food and feed
– 101	Meat and meat products
– ...	...
– 103	Fruit and vegetable products
– 1031	Processed and preserved potatoes
— 1031 11	Frozen potatoes
— ...	...
— 1031 14	Prepared or preserved potatoes
— ...	...
— 1031 14 603	Potato crisps and sticks
— 1031 14 605	Potato salad, no mayonnaise-based dressing
— ...	...
– 1032	Fruit and vegetable juices, unfermented and without added spirits

¹Codes are based on the GP 2009 product classification, which is a German derivative of the EU-wide PRODCOM classification.

²Since firms’ product portfolios do not vary much in the short run, year-on-year comparisons are not overly insightful.

³The 5-digit SIC codes that Bernard et al. (2010) use, for example, distinguish between 1,500 different products.

Unless noted otherwise, we operationalise output as sales value of production.⁴ Note that we do not use data on output quantities in our analyses. This is because the aggregation of quantities is conceptually problematic if products are measured in different units or of different quality.⁵

We also observe total energy expenditures for a rolling subsample of firms that are part of the cost structure survey (*Kostenstrukturhebung*, KSE). This allows us to compute the firm-level average energy price for these firms by dividing their energy expenditure by their energy consumption. The expenditure data are available for roughly 45% of firms each year. Every four years, a new stratified sample is drawn. The sample is chosen to represent firms of all sectors and all employment sizes. Only firms with more than 500 employees are always surveyed.

Mapping energy use to products

To analyse product composition shifts towards more or less energy-intensive products, we develop a novel measure of product energy intensities. Obtaining such product-level intensities is challenging in standard production data as ours because we observe energy inputs only at the firm level. For multi-product firms, we have no means of knowing how much of the energy they consume flows into which product.

To circumvent this problem, large parts of the existing literature have restricted their attention to single product firms. By contrast, we use a straightforward statistical approach to estimating product-level energy intensities: We regress firm-level energy use on firm-product level output and use the resulting product-specific coefficients as proxies for product energy intensity. Specifically, we consider the linear projection

$$m_i = \sum_{j=1}^{N_J} r_{ij} \eta_j + \varepsilon_i \quad (1)$$

where m_i is firm i 's energy use, r_{ij} is firm i 's output of product j and N_J is the number of distinct products in the economy. Each η_j then represents the average increase in energy use associated with an additional EUR of output of product j – that is, the average energy intensity of the product.

We use data over our entire sample range from 2003 to 2017 and pool observations in three-year bins to reduce small-sample error and flatten out period-specific shocks, and then estimate Equation 1 to get the vector of product-specific energy intensity estimates $(\hat{\eta}_j)_1^{N_J}$.⁶ This procedure yields reasonable product-level energy intensity estimates $\hat{\eta}_j$ for 82% of all products across our sample years. For some products that are only produced by a small number of firms, the resulting $\hat{\eta}_j$ are negative, in which case we set them to missing. The resulting product energy intensity estimates have a high predictive power for firm-level energy use and match aggregate energy intensity well (see Appendix B for more detail). The key advantage of this regression-based approach is that it allows us to leverage information from the full sample of firms to construct our measure of product energy intensity. This sets our study apart from the vast majority of previous empirical studies of energy use (as well as input use in general) and firm performance, which typically restrict their analyses to single-product firms to avoid the host

⁴Small differences between the sales value of production and revenue arise, for example, due to firms building up stock or selling out pre-existing inventory in a year, or due to systematic price discounts the firm offers on the market.

⁵That is true in particular given our aggregation to 6-digit product codes. But even in the case of a more granular definition we can think of a Ferrari and a Volkswagen Polo that would still constitute one unit of the product 'car'.

⁶Solving such a regression is computationally challenging because our granular product definition means that it includes around 1,400 regressors. To obtain a computationally feasible estimator, we use the fact that few firms produce more than a handful of different products. This means that the feature matrix $R = [r_1', \dots, r_{N'}']'$ is high-dimensional but sparse, since most of its entries are zero. We therefore implement an OLS computation procedure based on the R language's implementation of the QR decomposition of sparse matrices that can solve this regression problem in reasonable time.

of issues that come with attributing firm-level inputs to firm-product-level outputs (e.g. Ganapati et al., 2020).

If we were to follow this approach, we could simply sum up the energy consumption of all single-product firms producing a given product, and divide it by the total output of this product from single-product firms to obtain product-level energy intensities. While this may be a helpful simplification in some contexts, it is not well suited if one wants to explain changes in aggregate production and energy consumption. First, it ignores the bulk of economic activity and energy use: Over our period of observation, the 36% of firms that had multiple products in their portfolio produced 68% of aggregate output and were responsible for 80% of industrial energy consumption. Second, it is likely to give us a substantially distorted picture of the actual energy intensity of production: In our data, the aggregate energy intensity of single-product firms is 59% lower than the aggregate energy intensity of multi-product firms and 32% lower than total aggregate energy intensity.⁷ Indeed, we see in Appendix B that our regression-based product energy intensities explain actual energy use far better than single-product-based product energy intensities. At the same time, our regression-based approach asymptotically nests the single-product-based approach when applied to the subset of single-product firms.⁸ What is more, roughly 36% of aggregate revenue comes from products that are only produced by three or fewer single-product firms, making it virtually impossible to estimate a reliable single-product-based energy intensity for them.

Compared to structural approaches to input attribution like the one pioneered by De Loecker et al. (2016), our method of obtaining product energy intensities has the advantage that it only relies on a minimum of assumptions. This does come at a cost: Unlike De Loecker et al. (2016), we cannot attribute inputs to outputs for individual firms, and our product energy intensity measures $\hat{\eta}_j$ do not reflect deep technological parameters but only the empirical average amount of energy used per unit of product j . But we do get robust proxies for product energy intensity that do not rely on strong structural assumptions like the functional form of the production function and the timing of input choices.⁹

2.2 Product choice, energy use and emissions

We use the data described above to examine how aggregate and firm-level energy prices, energy use, and product composition relate to one another. Specifically, we find that energy prices have both increased and become more dispersed between 2005 and 2017. Curiously, this has coincided with a widening wedge between the energy intensity of the product mix (i.e., energy intensity if all products were still produced at their 2005 average energy intensity) and actual energy intensity.


Divergence between composition-induced and actual energy intensity

We observe that energy prices paid by German manufacturing firms increased substantially over our period of observation, but this increase was highly heterogeneous. As we can see in panel (a) of Figure 1, the distribution of energy prices in 2011 and 2017 shifted to the right compared to 2005. However, the dispersion of firm-level energy prices increased substantially over the same period. Panel (b) suggests that it was in particular firms with low energy consumption that saw their prices increase, with large consumers maintaining their much lower price levels: The unweighted average of firm-level

⁷ Across all years of observation, we observe an average energy intensity of 0.73 EUR/kWh among multi-product firms but only 0.42 EUR/kWh among single-product firms. Due to the large share of multi-product firms in overall energy consumption, total energy intensity across all firms is much closer to the energy intensity of multi-product firms at 0.63 EUR/kWh.

⁸ See Appendix B for a formal proof.

⁹ Even De Loecker et al. (2016) need a sufficiently large number of single-product firms that produce a product to be able to include it in their analysis, because they estimate product-level production functions off of single-product firms in a first step.



(a) Energy price dispersion 2005, 2011 and 2017

(b) Average energy prices, 2003-2017

Figure 1: Development of energy prices between 2005 and 2017

Note: Energy prices are based on the rolling subsample of around 11,000 manufacturing firms each year for which total energy expenditure is available.

Source: AFiD manufacturing census (see Appendix A)

energy prices increased substantially between 2003 and 2017.¹⁰ The consumption-weighted average price, which puts much higher weight on large consumers, increased somewhat up to 2008 but fell back to its original level over the subsequent years.

Both the extent and the direction of energy price dispersion we observe are in line with what we know about the institutional and regulatory landscape in our period of observation: The energy cost of a firm depends on its fuel mix, its choice of retailer, the quantities consumed, its location (network charges), and the taxes and surcharges it is subject to. Large users systematically face lower prices after taxes and surcharges. There is evidence of Ramsey pricing, e.g. for electricity (see also Davis et al., 2013 for the US case), which leads to lower prices for users with more elastic demand. Large users also tend to be connected to the grid at higher levels which lowers network charges for gas and electricity. Moreover, large users benefit from exemption schemes and relief measures designed to level the playing field in international markets. In Germany there are several such schemes where consumption levels play a role for eligibility: Reduced network charges, partial exemptions from the renewable energy surcharge, and electricity price compensation for ETS induced electricity price increases are the largest of these. These schemes all tend to favor the same type of industrial firm predominantly in energy intensive and trade-exposed sectors.¹¹

Secondly, we can observe that – in spite of rising energy prices – aggregate energy intensity in German manufacturing has *increased* over our period of observation (Figure 2). Over the same period of time, however, the sector's product mix became less energy-intensive: On aggregate, firms started producing relatively more of products with a low baseline energy intensity, and relatively less of products with a high baseline energy intensity.

This echoes previous results that Rottner and von Graevenitz (2024) find for German manufacturing emissions using the workhorse decomposition method developed by Levinson (2009, 2015) but still represents an economic puzzle. After all, profit-maximising firms should substitute away from an input

¹⁰The average price increased from around 12 ct/kWh in 2003 to almost 16 ct/kWh in 2014, before stabilising at 14 ct/kWh from 2016 on. The slight drop in prices after 2011 is most likely attributable to the substantial fall in industrial natural gas prices over the same period, which dropped from 3.61 ct/kWh in 2011 to 2.55 ct/kWh in 2017 (Destatis, 2023).

¹¹Past research has used these sources of variation to estimate causal effects of electricity prices on plant performance, e.g. Gerster and Lamp (2024); von Graevenitz and Rottner (2024). See these papers for more details on the pricing of electricity in Germany.




Figure 2: Actual energy intensity and energy intensity of the product mix in German manufacturing, 2003-2017

Note: Actual energy intensity is calculated as total energy consumption divided by total output in each year. The energy intensity of the product mix is calculated by taking product energy intensities from 2005 and weighting them with current product shares in total output each year.

Source: AFiD manufacturing census (see Appendix A)

when its price increases – or at least not increase their use of this input. The fact that this fundamental prediction does not hold on aggregate suggests that there are counteracting forces at play.

What is even more puzzling is that the energy intensity of the product mix develops in the opposite direction of the actual energy intensity. This is in line with the notion that an increase in the price of an input should lead firms to shift away from products that require more of this input (which follows, for example, from the product choice model that Boehm et al., 2022, develop). But it suggests that manufacturing firms shifted towards less energy-intensive products while at the same time producing them in a more energy-intensive way.

3 Theoretical model

To explain our empirical findings from Section 2, we propose a model of product portfolio choice with labour and energy as input factors. Firms have a core competence in one product but may opt to produce additional products. Such products are more expensive for a firm to produce the more they differ from its core product. Firms differ in their productivity and face different energy prices, while products differ in their energy intensity. A firm’s decision to expand into an additional product market is based on the interaction of firm- and product-specific features: Firms that have a high productivity in producing a specific product will be much more likely to enter the market for that product. And firms with high energy prices will struggle to profitably produce any good, especially those that are energy-intensive.

Our model is similar in spirit to the canonical multi-product models from the trade literature (Bernard et al., 2011; Eckel & Neary, 2010; Mayer et al., 2014) and in particular to the model developed by Barrows and Ollivier (2018). It expands on these models in two main ways: by allowing for firm-specific energy prices; and by allowing for heterogeneity in energy intensity between products.

To our knowledge, Barrows and Ollivier (2018) were the first to include energy as a second factor of production into a Melitz-type multi-product model. We build on their approach but allow energy prices to differ between firms. This has two advantages: First, it allows us to reflect the economic consequences of heterogeneous energy prices while nesting their simpler model and incorporates the key empirical finding that energy prices are heterogeneous between firms. Second, it allows us to make predictions about the effect of policies that change the energy prices of some firms more than others.

This is empirically relevant since it mirrors exactly how many industrial, climate and energy policies work in reality.

Unlike almost all previous multi-product models, we explicitly model heterogeneity between products. This allows us to derive *product-side* explanations for firms' product choice alongside *firm-side* explanations. Most importantly, we allow for products to have different baseline energy intensities. This is highly intuitive: Producing 1 000 EUR worth of steel will require much more energy than printing 1 000 EUR worth of books. Previous models of product choice – including Barrows and Ollivier (2018) – abstract from product heterogeneity, but it is crucial to explain the reshuffling patterns we observe in the data.

We discuss our fundamental model setup in Section 3.1. Building on this, we illustrate the role that reshuffling of producers plays for energy intensity based on a stylised energy price distribution in Section 3.2. After providing some empirical evidence of the underlying model mechanisms in Section 4, we proceed with a simple quantification based on a more realistic energy price distribution among German manufacturing firms in Section 5.

3.1 Model setup

Consumers love variety

We keep the demand side simple by assuming a representative consumer who demands different products in fixed proportions and exhibits a love for different varieties of each product. This gives producers some market power in the niche they occupy with their variety, which leads to them acting as monopolistic competitors on the market for each product.

Fundamentally, we assume that there is a continuum $I = [0, 1]$ of producers and a finite set of products J . Each producer i may or may not decide to produce her own variety of j . This means that, ultimately, a subset $I_j \subset [0, 1]$ of firm-specific varieties of product j will be produced in the economy.

The representative consumer derives utility

$$U(\mathbf{q}) := q_0 + \sum_{j \in J} \gamma \left(\underbrace{\left[\int_{i \in I} q_{ij}^{\frac{\sigma-1}{\sigma}} di \right]^{\frac{\sigma}{\sigma-1}}}_{=: U_j} \right)^{1-\delta} \quad (2)$$

from bundle $\mathbf{q} = (q_{ij})_{j \in J, i \in I}$. The inner term – which we denote by U_j – represents the CES utility the consumer derives from the aggregate of all varieties they consume of product j . σ is the elasticity of substitution between varieties of a product which we assume to be greater than one as is standard in the literature.

The outer term corresponds to a quasi-linear utility function over the utilities derived from each additional good: The consumer can derive linear utility from an outside option q_0 which can be thought of as representing saving and consumption of non-manufacturing goods and services. Product-level utilities U_j enter the consumer's aggregate utility function as additively separable concave terms. Intuitively, this means that the consumer trades off each individual manufacturing good against the outside option q_0 . If the aggregate price of one good j increases, the consumer will consume less of that good but will not reduce her consumption of other goods in order to be able to afford more of good j . γ represents the taste of the consumer for a good and $\delta \in (0, 1)$ governs how quickly the consumer's marginal utility from each manufacturing product decreases. For ease of exposition, we have imposed that γ , δ and σ are identical across products, but this is not essential to our conclusions.

