Make or buy? Allocative energy efficiency in production networks

Carole Marullaz* University of Zurich

Giacomo Romanini† Banca d'Italia

October 16, 2026

Updated frequently - Click here for the most recent version

Abstract

We develop a model of make-or-buy decisions in which firms choose between in-house production and alternative suppliers. Producer selection trades off efficiency in the use of production factors against relationship-specific productivity gains from outsourcing. When energy prices do not reflect the social cost of emissions, firms give too little weight to energy efficiency, distorting who produces each input – a new extensive margin of energy-use misallocation. The effects on aggregate energy intensity and welfare depend on the joint distribution of energy efficiency and outsourcing productivity along supply chains, as choosing a supplier also selects its own upstream producers. Moreover, a firm's observed energy intensity confounds its energy efficiency with which inputs it produces in-house. We therefore embed make-or-buy decisions in a sequential production network calibrated to Italian firm-to-firm transaction data and develop a method to infer each firm's unobserved in-house input set. When emissions are priced at their social cost, changes in producer selection reduce aggregate energy intensity by 12%, a fifth of the total gains.

*Corresponding author. carole.marullaz@econ.uzh.ch. Carole Marullaz gratefully acknowledges the support provided by the University of Zurich's Research Priority Program 'Equality of Opportunity'.

†giacomo.romanini@bancaditalia.it. The views expressed in this project do not necessarily reflect those of Banca d'Italia or the Eurosystem.

1 Introduction

Whether production is allocated to the most efficient producers is a central determinant of aggregate outcomes, from productivity to emissions. A large literature studies how distortions misallocate resources across active sellers, but typically takes as given who produces what and for whom. In reality, misallocation can also occur at the extensive margin, when firms decide whether to produce an input in-house or buy it, and from which supplier. This makes allocative efficiency a question of producer selection: which firm, among all potential producers, actually produces each input for each firm using it. We develop a unified framework to study the allocative role of this selection by comparing all potential producers jointly, and apply it to energy, for which this margin is particularly relevant. Most of the energy used in manufacturing goes into producing intermediate inputs (Shapiro, 2021), so make-or-buy decisions, by determining which firm produces each input, allocate energy demand across producers.[1] This allocation matters because comparable firms are highly heterogeneous in energy efficiency (Lyubich et al., 2018; De Lyon & Dechezleprêtre, 2025): the production of the same input requires more or less energy depending on who produces it. Moreover, the emission externality distorts energy prices absent any policy intervention.

Firms select their producers to minimize costs, but the most energy-efficient producer of an input is not necessarily the cheapest. Producer selection therefore involves a trade-off between energy efficiency and other cost advantages, and the price of energy sets its terms. When this price does not reflect the social cost of emissions, firms undervalue energy efficiency, so that their decisions might differ from efficient ones and misallocate energy use across potential producers at the extensive margin. Optimal carbon pricing can therefore save energy not only at the intensive margin, as producers use less energy and cleaner ones gain market share, but also by shifting producer selection toward more energy-efficient potential producers.

Despite a rich literature on firm boundaries (e.g., Antràs and Rossi-Hansberg (2009) and Boehm and Oberfield (2020)), we lack both a theoretical framework connecting make-or-buy decisions to the allocation of energy, and empirical evidence on whether firms select energy-efficient producers. We fill this gap on both fronts. Theoretically, we formalize this trade-off in a model of firms' make-or-buy decisions, in which producers differ in energy efficiency and in outsourcing capabilities, the relationship-specific productivity of outsourcing relative to in-house production. We show that the magnitude of the resulting energy-use misallocation from producer selection depends on their joint distribution, which determines both the impact and the frequency of inefficient decisions. Empirically, we use the universe of Italian firm-to-firm transactions and customs records to document how producer selection relates to producers' energy efficiency and to quantify the environmental costs of producer selection and the corresponding gains from carbon pricing.

Doing so requires identifying (i) which firm *actually* produces each input, which is unobserved whenever the input is produced in-house, and (ii) which firm *should* optimally produce it, which requires com-

[1] We use here make-or-buy decision in a broad sense as the decision to produce in-house or to outsource and, in the latter case, to which supplier. Together, these decisions determine producer selection.

paring all potential producers by their energy efficiency and outsourcing capabilities. Firm boundaries complicate the second step as well. Figure 1 plots Italian firms' intermediate expenditure shares, a proxy for outsourcing, against their energy shares relative to their suppliers'. Firms that outsource more are systematically less energy intensive than their suppliers, suggesting substantial scope for energy-use misallocation through producer selection. However, outsourcing moves the energy used at outsourced stages outside the buyer's boundary, so a firm may appear energy intensive simply because it produces energy-demanding inputs in-house, even if it does so more efficiently than its potential suppliers. Recovering firms' effective energy efficiency therefore requires explicitly accounting for make-or-buy decisions.

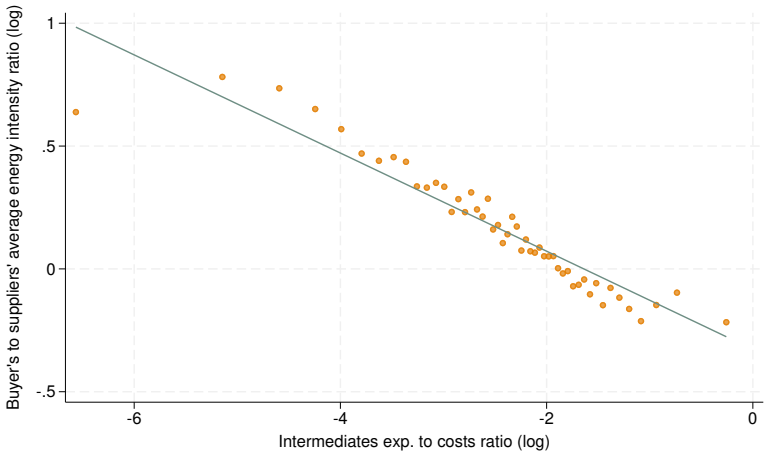


Figure 1: Outsourcing and buyer's relative energy intensity

Notes: The figure plots the relation between energy intensity, measured by energy expenditures per sales (in logarithm) and intermediate input expenditures to total costs (in logarithm). On the $y-$axis is buyer's energy intensity relative to the average energy intensity of its suppliers. The suppliers' average energy intensity is computed weighting by transaction shares. The relationship is measured controlling for product (6-digit), for the sample of manufacturing Italian firms in 2019.

We thus embed our model in a sequential production network as in Boehm and Oberfield (2023), where every firm faces a make-or-buy decision for each input. We depart from their framework in characterizing producers along two dimensions of technological heterogeneity that we must identify in the data. Because outsourcing an input leaves its upstream production to the supplier's own make-or-buy decisions, comparing potential producers requires recovering these two dimensions for every firm along the supply chain, not for the buyer and its suppliers in isolation. To do so, we calibrate the model exactly, in the spirit of Di Giovanni et al. (2024): each model firm is an actual Italian manufacturing firm, and the calibrated model reproduces every make-or-buy decision we recover from the data, without parametric restrictions on how energy efficiency and outsourcing capabilities are jointly distributed across firms. Our counterfactuals can then trace how each of these decisions responds to carbon pricing.

Central to the calibration is a novel method to infer each firm's unobserved in-house input set. Exploiting the sequential production structure, we infer that all required inputs lying downstream of a firm's purchased inputs must have been produced in-house.[2] For instance, from the universe of firms' transactions,

[2]We first use the universe of firm-to-firm transactions to construct the national, product-level analogue of the total requirements matrix from standard input–output tables, following Boehm and Oberfield (2020), Rachapalli (2021), and Boehm and Oberfield (2023), and use it to build product "recipes", i.e., all the crucial inputs required to produce that product, as a directed acyclic graph to capture the sequentiality of production. The recipes also determine the comparison set: they give us all buyer-input pairs for which a make-or-buy decision can occur and, for each of these pairs, the set of

we conclude that mortar production requires cement, and that cement production requires limestone. If we then observe a mortar producer buying limestone but not cement, we can infer that this producer makes its cement in-house. Beyond answering (i), the in-house set also addresses the confound in (ii), allowing us to recover energy efficiency net of firm boundaries.

Comparing actual to potential producers, we find substantial energy efficiency gaps: for roughly 90% of the inputs for which we observe a make-or-buy decision, the selected producer is at most half as energy-efficient as the most energy-efficient alternative available. However, these gaps need not reflect misallocation, as they may be offset by gains in outsourcing capability. We thus complete the identification by recovering relationship-specific outsourcing capabilities through a revealed preference argument. Intuitively, if a firm produces an input in-house, its outsourcing capability with every potential supplier of that input must be low enough to rationalize that choice. Each observed make-or-buy decision therefore bounds the outsourcing capabilities of the pairs involved, and we use these bounds to discipline the distribution from which they are drawn.

Equipped with these objects, we characterize the efficient allocation and quantify the welfare loss from energy-use misallocation in Italian manufacturing for 2019. Pricing emissions at the social cost of carbon restores the efficient allocation, lowers aggregate energy intensity by 47.6% and raises welfare by 2.10%. Correcting the misallocation from producer selection accounts for 19.1% and 10.7% of these gains. These contributions rest on about a fifth of the firms changing producer for at least one input, and this corresponds to switching only 3.5% of all decisions. The extensive margin operates through the impact of switches more than their frequency, because the pairs that move first are those where reallocation saves the most energy at the least cost. Producer selection is thus a margin of adjustment available to policy that delivers large emission reductions even when no firm improves its own technology. It also matters for evaluating these policies. For a 10% increase in the energy price, it raises the elasticity of substitution between energy and labor from 0.87 to 1.12, across the unit value typically assumed in macro-climate models, so that energy and labor become aggregate substitutes. Finally, we use our model to measure the effects of sector-specific carbon prices, showing that firms respond by selecting producers outside the covered sectors, a domestic form of carbon leakage that erodes the gains from the policy.

Relation to existing literature. Our paper builds on the seminal between-firm misallocation literature initiated by Restuccia and Rogerson (2008) and Hsieh and Klenow (2009), which quantifies the aggregate allocative consequences of distortions when firms are heterogeneous. Closest to our perspective are Klenow et al. (2024) and Kim (2025), who both show that a carbon tax can generate productivity gains by shifting demand toward less fossil-fuel-intensive firms, because, as they document, these firms also have higher marginal revenue productivity. We extend this logic in two directions. First, we directly quantify how suboptimal energy prices distort the allocation of energy use across producers at the extensive margin. Second, by explicitly incorporating in-house production, cross-firm reallocation operates across all producers of a given input, including the downstream firm itself, and we are able to identify

all potential producers, namely the buyer itself and all the external suppliers of that input.

energy efficiency net of firm boundaries.

We also speak to the literature studying how make-or-buy decisions might lead to misallocation and their macroeconomic consequences. Fally and Hillberry (2018), Chor and Ma (2020), Boehm and Oberfield (2020), and Boehm (2022) study the extent to which transaction costs and contracting frictions shape input sourcing patterns and aggregate productivity. We instead show how an aggregate distortion, the energy externality, also alters make-or-buy decisions. Characterizing every potential producer of each input by its efficiency, we can quantify the resulting misallocation jointly across sellers and along firm boundaries.[3]

Despite the importance of intermediate input production in firms' total energy demand, studying environmental distortions through the lens of firm boundaries is, to our knowledge, novel. The theoretical trade-environment literature has established that intermediate input outsourcing acts as a substitute for domestic energy use, and that firms facing high energy prices tend to outsource energy-intensive inputs abroad (Cole et al., 2014; Schenker et al., 2018; Bolz et al., 2024).[4] Yet empirically, energy has been found to play only a modest role in international outsourcing decisions as its price has remained low (Cherniwchan & Taylor, 2022; Dechezleprêtre et al., 2022; Levinson, 2023), consistent with our model's prediction that low energy prices translate into under-weighing energy in make-or-buy decisions. However, this literature abstracts from two features we explicitly incorporate: producer-level heterogeneity in energy efficiency, and in-house production as a concrete alternative to outsourcing. Because more upstream inputs require more energy to produce (Shapiro, 2021), much of a product's energy is used at upstream stages, and firm boundaries determine which producer performs them. We are the first to empirically characterize firms' energy efficiency and boundaries jointly.

Beyond quantifying misallocation, we evaluate the aggregate response of energy use to carbon pricing, which is hard to measure directly. Macro designs (e.g., synthetic control methods) capture it in reduced form, but their results hinge on the specific policy studied (Andersson, 2019; Metcalf & Stock, 2020; Bernard & Kichian, 2021). High-frequency estimates (Känzig, 2023; Ortubai et al., 2025) are less policy-specific, but trace responses over horizons that may be too short for reallocation through make-or-buy decisions, which are sticky (Antràs & Chor, 2013; Monarch, 2014; Martin et al., 2023). Firm-level studies credibly exploit variation across firms (Martin et al., 2014; Dechezleprêtre et al., 2019; Fontagné et al., 2024), but micro elasticities need not aggregate when production reallocates across firms (Oberfield & Raval, 2021). We bridge the two by building the aggregate response to a permanent price change from the micro level, combining firm-level substitution with reallocation across producers at both the intensive

[3]This literature typically focuses on contexts in which one mode, in-house production or outsourcing, is systematically more efficient than the other. When producers are compared on energy efficiency, we find that no such ranking holds: misallocation can run in either direction, through excessive in-house production or excessive outsourcing.

[4]Measuring carbon leakage abroad is a key focus of this literature, but make-or-buy decisions involve mostly domestic partners (Fort, 2017; Boehm & Oberfield, 2020; Fort, 2023). In our data, 87% of transactions are domestic.

and the extensive margin. The contribution of reallocation depends on the heterogeneity across producers, which our joint characterization of firms' energy efficiency and boundaries allows us to discipline.

To compare actual to potential producers in terms of effective energy efficiency, our quantitative exercise bridges the endogenous production-network literature (e.g., Acemoglu and Azar (2020), Bernard et al. (2022), Dhyne et al. (2023), and Arkolakis et al. (2025)) and the literature on firm boundaries (e.g., Grossman and Helpman (2002), Antràs and Chor (2013), and Alfaro et al. (2019)), which usually studies make-or-buy decisions at a single stage. Following Boehm and Oberfield (2023), we embed make-or-buy decisions in a sequential production network. In their framework, firm heterogeneity is specified to deliver tractable aggregate outcomes, and in-house production is disciplined by its extent, as in prior work measuring firms' vertical span (e.g., Antràs et al. (2012), Alfaro et al. (2019), and Boehm and Oberfield (2020)). In ours, heterogeneity is what we seek to measure. We therefore identify, input by input, which stages each firm produces in-house, building on the product-level sequential input-output structure of Boehm and Oberfield (2020) and Rachapalli (2021).[5] Beyond our core analysis, the approach also reveals substantial heterogeneity in make-or-buy choices: firms differ not only in the number of stages they internalize but also in which specific inputs they produce, with potential implications for input-specific policies (e.g., industrial policy, tariffs), whose incidence varies with firms' in-house portfolios.

Outline. The coming section clarifies the concept of energy-use misallocation across potential producers using a simple two-stage model. We describe the quantitative model in Section 3, and, in Section 4, develop a strategy to infer the set of inputs produced in-house and calibrate the effective energy efficiency and the distribution of outsourcing capabilities of each firm from the Italian data. Section 5 quantifies energy-use misallocation and discusses policy applications, and Section 6 concludes.

2 Make-or-buy decisions and the allocation of energy use

We introduce a conceptual framework to illustrate how make-or-buy decisions can induce energy-use misallocation across producers. Because firms can produce some of their inputs in-house, the resources used to produce a given input are allocated not only across its sellers, but across all its potential producers, including the downstream firms that use it. Make-or-buy decisions make this allocation endogenous, and thus subject to the distortion that the emission externality imposes on energy prices.[6]

2.1 Environment

We study a static economy with a fixed number of producers and constant technologies, in which production takes place in two stages: intermediate inputs are produced from energy and labor, and final

[5] In contemporaneous work in progress, Malik (2026) uses a similar algorithm to characterize Indian firms' organizational structure and their responses to tariffs.

[6] The logic extends to any setting in which potential producers differ along two dimensions whose ranking does not exactly agree, and that a distortion weighs differently, as formalized in Appendix A.1.

products are produced from intermediate inputs.

Emission externality and welfare. A representative household supplies labor inelastically, receives utility from the consumption of a good Y, which aggregates all final products, and disutility from energy:[7] $U(Y, E)$, where $\frac{\partial U}{\partial E} < 0$ and the Inada conditions on Y hold. The energy externality creates a distortion so, absent any policy intervention, energy is underpriced because firms bear only its private cost and do not internalize the disutility their emissions impose. Therefore, similar to spatial equilibrium models (e.g., Donald et al. (2025)), but different from the standard misallocation literature (e.g., Hsieh and Klenow (2009) and Baqaee and Farhi (2020)), output gains and welfare gains are different in this model. In fact, because there are no other distortions but the energy externality, output will be maximized.