The representative consumer solves the utility maximisation problem

$$\begin{aligned} \max_{\mathbf{q}} U(\mathbf{q}) \quad \text{s.t.} \quad & Y \geq q_0 + \sum_{j \in J} \int_{i \in I} p_{ij} q_{ij} di, \\ & q_0 \geq 0, \quad q_{ij} \geq 0 \quad \forall i, j \end{aligned} \quad (3)$$

where p_{ij} is the price of firm i 's variety of product j (relative to the outside option). On top of this, we assume that, given the set of prices $P = (p_{ij})_{i \in I, j \in J}$, consumer income Y is sufficiently large that the consumer does not spend all their income on manufacturing goods. This corresponds to assuming that

$$Y > \sum_j [(1 - \delta)\gamma]^{\frac{1}{\delta}} P_j^{1 - \frac{1}{\delta}} \quad (4)$$

where $P_j = \left[\int_{i \in I} p_{ij}^{1 - \sigma} \right]^{\frac{1}{1 - \sigma}}$ is the CES price aggregator.¹²

Under these assumptions, consumer demand for variety i of product j is

$$q_{ij}^* = [(1 - \delta)\gamma]^{\frac{1}{\delta}} P_j^{\sigma - \frac{1}{\delta}} p_{ij}^{-\sigma}. \quad (5)$$

As under standard CES utility, the demand for product variety ij decreases in its own price with a constant elasticity of σ . At the same time, the demand for product variety ij increases in the aggregate price index of product j with an elasticity of $\sigma - \frac{1}{\delta} > 0$. This elasticity results from the interplay of two different mechanisms: substitution between varieties, and substitution between the product and the numéraire good, i.e. saving or consumption outside the manufacturing sector.

If the aggregate price index P_j is high, the demand for product variety ij is higher *ceteris paribus* because it is more expensive for consumers to substitute variety i with other varieties (hence the σ in the elasticity term). At the same time, if P_j is high, consumers will buy less of product j because marginal utility is decreasing in aggregate consumption of the product. This effect is less pronounced if δ is close to 1, because then marginal utility from product j is decreasing so quickly that even large changes in the product price index only translate into moderate changes in demand. If δ is close to 0, on the other hand, the effect is potentially very large because the consumer is more indifferent between each product and the numéraire and small increases in the price index of product j may drive her away from the product (hence the $-\frac{1}{\delta}$ in the elasticity term).

We assume for the remainder of the paper that $\sigma > \frac{1}{\delta}$. That is, we assume that the substitution effect between varieties of a product dominates the substitution effect between the product and the numéraire, so that an increase in the aggregate price of product j will always increase the demand for product variety ij . This is intuitively plausible for most manufacturing goods: If all producers bar one increased their prices, we would expect demand to flock towards the producer who kept her prices constant instead of consumers just turning their back to the product as a whole.

Firms charge constant markup over unit costs

We assume that firms maximise profits given their production technology and consumer demand. Specifically, we assume that each firm i can produce its variety of product j from labour and energy using CES production technology

$$q_{ij}(l_{ij}, m_{ij}) = \varphi_{ij} [\alpha_j^l l_{ij}^\rho + \alpha_j^m m_{ij}^\rho]^{\frac{1}{\rho}}. \quad (6)$$

Here, φ_{ij} is a firm–product specific total factor productivity term which we will discuss in more detail below. l_{ij} and m_{ij} are the amounts of labour and energy that the firm spends on producing its

¹²We show below that such a level of income Y must exist for any equilibrium price set P^* , although it cannot be determined algebraically.

variety of the product. α_j^l and α_j^m are product-specific baseline labour and energy intensities which we normalise to $\alpha_j^l + \alpha_j^m = 1$ for each product. (Steel, for example, would have a very high α_j^m , whereas bookprinting would have a much lower α_j^m and, correspondingly, a higher α_j^l .) ρ determines the degree to which labour and energy can be substituted in the production process: The limiting cases of $\rho = 1$ and $\rho \rightarrow -\infty$ indicate perfect substitutability and perfect complementarity.¹³ We assume that $\rho < 0$, i.e. some degree of complementarity between labour and energy.

We assume that each firm has to pay the same wage w per effective unit of labour but draws a specific energy price τ_i . This results in the following cost minimisation problem

$$\min_{l,m} wl + \tau_i m \quad \text{s.t.} \quad q_{ij}(l_{ij}, m_{ij}) \geq q_{ij}. \quad (7)$$

Its solution tells us that firms produce at a firm-product specific constant unit cost

$$c_{ij} = \frac{1}{\varphi_{ij}} \left[(1 - \alpha_j^m)^{-\frac{1}{\rho-1}} w^{\frac{\rho}{\rho-1}} + (\alpha_j^m)^{-\frac{1}{\rho-1}} \tau_i^{\frac{\rho}{\rho-1}} \right]^{\frac{\rho-1}{\rho}} = \frac{1}{\varphi_{ij}} \tilde{c}_j(\tau_i) \quad (8)$$

where $\tilde{c}_j(\tau_i)$ represents the cost net of productivity of a firm with energy price τ_i that wants to produce one unit of its variety of product j .

The unit cost of production depends on the interaction of firm- and product-specific factors. It decreases in firm-product-specific productivity φ_{ij} and increases in wage w and energy price τ_i . The cost impact of a higher energy price is more pronounced if the product has a relatively higher baseline energy intensity α_j^m whereas the cost impact of a higher wage is more pronounced if the product has a relatively higher baseline labour intensity. High energy prices matter more if ρ is low, because then the firm cannot compensate for its higher energy prices by substituting energy with labour.

Conditional on entering the market for product j , firm i maximises its profits by setting

$$p_{ij}^* = \arg \max_p (p - c_{ij}) [(1 - \delta)\gamma^{1-\delta}]^{\frac{1}{\delta}} P_j^{\sigma-\frac{1}{\delta}} p_{ij}^{-\sigma}. \quad (9)$$

Since there is a continuum of firms, each individual firm takes the aggregate price index P_j as given and sets the price for its variety of j to maximise its profits in monopolistic competition with all other firms. Its optimal conditional profits are therefore given by

$$\pi_{ij}^*(P_j) = AP_j^{\sigma-\frac{1}{\delta}} c_{ij}^{1-\sigma} \quad (10)$$

with $A = \frac{(\sigma-1)^{\sigma-1}}{\sigma^\sigma} [(1 - \delta)\gamma]^{\frac{1}{\delta}}$ being a profit multiplier that only depends on model parameters.

Firms draw individual productivities for each product

We assume that each firm draws its productivity for each product independently from a Pareto distribution. This generates a hierarchy of products for each firm. In analogy to the typical model setup in the multi-product firm literature (e.g. Mayer et al., 2014, Nocke & Yeaple, 2014, Barrows & Ollivier, 2018), we can conceive of the product for which a firm is most productive as the firm's core product, and products for which the firm is less productive as products that are increasingly far away from its core competency.¹⁴

¹³Specifically, the elasticity of substitution between labour and energy is $\frac{1}{1-\rho}$.

¹⁴Standard models from the multi-product literature typically assume that each firm's productivity decreases deterministically with a product's distance from the firm's core competency. This is a helpful assumption when one considers homogeneous products but quickly becomes untractable when analysing heterogeneous products. In a heterogeneous-product model, interactions between a product's properties and its distance to other products immediately generate high-dimensional interactions that are hard to track analytically.

Specifically, each firm i draws its productivity for product j from a Pareto distribution with cdf G with dispersion parameter k and minimum value 1. We make the standard assumption that $k > \sigma - 1$, that is, we put an upper bound on productivity dispersion (a value of $k \rightarrow \infty$ would imply that all firms have a productivity of 1, whereas a value of $k \rightarrow 0$ would imply that each level of productivity between 0 and ∞ is equally likely). As we show in Appendix C.1, this assumption is necessary to ensure that we have a positive equilibrium price index for each product. Intuitively, the underlying mechanism works as follows: If productivity is extremely dispersed (i.e. we have a verly low k), there will be a non-negligible share of firms with extremely high productivities which can produce the product almost for free, drawing demand away from any firm which produces a variety of the product at a positive price and preventing the formation of an equilibrium in which the product is produced at all.

Firms draw energy prices

We assume that firms draw their energy price τ_i from a distribution H with strictly positive support and $\mathbb{E}[\tau_i^{1-\sigma}] < \infty$. We assume a stylised distribution of τ_i in Section 3.2 to derive theoretical predictions and quantitatively investigate the effects of a more realistic energy price distribution in Section 5.

We assume that $\tau_i \perp \varphi_{ij}$, i.e. that a firm's productivity does not affect its energy price *ex ante*. Importantly, this does not rule out a relationship between energy price and productivity *conditional on entry*. Indeed, it is a key prediction of our model that firms need to be either very productive or have a low energy price in order to succeed in the market, which introduces a positive relationship between energy price and productivity conditional on entering the market.

If we consider our model to represent a snapshot of the German manufacturing sector at a given point in time, there are valid reasons to suspect both a positive and a negative *ex ante* relationship between energy prices and productivity. On the one hand, energy prices may be higher for less productive firms because they are generally smaller and have less market power on factor markets, and generally buy smaller quantities of energy. What is more, strategic sourcing of affordable energy may correlate with general management quality and hence with overall productivity. This would rationalise a negative relationship between φ_{ij} and τ_i . On the other hand, energy prices may be lower for less productive firms due to dynamic selection effects that our static model cannot capture. Firstly, firms with a low energy price may have been under less pressure to engage in productivity-enhancing investments (Colmer et al., 2024; Hawkins-Pierot & Wagner, 2024). Secondly, existing firms with low productivity may have been more likely to be targeted by subsidies in order to keep them alive. Keeping firms from shutting down has in fact been an important rationale of German industrial policy in the face of economic pressure for the past decades, as witnessed by a steep rise in the share of 'zombie firms' over our period of observation (although the share of zombie firms remains small in absolute terms; Albuquerque & Iyer, 2023). And thirdly, firms with a high productivity may simply have been under less pressure to source cheap energy than firms that are scraping to get by. By assuming that productivity φ_{ij} and energy price τ_i are independent, we take a middle ground between these different potential states of the world, which has the advantage that we can attribute any resulting relationship between energy price and productivity purely to our model mechanisms.

Through the lens of the allocative efficiency literature (e.g. Ruzic & Ho, 2023 or Hsieh & Klenow, 2009), our assumption of independence between τ_i and φ_{ij} implies that different energy prices are *ex ante* allocatively neutral: Lower energy prices increase firms' share in overall production (and input use), but these firms are not systematically more or less productive than others. As we discuss in more detail in Section 6, a negative correlation between τ_i and φ_{ij} would mean that energy price dispersion helps combat misallocation because more productive firms would have access to cheaper energy and be able to increase their share in production, and vice versa.

Product market entry and equilibrium

We assume that entering a product market requires an irrecoverable upfront investment of F . Firm i considers each product market j separately and decides to enter it if its profits from that product market, π_{ij}^* , exceed entry costs F .¹⁵ This defines a cost threshold

$$\hat{c}_j(P_j) = \left[\frac{A}{F} P_j^{\sigma - \frac{1}{\delta}} \right]^{\frac{1}{\sigma - 1}}. \quad (11)$$

All firms with cost $c_{ij} < \hat{c}_j$ enter the market for product j , while all firms with costs above $\hat{c}_j$ stay out of the market and supply $q_{ij} = 0$.

This cost threshold depends on the CES price index for product j , P_j : If the general price level for product j is high, firms can profitably enter the product market even if they have relatively high production costs. P_j , however, depends itself on the cost cutoff $\hat{c}_j$ because it reflects aggregate product variety prices of those firms that end up producing product j . Defining

$$I_j(P_j) = \{i \in I \mid c_{ij} < \hat{c}_j(P_j)\} \quad (12)$$

as the set of firms that would enter the market for product j given price index P_j , the equilibrium condition for the market for product j is given by

$$P_j^* = \left[\int_{i \in I_j(P_j^*)} p_{ij}^{1-\sigma} di \right]^{\frac{1}{1-\sigma}}. \quad (13)$$

As we show in Appendix C.1, this pins down a unique equilibrium with $P_j^* < \infty$. Intuitively, the main mechanism behind this equilibrium is the following: If the CES price index for product j increases (i.e. if the left-hand side of Equation 13 increases), more firms will enter the market for the product. However, the more firms enter the market, the cheaper it gets for the consumer to obtain one unit of utility from product j because she can spread her consumption across more varieties (which puts more competitive pressure on each individual competitive monopolist in the market). Since P_j reflects the cost of obtaining one unit of utility from product j , this decreases the CES price index (i.e. the right-hand side of Equation 13).

Given our assumption that firm-product-specific productivities are Pareto($k, 1$)-distributed, the equilibrium price index is given by

$$P_j^* = \left[\frac{k - (\sigma - 1)}{k} \left(\frac{\sigma}{\sigma - 1} \right)^{\sigma - 1} \left(\frac{F}{A} \right)^{\frac{k}{\sigma - 1} - 1} \left(\int_0^\infty \tilde{c}_j(\tau)^{-k} dH(\tau) \right)^{-1} \right]^{\frac{1}{\xi k + (1 - \xi)(\sigma - 1)}} \quad (14)$$

where we have defined

$$\xi := \frac{\sigma - \frac{1}{\delta}}{\sigma - 1} \in (0, 1), \quad (15)$$

for notational convenience.¹⁶ Here, ξ represents the strength of the demand response to an increase in the CES price index P_j in our world with decreasing marginal utility from each composite product

¹⁵Strictly speaking, we also assume that firms know the distribution of firm-product productivities φ_{ij} and energy prices τ_i and believe that all other firms are maximising profits as well, so they can anticipate the resulting equilibrium price index P_j . Without this additional assumption, firms would not be able to determine their potential profits π_{ij}^* before entering.

¹⁶As we show in Appendix C.1, the integral $\int_0^\infty \tilde{c}_j(\tau)^{-k} dH(\tau)$ must be finite by our condition that $\int_0^\infty \tau^{1-\sigma} dH(\tau) < \infty$. Also note that this closed-form expression for P_j^* allows us to specify our assumption from Equation 4: If we assume

j relative to the strength of the demand response in a counterfactual world with constant marginal utility from each composite product.

The explicit expression for P_j^* tells us that the equilibrium CES price index of product j is increasing in net-of-productivity costs $\tilde{c}_j(\tau)$: If all firms' costs uniformly increase by 1%, the equilibrium price index increases by $\frac{k}{\xi k + (1-\xi)(\sigma-1)} > 1\%$. This is because, in response to a 1% cost shock, all firms increase the prices of their varieties p_{ij} by 1%. Ceteris paribus, this increases the CES price index of product j by 1% as well. However, since the consumer's marginal utility from product j is decreasing, she will only increase her spending on the product by $\xi\%$, where $\xi < 1$. That means that each firm's profits actually decrease by $(1 - \xi)\%$. This drives some firms that were just breaking even before out of the market for product j . This, in turn, reduces competitive pressure, which increases the CES price index beyond the mechanical 1% increase.