Intermediate input production. Intermediate inputs $\omega \in \Omega$ are all produced with a CRS production function $F_\omega(e, l; \nu)$ that combines energy e and labor l, and in which ν is the effective energy efficiency (defined formally below) characterizing the producer. F_ω is differentiable, increasing and concave in both energy and labor, and can vary across inputs ω, but its functional form is common to all producers, whose technologies differ only through ν. This symmetry in input production allows us to directly compare producers in terms of their effective energy efficiency. With factor prices p_e and w, the corresponding dual unit cost function is $c_\omega(p_e, w; \nu)$.

Potential producers. There are two types of perfectly competitive firms, and both can produce ω: downstream firms, indexed by $j \in \mathcal{S}$, and upstream firms, indexed by $i \in \mathcal{S}_\omega$. Downstream firms sell their final product j directly to the household facing demand $y_j = D(p_j)$, which they take as given. The production of $y_j = \mathrm{F}(\{x_j(\omega)\}_{\omega \in \Omega})$ combines all intermediate inputs $\omega \in \Omega$, and $\mathrm{F}(.)$ exhibits CRS and is separable, differentiable, increasing and concave in all its arguments. Downstream firms can either produce each intermediate input ω in-house using F_ω and their own effective energy efficiency ν_j, or outsource their production to one upstream supplier selected from the set $\mathcal{S}_\omega$. Upstream suppliers similarly produce with F_ω but use their own effective energy efficiency ν_i, produce only one input ω, and cannot outsource production.

Energy efficiency. ν_k corresponds to the effective energy efficiency of firm $k \in \mathcal{S} \cup \bigcup_{\omega \in \Omega} \mathcal{S}_\omega$. Like any efficiency parameter, a higher ν_k lowers unit cost; what makes it specific to energy is that its interactions with e and l are not symmetric: a higher ν_k reduces the energy intensity of input production. By Shephard's lemma, input energy intensity is $\frac{e_\omega(p_e, w; \nu_k)}{c_\omega(p_e, w; \nu_k) F_\omega(e, l; \nu_k)} = \frac{\partial \log c_\omega(p_e, w; \nu_k)}{\partial p_e}$, the price-derivative of the log unit cost under CRS. It therefore falls in ν_k when the log unit cost function exhibits decreasing differences in (p_e, ν). We take those properties as our definition.

[7] To simplify exposition, we talk about energy externality to refer to the emission externality and ignore differences in fuel mixes. We focus on the energy used for production and abstract from transport emissions, as it accounts for less than 5% of intermediates' total greenhouse gas emissions (Le Moigne et al., 2024).

Definition 1 (Effective energy efficiency)**.** *ν_k characterizes the effective energy efficiency of firm k, in the sense that it* (i) *decreases its input energy intensity $\frac{e_\omega(\nu_k)}{c_\omega(\nu_k)F_\omega(\nu_k)}$ and* (ii) *decreases its input unit cost $c_\omega(\nu_k)$, $\forall\omega$. Formally,*

$$\text{(i)}\ \frac{\partial^2 \log c_\omega(p_e, w; \nu_k)}{\partial p_e\, \partial \nu_k} < 0 \qquad \textit{and} \qquad \text{(ii)}\ \frac{\partial c_\omega(p_e, w; \nu_k)}{\partial \nu_k} < 0, \qquad \forall\omega. \tag{1}$$

For instance, in a Cobb-Douglas production function, which we will use to illustrate our main mechanisms, this definition maps into $\frac{1}{A(\nu_k)} e^{1-\nu_k} l^{\nu_k}$: $\frac{e(p_e,w;\nu_k)}{c_\omega(p_e,w;\nu_k)F_\omega(e_k,l_k;\nu_k)} = \frac{1-\nu_k}{p_e}$ decreases with ν_k $\forall\omega$, and (ii) holds whenever $p_e > w$. In a CES framework with complementary factors, which we will use in our quantitative analysis, ν corresponds to the energy-biased technology. With substitutable factors, ν_k^{-1} is the energy-saving technology.

Outsourcing capabilities. Each relationship between a downstream buyer j and a potential supplier $i \in \mathcal{S}_\omega$ is characterized by a Hicks-neutral parameter $q_{ij} \in (0, \infty)$, which we term outsourcing capabilities, measured relative to in-house production, normalized to $q_{jj} = 1$.[8] It reflects any relationship-specific factor that makes outsourced inputs differ from in-house inputs, except for energy-related factors. For instance, if outsourcing to supplier i generates specialization gains for firm j, $q_{ij} > 1$. In contrast, if outsourcing to i implies a loss of control over input quality for firm j, $q_{ij} < 1$. To focus on the environmental externality, we abstract from any other distortion and interpret q_{ij} as a technological wedge, borrowing from the terminology of Bergquist et al. (2026), though our framework could equally be used to study the implications of a distortive wedge or a trade policy.

Formally, this means that for all inputs $\omega \in \Omega$ used by downstream firm j:

$$x_j(\omega) = \begin{cases} F_\omega\left(e_{\omega j}, l_{\omega j}; \nu_j\right) & \text{if } \omega \text{ is in-house} \\ q_{ij}\, F_\omega\left(e_{ij}, l_{ij}; \nu_i\right) & \text{if } \omega \text{ is outsourced to } i \in \mathcal{S}_\omega \end{cases} \tag{2}$$

2.2 Equilibrium decisions

Given the make-or-buy allocation, all firms choose their input demand separately to minimize costs, taking input prices as given. Perfect competition in the upstream sector implies marginal cost pricing, so the effective unit cost to firm j of input ω produced by $k \in \mathcal{S}_\omega \cup \{j\}$ is $c_\omega\left(p_e, w; \nu_k\right)/q_{kj}$, with $q_{jj} = 1$. Under CRS, the make-or-buy decision for one input does not depend on those for other inputs, so downstream firms source each input from the producer offering the lowest effective unit cost:

$$k_j(\omega; p_e) = \arg \min_{k \in \mathcal{S}_\omega \cup \{j\}} \frac{c_\omega\left(p_e, w; \nu_k\right)}{q_{kj}} \tag{3}$$

[8] We isolate the effective energy efficiency to characterize firms through this dimension, but make-or-buy decisions are far from being based on energy considerations alone (Dechezleprêtre et al., 2022; Levinson, 2023). We observe that comparable firms exhibit substantial variation in their input mixes, which can be attributed to multiple factors, from specialization patterns to trade costs (Antràs, 2020; Boudreau et al., 2023). This requires a second primitive - our outsourcing capability - that aggregates these idiosyncratic, non-energy factors. It rationalizes why firms may not necessarily select the most energy-efficient production process, thereby inducing energy misallocation.

Each input of each buyer is thus assigned to a single producer: in-house production and outsourcing are mutually exclusive, so make-or-buy decisions operate at the extensive margin and determine how aggregate production is allocated across producers.

Separating the firm itself from its external suppliers, (3) implies in-house production if

$$\underbrace{c_\omega\left(p_e, w; \nu_j\right) < \underbrace{\min_{i \in \mathcal{S}_\omega} \frac{c_\omega\left(p_e, w; \nu_i\right)}{q_{ij}}}_{\text{Supplier selection}}}_{\text{Producer selection}}$$

Make-or-buy thus extends supplier selection to the firm itself. Both select the producer with the lowest effective unit cost, but they are reflected by different objects in the data: supplier selection translates into variations in sales share across sellers, while in-house production is observed in the buyer's input shares.

2.3 Energy-use Misallocation

We now turn to the energy distortion and the induced misallocation. To obtain tractable results, we suppose that there is only one upstream supplier per input ($\mathcal{S}_\omega = \{i\}$) and focus on the decision between in-house production and outsourcing to i.[9] Define the in-house net advantage as

$$\Delta_{ij}(\omega; p_e) \equiv \log c_\omega(p_e, w; \nu_i) - \log c_\omega(p_e, w; \nu_j) - \log(q_{ij}) \tag{4}$$

Then firm j produces in-house if and only if $\Delta_{ij}(\omega; p_e) > 0$. The primitives $\{\nu_j, \nu_i\}$ and q_{ij} enter (4) asymmetrically. Outsourcing capabilities are Hicks-neutral: q_{ij} scales the cost of all factors, so it shifts Δ_{ij} by a constant, independent of p_e. Energy efficiency affects energy costs only, so its contribution to Δ_{ij}, through $\log c_\omega$, varies with p_e. The energy price therefore sets how much weight the decision places on energy efficiency relative to outsourcing capabilities: the higher is p_e, the larger the cost advantage of the more energy-efficient producer. This is why a uniform energy price change can affect the make-or-buy decision, and why a common distortion can induce between-producer misallocation.