Energy intensity depends on interaction of product features and energy prices

To assess the role of energy price changes and product choice for energy intensity, let us define the energy intensity of product variety ij in revenue terms as

$$\eta_{ij} := \frac{m_{ij}}{r_{ij}} = \frac{\sigma - 1}{\sigma} \frac{s_{ij}}{\tau_i} \quad (16)$$

where

$$s_{ij} := \frac{(\alpha_j^m)^{-\frac{1}{\rho-1}} \tau_i^{\frac{\rho}{\rho-1}}}{(1 - \alpha_j^m)^{-\frac{1}{\rho-1}} w^{\frac{\rho}{\rho-1}} + (\alpha_j^m)^{-\frac{1}{\rho-1}} \tau_i^{\frac{\rho}{\rho-1}}} \quad (17)$$

represents the share of unit energy costs in total unit costs.¹⁷

Lemma 1: Revenue-based energy intensity η_{ij} is decreasing in firm energy price τ_i and increasing in baseline product energy intensity α_j^m .

Proof: Simple algebra reveals that

$$\frac{ds_{ij}}{d\alpha_j^m} = \frac{-1}{\rho} \frac{s_{ij}(1 - s_{ij})}{\alpha_j^m(1 - \alpha_j^m)} > 0 \quad (18)$$

and

$$\frac{ds_{ij}}{d\tau_i} = \frac{\rho}{\rho - 1} \frac{1}{\tau_i} s_{ij}(1 - s_{ij}) > 0. \quad (19)$$

To see that η_{ij} is decreasing in τ_i , consider the derivative

$$\frac{d\eta_{ij}}{d\tau_i} = \eta_{ij} \frac{1}{\tau_i} \cdot \underbrace{\left[\frac{\rho}{\rho - 1} (1 - s_{ij}) - 1 \right]}_{\Delta}. \quad (20)$$

Clearly, $\frac{\rho}{\rho-1} < 1$ by $\rho < 0$ and $s_{ij} < 1$ by our considerations above. But this implies that $\frac{\rho}{\rho-1}(1 - s_{ij}) < 1$ and hence $\Delta < 0$, so $\frac{d\eta_{ij}}{d\tau_i} < 0$. $\square$

that $Y > \sum_j [(1 - \delta)\gamma^{1-\delta}]^{\frac{1}{\delta}} (P_j^*)^{1-\frac{1}{\delta}}$, the consumer's income must be sufficiently large for her not to spend all her money on manufacturing goods under equilibrium prices.

¹⁷Specifically, $s_{ij}^{\frac{\rho}{\rho-1}}$ is the ratio between the unit cost of variety ij in a counterfactual world where energy is the only costly input relative to the actual unit cost.

Lemma 1 tells us that η_{ij} is decreasing in the firm's energy price τ_i and increasing in the product's baseline energy intensity α_j^m , which is intuitively straightforward. Since firms can at least partially substitute energy with labour in the production process (because $\rho > -\infty$), higher energy prices will translate into lower energy intensity. We can think of this mechanism as representing both behavioural changes and investments into energy-saving technologies in our static model. At the same time, products with a higher baseline energy intensity will require more energy even after accounting for potential input substitution patterns. This represents fundamental properties of the production process: For example, to produce a ton of steel, a significant amount of energy is required to extract iron from iron ore and then oxidate it, and a firm can only get so far by investing into more energy-efficient machinery or implementing energy saving measures.

Our empirical focus is on revenue-based energy intensity because this is the outcome we can observe in the data.¹⁸ However, to get a clearer picture of the underlying mechanisms, it is instructive to look at quantity-based 'physical' energy intensity

$$\tilde{\eta}_{ij} := \frac{m_{ij}}{q_{ij}} = \frac{1}{\varphi_{ij}} \left(\frac{\alpha_j^m}{s_{ij}} \right)^{-\frac{1}{\rho}} \quad (21)$$

as well in the theoretical part. The most striking difference compared to revenue-based energy intensity η_{ij} is that $\tilde{\eta}_{ij}$ depends positively on $\left(\frac{\alpha_j^m}{s_{ij}} \right)^{-\frac{1}{\rho}}$ rather than on $\frac{s_{ij}}{\tau_i}$. This term is still decreasing in the energy price τ_i . But unlike our revenue-based energy intensity measure, it does not include the additional intensity-reducing effect that comes from the fact that an increase in energy costs will translate into higher overall costs which increase revenue because they are passed through to the consumer.

The second major difference between physical energy intensity $\tilde{\eta}_{ij}$ and revenue-based energy intensity q_{ij} is that physical energy intensity negatively depends on total factor productivity φ_{ij} . This reflects the fact that firms with a high TFP require less of each input to produce the same physical output. However, since firms in our model charge a constant markup on each product, these efficiency gains are fully passed through to the consumer. This, of course, is a simplifying assumption. In reality, firms will only partially pass through efficiency gains to consumers on most markets, and hence revenue-based energy intensity will to some extent depend on productivity as well (Ganapati et al., 2020). But since this channel is not essential to the reshuffling phenomenon we attempt to explain and substantially increases the mathematical complexity of our model, we follow a significant strand of the product choice literature (e.g. Bernard et al., 2011; Nocke & Yeaple, 2014) in assuming mathematically tractable CES utility which results in iso-elastic demand and constant multiplicative mark-ups.

Finally, we also define aggregates of the energy intensity measures.

$$\eta_j = \frac{M_j}{R_j} = \frac{\int_{i \in I_j} \eta_{ij} r_{ij} di}{\int_{i \in I_j} r_{ij} di} \quad (22)$$

with $M_j = \int_{i \in I_j} m_{ij} di$ and $R_j = \int_{i \in I_j} r_{ij} di$ is the aggregate energy intensity of product j , and

$$\eta = \frac{M}{R} = \frac{\sum_{j \in J} M_j}{\sum_{j \in J} R_j} = \sum_{j \in J} \eta_j \frac{R_j}{R} \quad (23)$$

¹⁸While quantity-based energy intensity measures are in principle available in our data, they are difficult to interpret since products can be quite heterogeneous even within six-digit categories.

is the aggregate energy intensity of the entire manufacturing sector, with $R = \sum_{j \in J} R_j$ and $M = \sum_{j \in J} M_j$.

3.2 Illustration under stylised energy price distribution

To illustrate how our hypothesised reshuffling mechanism can increase the energy intensity of production even while the product mix becomes cleaner, we run through our model under a stylised energy price distribution in this section. This serves to highlight the key mechanisms at play. In Section 5, we will quantitatively assess how our model behaves under a more realistic distribution of energy prices, for which theoretical statements are much harder to derive.

To illustrate how the ‘reshuffling’ mechanism works in principle, consider an economy as we have described it in Section 3.1 in which the energy price distribution H is binary: Firms pay low energy price $\underline{\tau}$ with probability p and high energy price $\bar{\tau} = \nu \underline{\tau}$, $\nu > 1$, with probability $1 - p$. Trivially, this distribution of τ_i satisfies the regularity condition that $\mathbb{E}[\tau_i^{1-\sigma}] < \infty$.

As derived in Appendix C.1, we can write the average revenue from product j among the group of firms with energy price τ as

$$R_j(P_j, \tau) := \frac{k}{k - (\sigma - 1)} A^k F^{-(\frac{k}{\sigma-1}-1)} \tilde{c}_j(\tau)^{-k} P_j^{\xi k}. \quad (24)$$

Similarly, since by Equation 16, η_{ij} only depends on product parameters and the firm energy price τ_i , each firm with the same energy price τ will produce product j with the same energy intensity

$$\eta_j(\tau) = \frac{\sigma - 1}{\sigma} \frac{s_j(\tau)}{\tau}, \quad (25)$$

with

$$s_j(\tau) := \frac{(\alpha_j^m)^{-\frac{1}{\rho-1}} \tau^{\frac{\rho}{\rho-1}}}{(1 - \alpha_j^m)^{-\frac{1}{\rho-1}} w^{\frac{\rho}{\rho-1}} + (\alpha_j^m)^{-\frac{1}{\rho-1}} \tau^{\frac{\rho}{\rho-1}}} \quad (26)$$

being the share of energy costs in total costs for any firm with energy price τ when producing product j . Of course, this does not affect the applicability of Lemma 1, and in particular the fact that $\eta_j(\tau)$ is decreasing in τ .

Under our binary energy distribution, we can then define the aggregate energy intensity with which product j is produced in terms of the distributional parameters ν and $\underline{\tau}$:

$$\eta_j(\nu, \underline{\tau}) = \theta_j(\nu, \underline{\tau}) \eta_j(\underline{\tau}) + (1 - \theta_j(\nu, \underline{\tau})) \eta_j(\nu \underline{\tau}). \quad (27)$$

Here, $\theta_j(\nu, \underline{\tau}) = \frac{p R_j(P_j, \underline{\tau})}{p R_j(P_j, \underline{\tau}) + (1-p) R_j(P_j, \nu \underline{\tau})}$ is the share of total revenue from product j that is produced by the group of low energy-price firms. Note that P_j , on which the average revenues of both groups depend, cancels out of the expression for the revenue share θ_j by Equation 24. This is due to our assumption that productivity is Pareto-distributed, which implies that the reduction in the mass of active firms in the market due to a price increase is proportional to the original mass of active firms. If we now consider a marginal increase in the energy price of the high-price firms – that is, a marginal increase in ν –, it is straightforward to show that

$$\frac{d\eta_j}{d\nu} = \underbrace{(\eta_j(\underline{\tau}) - \eta_j(\nu \underline{\tau})) \frac{d\theta_j}{d\nu}}_{\text{reshuffling effect}} + \underbrace{(1 - \theta_j) \frac{d\eta_j(\nu \underline{\tau})}{d\nu}}_{\text{technique effect}} \underline{\tau}. \quad (28)$$

Intuitively, this means the following: If the energy price of the high-price firms increases, this induces them to produce in a more energy-efficient way. This is the technique effect and it reduces the energy intensity of product j by Lemma 1. However, an energy price increase for the high-price firms also has a composition effect on aggregate energy intensity at the product level because it affects the revenue share of low-price firms in the market for product j . Since firms with a low energy price produce the same product in a more energy-intensive way, this reshuffling effect *increases* observed product-level energy intensity if and only if $\frac{d\theta_j}{d\nu} > 0$, that is, if the increase in ν increases the revenue share of the low-energy-price group.

Using our expression for R_j from Equation 24, we can see that

$$\frac{d\theta_j}{d\nu} = k\theta_j(1 - \theta_j)\frac{s_j(\nu_{\underline{T}})}{\nu} > 0, \quad (29)$$

that is, an increase in the energy price of high-price firms increases the market share of low-price firms in the market for any product j . This result is quite intuitive: Given an increase in the energy price of high-price firms, these firms will have to increase their prices and lose market shares. What is more, because they become less profitable, some high-price firms close to the extensive margin will choose to exit the market for product j altogether, which further reduces their market share.¹⁹

Taken together, this means that the observed energy intensity of product j changes by

$$\frac{d\eta_j}{d\nu} = \frac{(1 - \theta_j)\eta_j(\nu_{\underline{T}})}{\nu} \left[\underbrace{\frac{\rho}{\rho - 1}(1 - s_j(\nu_{\underline{T}})) - 1}_{<0} + \underbrace{(s_j(\nu_{\underline{T}})\nu - s_j(\nu_{\underline{T}}))k\theta_j}_{>0} \right]. \quad (30)$$

The sign of this total effect can be positive or negative depending on the the product's energy intensity α_j^m , the baseline dispersion of energy prices – given by ν in our stylised distribution – and the exact parametrisation of the model.²⁰ That means we cannot analytically determine whether the reshuffling effect or the technique effect will dominate for a given product. But Figure 3 shows that it certainly is possible that the reshuffling effect will be sufficiently large that an increase in the energy price for a subset of firms leads to an overall *increase* in energy intensity for some products.

Broadening our view across all products to look at aggregate energy intensity η , we can see that

$$\frac{d\eta}{d\nu} = \underbrace{\sum_{j \in J} \frac{d\eta_j}{d\nu} \frac{R_j}{R}}_{\text{avg. within-product effect}} + \underbrace{\frac{1}{R} \left(\sum_{j \in J} \eta_j \frac{dR_j}{d\nu} - \eta \sum_{j \in J} \frac{dR_j}{d\nu} \right)}_{\text{product composition effect}}. \quad (31)$$

This means that the response of aggregate energy intensity to an increase in ν depends on two additively separable effects: the average within-product effect and the product composition effect. The average within-product effect is simply the weighted sum of the effects of an increase in ν on the energy intensity of each individual product that we have discussed at length above. Its sign is indeterminate, as it nests the energy demand reducing technique effect and the energy demand increasing reshuffling effect. If the within-product effect is energy demand increasing for some products and energy demand reducing for others, products with higher total revenue will have a higher weight.

¹⁹In general, a third effect will be at work as well: An increase in ν increases the equilibrium price P_j^* because it increases the cost of a non-zero mass of firms. In principle, such an increase in P_j^* may affect the revenue of high-energy-price firms and low-energy-price firms differently, leading to an orthogonal shift in θ_j . But by our assumption of Pareto-distributed productivities, θ_j does not depend on P_j , so this channel is shut down.

²⁰For example, a lower productivity dispersion (i.e. a higher k), a stronger complementarity between energy and labour (i.e. a lower ρ) or a larger gap between $s_j(\nu_{\underline{T}})$ and $s_j(\nu_{\underline{T}})$ increase $\frac{d\eta_j}{d\nu}$.




Figure 3: Elasticity of product-level energy intensity with respect to ν

Note: Each line represents the elasticity of observed product-level energy intensity with respect to ν , $\frac{d\eta_j}{d\nu} \frac{\nu}{\eta_j}$, for a different baseline level of ν . It is based on $p = 0.3$ and $\tau = 1$, with all other parameters corresponding to their values in the main quantification as listed in Appendix D.

The product composition effect reflects the fact that a change in ν may affect the shares of products with a different energy intensity in different ways. It is energy-decreasing if and only if the empirical covariance between η_j and $\frac{dR_j}{d\nu}$ is negative, that is: If the revenue share of products with a high energy intensity decreases in response to an increase in ν , and vice versa.

Since products only differ in their baseline energy intensity α_j^m , and since we know by Lemma 1 that η_j is strictly increasing in α_j^m , this boils down to the question whether or not $\frac{dR_j}{d\nu}$ is increasing in α_j^m , that is, whether or not $\frac{d^2 R_j}{d\alpha_j^m d\nu} > 0$. Similarly to the case of the within-product effect, it turns out that this depends on the product's energy intensity, the baseline level of ν and other model parameters. Specifically, some algebra shows that, in equilibrium,

$$\frac{dR_j}{d\nu} = -k \underbrace{\frac{(1-\xi)(\sigma-1)}{(1-\xi)(\sigma-1) + \xi k}}_{\in(0,1)} \frac{R_j}{\nu} (1-\theta_j) s_j(\nu \underline{\tau}) < 0. \quad (32)$$

That means that, as one would expect, aggregate revenue from each product decreases when ν increases. However, whether or not this decrease is weaker or stronger among products with a high baseline energy intensity α_j^m crucially depends on the $(1-\theta_j)s_j(\nu \underline{\tau})$ term. That term reflects two opposing forces with respect to the role of α_j^m : On the one hand, products with a higher α_j^m are simply produced less by firms with a high energy price. Hence, $(1-\theta_j)$ is relatively lower for such products and the increase in ν affects them less. On the other hand, the adverse effects of an increase in ν are stronger among products for which the share of energy in total costs, $s_j(\nu \underline{\tau})$, is high. Trivially, this is the case for products with a higher α_j^m .

Analytically, the sign of $\frac{d^2 R_j}{d\alpha_j^m d\nu}$ is indeterminate, as it is a complex non-linear and non-monotonic function of product energy intensity α_j^m and the energy price distribution. Figure 4 highlights that it indeed strongly depends on the energy price distribution how $\frac{dR_j}{d\nu}$ evolves with α_j^m . This implies that the sign of the composition effect will depend on the empirical distribution of product baseline energy intensities as well.

Putting everything together, we conclude that a non-uniform increase in the energy price can indeed explain the puzzling pattern we observe in the data: An asymmetric price increase can lead to a decrease in the energy intensity of the product mix – through the product composition effect – paired




Figure 4: Elasticity of aggregate product revenue with respect to ν

Note: Each line represents the elasticity of aggregate product revenue with respect to ν , $\frac{dR_j}{d\nu} \frac{\nu}{R_j}$, for a different baseline level of ν . It is based on $p = 0.3$ and $\bar{\tau} = 1$, with all other parameters corresponding to their values in the main quantification as listed in Appendix D.

with an increase in actual energy intensity through the reshuffling effect. Whether or not it will produce these results depends on the exact parameters of the model.

One stylised setting in which our framework trivially produces exactly such a result is a world with two products – product A which requires no energy at all and product B which requires some energy –, a high baseline level of ν and a high share p of firms with a persistent low energy price. Then, the product composition effect is clearly energy-reducing because the revenue from product A is unaffected by an increase in ν whereas revenue from product B shrinks. The energy-reducing technique effect is rather small because only few firms are actually affected by the increase in ν . But at the same time, the energy-increasing reshuffling effect is quite strong because $\eta_B(\bar{\tau}) - \eta_B(\nu\bar{\tau})$ is large.


4 Evidence of model mechanisms

Suggestive evidence at the product, firm and firm-product level provides support to several key mechanisms posited by our model as we have described it in Section 3. In particular, it supports our reshuffling hypothesis that producers with a low energy price partially replace producers with a higher energy price in the markets for energy-intensive goods, which may result in increasing energy intensity.


First of all, Figure 5 shows that producer reshuffling is more pronounced for more energy intensive products. From panel (a), we can see that both add and drop rates – that is, the share of firms that add or drop the product over a given time period relative to the total number of firms that produce it – are higher for more energy-intensive products. That means that energy-intensive products are dropped more often than others (relative to the total number of firms that produce them), just as one would expect under rising energy prices. But they are also *added* more often than others – indicating that new producers may fill part of the gap created by old producers moving out of the product market.

As can be seen from panel (b) of Figure 5, this is also reflected in a higher standard deviation of firm-product level growth rates among more energy intensive products. That means that we also find evidence of reshuffling when we look at both extensive-margin changes (i.e. when a firm adds or drops a product) and intensive-margin changes (i.e. when a firm increases or decreases its output of a product) combined.

In line with our hypothesised reshuffling mechanism, the firms that are adding more energy-intensive products tend to be those with a lower energy price, and vice versa for the case of product dropping.



(a) Product add and drop rates by energy intensity



(b) Standard deviation of symmetric growth rate by energy intensity

Figure 5: Reshuffling among producers of most energy-intensive products

Note: Each point in the binned scatter plots represents around 100 product-period combinations. Each underlying data point represents one product over a six-year interval from 2005-2011 or 2011-2017. Product add and drop rates are computed as $\frac{\# \text{ of producing firms adding/dropping the product}}{\# \text{ of producers in base period} + \# \text{ of producing firms in final period}}$ in line with Bernard et al. (2010). Firm-product level symmetric growth rates are computed in terms of revenue r_{ijt} as $\frac{r_{ij,t+1} - r_{ijt}}{r_{ij,t+1} + r_{ijt}}$.

Source: AFiD manufacturing census (see Appendix A)

As Panel (a) of Figure 6 shows, the probability that a firm adds a specific product decreases in the firm's energy price. This relationship between energy price and adding probability plays almost no role among products with a low energy intensity but is strongly pronounced among products with high energy intensities. For example, the predicted probability that a firm with a low energy price of 5 ct/kWh adds a product with a high energy intensity of 4 kWh/EUR is more than seven times as high as the probability that a firm with a high energy price of 25 ct/kWh adds the same product.

Panel (b) of Figure 6 documents the reverse tendency for product dropping: The probability that a firm with a high energy price drops a product is much higher for products with a high energy intensity than for products with a low energy intensity. For example, for a firm with a high energy price of 25 ct/kWh, the predicted probability of dropping a product with a high energy intensity of 4 kWh/EUR is 60 percent higher than the predicted probability of dropping a product with a low energy intensity of 0.5 kWh/EUR. Conversely, firms with a lower energy price are much less likely to drop energy-intensive products.

Unexpectedly, though, firms with a low energy price are more likely than others *in absolute terms* to drop products with a low energy intensity. This seems counterintuitive at first because a lower energy price is – at least through the lens of our model – an unequivocal benefit: It reduces production costs for both energy-intensive and less energy-intensive products, even though cost reductions are much larger for energy-intensive products. A large part of this relationship can be explained by the simple fact that we only include surviving firms in the regressions underlying Figure 6. While this gives us clean product adding and dropping effects that are not affected by firm entry and exit, it implies that we mostly observe product dropping among multi-product firms, as product switches among single-product firms are rare. But firms with a low energy price are more likely to be multi-product firms, so we are more likely to observe them drop a product.²¹

²¹The fact that firms with a low energy price will, on average, produce more different products follows directly from our model (since productivity and energy prices are independent from one another, and the productivity cutoff for product market entry is increasing in the energy price). We also document it empirically in Figure 11b in Appendix E.




Figure 6: Predicted adding and dropping probabilities of products by product energy intensity and firm energy price

Note: Visualisations are based on a regression of adding and dropping dummies on log firm energy price, product energy intensity and the interaction of the two. We report the regression results in Table 5 in Appendix E. The adding probabilities in Panel (a) are based on the sample of 18 007 162 hypothetical firm-product combinations over the two six-year periods 2005-2011 and 2011-2017 that are not produced in the respective period’s base year (and so could potentially be added). The dropping probabilities in Panel (b) are based on the sample of 35 306 firm-product combinations over the two six-year periods 2005-2011 and 2011-2017 that are produced in the respective period’s base year (and so could potentially be dropped). The sample is restricted to firms that survive over the respective six-year period. The graphs’ x and y axes represent the empirical range of product energy intensities and firm energy prices, respectively.

Source: AFiD manufacturing census (see Appendix A)

On top of these results on energy-price-related product adding and dropping, Figure 7 documents a positive correlation between a firm’s energy price and its total factor productivity, conditional on output. This is evidence of exactly the selection mechanism that our model posits: Both a low energy price and a high productivity reduce firms’ production costs, increasing their output and profits. Hence, at any given level of output, we expect a positive relationship between energy price and productivity – even if energy prices and productivity are *ex ante* independent. Specifically, our model predicts that

$$\tau_i = \left(\frac{1}{\alpha_j^m} \right)^{-\frac{1}{\rho}} \left[\left(\frac{\varphi_{ij}}{r_{ij}^{\frac{1}{\sigma-1}}} A^{\frac{1}{\sigma-1}} (P_j^*)^\xi \right)^{\frac{\rho}{\rho-1}} - (1 - \alpha_j^m)^{-\frac{1}{\rho-1}} w^{\frac{\rho}{\rho-1}} \right]^{\frac{\rho-1}{\rho}} \quad (33)$$

for each firm i and product j . This means that the firm energy price τ_i is positively correlated with $\frac{\varphi_{ij}}{r_{ij}^{\frac{1}{\sigma-1}}}$. Since firm-product revenue r_{ij} is a function of $\varphi_{ij}^{\sigma-1}$, φ_{ij} cancels out of the right-hand side of the equation in the cross section.²² However, once we condition on a specific level of revenue, our model predicts a clear positive relationship between τ_i and φ_{ij} . While we do not observe firm-product specific productivity φ_{ij} in the data – because we cannot attribute the inputs of multi-product firms to their different outputs – this result generalises to the relationship between firm-level productivity and the energy price.²³

In fact, when adding ‘2-5 products’ and ‘>5 products’ dummies to our regression specification in Table 5, the negative energy price coefficient drops from -0.2 to -0.11 .

²²See Equation 54 in Appendix C.1 for the exact expression for optimal r_{ij} .

²³The firm-level productivity we observe in the data is, effectively, a weighted average of the firm’s specific productivities in producing the different products it has in its portfolio.




Figure 7: Positive correlation between energy price and TFP


Note: Bin scatter plot based on $N = 34\,915$ firm-year combinations in 2005, 2011 and 2017. Both x and y axis are net of a linear control term for log firm revenue. TFPR is calculated using median cost shares of capital, labour, materials and energy at the 4-digit industry level as proxies for the respective input elasticities as described in De Loecker and Syverson (2021). The regression coefficient underlying the slope of the fitted line is 0.07, measured with a standard error of 0.01. Similar results hold for most other productivity measures, as documented in Table 7 in Appendix E.

Source: AFiD manufacturing census (see Appendix A)

Finally, the data also suggest that firms with a lower energy price are more energy intensive (Figure 8). This holds both unconditionally (as depicted in panel (a) of the figure) and conditional on firms' exact product mix (as depicted in panel (b)). To arrive at the latter, we first compute the energy intensity of the firm's product mix analogously to the energy intensity of the entire manufacturing sector's product mix in Figure 2. If a firm's actual energy intensity is higher than the energy intensity of its product mix, this means that it requires more energy to produce its exact combination of outputs than the average firm would need. This quantity – which we have dubbed the actual–predicted energy ratio – allows us to compare the energy intensity of firms conditional on their potentially very diverse product mixes. Of course, the relationship depicted in Figure 8 does not represent the causal effect of energy prices on energy intensity: At least some degree of reverse causality is to be expected because, typically, firms with a higher energy intensity will use more energy and for this reason get quantity discounts from their energy providers. But it shows that there is a clear negative relationship between energy intensity and the energy price even after accounting for the fact that firms with a low energy price may select into more energy-intensive products and industries. A causal effect from higher energy prices to lower energy intensity is then both in line with fundamental economic intuition and existing empirical evidence (Gerster & Lamp, 2024; Marin & Vona, 2021).

5 Model quantification

As we have illustrated in Section 3, heterogeneous energy price changes can affect actual energy intensity and the energy intensity of the product mix in different ways due to a reshuffling among the producers of energy-intensive products in particular. To get an idea of the magnitude of the effect that this reshuffling could have had in German manufacturing over our period of observation, we walk through a simple and preliminary quantification of our model in this section. We simulate our model with 75,000 firms and 100 products over three periods that roughly represent our main years of



(a) Energy intensity

(b) Actual / predicted energy use

Figure 8: Relationship between energy price and energy intensity

Note: Bin scatter plot based on $N = 35\,301$ firm-year combinations in 2005, 2011 and 2017. The ratio of actual to predicted energy use is calculated as $\frac{m_i}{\sum_{j=1}^{N_j} \hat{\eta}_j r_{ij}}$, where m_i represents firm i 's energy consumption, $\hat{\eta}_j$ is the average energy intensity of product j as described in Section 2.1 and r_{ij} is firm i 's output of product j .

Source: AFiD manufacturing census (see Appendix A)

observation: 2005, 2011 and 2017.²⁴

In contrast to the stylised setup in Section 3.2, we plug a more realistic distribution of energy prices, H , into our model in this section. To roughly mimic the empirical energy price distribution, we assume that energy prices are distributed log-normally with log mean $\mu_{\tau,1}$ and log standard deviation $\sigma_{\tau,1}$ in the first period. For each subsequent period $t > 1$, firm i 's energy price τ_i is updated according to the law of motion

$$\ln \tau_{it} = \ln \tau_{i,t-1} + \beta_{\tau,t} (\ln \tau_{t-1} - \overline{\ln \tau_{i,t-1}}) + \ln \hat{\tau}_{it}, \quad (34)$$

where $\hat{\tau}_{it}$ is a random shock drawn from a normal distribution with mean $\hat{\mu}_{\tau,t}$ and standard deviation $\hat{\sigma}_{\tau,t}$, and $\overline{\ln \tau_t}$ is the average log energy price in period t . If the shock distribution has a higher mean, the shock mainly results in an increased average energy price, and if it has a higher standard deviation, the shock mainly results in an increased energy price dispersion (although mean and variance of the energy price distribution positively depend on each other due to its log-normal structure). A higher $\beta_{\tau,t}$ implies that firms with a higher baseline energy price will be more likely to have their energy price increase, which results in a higher energy price dispersion on aggregate.²⁵

To reproduce the development of the empirical energy price distribution from 2005 to 2017, we set the baseline parameters $\mu_{\tau,1} = 0$ and $\sigma_{\tau,1} = 0.4$ and pick law of motion parameters $(\hat{\mu}_{\tau,t}, \hat{\sigma}_{\tau,t}, \beta_{\tau,t})_{t=2,3}$ that yield the energy price distribution depicted in Panel (b) of Figure 9.²⁶ Calibration choices with respect to other model components are described in more detail in Appendix D.

Table 2 shows the results of the corresponding model simulation. Given our choice of parameters and our approximation of the true energy price distribution, we see that the energy intensity of the product mix inversely follows the average energy price: It slightly decreases from period 1 to period 2, when

²⁴We have chosen the number of firms to get the right order of magnitude of firms entering the market in our resulting simulations. The number of products is a compromise between how granular we are able to represent real-world product heterogeneity versus how computationally intensive it is to solve the model.

²⁵Specifically, since the right-hand side is a linear combination of normal distributions, the log of the updated energy price distribution will be normally distributed as well, with $\mu_{\tau,t} = \mu_{\tau,t-1} + \hat{\mu}_{\tau,t}$ and $\sigma_{\tau,t} = \sqrt{(\beta^2 + \beta + 1)\sigma_{\tau,t}^2 + \hat{\sigma}_{\tau,t}^2}$.

²⁶The exact parameters to produce these distributions are $\hat{\mu}_{\tau,2} = 0.6$, $\hat{\sigma}_{\tau,2} = 0$, $\beta_{\tau,2} = 0$, $\hat{\mu}_{\tau,3} = 0$, $\hat{\sigma}_{\tau,3} = 0.2$ and $\hat{\beta}_{\tau,3} = 0.2$.




Figure 9: Actual distribution of energy prices vs distribution used in simulation

Note: Panel (b) represents kernel density estimates of the distribution of energy price draws among the 75,000 firms in the population of our simulated model world. Panel (a) is identical to Figure 1a and just repeated here for ease of comparison with panel (b).