Lemma 1 (Monotone selection in energy price)**.** *Under Definition 1(i), if $\nu_j > (<)\ \nu_i$, then $\Delta_{ij}(\omega; p_e)$ is strictly increasing (decreasing) in p_e. Consequently, there exists* at most *one threshold $\bar{p}_e$ such that firm j produces in-house (outsources) for $p_e \geq \bar{p}_e$.*

Proof. (sketch)

$$\frac{d\Delta_{ij}}{dp_e} = \frac{\partial \log c_\omega(p_e, w; \nu_i)}{\partial p_e} - \frac{\partial \log c_\omega(p_e, w; \nu_j)}{\partial p_e}$$

The Hicks-neutral q_{ij} drops out, and by decreasing differences, $\frac{d\Delta_{ij}}{dp_e} > (<)\ 0$ if $\nu_j > (<)\ \nu_i$. □

Selection is therefore monotone in p_e: raising the energy price never leads a firm to switch toward a less energy-efficient producer. Because the choice is discrete, however, the shifting weights change the selected producer only when the ranking of effective costs reverses, which happens at most once, at $\bar{p}_e$.

[9] We reintroduce multiple external suppliers in the quantitative analysis. Most of the results we derive apply equally to supplier selection alone, and could therefore be used to study that margin in isolation.

Lemma 2 (Reallocation requires a trade-off)**.** *Under Definition 1, a change in the energy price can alter a pair's make-or-buy decision only if energy efficiency and outsourcing capability favor opposite producers, i.e., if $\nu_j > \nu_i$ and $q_{ij} > 1$, or $\nu_j < \nu_i$ and $q_{ij} < 1$. Otherwise, the same producer is weakly cheaper at every p_e, and no price change induces a switch.*

Under a trade-off, selection is governed by *relative* advantage: the more energy-efficient producer is chosen if and only if its absolute advantage in energy use outweighs its absolute disadvantage in pairwise total-factor use, captured by the outsourcing capability q_{ij}. The energy price sets the terms of this comparison, so a price that fails to reflect the social cost of emissions distorts it.

Inefficient Allocation. Inefficiencies arise because firms evaluate costs at the prevailing energy price p_e, which does not internalize the environmental externality. In Appendix A.2.1, we derive the optimal allocation and show that it can be decentralized as a competitive equilibrium with a Pigouvian tax on energy such that the energy price corresponds to the socially optimal price, p_e^{SCC}, which reflects both the direct utility cost of energy and its effect on aggregate output. Energy-use misallocation can thus be measured by the welfare gains of moving to p_e^{SCC}. The wedge between p_e and p_e^{SCC} distorts both the intensive and the extensive margins:

Proposition 1. *If $p_e < p_e^{SCC}$, the competitive equilibrium is inefficient relative to the efficient allocation that evaluates costs at p_e^{SCC}.*

- ***Factor substitution (intensive margin, within producers):*** *each producer uses too much energy relative to labor.*
- ***Product substitution (intensive margin, across producers):*** *for a given make-or-buy allocation, households and downstream firms demand relatively too much of the output of energy-inefficient producers.*
- ***Producer selection (extensive margin, across producers):*** *there may be firm-input pairs for which the privately optimal make-or-buy decision (3) differs from the socially optimal one:*

$$\arg \min_{k \in \mathcal{S}_\omega \cup \{j\}} \frac{c_\omega\left(p_e, w; \nu_k\right)}{q_{kj}} \neq \arg \min_{k \in \mathcal{S}_\omega \cup \{j\}} \frac{c_\omega\left(p_e^{SCC}, w; \nu_k\right)}{q_{kj}} \tag{5}$$

The intensive-margin distortions correspond to the standard misallocation from factor price distortions within (Metcalf et al., 2003; Fullerton & Heutel, 2017) and across (Hsieh & Klenow, 2009) firms. The extensive-margin distortion arises from endogenous producer selection. By Lemma 1, whenever a pair's decision differs at the two prices, production is allocated to the less energy-efficient producer in equilibrium, which is privately cost-minimizing, but socially too energy intensive.

Producer selection responds to the energy price, so a change in the price reallocates production across producers, changing aggregate output and environmental damages. Because p_e^{SCC}, and hence the corresponding optimal carbon tax, is determined by these very quantities, it cannot be disciplined from

intensive-margin responses alone.[10] Evaluating the effect of moving to p_e^{SCC} likewise requires measuring the reallocation it induces along all three margins of Proposition 1. In our quantitative analysis, we measure their respective contribution.

2.4 Energy intensity

The input-level energy intensity of producing at k for buyer j (i.e., the *quantity* of energy used by k to produce that input relative to its total value for downstream firm j) is

$$\sigma_\omega(p_e;\nu_k) \equiv \frac{e_\omega(p_e;\nu_k)}{(c_\omega(p_e;\nu_k)/q_{kj})x_{kj}(\omega)} = \frac{\partial \log c_\omega(p_e;\nu_k)}{\partial p_e} \quad \forall j \tag{6}$$

Where we now drop the wage index to simplify notation. Thus, the embedded quantity of energy used to produce one monetary unit of input ω for firm j, $\sigma_\omega(p_e;\nu_{k_j(\omega;p_e)})$, is determined, at the intensive margin, by the effective energy efficiency of the producer (ν_k) and the factor prices (p_e). With heterogeneous effective energy efficiency, it is additionally determined, at the extensive margin, by the allocation of input production ($k_j(\omega;p_e)$), through the make-or-buy decision.

The downstream firm *measured* energy intensity is

$$\begin{aligned}\frac{e_j}{c_j y_j} &= \frac{\sum_{\omega\in\Omega} e_\omega(p_e;\nu_j)\mathbf{1}\{k_j(\omega;p_e)=j\}}{c_j y_j} \\ &= \sum_{\omega\in\Omega}\sigma_\omega(p_e;\nu_j)\frac{c_\omega(p_e;\nu_j)x_j(\omega)}{c_j y_j}\mathbf{1}\{k_j(\omega;p_e)=j\} = \sum_{\omega\in\mathcal{I}_j}\sigma_\omega(p_e;\nu_j)\frac{c_\omega(p_e;\nu_j)x_j(\omega)}{c_j y_j}\end{aligned} \tag{7}$$

Where $\mathcal{I}_j$ is the set of inputs produced in-house and $c_j \equiv c_j\left(\left\{c_\omega(p_e;\nu_{k_j(\omega;p_e)})/q_{k_j(\omega;p_e)j}\right\}_{\omega\in\Omega}\right)$ is firm j's equilibrium unit cost of production. The empirical energy intensity of the firm similarly depends on its own effective energy efficiency (and factor prices) at the intensive margin, while the make-or-buy decision drives the extensive margin through the set of inputs kept in-house. Analogously to Figure 1, Equations (6) and (7) highlight the two key features of the make-or-buy decision that our project takes into account to connect it to aggregate energy intensity. First, Equation (6) establishes its *allocative role*, the quantification of which requires characterizing potential producers by their effective energy efficiency. But make-or-buy decisions also have a *measurement role*, made salient by Equation (7): for a given effective energy efficiency ν_j, observed firm-level energy intensity mechanically rises with the share of inputs produced in-house, since in-house production brings the associated energy use inside the firm's boundary, whereas outsourcing shifts it to suppliers. Observed firm-level energy intensity is therefore not

[10] Under specific restrictions, this dependence disappears: Golosov et al. (2014) show that, when damages are proportional to output and marginal utility is inversely proportional to output, the optimal carbon tax does not depend on the optimal allocation. Even then, one still needs to know the optimal allocation to evaluate the effect of the tax on output and welfare. In appendix A.2.4, we show how our framework aggregates to a static version of the sectoral representation used in their canonical framework, highlighting that make-or-buy decisions shape sectoral technologies, typically treated as primitives, by endogenizing the underlying distribution of producers.

a sufficient statistic for effective energy efficiency.[11] Normalizing by value added instead of costs removes this span effect but leaves a composition effect: firm energy intensity becomes a cost-weighted average of the input-level energy intensities σ_ω over the inputs produced in-house, and thus depends on which inputs the firm makes unless σ_ω is constant across inputs, which the data reject.[12] We now discuss the first role in more details, while our quantitative model is built to deal with the second one.

Embedded energy intensity. Consider an increase in the energy price from p_e to $p'_e > p_e$. The total change in energy intensity of producing ω for firm j, can be decomposed as

$$\Delta\sigma_\omega(p_e;\nu_{k_j(\omega;p_e)}) = \underbrace{\left[\sigma_\omega(p'_e;\nu_{k_j(\omega;p_e)}) - \sigma_\omega(p_e;\nu_{k_j(\omega;p_e)})\right]}_{\text{Intensive margin}} + \underbrace{\left[\sigma_\omega(p'_e;\nu_{k_j(\omega;p'_e)}) - \sigma_\omega(p'_e;\nu_{k_j(\omega;p_e)})\right]}_{\text{Extensive margin}}. \quad (8)$$

The first term captures the within-producer response, driven by factor substitution and scale effects: holding fixed the identity of the producer, at a higher p_e, firms use less energy relative to labor:

$$\frac{\partial\sigma_\omega(p_e;\nu)}{\partial p_e}\Big|_\nu < 0,$$

so the intensive margin weakly reduces energy intensity. The second term captures the make-or-buy reallocation effect. Suppose that an increase in p_e induces a switch ($k_j(\omega;p_e) \neq k_j(\omega;p'_e)$). Under Definition 1,

$$\frac{\partial\sigma_\omega(p_e;\nu)}{\partial\nu} \leq 0,$$

and Lemma 1 implies $\nu_{k_j(\omega;p'_e)} > \nu_{k_j(\omega;p_e)}$. Hence, the extensive margin reduces energy intensity through reallocation toward more energy-efficient producers:

$$\sigma_\omega(p'_e;\nu_{k_j(\omega;p'_e)}) - \sigma_\omega(p'_e;\nu_{k_j(\omega;p_e)}) \leq 0,$$

Aggregate Energy Intensity. These results hold at the firm-input level, but the outcomes that matter for policy are aggregate: welfare, aggregate energy use and output. We use aggregate energy intensity, a key driver of carbon emissions (Peters et al., 2017; Le Quéré et al., 2019), to illustrate what moving from one level to the other involves. Aggregating over inputs and downstream firms:

$$\begin{aligned}\frac{E}{PY} &= \frac{1}{PY}\int_{j\in\mathcal{S}}\sum_{\omega\in\Omega} e_j dj = \int_{j\in\mathcal{S}}\frac{c_j y_j}{PY}\left(\sum_{\omega\in\mathcal{I}_j}\frac{c_\omega x_j(\omega)}{c_j y_j}\sigma_\omega(p_e;\nu_j) + \sum_{\omega\in\mathcal{O}_j}\frac{(c_\omega/q_{k_j(\omega,p_e)j})x_j(\omega)}{c_j y_j}\sigma_\omega(p_e;\nu_{k_j(\omega;p_e)})\right)dj \\ &= \int_{j\in\mathcal{S}}\frac{c_j y_j}{PY}\sum_{\omega\in\Omega}\frac{c_\omega x_j(\omega)}{c_j y_j}\sigma_\omega(p_e;\nu_{k_j(\omega;p_e)})dj\end{aligned}$$

[11] Interestingly, the measurement role of firm boundaries implies that a higher energy price might *compress* the cross-seller dispersion in observed firm-level energy intensity, as more energy-efficient firms bring production in-house (see Appendix B.6.3 and Figure B.5).

[12] The correlation of this composition effect with ν_j is ambiguous in sign. If energy efficiency lowers in-house costs proportionally more for energy-intensive inputs, as in a CES, firms more energy-efficient than their suppliers produce their most energy-intensive inputs in-house, which raises their measured intensity, but higher efficiency then brings less energy-intensive inputs in-house at the margin, which lowers it. The composition effect can thus attenuate or amplify measured efficiency differences across firms (Appendix A.2.2).

The pair-level mechanism governs σ_ω, but aggregate intensity also depends on the shares $c_j y_j / PY$ and $\frac{(c_\omega(p_e,w;\nu_{k_j(\omega;p_e)})/q_{k_j(\omega;p_e)\,j})\,x_j(\omega)}{c_j y_j}$, which respond to p_e through the product-demand channel of Proposition 1 and shift further as make-or-buy switches change firms' unit costs. How they respond depends on substitution elasticities and on how switches reorder firms' costs, which ν alone does not determine, as we show formally in Appendix A.2.3. The aggregate effect of $p_e < p_e^{SCC}$ is therefore a quantitative question: measuring it requires more structure on production, which we impose in Section 3, and identifying the distribution of (ν, q), which we do in Section 4.

2.5 Cobb-Douglas Illustration

Before turning to that structure, we briefly illustrate the results above in the most tractable case, Cobb-Douglas input production ($F_\omega = \frac{1}{A(\nu)} e^{1-\nu} l^\nu$, fully characterized in Appendix A.2.4), in which both Δ_{ij} and σ_ω are linear in ν. We use it for illustration only; the quantitative model relies on a CES specification.

First, the in-house net advantage (4) becomes:

$$\Delta_{ij}(\omega; p_e) = \underbrace{-\log(q_{ij})}_{\equiv \Delta \log q} + \log\left(\frac{p_e}{w}\right) \underbrace{(\nu_j - \nu_i)}_{\equiv \Delta\nu} \quad \Rightarrow \quad \frac{d\Delta_{ij}}{dp_e} = \frac{\nu_j - \nu_i}{p_e} \geq (\leq)\, 0 \;\text{ if } \nu_j \geq (\leq)\, \nu_i$$

The q-gap $\Delta \log q \equiv \log(q_{jj}) - \log(q_{ij}) = -\log(q_{ij})$ enters with unit weight and the ν-gap $\Delta\nu$ with weight $\log(p_e/w)$. This makes explicit the asymmetry behind Lemma 1, illustrated in closed-form in the right-hand side.

Consider an energy price increase from p_e to $p_e' > p_e > w$. Since a switch from outsourcing to in-house ($\mathcal{O}_j \to \mathcal{I}_j$) requires $\Delta_{ij}(p_e) \leq 0$ and $\Delta_{ij}(p_e') \geq 0$, while a switch from in-house to outsourcing ($\mathcal{I}_j \to \mathcal{O}_j$) requires $\Delta_{ij}(p_e) \geq 0$ and $\Delta_{ij}(p_e') \leq 0$,

$$[\mathcal{O}_j \to \mathcal{I}_j] \qquad \log\left(\frac{p_e'}{w}\right) \geq \frac{-\Delta \log q}{\underbrace{\Delta\nu}_{\geq 0 \text{ by Lemma 1}}} \geq \log\left(\frac{p_e}{w}\right) \geq 0 \tag{9a}$$

$$[\mathcal{I}_j \to \mathcal{O}_j] \qquad -\log\left(\frac{p_e'}{w}\right) \leq \frac{\Delta \log q}{\underbrace{\Delta\nu}_{\leq 0 \text{ by Lemma 1}}} \leq -\log\left(\frac{p_e}{w}\right) \leq 0 \tag{9b}$$

This is the trade-off of Lemma 2 in closed form: the inequalities hold only when $\Delta \log q$ and $\Delta\nu$ carry opposite signs. Concretely, an energy-efficient firm j ($\Delta\nu > 0$) that nonetheless outsources at p_e must have high outsourcing capabilities ($\Delta \log q \leq -\log(p_e/w)\,\Delta\nu < 0$); it switches in-house as the price rises to p_e' only if that advantage is not too strong, $\Delta \log q \geq -\log(p_e'/w)\,\Delta\nu$, so that whether a pair switches depends on where its $(\Delta \log q, \Delta\nu)$ lies. Figure 2 specializes these thresholds to the environmental distortion, plotting them at the market price (p_e^{mkt}) and at the social cost (p_e^{SCC}). The shaded areas between the two thresholds, which lie exclusively in the trade-off regions, contain the inefficient allocations for which the privately optimal decision differs from the planner's.

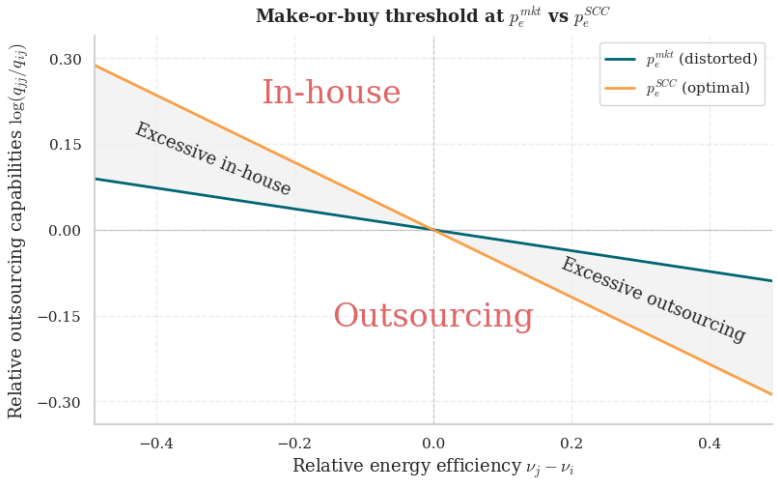


Figure 2: Outsourcing capabilities versus effective energy efficiency

Notes: Simulated make-or-buy thresholds at the equilibrium (p_e^{mkt}) and the optimal (p_e^{SCC}) energy price for different relative outsourcing capabilities ($\Delta \log q$) and relative energy efficiency ($\Delta\nu$) combinations, under a Cobb-Douglas production function and $w < p_e^{mkt} < p_e^{SCC}$. An input characterized by a pair above (below) the line is produced in-house (outsourced). The shaded areas contain the pairs for which the equilibrium and optimal allocations differ.