Table 2: Energy price dispersion and energy intensity in simulated model

t	$\bar{\tau}_{\text{weighted}}$	$\bar{\tau}_{\text{unweighted}}$	σ_{τ}	η (product mix)	η (actual)
1	0.94	0.84	0.39	100.00%	100.00%
2	1.67	1.49	0.69	99.77%	68.10%
3	1.57	1.31	0.85	99.81%	75.01%

Note: $\bar{\tau}_{\text{weighted}}$ represents the consumption-weighted average of the energy prices of all firms active in the market. $\bar{\tau}_{\text{unweighted}}$, conversely, represents the unweighted average of the same energy prices. σ_{τ} represents the (unweighted) standard deviation of energy prices. Predicted η denotes aggregate (revenue-based) energy intensity if each product was produced at its baseline energy intensity from $t = 1$. Actual η denotes actual (revenue-based) energy intensity.

the average energy price rises strongly, and it minimally increases from period 2 to period 3, when the average energy price slightly falls. This is the expected result: Energy price increases make energy-intensive products relatively more expensive to produce and hence less profitable, so firms should disproportionately reduce their output of them.

Overall, though, the energy intensity of the product mix is quite persistent despite substantial price changes. This is the reshuffling mechanism at play: As energy prices increase on average, the comparative advantage of those firms with a relatively low energy price increases. Many firms reduce their output of energy-intensive products, but the producers with the relatively lowest energy prices increase their output, so in sum, production does not decrease by that much.

By contrast, actual energy intensity changes substantially between the three periods. That means: *What* is produced does not change by that much, but *how* it is produced changes a lot. From period 1 to period 2, total energy intensity drops by 32%. This is because the effect of the energy price increase across the board (which causes all firms to reduce their energy intensity) dominates the effect of the increase in the dispersion of energy prices (which leads to a reallocation towards firms with a relatively lower energy price). From period 2 to period 3, on the other hand, total energy intensity increases by 7%. This is in part the result of a slight decrease in the average energy price and in part due to the increase in energy price dispersion.

While these simulation results are internally consistent and do highlight that reshuffling due to increasing energy price dispersion can drive a wedge between actual energy intensity and the energy

intensity of the product mix, they are not able to explain the actual development of energy intensity in German manufacturing. Recall from Figure 2 that the energy intensity of the product mix fell slightly but consistently between 2005 and 2017. Actual energy intensity fluctuated substantially around a slight upward trend between 2005 and 2011 and then increased substantially between 2011 and 2017. Our model results do offer a convincing explanation how increasing energy price dispersion may have contributed to the increase in actual energy intensity relative to the energy intensity of the product mix, in particular over the period from 2011 to 2017. But additional factors that we do not include in our simple model would be required to account for the overall increase in energy intensity despite rising average energy prices. One reason might be that the cost of other inputs has increased as well over our period of observation, which could explain why an increase in the average energy price did not lead to the aggregate shift away from energy that our model predicts. Alternatively, a shift in production technology might be at play if the increased availability of automation technologies introduced an ‘energy-biased technological change’. Investigating these mechanisms in more detail is an important avenue for future research.

6 Conclusion

In this study, we have shown that an increase in the dispersion of energy prices can lead to an increase in the aggregate energy intensity of production, and that this can happen independently of the energy intensity of the product mix. Using administrative microdata, we have shown that the energy prices paid by German manufacturing firms have become substantially more dispersed between 2005 and 2017. And that this has gone along with a reshuffling among the producers of more energy-intensive products: These were more likely to be added by firms with a low energy price and (relatively) more likely to be dropped by firms with a high energy price. We have developed a model that explains how an increased dispersion of energy prices leads to such a reshuffling and that this reshuffling increases aggregate energy intensity because firms with a low energy price tend to produce the same goods in a more energy-intensive way than firms with a high energy price. In a simple and preliminary quantification, we have been able to confirm the relevance of this mechanism, but have observed that, in our model, it is dominated by the energy-saving effects of rising average energy prices.

Our findings have two main implications for policy. First, they demonstrate that, in order to reduce energy consumption and emissions, it is not enough to focus on a shift towards less energy-intensive products alone. As we have seen empirically in the case of the German manufacturing sector, and illustrated in our model, it is possible that energy intensity rises even though the manufacturing sector as a whole shifts towards less energy-intensive products. The crucial policy lesson here is that policies that aim at making less energy-intensive products more attractive without strictly penalising the production of more energy-intensive products do lead to a shift towards less energy-intensive products. But they can make the production of highly energy-intensive products even more energy-intensive. This is because policies that induce a shift towards less energy-intensive products reduce competitive pressure in the markets for more energy-intensive products. Exactly those firms with the highest comparative advantage in the production of energy-intensive products will benefit the most from this reduced pressure and step up their production. These firms, however, are often the ones with the lowest effective energy or emissions prices. Consequently, they produce the same goods in a more energy-intensive way and counteract the beneficial effects of the increased production of clean goods. The second main policy implication is related to the first: It is that allocation – that is, who produces how much of what – matters for energy consumption. Policies like the compensation of energy-intensive producers in trade-exposed industries may aim to ensure a level playing field internationally, where many countries can offer cheaper and less heavily regulated access to energy. But they tilt the domestic playing field towards the compensated firms.

This is a highly relevant aspect of such policies that public debate has often neglected so far. The allocative efficiency literature has long found that the inefficient allocation of resources (away from the firm that can put them to their most productive use) is a significant barrier to efficient production that may cost a substantial share of GDP (Hopenhayn, 2014; Hsieh & Klenow, 2009; Ruzic & Ho, 2023). Our findings demonstrate that dispersion in energy prices has allocative effects as well: It leads to an allocation of the production of the most energy-intensive products to the firms with the lowest energy price and the least energy-efficient technology. This contributes to an increase in aggregate energy intensity. Whether it exacerbates or mitigates the problem of allocative efficiency depends on who the firms are that benefit from lower energy prices: If the institutional and regulatory conditions in the industrial energy market make it easier for more productive firms to source cheap energy, this may lead to increased energy intensity but increase allocative efficiency because it causes an allocation of production towards more productive firms. If the reverse is true – i.e. if less productive firms have access to cheaper energy, on average – the dispersion of energy prices makes allocative inefficiencies worse.

Bibliography

- Abeberese, A. B. (2017). Electricity Cost and Firm Performance: Evidence from India. *The Review of Economics and Statistics*, 99(5), 839–852. https://doi.org/10.1162/REST_a_00641
- Akerberg, D. A., Caves, K., & Frazer, G. (2015). Identification Properties of Recent Production Function Estimators. *Econometrica*, 83(6), 2411–2451. <https://doi.org/10.3982/ECTA13408>
- Albuquerque, B., & Iyer, R. (2023). *The Rise of the Walking Dead: Zombie Firms Around the World* (Number WP/23/125). https://www.imf.org/en/Publications/WP/Issues/2023/06/16/The-Rise-of-the-Walking-Dead-Zombie-Firms-Around-the-World-534866?utm_source=chatgpt.com
- Aldy, J. E., & Pizer, W. A. (2015). The Competitiveness Impacts of Climate Change Mitigation Policies. *Journal of the Association of Environmental and Resource Economists*, 2(4), 565–595. <https://doi.org/10.1086/683305>
- Barrows, G., & Ollivier, H. (2018). Cleaner Firms or Cleaner Products? How Product Mix Shapes Emission Intensity from Manufacturing. *Journal of Environmental Economics and Management*, 88, 134–158. <https://doi.org/10.1016/j.jeem.2017.10.008>
- Barrows, G., Calel, R., Jégard, M., & Ollivier, H. (2024, May 13). *Equilibrium Effects of Carbon Policy*. Mannheim Conference on Energy and the Environment.
- Basaglia, P., Behr, S. M., & Drupp, M. A. (2024). *De-Fueling Externalities: Causal Effects of Fuel Taxation and Mediating Mechanisms for Reducing Climate and Pollution Costs* (Number 10508). <https://doi.org/10.2139/ssrn.4477996>
- Bernard, A. B., Redding, S. J., & Schott, P. K. (2011). Multiproduct Firms and Trade Liberalization. *The Quarterly Journal of Economics*, 126(3), 1271–1318. <https://doi.org/10.1093/qje/qjr021>
- Bernard, A. B., Redding, S. J., & Schott, P. K. (2010). Multiple-Product Firms and Product Switching. *American Economic Review*, 100(1), 70–97. <https://doi.org/10.1257/aer.100.1.70>
- Boehm, J., Dhingra, S., & Morrow, J. (2022). The Comparative Advantage of Firms. *Journal of Political Economy*, 130(12), 3025–3100. <https://doi.org/10.1086/720630>
- Bretschger, L., & Jo, A. (2024). Complementarity between Labor and Energy: A Firm-Level Analysis. *Journal of Environmental Economics and Management*, 124, 102934. <https://doi.org/10.1016/j.jeem.2024.102934>
- Calel, R., & Dechezleprêtre, A. (2016). Environmental Policy and Directed Technological Change: Evidence from the European Carbon Market. *The Review of Economics and Statistics*, 98(1), 173–191. https://doi.org/10.1162/REST_a_00470

- Colmer, J., Martin, R., Muûls, M., & Wagner, U. J. (2024). Does Pricing Carbon Mitigate Climate Change? Firm-Level Evidence from the European Union Emissions Trading System. *The Review of Economic Studies*. <https://doi.org/10.1093/restud/rdae055>
- Davis, S. J., Grim, C., Haltiwanger, J., & Streitwieser, M. (2013). Electricity Unit Value Prices and Purchase Quantities: U.S. Manufacturing Plants, 1963–2000. *The Review of Economics and Statistics*, 95(4), 1150–1165. https://doi.org/10.1162/REST_a_00309
- De Loecker, J., & Syverson, C. (2021). An Industrial Organization Perspective on Productivity. In *Handbook of Industrial Organization: Vol. 4. Handbook of Industrial Organization* (pp. 141–223). Elsevier. <https://doi.org/10.1016/bs.hesind.2021.11.003>
- De Loecker, J., Goldberg, P. K., Khandelwal, A. K., & Pavcnik, N. (2016). Prices, Markups, and Trade Reform. *Econometrica*, 84(2), 445–510. <https://doi.org/10.3982/ECTA11042>
- Destatis. (2023). *Preise: Daten zur Energiepreisentwicklung. Lange Reihen von Januar 2005 bis Januar 2023* (Number 5619001231014). https://www.destatis.de/DE/Themen/Wirtschaft/Preise/Erdgas-Strom-Durchschnittspreise/_inhalt.html#_a9l5v6d6y
- Eckel, C., & Neary, J. P. (2010). Multi-Product Firms and Flexible Manufacturing in the Global Economy. *Review of Economic Studies*, 77(1), 188–217. <https://doi.org/10.1111/j.1467-937X.2009.00573.x>
- Elliott, R., Sun, P., & Zhu, T. (2019). Electricity Prices and Industry Switching: Evidence from Chinese Manufacturing Firms. *Energy Economics*, 78, 567–588. <https://doi.org/10.1016/j.eneco.2018.11.029>
- Elrod, A. A., & Malik, A. S. (2017). The Effect of Environmental Regulation on Plant-Level Product Mix: A Study of EPA's Cluster Rule. *Journal of Environmental Economics and Management*, 83, 164–184. <https://doi.org/10.1016/j.jeem.2017.03.002>
- Elrod, A. A., & Malik, A. S. (2019). The Effect of County Non-Attainment Status on the Product Mix of Plants in the Pulp, Paper, and Paperboard Industries. *Journal of Environmental Economics and Policy*, 8(3), 283–300. <https://doi.org/10.1080/21606544.2019.1569561>
- Flues, F. S., & Lutz, B. J. (2015). *The Effect of Electricity Taxation on the German Manufacturing Sector: A Regression Discontinuity Approach* (Numbers 15–13). <https://doi.org/10.2139/ssrn.2580742>
- Ganapati, S., Shapiro, J. S., & Walker, R. (2020). Energy Cost Pass-Through in US Manufacturing: Estimates and Implications for Carbon Taxes. *American Economic Journal: Applied Economics*, 12(2), 303–342. <https://doi.org/10.1257/app.20180474>
- Gerster, A., & Lamp, S. (2024). Energy Tax Exemptions and Industrial Production. *The Economic Journal*, 134(663), 2803–2834. <https://doi.org/10.1093/ej/ueae048>
- Goldberg, P. K., Khandelwal, A. K., Pavcnik, N., & Topalova, P. (2010). Multiproduct Firms and Product Turnover in the Developing World: Evidence from India. *The Review of Economics and Statistics*, 92(4), 1042–1049. <https://www.jstor.org/stable/40985812>
- Hawkins-Pierot, J. T., & Wagner, K. R. H. (2024). *Technology Lock-In and Costs of Delayed Climate Policy*. http://www.krhwagner.com/papers/carbon_lockin.pdf
- Hopenhayn, H. A. (2014). Firms, Misallocation, and Aggregate Productivity: A Review. *Annual Review of Economics*, 6, 735–770. <https://doi.org/10.1146/annurev-economics-082912-110223>
- Hsieh, C.-T., & Klenow, P. J. (2009). Misallocation and Manufacturing TFP in China and India*. *The Quarterly Journal of Economics*, 124(4), 1403–1448. <https://doi.org/10.1162/qjec.2009.124.4.1403>
- Levinson, A. (2009). Technology, International Trade, and Pollution from US Manufacturing. *American Economic Review*, 99(5), 2177–2192. <https://doi.org/10.1257/aer.99.5.2177>
- Levinson, A. (2015). A Direct Estimate of the Technique Effect: Changes in the Pollution Intensity of US Manufacturing, 1990–2008. *Journal of the Association of Environmental and Resource Economists*, 2(1), 43–56. <https://doi.org/10.1086/680039>