Finally, since $\sigma_\omega(p_e;\nu_k) = \frac{1-\nu_k}{p_e}$, decomposition (8) becomes, as illustrated in Figure A.1,

$$\sigma_\omega(p_e';\nu_{k_j(\omega;p_e')}) - \sigma_\omega(p_e;\nu_{k_j(\omega;p_e)}) = \underbrace{(1-\nu_{k_j(\omega;p_e)})\left(\frac{1}{p_e'}-\frac{1}{p_e}\right)}_{\text{Int. margin } \leq 0} + \underbrace{\frac{\nu_{k_j(\omega;p_e)}-\nu_{k_j(\omega;p_e')}}{p_e'}}_{\text{Ext. margin } \leq 0 \text{ (Lemma 1)}} \tag{10}$$

Moving to aggregate outcomes, we additionally assume that intermediate inputs within firms and final products in household demand are aggregated Cobb-Douglas, which fixes input and sales shares (Appendix A.2.3).[13] In that case, Figure 2 and this decomposition together characterize the extensive-margin contribution. It depends on a switching pair's impact $\frac{|\Delta\nu|}{p_e^{SCC}}$, the second term in 10 and the pair's horizontal position in the figure, and on the frequency of switches, which is the mass of pairs between the two thresholds (i.e., (9a)-(9b) evaluated at $p_e = p_e^{mkt}$ and $p_e' = p_e^{SCC}$). The band widens with $|\Delta\nu|$, and so does the impact of each switch, but whether more mass actually falls in it depends on how q-gaps co-vary with ν-gaps: the joint distribution of $(\Delta \log q, \Delta\nu)$ governs extensive-margin misallocation. In Appendix A.2.4, we make illustrative distributional assumptions on this joint distribution to derive closed-form solutions, and show that the same ν-gap distribution can generate very different magnitudes of misallocation depending on the sign and shape of this relation and on the noise around it.

The main takeaway of this section is that make-or-buy decisions govern how energy use is allocated across firms and therefore shape aggregate energy intensity and welfare. When energy prices are below the social cost, firms place insufficient weight on relative energy efficiency in these decisions, leading to cross-producer energy-use misallocation at the extensive margin. The direction of this misallocation depends on (ν_j, ν_i), while its magnitude depends on the joint distribution of (ν_j, ν_i, q_{ij}) - primitives

[13] Cobb-Douglas aggregation fixes expenditure shares regardless of relative prices. The product-demand channel of Proposition 1 still distorts relative quantities, as buyers demand relatively more of the output of energy-inefficient producers, but it no longer reallocates expenditure and therefore leaves aggregate energy intensity unaffected. Producer selection, by contrast, still moves it, since it changes which producer's energy intensity enters each fixed share. Potential producers are perfect substitutes in the make-or-buy decisions, but that alone would not be sufficient: a common change in p_e moves their ranking because energy efficiency is not Hicks-neutral, and therefore shifts producers' costs by different proportions.

that are not directly observable because empirical firm-level energy intensity reflects both production technology and firm boundaries. The aggregate effect of distorted energy prices combines the extensive margin with the within-firm and product-demand channels of Proposition 1, which the extensive margin itself feeds back into through equilibrium prices and costs. Assessing it requires a structural approach, to which we now turn.

3 A quantitative model of make-or-buy decisions in a production network

Quantifying the misallocation induced by distorted energy prices is an ex-post analysis, in the sense that there is no sufficient statistic for it: the extensive margin depends on the joint distribution of (ν, q) across all potential pairs rather than on observed shares only. It therefore requires knowledge of both the observed equilibrium allocation and the counterfactual efficient one. Since no empirical "efficient" benchmark exists for the Italian economy, we impose parametric structure on our framework to directly recover this distribution from the data and construct the counterfactual.

Endogenous network models determine which firm-pairs trade, but usually take firm boundaries as given (e.g., Huneeus (2018), Bernard et al. (2022), Dhyne et al. (2023), and Arkolakis et al. (2025));[14] the firm-boundaries literature explicitly endogenizes make-or-buy decisions but typically at a single stage (e.g., Grossman and Helpman (2002), Antràs and Chor (2013), and Alfaro et al. (2019)). In contrast, our quantitative model needs to characterize all alternative producers in terms of effective energy efficiency and unit costs at different energy prices, taking into account that make-or-buy decisions can happen at any intermediate stage, while remaining numerically solvable. Therefore, we embed our model of make-or-buy decisions and supplier selection into a sequential production network in the spirit of Boehm and Oberfield (2023). A manufacturing supply chain can be composed of multiple stages (vertical), and each stage might use multiple inputs (horizontal). Each manufacturing firm, regardless of its position in the supply chain, can produce its intermediate inputs in-house or outsource them to other firms, meaning that their producer selection decisions jointly shape the production network.[15] Below, we lay out the main components, while additional details can be found in Appendix A.3.

[14] Make-or-buy decisions could in principle be included in these models by treating in-house production as the outside option to trading with a supplier, as in Eaton et al. (2026). However, they model in-house production with a process that systematically differs from outsourced production in terms of input requirements, preventing a direct comparison of energy efficiency across downstream and upstream firms. More generally, the asymmetry between potential external suppliers and in-house production generates challenges that make these already complex frameworks less suitable for analyzing make-or-buy decisions.

[15] Our model is nested in the general frameworks of Acemoglu and Azar (2020) and Baqaee and Farhi (2020) describing the effects of distortions in a production network, since make-or-buy decisions can be reinterpreted as selection across perfectly substitutable potential producers, including a fictitious one corresponding to the firm itself. However, their non-parametric results require observing all allocations under comparison, which we do not. By specifically modeling the primitives governing make-or-buy decisions, ν and q, and identifying them in the data, we can evaluate the network at any (distorted) energy price and quantify the resulting misallocation.

3.1 Environment

Household. We consider a closed economy.[16] A representative household supplies L units of labor inelastically, receives utility from consumption and disutility from energy. The household consumes a numeraire good X and a manufacturing good Y, with quasi-linear preferences: $U(X,Y,E) = X_c + \frac{Y^\alpha}{\alpha} - \Upsilon E$, where Υ represents the environmental damages. This linear specification implies that the optimal carbon tax is independent of the underlying allocation and exactly equal to Υ.[17]Energy is produced competitively using the numeraire good with technology $E = \frac{X_e}{\chi}$. The numeraire good is solely produced with labor ($X = X_c + X_e = L_x$). The manufacturing good Y aggregates the different products $\omega \in \Omega$ sold by all the manufacturing firms: $Y = \left(\sum_{\omega\in\Omega} \delta_\omega^{\frac{1}{\eta}} Y_\omega^{\frac{\eta-1}{\eta}}\right)^{\frac{\eta}{\eta-1}}$, where δ_ω is a product demand shifter, and where the product bundle for product ω aggregates the production of all the different firms in the set $\mathcal{S}_\omega$ that sell it, $Y_\omega = \left(\int_{j\in\mathcal{S}_\omega} y_j(\omega)^{\frac{\varepsilon-1}{\varepsilon}} dj\right)^{\frac{\varepsilon}{\varepsilon-1}}$. Therefore, firm j selling product ω faces demand $p_j(\omega) = \delta_\omega^{\frac{1}{\eta}} P \left(\frac{Y}{Y_\omega}\right)^{\frac{1}{\eta}} \left(\frac{Y_\omega}{y_j(\omega)}\right)^{\frac{1}{\varepsilon}}$, the product price index is $P_\omega = \left(\int_{j\in\mathcal{S}_\omega} p_j(\omega)^{1-\varepsilon} dj\right)^{\frac{1}{1-\varepsilon}}$ and the aggregate price index is $P = \left(\sum_{\omega\in\Omega} \delta_\omega P_\omega^{1-\eta}\right)^{\frac{1}{1-\eta}}$.

Labor markets are perfectly competitive, thus the wage is equalized across all sectors. As in Boehm and Oberfield (2023), manufacturing firms are perfectly competitive when selling to other firms but monopolistically competitive when selling to the household, charging a markup $\frac{\varepsilon}{\varepsilon-1}$ over marginal cost. The household faces the same markup on the numeraire good, whose price we normalize to $P_x = 1$, so that all consumer prices are scaled by a common factor. Energy production, by contrast, uses the numeraire good at marginal cost, leaving the energy price free of markups. Profits (Π_X, Π_Y) are rebated lump-sum, meaning that markups do not distort neither the allocation nor the scale of production, and the environmental externality remains the only source of distortion. If the government imposes a carbon tax τ on energy, its revenue $\tilde{T}$ is also rebated lump-sum to the household, whose budget constraint satisfies $P_x X_c + PY = wL + \tilde{T} + \Pi_X + \Pi_Y$. These functional form assumptions imply that

$$X = L_x = L - L_m \qquad P_x = \frac{\varepsilon}{\varepsilon - 1} w = 1 \qquad p_e = \chi w + \tau \qquad PY = Y^\alpha$$

Manufacturing production is therefore determined by the supply side, while the numeraire good absorbs the labor adjustments.

Inputs production. Supply chains can be interpreted as a tree, as illustrated by Figure 3. Each manufacturing input $\omega \in \Omega$ is produced by combining energy and labor with a set $\mathcal{U}_\omega^d \subset \Omega$ of direct inputs produced at the previous stage. This sequentiality implies that the production of ω draws directly on $\mathcal{U}_\omega^d$ and, since each of these inputs is itself produced the same way, indirectly on their own direct inputs, and so on down to the most upstream stage. We denote by Ω_ω the recipe of ω: the set of all direct and

[16] In our empirical analysis, we allow manufacturing firms to purchase and sell intermediate inputs abroad, but impose balanced trade.

[17] Static linear damages are a reasonable approximation for Italian manufacturing emissions, which are too small to affect marginal damages. A counterfactual with non-linear damages, relevant at the global scale, shows how extensive-margin reallocation affects the optimal tax once it depends on the allocation.

indirect inputs that enter its production at each stage.[18] There is no round-about production, $\omega \notin \Omega_\omega$, but an input can enter a recipe through several branches of the tree. We call each such occurrence a production stage: the production of a given input at a given position in the recipe. With a slight abuse of notation, we index both a stage and the input it produces by ω, as all characteristics of an input, such as its production parameters and its set of suppliers $\mathcal{S}_\omega$, are common to all the stages producing it, which differ only in their position in the recipe. Formally, the production of input ω by producer j takes the form:

$$x_j(\omega) = \left((\alpha_\omega^{el})^{1-\gamma} \left(\beta_\omega^{1-\rho_\omega} (\nu_j e_j(\omega))^{\rho_\omega} + (1-\beta_\omega)^{1-\rho_\omega} l_j(\omega)^{\rho_\omega} \right)^{\frac{\gamma}{\rho_\omega}} + \sum_{\hat{\omega} \in \mathcal{U}_\omega^d} (\alpha_\omega^{\hat{\omega}})^{1-\gamma} (q_{ij} x_{ij}(\hat{\omega}))^{\gamma} \right)^{\frac{1}{\gamma}} \tag{11}$$

Where α_ω^{el} and $\alpha_\omega^{\hat{\omega}}$ represent respectively the requirement share of the primary factors and the direct upstream inputs in the production of ω. We normalize $\alpha_\omega^{el} + \sum_{\hat{\omega} \in \mathcal{U}_\omega^d} \alpha_\omega^{\hat{\omega}} = 1 \ \forall \omega$, ruling out any love-of-variety effect. $\frac{1}{1-\rho_\omega}$ is the elasticity of substitution between energy and labor in the production of ω, and β_ω represents the share of energy needed in the production of input ω. $\frac{1}{1-\gamma}$ is the elasticity of substitution between intermediate inputs. As in the two-stage model, $\nu_j \in (0, \infty)$ is producer j's effective energy efficiency, which takes the form of an energy-biased technology under our CES specification, and $q_{ij} x_{ij}(\hat{\omega})$ represents the quality-adjusted volume of input $\hat{\omega}$ that is used by j. The outsourcing capability $q_{ij} \in (0, \infty)$ depends on which firm i produces input $\hat{\omega}$ for firm j, and we again normalize $q_{jj} = 1$ if the input is produced in-house by firm j.

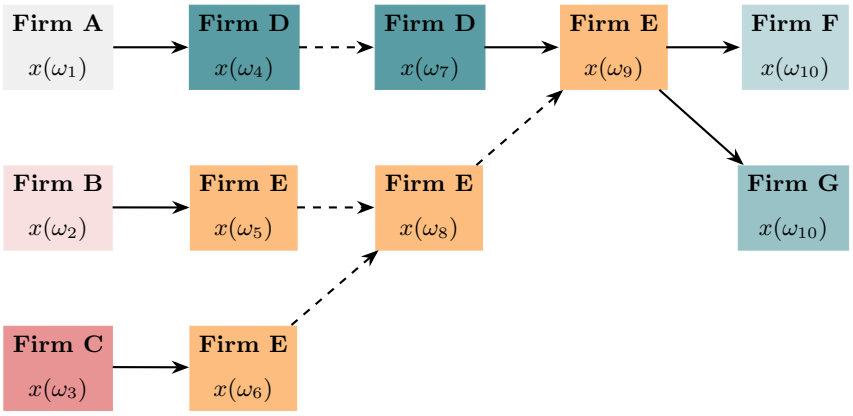


Figure 3: Illustration of the production network

Notes: Each box is a production stage: a firm producing input ω, which combines energy, labor, and its direct inputs (Equation (11)). Arrows point from an input to the stage that uses it. Solid arrows are market transactions between firms (outsourcing), while dashed arrows link stages performed within the same firm (in-house production); boxes of the same color belong to the same firm. Each firm sells only the output of its last stage, its core product, both to other firms and to the household (not shown). In this example, firms F and G buy ω_9 from firm E, which produces ω_5, ω_6 and ω_8 in-house and buys ω_2, ω_3 and ω_7 from firms B, C and D.

Firm production. Firms produce inputs according to the production function (11). Firms sell only one input, both to other firms and to the household, which we call their *core* product.[19] The recipe of this

[18] Our *products'* recipes differ from the *firms'* recipes of Boehm and Oberfield (2020), which correspond to the set of inputs a firm purchases.

[19] When mapping the model to the data, we include multi-product firms by assuming that they aggregate their core products through a CES with elasticity of substitution $\frac{1}{1-\gamma}$. This implies that, in effect, they act, for each product

core product determines which stages its production involves. All firms have to perform the last stage of their production in-house: a firm $j \in \mathcal{S}_\omega$ selling ω cannot outsource the production of ω. This last stage is characterized by a firm-specific TFP ψ_j.[20] In contrast, the firm faces a make-or-buy decision for each direct input $\hat{\omega} \in \mathcal{U}^d_\omega$: it can buy it from one supplier $i \in \mathcal{S}_{\hat{\omega}}$, or produce it in-house according to (11), in which case the same decision arises for each of $\hat{\omega}$'s own direct inputs $\tilde{\omega} \in \mathcal{U}^d_{\hat{\omega}}$. Make-or-buy decisions thus proceed sequentially upstream from the core product, and a firm performs only consecutive stages in-house: no stage upstream of an outsourced input is performed by the firm.[21]

3.2 Equilibrium

Timing and equilibrium. The model is static and technologies are fixed. Firms observe their effective energy efficiency and the prices offered by all suppliers, and draw their relationship-specific outsourcing capabilities q_{ij} from a distribution $\mathbf{Q}_j(q)$ for each potential supplier $i \in \mathcal{S}_\omega$ and each input ω in their core product's recipe. With this information, they decide for each input that can be outsourced, whether to produce it in-house or to outsource it to their preferred supplier among the set of potential suppliers. Then, they choose their input demand.[22] After that, all firms produce, markets clear, and the household consumes. The equilibrium is formally defined in Definition 2.