- Liu, R. (2010). Import Competition and Firm Refocusing. *Canadian Journal of Economics/Revue canadienne d'économie*, 43(2), 440–466. <https://doi.org/10.1111/j.1540-5982.2010.01579.x>
- Ma, Y., Tang, H., & Zhang, Y. (2014). Factor Intensity, Product Switching, and Productivity: Evidence from Chinese Exporters. *Journal of International Economics*, 92(2), 349–362. <https://doi.org/10.1016/j.jinteco.2013.11.003>
- Marin, G., & Vona, F. (2021). The Impact of Energy Prices on Socioeconomic and Environmental Performance: Evidence from French Manufacturing Establishments, 1997–2015. *European Economic Review*, 135, 103739. <https://doi.org/10.1016/j.eurocorev.2021.103739>
- Mayer, T., Melitz, M. J., & Ottaviano, G. I. P. (2014). Market Size, Competition, and the Product Mix of Exporters. *American Economic Review*, 104(2), 495–536. <https://doi.org/10.1257/aer.104.2.495>
- Mayr-Dorn, K. (2024). *How Effective Is Emissions Pricing? The Role of Firm-Product Level Adjustment* (Number 2410). <http://www.economics.jku.at/papers/2024/wp2410.pdf>
- Navarro, L. (2012). Plant Level Evidence on Product Mix Changes in Chilean Manufacturing. *The Journal of International Trade & Economic Development*, 21(2), 165–195. <https://doi.org/10.1080/09638191003710397>
- Nocke, V., & Yeaple, S. (2014). Globalization and Multiproduct Firms. *International Economic Review*, 55(4), 993–1018. <https://doi.org/10.1111/iere.12080>
- Pan, X., Guo, S., Xu, H., Tian, M., Pan, X., & Chu, J. (2022). China's Carbon Intensity Factor Decomposition and Carbon Emission Decoupling Analysis. *Energy*, 239, 122175. <https://doi.org/10.1016/j.energy.2021.122175>
- Qiu, L. D., & Zhou, W. (2013). Multiproduct Firms and Scope Adjustment in Globalization. *Journal of International Economics*, 91(1), 142–153. <https://doi.org/10.1016/j.jinteco.2013.04.006>
- Ritchie, H., & Roser, M. (2023). *Sector by Sector: Where Do Global Greenhouse Gas Emissions Come From?* <https://ourworldindata.org/ghg-emissions-by-sector>
- Rottner, E., & von Graevenitz, K. (2024). What Drives Carbon Emissions in German Manufacturing: Scale, Technique or Composition?. *Environmental and Resource Economics*, 87(9), 2521–2542. <https://doi.org/10.1007/s10640-024-00894-7>
- Ruzic, D., & Ho, S.-J. (2023). Returns to Scale, Productivity, Measurement, and Trends in U.S. Manufacturing Misallocation. *Review of Economics and Statistics*, 105(5), 1287–1303. https://doi.org/10.1162/rest_a_01121
- Shapiro, J. S., & Walker, R. (2018). Why Is Pollution from US Manufacturing Declining? The Roles of Environmental Regulation, Productivity, and Trade. *American Economic Review*, 108(12), 3814–3854. <https://doi.org/10.1257/aer.20151272>
- von Graevenitz, K., & Rottner, E. (2024). *Do Manufacturing Plants Respond to Exogenous Changes in Electricity Prices? Evidence From Administrative Micro-Data* (Numbers 22–38). <https://doi.org/10.2139/ssrn.4251564>
- von Graevenitz, K., Rottner, E., & Richter, P. (2024). *Is Germany Becoming the European Pollution Haven?* (Numbers 23–69). <https://doi.org/10.2139/ssrn.4690103>
- Wolverton, A., Shadbegian, R., & Gray, W. B. (2022). *The U.S. Manufacturing Sector's Response to Higher Electricity Prices: Evidence from State-Level Renewable Portfolio Standards* (Number 30502). <https://doi.org/10.3386/w30502>
- Wooldridge, J. M. (2009). On Estimating Firm-Level Production Functions Using Proxy Variables to Control for Unobservables. *Economics Letters*, 104(3), 112–114. <https://doi.org/10.1016/j.econlet.2009.04.026>

Zhang, W., Xu, H., & Xu, Y. (2023). Does Stronger Environmental Regulation Promote Firms' Export Sophistication? A Quasi-Natural Experiment Based on Sewage Charges Standard Reform in China. *Sustainability*, 15(11), 9023. <https://doi.org/10.3390/su15119023>

A Exact data versions used

For our analyses, we use the following modules of the German AFiD manufacturing census, provided by the Research Data Centres of the Federal Statistical Office and the Statistical Offices of the Federal States ('Forschungsdatenzentrum des Bundes und der Länder'):

- for information on firm employment, the panel of industrial establishments ('AFiD-Panel Industriebetriebe') with DOI 10.21242/42111.2017.00.01.1.1.0;
- for information on firm-product-level sales, the production survey ('AFiD-Modul Produkte') with DOI 10.21242/42131.2017.00.03.1.1.0;
- for information on firm-level energy use, the energy use survey ('AFiD-Modul Energieverwendung') with DOI 10.21242/43531.2017.00.03.1.1.0;
- for information on energy and materials expenditures, the cost structure panel ('Panel der Kostenstrukturhebung im Bereich Verarbeitendes Gewerbe, Bergbau und Gewinnung von Steinen und Erden') with DOI 10.21242/42251.2017.00.01.1.1.0.

B Product-level energy intensities

Lemma 2 shows that the regression-based and the single-product-based estimator of average product energy intensity are closely related when we consider a sample of single-product firms only. The key difference between the two of them is that the regression-based estimator puts a slightly higher weight on the energy intensity of large firms but this evens out asymptotically, so both estimators converge to η_j , that is, the average amount of energy it takes a firm to generate one more unit of revenue from product j .

Lemma 2: The two estimators for average product energy intensity, $\hat{\eta}_j$ and $\hat{\eta}_j^{\text{SP}}$ are asymptotically identical in a sample of single-product firms.

Proof: To see that, define η_j as the coefficient from a linear projection of firm energy use on output,

$$m_{ij} = \eta_j r_{ij} + \varepsilon_{ij}, \quad (35)$$

where r_{ij} is firm i 's output of its only product j , $m_{ij} = m_i$ is the amount of energy firm i uses to produce j , and $\text{Cov}(\varepsilon_{ij}, r_{ij}) = \mathbb{E}[\varepsilon_{ij}] = 0$ by construction of the linear projection.

It is straightforward to show that the OLS estimator $\hat{\eta}_j$ in a regression

$$m_{ij} = \sum_{j=1}^{N_J} r_{ij} \eta_j + \varepsilon_{ij} \quad (36)$$

is defined as

$$\hat{\eta}_j = \frac{\sum_{j=1}^{N_J} r_{ij} m_{ij}}{\sum_{j=1}^{N_J} r_{ij}^2} = \eta_j + \frac{\sum_{j=1}^{N_J} r_{ij} \varepsilon_{ij}}{\sum_{j=1}^{N_J} r_{ij}^2}. \quad (37)$$

Since $\text{Cov}(\varepsilon_{ij}, r_{ij}) = 0$ and $\mathbb{E}[\varepsilon_{ij}] = 0$, clearly $\hat{\eta}_j \xrightarrow{p} \eta_j$ by the law of large numbers and the continuous mapping theorem.

By contrast, the single-product based product energy intensity estimator $\hat{\eta}_j^{\text{SP}}$ is defined as

$$\hat{\eta}_j^{\text{SP}} = \frac{\sum_{j=1}^{N_J} m_{ij}}{\sum_{j=1}^{N_J} r_{ij}} = \eta_j + \frac{\sum_{j=1}^{N_J} \varepsilon_{ij}}{\sum_{j=1}^{N_J} r_{ij}}. \quad (38)$$

Trivially, it also holds that $\hat{\eta}_j^{\text{SP}} \xrightarrow{p} \eta_j$ by $\mathbb{E}[\varepsilon_{ij}] = 0$. □

However, restricting our attention to single-product firms ignores a substantial part of production and energy consumption. We have explained in Section 2.1 that multi-product firms are substantially more energy-intensive than single-product firms and make up the bulk of aggregate production and energy consumption. This is reflected in the performance of the two product energy intensity measures because many products – and in particular the most energy-intensive ones – are not produced by single-product firms at all. When we compare the coverage of the two measures, we see that products representing 23% of total output have only a regression-based energy intensity measure, whereas just 5% have only a single-product-based measure (Table 3).

When we shift our glance to the firm level, we can see that, based on our product energy intensities, we can predict the energy use of firms covering 63% of output and 49% of energy use when using our regression-based measure (compared to 61% of output and 39% of energy use when using the single-product based measure). The high share of unpredictable output and energy use at the firm level comes from the fact that we can only predict the energy use of firms for all of whose products we can measure product energy intensity. However, among the firms with predictable energy use,

the correlation between actual energy use and the energy use predicted by regression-based product intensity measures is very high at 0.88.

Table 3: Coverage of regression-based and single-product-based product energy intensity measures

	both measures	reg only	SP only	none
	(1)	(2)	(3)	(4)
Share of products	59%	23%	5%	13%
Share of output	59%	23%	5%	13%
Output share of firms with predictable energy use	54%	9%	7%	28%
Energy share of firms with predictable energy use	36%	13%	3%	47%

Note: The share of products and share of output are based on a sample of 3 582 product-year combinations across our focal years of 2005, 2011 and 2017. The shares of firms are based on 85 445 firm-year combinations across the same focal years. Column (1) represents the share of firms or products for which both a regression-based and a single-product-based energy intensity measure is available. Column (2) and column (3) represent the shares of firms or products for which either a regression-based or a single-product-based energy intensity measure is available, respectively. Column (4) represents the share of firms or products for which neither a regression-based nor a single-product-based energy intensity is available.

Source: AFiD manufacturing census (see Appendix A)

C Model proofs and derivations

C.1 Fundamental model behaviour

The consumer's utility maximisation problem

The Karush–Kuhn–Tucker conditions of the consumer's utility maximisation problem defined by Equation 3 are

$$\left. \frac{\partial U}{\partial q_0} \right|_{\mathbf{q}^*} = 1 = \lambda^* + \mu_0^* \quad (39)$$

$$\left. \frac{\partial U}{\partial q_{ij}} \right|_{\mathbf{q}^*} = \left. \frac{\partial U}{\partial U_j} \right|_{\mathbf{q}^*} \cdot \left. \frac{\partial U_j}{\partial q_{ij}} \right|_{\mathbf{q}^*} = \lambda^* p_{ij} + \mu_{ij}^* \quad \forall i, j \quad (40)$$

$$\mu_0^* \geq 0 \quad \mu_{ij}^* \geq 0 \quad \forall i, j \quad \mu_0^* q_0^* = 0 \quad \mu_{ij}^* q_{ij}^* = 0 \quad \forall i, j, \quad (41)$$

where μ_0 and μ_{ij} are the Lagrangian multipliers corresponding to the non-negativity constraints on q_0 and q_{ij} , respectively, and λ is the Lagrangian multiplier corresponding to the budget constraint. Consider a consumption bundle $\mathbf{q}^*$ that satisfies the Karush–Kuhn–Tucker conditions and maximises the consumer's utility. First, note that the U is strictly increasing in q_0 and each q_{ij} , which implies that the budget constraint has to be binding for $\mathbf{q}^*$. Second, spelling out Equation 40 tells us that

$$(1 - \delta_j) \gamma_j (U_j(\mathbf{q}_j^*))^{\frac{1 - \sigma_j \delta_j}{\sigma_j}} (q_{ij}^*)^{-\frac{1}{\sigma_j}} = \lambda^* p_{ij} + \mu_{ij}^* \quad \forall i, j, \quad (42)$$

from which we can see that, for any finite price p_{ij} , optimal consumption q_{ij}^* of variety i of product j has to be strictly positive. Otherwise, the consumer could marginally reduce her consumption of any other variety $i'j'$ with $q_{i'j'}^* > 0$ (resulting in a finite marginal reduction in utility), and marginally increase her consumption of variety ij (resulting in an infinite marginal increase in utility) to increase her utility. By the complementary slackness conditions, this implies that $\mu_{ij}^* = 0 \quad \forall i, j$.²⁷

Let us assume for now that our consumption bundle $\mathbf{q}^*$ features $q_0^* > 0$, i.e. positive consumption of the numéraire good. (We will verify below that this is feasible and show that indeed no alternative consumption bundle $\mathbf{q}'$ with $q'_0 = 0$ can be optimal.) This implies that $\mu_0^* = 0$ and hence – by Equation 39 – that $\lambda^* = 1$.

Then, we can simplify Equation 42 to the demand function

$$q_{ij}^* = ((1 - \delta_j) \gamma_j)^{\sigma_j} U_j(\mathbf{q}_j^*)^{1 - \delta_j \sigma_j} p_{ij}^{-\sigma_j} \quad \forall i, j. \quad (43)$$

We can transform and integrate both sides of Equation 43 over the entire continuum of firms to solve for optimal product utility $U_j(\mathbf{q}_j^*)$ in terms of the standard CES price aggregator $P_j := \left[\int_{i \in I_j} p_{ij}^{1 - \sigma} di \right]^{\frac{1}{1 - \sigma}}$. This reveals that

$$U_j(\mathbf{q}_j^*) = [(1 - \delta_j) \gamma_j]^{\frac{1}{\delta_j}} P_j^{-\frac{1}{\delta_j}} \quad \forall j, \quad (44)$$

²⁷We will see that, in equilibrium, firm i may decide not to supply its variety of j at all because the entry costs into the product market outweigh its profits. This imposes $q_{ij}^* = 0$, which in turn means that the equilibrium price p_{ij}^* cannot be finite. Strictly speaking, infinite prices do not exist, so neither does an equilibrium on the market for variety i of product j . However, we can formally think of an equilibrium on the market for ij as the limiting point of a sequence $(p_{ij,n}^*, q_{ij,n}^*)$ where $q_{ij,n}^* \rightarrow 0$ and $p_{ij,n}^* \rightarrow \infty$ as $n \rightarrow \infty$. In slightly informal notation, we denote this by writing $q_{ij}^* = 0$ and $p_{ij}^* = \infty$.

which we can use to derive the set of final demand functions for each variety of each product,

$$q_{ij}^* = [(1 - \delta_j)\gamma_j]^{\frac{1}{\delta_j}} P_j^{\frac{\sigma_j \delta_j - 1}{\delta_j}} p_{ij}^{-\sigma_j} \quad \forall i, j. \quad (45)$$

Using these demand functions, we can show that total consumer expenditure on manufacturing goods is

$$\sum_{j \in J} \int_{i \in I} p_{ij} q_{ij}^* di = \sum_{j \in J} [(1 - \delta_j)\gamma_j]^{\frac{1}{\delta_j}} P_j^{\frac{\delta_j - 1}{\delta_j}} < Y, \quad (46)$$

where the inequality holds by our assumption from Equation 4. Since we have established above that the budget constraint needs to hold with equality if $\mathbf{q}^*$ maximises the consumer's utility, this implies that

$$q_0^* = Y - \sum_{j \in J} \int_{i \in I} p_{ij} q_{ij}^* di > 0, \quad (47)$$

as we have assumed above.

We have seen that, given an income Y and a set of prices P that satisfy our assumption from Equation 4, the Karush–Kuhn–Tucker conditions pin down a unique optimal solution $\mathbf{q}^*$ among all consumption bundles with $q_0 > 0$. We can show on top of this that no consumption bundle with $q_0 = 0$ can be optimal.

Consider an optimal consumption bundle $\mathbf{q}'$ and suppose for contradiction that $q'_0 = 0$. Then, $\mu'_0 \geq 0$ and $\lambda' \leq 1$ by Equation 39. Since optimality of $\mathbf{q}'$ implies that $\mu'_{ij} = 0$, Equation 42 requires that

$$(1 - \delta_j)\gamma_j (U_j(\mathbf{q}'))^{\frac{1 - \sigma_j \delta_j}{\sigma_j}} (q'_{ij})^{-\frac{1}{\sigma_j}} = \lambda' p_{ij} \leq p_{ij} \quad \forall i, j. \quad (48)$$

Now, if $\lambda' = 1$, it is clear that $q'_{ij} = q_{ij}^* \quad \forall i, j$, which implies by the bindingness of the budget constraint that $q'_0 > 0$, which is a contradiction. If, on the other hand, $\lambda' < 1$, the inequality from Equation 48 becomes strict, which means that $q'_{ij} > q_{ij}^* \quad \forall i, j$. This means that, for any ij , the marginal utility from another unit of q_{ij} at consumption level q'_{ij} is below the price p_{ij} . At the same time, the marginal utility from consuming one more unit of the numéraire is 1, so the consumer can translate any budget not spent on ij one for one into utility. But then the consumer can increase her utility by marginally reducing her consumption of ij and spending the freed budget on the numéraire. Hence, $\mathbf{q}'$ cannot be optimal, which is a contradiction.

The firm's profit maximisation problem

The cost minimisation problem specified in Equation 7 consists of a linear objective function and a strictly convex constraint, so the first order conditions of the corresponding Lagrangian uniquely pin down the minimising amounts of labour and energy required to produce q_{ij} units of output. Specifically, these are

$$l_{ij}^*(q_{ij}) = \frac{1}{\varphi_{ij}} \left(\frac{\alpha_j^l}{w} \right)^{-\frac{1}{\rho-1}} \left[(\alpha_j^l)^{-\frac{1}{\rho-1}} w^{\frac{\rho}{\rho-1}} + (\alpha_j^m)^{-\frac{1}{\rho-1}} \tau_i^{\frac{\rho}{\rho-1}} \right]^{-\frac{1}{\rho}} q_{ij} \quad (49)$$

$$m_{ij}^*(q_{ij}) = \frac{1}{\varphi_{ij}} \left(\frac{\alpha_j^m}{\tau_i} \right)^{-\frac{1}{\rho-1}} \left[(\alpha_j^l)^{-\frac{1}{\rho-1}} w^{\frac{\rho}{\rho-1}} + (\alpha_j^m)^{-\frac{1}{\rho-1}} \tau_i^{\frac{\rho}{\rho-1}} \right]^{-\frac{1}{\rho}} q_{ij}, \quad (50)$$

which implies constant firm-product specific unit costs

$$c_{ij} = \frac{1}{\varphi_{ij}} \left[(\alpha_j^l)^{-\frac{1}{\rho-1}} w^{\frac{\rho}{\rho-1}} + (\alpha_j^m)^{-\frac{1}{\rho-1}} \tau_i^{\frac{\rho}{\rho-1}} \right]^{\frac{\rho-1}{\rho}}. \quad (51)$$

The firm's profit maximisation problem specified in Equation 9 consists of maximising a strictly concave function without constraints, so the first-order conditions pin down the unique maximiser. As in a standard CES demand case, the profit-maximising price is

$$p_{ij}^* = \frac{\sigma_j}{\sigma_j - 1} c_{ij}, \quad (52)$$

which results in unit sales of

$$q_{ij}^* = \left(\frac{\sigma_j - 1}{\sigma_j} \right)^{\sigma_j} [(1 - \delta_j) \gamma_j]^{\frac{1}{\delta_j}} P_j^{\sigma_j - \frac{1}{\delta_j}} c_{ij}^{-\sigma_j} \quad (53)$$

and revenue

$$r_{ij}^* = \left(\frac{\sigma_j - 1}{\sigma_j} \right)^{\sigma_j - 1} [(1 - \delta_j) \gamma_j]^{\frac{1}{\delta_j}} P_j^{\sigma_j - \frac{1}{\delta_j}} c_{ij}^{1 - \sigma_j}. \quad (54)$$

Optimal profits are then

$$\pi_{ij}^* = \frac{(\sigma_j - 1)^{\sigma_j - 1}}{\sigma_j^{\sigma_j}} [(1 - \delta_j) \gamma_j]^{\frac{1}{\delta_j}} P_j^{\sigma_j - \frac{1}{\delta_j}} c_{ij}^{1 - \sigma_j}. \quad (55)$$

Distributions and the CES price index

Firms draw their productivities for each product independently from a Pareto($k, 1$) distribution with cdf

$$G(\varphi_{ij}) = 1 - \varphi_{ij}^{-k} \quad (56)$$

and pdf

$$g(\varphi_{ij}) = k \varphi_{ij}^{-k-1} \quad (57)$$

over a support of $[1, \infty)$. Obviously, this means that firm-product-specific productivities for each product are continuously distributed with a strictly decreasing pdf over the entire support.

Firms draw their energy prices independently from any firm-product productivities from a distribution with cdf H and strictly positive support. Hence, for any firm i and product j ,

$$\tau_i \mid \varphi_{ij} \sim H. \quad (58)$$

We only impose the regularity condition that $\mathbb{E}[\tau_i^{1-\sigma_j}] = \int_0^\infty \tau^{1-\sigma_j} dH(\tau) < \infty$ for all products j . This is a sufficient condition for

$$\mathbb{E}[\tilde{c}_j(\tau_i)^{1-\sigma_j}] = \int_0^\infty \left(\left[(\alpha_j^l)^{-\frac{1}{\rho-1}} w^{\frac{\rho}{\rho-1}} + (\alpha_j^m)^{-\frac{1}{\rho-1}} \tau^{\frac{\rho}{\rho-1}} \right]^{\frac{\rho-1}{\rho}} \right)^{1-\sigma_j} dH(\tau) < \infty \quad \forall j. \quad (59)$$

We can then rewrite the equilibrium condition from Equation 13 as

$$P_j^* = \left[\int_0^\infty \int_{\hat{\varphi}_j(P_j^*, \tau)}^\infty p_j(\varphi, \tau)^{1-\sigma_j} g(\varphi) d\varphi dH(\tau) \right]^{\frac{1}{1-\sigma_j}}, \quad (60)$$

where

$$\hat{\varphi}_j(P_j, \tau) = \left(\frac{F}{A_j} \right)^{\frac{1}{\sigma_j-1}} P_j^{-\xi_j} \tilde{c}_j(\tau) \quad (61)$$

is the productivity threshold above which a firm with energy price τ enters the market for product j given aggregate price level P_j with

$$\xi_j := \frac{\sigma_j - \frac{1}{\delta_j}}{\sigma_j - 1} \in (0, 1) \quad (62)$$

analogous to the definition of ξ in Equation 15 in the main text.

Using our CES pricing condition from Equation 52, $p_j(\varphi, \tau)$ is defined as

$$p_j(\varphi, \tau) = \frac{\sigma_j}{\sigma_j - 1} \frac{1}{\varphi} \tilde{c}_j(\tau), \quad (63)$$

that is, as the price that a firm with energy price τ and firm-product-specific productivity φ charges for its variety of product j . Using the definition of $g(\varphi) = k\varphi^{-k-1}$ and our formula for $p_j(\varphi, \tau)$, we can integrate over our productivity distribution and further simplify the equilibrium condition to

$$P_j^* = \frac{\sigma_j}{\sigma_j - 1} \left[\frac{k}{k - (\sigma_j - 1)} \int_0^\infty \hat{\varphi}_j(P_j^*, \tau)^{\sigma_j-1-k} \cdot \tilde{c}_j(\tau)^{1-\sigma_j} dH(\tau) \right]^{\frac{1}{1-\sigma_j}}. \quad (64)$$

Note that integration is possible because our assumption that $k > \sigma_j - 1 \forall j$ implies that $\int_{\underline{\varphi}}^\infty \varphi^{\sigma_j-k-1} d\varphi$ is finite. We can then use the definition of $\hat{\varphi}$ to obtain the expression for the equilibrium price given in Equation 14 in the main text,

$$P_j^* = \left[\frac{k - (\sigma_j - 1)}{k} \left(\frac{\sigma_j}{\sigma_j - 1} \right)^{\sigma_j-1} \left(\frac{F}{A_j} \right)^{\frac{k}{\sigma_j-1}-1} \left(\int_0^\infty \tilde{c}_j(\tau)^{-k} dH(\tau) \right)^{-1} \right]^{\frac{1}{\xi_j k + (1-\xi_j)(\sigma_j-1)}} \quad (65)$$

Note that the integral on the right-hand side must exist and be finite by $k > \sigma_j - 1$: This implies that $\tilde{c}_j(\tau)^{-k}$ is asymptotically dominated by $\tilde{c}_j(\tau)^{1-\sigma_j}$. Since $\int_0^\infty \tilde{c}_j(\tau)^{1-\sigma_j} dH(\tau)$ is finite by assumption, the same must be true for $\int_0^\infty \tilde{c}_j(\tau)^{-k} dH(\tau)$.

Existence and uniqueness of product market equilibrium

To show that a unique equilibrium for each product market must exist, we first need to prove that the distribution of firm-specific unit cost c_{ij} is continuous under an arbitrary distribution of energy prices, H .

Lemma 3: The distribution of firm-specific unit costs c_{ij} for producing product j is continuous over its entire support $[0, \bar{c}_j)$ with $\bar{c}_j > 0$ and $\bar{c}_j \rightarrow \infty$ if the support of the energy price distribution H has no upper limit.

Proof: By our assumptions in Section 3.1, firm-product productivity φ_{ij} is distributed according to a continuous distribution with cdf G over a support range $[1, \infty)$. Energy prices τ_i are

distributed according to an arbitrary distribution H with strictly positive support, and firm-product specific costs are given by

$$c_{ij} = \frac{1}{\varphi_{ij}} \left[(\alpha_j^l)^{-\frac{1}{\rho-1}} w^{\frac{\rho}{\rho-1}} + (\alpha_j^m)^{-\frac{1}{\rho-1}} \tau_i^{\frac{\rho}{\rho-1}} \right]^{\frac{\rho-1}{\rho}} = \frac{1}{\varphi_{ij}} \tilde{c}_j(\tau_i). \quad (66)$$

Let us define the distribution of firm-product specific costs for product j as Z_j such that $c_{ij} \sim Z_j$. We begin by showing that Z_j must be continuous. This is equivalent to saying that $\mathbb{P}(c_{ij} = c) = 0 \forall c \in [0, \bar{c})$, i.e. that the distribution does not have any mass point.

To see that this must be the case, consider any $c \in [0, \bar{c})$ and note that, by the law of total probability,

$$\begin{aligned} \mathbb{P}(c_{ij} = c) &= \int_1^\infty \mathbb{P}(c_{ij} = c \mid \varphi_{ij} = \varphi) dG(\varphi) \\ &= \int_1^\infty \mathbb{P} \left(\tau_i = \left[\left(\frac{1}{\alpha_j^m} \right)^{-\frac{1}{\rho-1}} (\varphi c)^{\frac{\rho}{\rho-1}} + \left(\frac{\alpha_j^l}{\alpha_j^m} \right)^{-\frac{1}{\rho-1}} w^{\frac{\rho}{\rho-1}} \right] \right) g(\varphi) d\varphi, \end{aligned} \quad (67)$$

where we have used the continuity of G to write $dG(\varphi)$ as $g(\varphi)d\varphi$. Now, if H has no mass points, then $\mathbb{P}(\tau_i = \tau) = 0$ for any τ and hence $\mathbb{P}(c_{ij} = c) = 0$ as well. If H has a non-empty set of mass points M_H , then for any $\tau \in M_H$, $0 < \mathbb{P}(\tau_i = \tau) < 1$, and $\mathbb{P}(\tau_i = \tau) = 0$ for any $\tau \notin M_H$. Trivially, the set of mass points has to be finite as otherwise the total probability mass of the distribution would be infinite. But a Riemann integral over a function with a finite set of finite non-zero points is still zero, so $\mathbb{P}(c_{ij} = c) = 0$ in this case as well.

Having established that Z_j is continuous, we can proceed to show that it must have support $[0, \bar{c})$. Trivially, H must have non-empty support $\mathcal{H}$, where we know by assumption that $\mathcal{H} \subset \mathbb{R}_+$. Using our expression for $\tilde{c}_j$ from above, this clearly means that the support of $\tilde{c}_j(\tau_{ij})$, $\mathcal{C}$, must be a strict subset of $\mathbb{R}_+$ as well. Now, as G has support $[1, \infty)$, the distribution of $\frac{1}{\varphi_{ij}}$ must have support $(0, 1]$. This clearly implies that the lower bound of the support of $c_{ij} = \frac{1}{\varphi_{ij}} \tilde{c}_j$ must be 0.

What is more, it implies that the support of c_{ij} must cover the entire range from 0 to

$$\bar{c} := \max_{\tau \in \mathcal{H}} \tilde{c}_j(\tau). \quad (68)$$

If $\sup \mathcal{H} = \infty$, this means that $\bar{c}_j \rightarrow \infty$ because $\lim_{\tau \rightarrow \infty} \tilde{c}_j(\tau) = \infty$. If $\sup \mathcal{H} = \bar{\tau}$ for some $\bar{\tau} \in \mathbb{R}_+$, $\bar{c}_j < \infty$. $\square$

Given Lemma 3, it is then straightforward to prove that each product market must have a unique equilibrium.

Lemma 4: The market for each product j is cleared by a unique equilibrium price index $P_j^* < \infty$.

Proof: We can use the cost cutoff $\hat{c}_j$, the continuity of the cost distribution established in Lemma 3 and firms' optimal pricing condition to rewrite the equilibrium condition in Equation 13 as

$$P_j^* = \frac{\sigma_j}{\sigma_j - 1} \left[\int_0^{\hat{c}_j} (P_j^*) c^{1-\sigma_j} f(c) dc \right]^{\frac{1}{1-\sigma_j}}, \quad (69)$$

where f is the pdf of the cost distribution.

From this equation, it becomes clear that a unique equilibrium price index $P_j^* < \infty$ must exist. Obviously, the left-hand side of the equation is continuous and strictly increasing in P_j^* , is zero

if $P_j^* = 0$ and tends to infinity as $P_j^* \rightarrow \infty$. The right-hand side of the equation is continuous in P_j^* as well because $\hat{c}_j$ is a continuous function of P_j^* . It tends to infinity as $P_j^* \rightarrow 0$, because then $\hat{c}_j \rightarrow 0$. It is weakly decreasing in P_j^* because $\hat{c}_j$ is strictly positive and is strictly increasing in P_j^* , and $f(c) > 0$ over the non-empty interval $[0, \hat{c}_j)$. And it converges to the minimum possible price index $\underline{P}_j := \frac{\sigma_j}{\sigma_j-1} \left[\int_{i \in [0,1]} c_{ij}^{1-\sigma_j} di \right]^{\frac{1}{1-\sigma_j}}$ as $P_j^* \rightarrow \infty$ because then $\hat{c}_j \rightarrow \infty$ and hence all firms will enter the market. $\square$

Revenue and energy intensity in equilibrium

We can write the conditional average revenue of firms with energy price τ from product j given aggregate price index P_j as

$$R_j(P_j, \tau) := A_j P_j^{\sigma_j - \frac{1}{\delta_j}} \tilde{c}_j(\tau)^{1-\sigma_j} \int_{\hat{\varphi}_j(P_j, \tau)}^{\infty} \varphi^{\sigma_j-1} g(\varphi) d\varphi \quad (70)$$

relying on our expression for optimal firm-level revenue from Equation 54. Using the fact that $g(\varphi) = k\varphi^{-k-1}$, we can rewrite this as

$$R_j(P_j, \tau) = \frac{k}{k - (\sigma_j - 1)} A_j P_j^{\sigma_j - \frac{1}{\delta_j}} \tilde{c}_j(\tau)^{1-\sigma_j} \hat{\varphi}_j(P_j, \tau)^{(\sigma_j-1)-k}, \quad (71)$$

which boils down to the expression given in Equation 24 in the main text,

$$R_j(P_j, \tau) := \frac{k}{k - (\sigma_j - 1)} A_j^k F^{-\left(\frac{k}{\sigma_j-1}-1\right)} \tilde{c}_j(\tau)^{-k} P_j^{\xi_j k}. \quad (72)$$

Aggregate revenue from product j conditional on product price index P_j is then given by

$$\begin{aligned} R_j(P_j) &:= \int_0^\infty R_j(P_j, \tau) dH(\tau) \\ &= \frac{k}{k - (\sigma_j - 1)} A_j^k F^{-\left(\frac{k}{\sigma_j-1}-1\right)} P_j^{\xi_j k} \int_0^\infty \tilde{c}_j(\tau)^{-k} dH(\tau). \end{aligned} \quad (73)$$

The energy intensity of product j is then

$$\eta_j := \frac{\int_0^\infty R_j(P_j, \tau) \frac{\sigma_j-1}{\sigma_j} \frac{s_j(\tau)}{\tau} dH(\tau)}{R_j(P_j)} = \frac{\sigma_j-1}{\sigma_j} \frac{\int_0^\infty \tilde{c}_j(\tau)^{-k} \frac{s_j(\tau)}{\tau} dH(\tau)}{\int_0^\infty \tilde{c}_j(\tau)^{-k} dH(\tau)} \quad (74)$$

and depends only on the cost structure of the product and the energy price distribution. We can rewrite it in a more intuitive form as

$$\eta_j = \frac{\sigma_j-1}{\sigma_j} \int_0^\infty \frac{s_j(\tau)}{\tau} d\Theta_j(\tau) \quad (75)$$

where

$$d\Theta_j(\tau) = \frac{\tilde{c}_j(\tau)^{-k} dH(\tau)}{\int_0^\infty \tilde{c}_j(\tau')^{-k} dH(\tau')} = \frac{R_j(P_j, \tau) dH(\tau)}{\int_0^\infty R_j(P_j, \tau') dH(\tau')} \quad (76)$$

is the ‘revenue share’ of firms with energy price τ in product j . (Of course, if H is a continuous distribution, it will technically be the revenue-weighted density, not a share.)

Plugging the equilibrium price index P_j^* into our expression for $R_j(P_j)$, we can see that equilibrium revenue from product j must be equal to

$$R_j^* = \left(\frac{\sigma_j}{\sigma_j - 1} \right)^{\zeta_j(\sigma_j - 1)} A_j^{k+1 - \frac{k\zeta_j}{\sigma_j - 1}} \left[\left(\frac{k}{k - \sigma_j - 1} \right) F^{-\left(\frac{k}{\sigma_j - 1} - 1\right)} \int_0^\infty \tilde{c}_j(\tau)^{-k} dH(\tau) \right]^{1 - \zeta_j} \quad (77)$$

with

$$\zeta_j := \frac{\xi_j k}{\xi_j k + (1 - \xi_j)(\sigma_j - 1)} \in (0, 1). \quad (78)$$

Comparative statics with respect to α_j^m

In the most general form, net-of-productivity costs $\tilde{c}_j(\tau)$ may increase or decrease given an infinitesimal increase in α_j^m , as

$$\frac{d\tilde{c}_j(\tau)}{d\alpha_j^m} = \frac{-1}{\rho} \tilde{c}_j(\tau) s_j(\tau) \frac{\frac{d\alpha_j^l}{d\alpha_j^m} (\alpha_j^l)^{-\frac{\rho}{\rho-1}} w^{\frac{\rho}{\rho-1}} + (\alpha_j^m)^{-\frac{\rho}{\rho-1}} \tau^{\frac{\rho}{\rho-1}}}{(\alpha_j^m)^{-\frac{1}{\rho-1}} \tau^{\frac{\rho}{\rho-1}}} \quad (79)$$

is positive if and only if

$$\left(\frac{\alpha_j^l}{\alpha_j^m} \frac{\tau}{w} \right)^{\frac{\rho}{\rho-1}} > -\frac{d\alpha_j^l}{d\alpha_j^m}. \quad (80)$$

This clearly is the case if α_j^l is fixed, in which case $\frac{d\alpha_j^l}{d\alpha_j^m} = 0$. However, since we assume that the cost shares add up to 1 for each product, i.e. that $\alpha_j^l = 1 - \alpha_j^m$,

$$\frac{d\tilde{c}_j(\tau)}{d\alpha_j^m} = \frac{-1}{\rho} \tilde{c}_j(\tau) s_j(\tau) \frac{1}{\alpha_j^m} \underbrace{\left[1 - \left(\frac{\alpha_j^m w}{(1 - \alpha_j^m) \tau} \right)^{\frac{\rho}{\rho-1}} \right]}_{=: z_j(\tau)}, \quad (81)$$

which clearly is positive if and only if

$$\tau > \frac{\alpha_j^m}{1 - \alpha_j^m} w. \quad (82)$$

Since equilibrium product revenue R_j^* depends negatively on aggregate costs, this means that an increase in α_j^m has an indeterminate effect on revenue, as

$$\frac{dR_j^*}{d\alpha_j^m} = -(1 - \zeta_j) k \frac{-1}{\rho} \frac{1}{\alpha_j^m} R_j^* \int_0^\infty s_j(\tau) z_j(\tau) d\Theta_j(\tau). \quad (83)$$

C.2 Results under stylised energy distribution

Comparative statics with respect to ν and α_j^m

Based on our definition of $R_j(P_j, \tau)$ from Equation 24 and our definition of θ_j , it is straightforward to show that

$$\theta_j := \frac{pR_j(P_j, \tau)}{pR_j(P_j, \tau) + (1-p)R_j(P_j, \nu\tau)} = \frac{p}{p + (1-p)\left(\frac{\hat{c}_j(\tau)}{\hat{c}_j(\nu\tau)}\right)^k} = d\Theta_j(\tau), \quad (84)$$

which is only a function of ν and τ . Simple application of the chain rule then yields the result from Equation 29 in the main text that

$$\frac{d\theta_j}{d\nu} = k\theta_j(1-\theta_j)\frac{s_j(\nu\tau)}{\nu}. \quad (85)$$

Similarly, some simple albeit tedious algebra reveals that

$$\frac{dR_j^*}{d\nu} = -k(1-\zeta_j)\frac{R_j^*}{\nu}(1-\theta_j)s_j(\nu\tau) \quad (86)$$

where $R_j^* = R_j(P_j^*)$ is equilibrium revenue from product j . For the cross derivative $\frac{d^2 R_j^*}{d\alpha_j^m d\nu}$, it helps to note that

$$\frac{d(1-\theta_j)}{d\alpha_j^m} = -\frac{d\theta_j}{d\alpha_j^m} = -\theta_j(1-\theta_j)\frac{-1}{\rho}\frac{1}{\alpha_j^m}k(s_j(\nu\tau)z_j(\nu\tau) - s_j(\tau)z_j(\tau)) \quad (87)$$

as well as

$$\frac{ds_j(\tau)}{d\alpha_j^m} = \frac{-1}{\rho-1}\frac{s(\tau)(1-s(\tau))}{\alpha_j^m(1-\alpha_j^m)} \quad (88)$$

and

$$\frac{dR_j^*}{d\alpha_j^m} = -k(1-\zeta_j)\frac{-1}{\rho}\frac{1}{\alpha_j^m}R_j^*[\theta_j s_j(\tau)z_j(\tau) + (1-\theta_j)s_j(\nu\tau)z_j(\nu\tau)]. \quad (89)$$

Taken together, this reveals – after some simple but tedious algebra – that

$$\begin{aligned} \frac{d^2 R_j^*}{d\nu d\alpha_j^m} = & -(1-\zeta_j)k\frac{-1}{\rho}\frac{1}{\alpha_j^m}\frac{R_j^*}{\nu}(1-\theta_j)s_j(\nu\tau) \cdot \\ & \left[k\zeta_j\theta_j s_j(\tau)z_j(\tau) - \left(k(1-\zeta_j(1-\theta_j)) + \frac{\rho}{\rho-1} \right) s_j(\nu\tau)z_j(\nu\tau) - \frac{\rho}{\rho-1} \right]. \end{aligned} \quad (90)$$

Clearly, this is negative – and so the change in revenue due to an increase in ν is decreasing in α_j^m – if and only if

$$k\zeta_j\theta_j s_j(\tau)z_j(\tau) > \left(k(1-\zeta_j(1-\theta_j)) + \frac{\rho}{\rho-1} \right) s_j(\nu\tau)z_j(\nu\tau) + \frac{\rho}{\rho-1}. \quad (91)$$

D Model calibration

To simulate the model, we have used the fundamental parameters described in Table 4.

Table 4: Fundamental parameters of simulated model

Parameter	Value	Comment
σ	4.5	As in Shapiro and Walker (2018)
γ	3	Calibrated to get realistic number of active firms
δ	0.33	Calibrated to get realistic number of active firms
F	1.5	Calibrated to get realistic number of active firms
k	5.4	Estimated from AFiD data (von Graevenitz et al., 2024)
μ	2	Calibrated to get realistic distribution of # of products per firm
ρ	-0.66	Taken from Bretschger and Jo (2024); corresponds to an elasticity of substitution of 0.6 between energy and labour
w	1	Normalised to 1

We have drawn product-level energy intensities α_j^m from a Beta(28, 400) distribution, resulting in the distribution of baseline energy intensities depicted in Figure 10.




Figure 10: Distribution of product-level energy intensities in simulated model

Table 5: Regressions of product adding and dropping on firm energy price and product energy intensity

	(1)	(2)
Dep. variable	Adding	Dropping
prod. ener. int. (kWh/EUR)	-5.5×10^{-5} ($< 2.7 \times 10^{-5}$)	0.018 (0.009)
log firm ener. price (EUR/kWh)	-1.5×10^{-5} ($< 0.7 \times 10^{-5}$)	-0.020 (0.004)
prod. ener. int. $\times$ log firm ener. price	-2.2×10^{-5} ($< 1.1 \times 10^{-5}$)	0.004 (0.003)
constant	10.2×10^{-5} ($< 5.1 \times 10^{-5}$)	0.038 (0.008)
<i>N</i>	18 007 162	35 306

Note: The regression in column (1) is based on the (pooled) universe of possible firm-product combinations over the two six year intervals 2005-2011 and 2011-2017 that are *not* produced in the interval's base year. The regression in column (2) is based on the (pooled) universe of firm-product combinations over the same two intervals that *are* produced in the interval's base year. All coefficients in column (1) are significant at the 5 percent level but exact standard errors are irrecoverable for data confidentiality reasons.

Source: AFiD manufacturing census (see Appendix A)

E Additional empirical results

The regression model underlying the results reported in Table 5 is

$$D_{ijt} = \alpha + \beta_1 \text{prodenerint}_{jt} + \beta_2 \ln \text{enerprice}_{it} + \beta_3 \text{prodenerint}_{jt} \times \ln \text{enerprice}_{it} + \varepsilon_{ijt}, \quad (92)$$

where D_{ijt} is a dummy that indicates whether or not firm i added or dropped product j between the period's base year t and the base year of the next period, t' .

Firm-product level regressions reveal that the higher add and drop rates among energy-intensive products stem from the fact that absolute adding and dropping probabilities decrease minimally in product energy intensity, but energy-intensive products are produced by far fewer firms. This can be seen in Table 6: Column (1) reveals that the probability that a firm produces a specific product decreases substantially in the product's energy intensity. Columns (2) and (3) make clear that the absolute probability of adding or dropping a product slightly decreases in its energy intensity as well, but at a much slower rate.²⁸


²⁸For example, the ratio of the probability that a firm adds a product to the baseline probability that a firm produces the same product is $\frac{11.5-1.3 \cdot 0.5}{200.9-30.7 \cdot 0.5} = 0.058$ for a product with a low energy intensity of 0.5 kWh/EUR, but $\frac{11.5-1.3 \cdot 3}{200.9-30.7 \cdot 3} = 0.070$ for a product with a rather high energy intensity of 3 kWh/EUR.

Table 6: Regressions of producing, adding and dropping on product energy intensity

	(1)	(2)	(3)
Dep. variable	Producing	Adding	Dropping
energy intensity (kWh/EUR)	-30.7×10^{-5} ($< 15.3 \times 10^{-5}$)	-1.3×10^{-5} ($< 0.6 \times 10^{-5}$)	-1.7×10^{-5} ($< 0.8 \times 10^{-5}$)
constant	200.9×10^{-5} ($< 100.4 \times 10^{-5}$)	11.5×10^{-5} ($< 5.7 \times 10^{-5}$)	14.6×10^{-5} ($< 7.3 \times 10^{-5}$)
<i>N</i>	42 581 251	42 581 251	42 581 251

Note: Regressions are based on the (pooled) universe of possible firm-product combinations over the two six year intervals 2005-2011 and 2011-2017. All coefficients are significant at the 5 percent level but exact standard errors are irrecoverable for data confidentiality reasons.

Source: AFiD manufacturing census (see Appendix A)

**Figure 11:** Relationship between firm size and energy price

Note: Bin scatter plot based on the subsample of $N = 35\,808$ firm-year combinations in 2005, 2011 and 2017 for which energy price data is available. All values are net of 4-digit main industry fixed effects.

Source: AFiD manufacturing census (see Appendix A)

Table 7: Regressions of energy price on productivity

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dep. variable	log energy price						
Prod. measure	OLS	CS	W2009	Labour	Energy	Capital	Materials
log productivity	0.123 (0.009)	0.072 (0.01)	−0.001 (0.001)	0.035 (0.007)	0.263 (0.004)	0.059 (0.004)	−0.045 (0.008)
log real sales value	−0.104 (0.004)	−0.100 (0.004)	−0.083 (0.003)	−0.099 (0.004)	−0.112 (0.003)	−0.093 (0.003)	−0.080 (0.003)
constant	−0.463 (0.057)	−0.530 (0.063)	−0.811 (0.051)	−0.957 (0.061)	−0.768 (0.044)	−0.717 (0.051)	−0.831 (0.052)
<i>N</i>	34 914	34 914	34 887	35 299	35 304	34 954	35 268

Note: Based on pooled firm-year observations from 2005, 2011 and 2017. Standard errors clustered at the firm level. All productivities are measured in terms of revenue. OLS refers to the residuals of a OLS regression of log revenue on the four factors of production (in log terms). CS refers to total factor productivity with individual factor productivities weighted by cost-share based elasticity estimates. W2009 refers to total factor productivity estimated as in Wooldridge (2009). Individual

factor productivities are calculated as $\log \frac{\text{revenue}}{\text{factor use}}$.
Source: AFiD manufacturing census (see Appendix A)



Download ZEW Discussion Papers:

<https://www.zew.de/en/publications/zew-discussion-papers>

or see:

<https://www.ssrn.com/link/ZEW-Ctr-Euro-Econ-Research.html>

<https://ideas.repec.org/s/zbw/zewdip.html>



IMPRINT

**ZEW – Leibniz-Zentrum für Europäische
Wirtschaftsforschung GmbH Mannheim**

ZEW – Leibniz Centre for European
Economic Research

L 7,1 · 68161 Mannheim · Germany

Phone +49 621 1235-01

info@zew.de · zew.de

Discussion Papers are intended to make results of ZEW research promptly available to other economists in order to encourage discussion and suggestions for revisions. The authors are solely responsible for the contents which do not necessarily represent the opinion of the ZEW.