Prices and marginal costs. Firms are perfectly competitive and price at marginal cost when selling to other firms. For each of its intermediate inputs ω, the buyer firm chooses the supplier that can offer the lowest quality-adjusted effective cost among the set of potential suppliers $\mathcal{S}_\omega$. Therefore, the cost of outsourcing input ω is $c_j^{\mathcal{O}}(\omega) = \min_{i \in \mathcal{S}_\omega} \frac{c_i(\omega)}{q_{ij}}$. In contrast, the cost of producing input ω in-house delivers the effective unit cost

$$c_j^{\mathcal{I}}(\omega) = \left(\alpha_\omega^{el} c_j^{el}(\omega)^{\frac{-\gamma}{1-\gamma}} + \sum_{\hat{\omega} \in \mathcal{U}^d_\omega} \alpha_\omega^{\hat{\omega}} c_j(\hat{\omega})^{\frac{-\gamma}{1-\gamma}} \right)^{\frac{-(1-\gamma)}{\gamma}}$$

Where $c_j^{el}(\omega) = \left(\beta_\omega \left(\frac{p_e}{\nu_j} \right)^{\frac{-\rho_\omega}{1-\rho_\omega}} + (1-\beta_\omega) w^{\frac{-\rho_\omega}{1-\rho_\omega}} \right)^{\frac{-(1-\rho_\omega)}{\rho_\omega}}$ reflects the costs of the primary factors, and $c_j(\hat{\omega})$ is the effective cost of the upstream input $\hat{\omega}$, that we now characterize. For each input that can be outsourced, the firm chooses separately whether to produce it in-house or to outsource it to minimize

separately, as a "fictitious" single-product firm, but each is characterized by the same firm-level primitives (see Appendix B.3.1).

[20] Adding firm heterogeneity through a firm-specific TFP allows us to separate the outsourcing capabilities q_{ij}, which then only collect outsourcing-specific factors, from suppliers' productivity, which are absorbed in the costs.

[21] Alfaro et al. (2019) show that firms integrate stages that are bunched together.

[22] Because the energy-labor nest differs across inputs, we model the cost minimization problem to choose factor demand for each in-house input separately. Thanks to the independence of irrelevant alternatives property of the CES, this is actually equivalent to assuming that firms choose their total demand for energy and labor and then choose how to allocate it across their different inputs.

its effective unit cost.[23] Therefore,

$$
\begin{aligned}
c_j(\hat{\omega}) &= \min\left\{c_j^{\mathcal{I}}(\hat{\omega}),\ c_j^{\mathcal{O}}(\hat{\omega})\right\} \\
&= \min\Bigg\{\left(\alpha_{\hat{\omega}}^{el}\left(\beta_{\hat{\omega}}\left(\frac{p_e}{\nu_j}\right)^{\frac{-\rho_{\hat{\omega}}}{1-\rho_{\hat{\omega}}}} + (1-\beta_{\hat{\omega}})w^{\frac{-\rho_{\hat{\omega}}}{1-\rho_{\hat{\omega}}}}\right)^{\frac{(1-\rho_{\hat{\omega}})\gamma}{(1-\gamma)\rho_{\hat{\omega}}}} + \sum_{\tilde{\omega}\in\mathcal{U}_{\hat{\omega}}^d}\alpha_{\hat{\omega}}^{\tilde{\omega}}c_j(\tilde{\omega})^{\frac{-\gamma}{1-\gamma}}\right)^{\frac{-(1-\gamma)}{\gamma}}, \\
&\qquad \min_{i\in\mathcal{S}_{\hat{\omega}}}\left(\frac{c_i(\hat{\omega})}{q_{ij}}\right)\Bigg\}
\end{aligned}
\tag{12}
$$

Firms take into account that, if they produce in-house, they have to pay for energy and labor, but also for the upstream stages, whose equilibrium cost embeds their own make-or-buy outcome. In contrast, if they outsource, they cannot perform in-house the stages located further upstream, whose costs are now embedded in the selected supplier's cost through its own make-or-buy decisions.

Finally, the effective unit cost of the core product ω^c of a firm is

$$
c_j(\omega^c) = \frac{1}{\psi_j}\left(\alpha_{\omega^c}^{el}\left(\beta_{\omega^c}\left(\frac{p_e}{\nu_j}\right)^{\frac{-\rho_{\omega^c}}{1-\rho_{\omega^c}}} + (1-\beta_{\omega^c})w^{\frac{-\rho_{\omega^c}}{1-\rho_{\omega^c}}}\right)^{\frac{(1-\rho_{\omega^c})\gamma}{(1-\gamma)\rho_{\omega^c}}} + \sum_{\hat{\omega}\in\mathcal{U}_{\omega^c}^d}\alpha_{\omega^c}^{\hat{\omega}}c_j(\hat{\omega})^{\frac{-\gamma}{1-\gamma}}\right)^{\frac{-(1-\gamma)}{\gamma}}
$$

We define $\mathcal{I}_j$ as the set of inputs that firm j produces in-house and $\mathcal{O}_j$ as the set of inputs it purchases. Since costs depend on the input but not on the stage at which it is used, a firm makes the same choice at every stage of a given input that it decides on, so each such input belongs to either $\mathcal{I}_j$ or $\mathcal{O}_j$. The in-house total requirement $\tilde{\alpha}_{\omega}^{\hat{\omega}}(j) \equiv \sum_{K\geq 0}\sum_{\omega_1,\ldots,\omega_K\in\mathcal{I}_j}\alpha_{\omega}^{\omega_1}\alpha_{\omega_1}^{\omega_2}\cdots\alpha_{\omega_K}^{\hat{\omega}}$ (with $\tilde{\alpha}_{\omega}^{\omega}(j)\equiv 1$) sums the product of direct requirements over all supply chains connecting ω to $\hat{\omega}$ whose intermediate stages j performs in-house, and therefore counts each stage at which $\hat{\omega}$ enters.

Equilibrium conditions. We summarize here the key equilibrium objects that we directly map to the data, while appendix A.3.1 describes the remaining equilibrium conditions and discusses the intensive and extensive margin effects on aggregate energy intensity.[24] The total expenditure shares of firm j (producing ω^c) for energy and labor are:

[23] Consistent with Miyauchi et al. (2018)'s empirical findings, the likelihood to find a suitable supplier, hence to outsource, increases with the mass of potential suppliers $\mathcal{S}_\omega$, but is uncorrelated with the mass of buyers.

[24] Our framework differs from Boehm and Oberfield (2023) by introducing two dimensions of technological heterogeneity that we must identify separately to compare our potential producers: effective energy efficiency and outsourcing capabilities. This prevents us from applying the distributional assumptions they use to solve for prices, and motivates the exact calibration we describe below.

$$
\begin{aligned}
\frac{p_e e_j}{c_j(\omega^c)x_j(\omega^c)} &= \psi_j^{\frac{\gamma}{1-\gamma}} \sum_{\hat{\omega}\in\mathcal{I}_j\cup\{\omega^c\}} \tilde{\alpha}^{\hat{\omega}}_{\omega^c}(j)\alpha^{el}_{\hat{\omega}}\beta_{\hat{\omega}} \left(\frac{c_j(\omega^c)}{c_j^{el}(\hat{\omega})}\right)^{\frac{\gamma}{1-\gamma}} c_j^{el}(\hat{\omega})^{\frac{\rho_{\hat{\omega}}}{1-\rho_{\hat{\omega}}}} \left(\frac{p_e}{\nu_j}\right)^{\frac{-\rho_{\hat{\omega}}}{1-\rho_{\hat{\omega}}}} \\
\frac{wl_j}{c_j(\omega^c)x_j(\omega^c)} &= \psi_j^{\frac{\gamma}{1-\gamma}} \sum_{\hat{\omega}\in\mathcal{I}_j\cup\{\omega^c\}} \tilde{\alpha}^{\omega}_{\omega^c}(j)\alpha^{el}_{\omega}(1-\beta_{\omega}) \left(\frac{c_j(\omega^c)}{c_j^{el}(\omega)}\right)^{\frac{\gamma}{1-\gamma}} c_j^{el}(\hat{\omega})^{\frac{\rho_{\hat{\omega}}}{1-\rho_{\hat{\omega}}}} w^{\frac{-\rho_{\hat{\omega}}}{1-\rho_{\hat{\omega}}}} \\
\Rightarrow \frac{p_e e_j}{wl_j} &= \frac{\sum_{\hat{\omega}\in\mathcal{I}_j\cup\{\omega^c\}} \tilde{\alpha}^{\hat{\omega}}_{\omega^c}(j)\alpha^{el}_{\hat{\omega}}\beta_{\hat{\omega}} c_j^{el}(\hat{\omega})^{\frac{\rho_{\hat{\omega}}-\gamma}{(1-\rho_{\hat{\omega}})(1-\gamma)}} \left(\frac{p_e}{\nu_j}\right)^{\frac{-\rho_{\hat{\omega}}}{1-\rho_{\hat{\omega}}}}}{\sum_{\hat{\omega}\in\mathcal{I}_j\cup\{\omega^c\}} \tilde{\alpha}^{\hat{\omega}}_{\omega^c}(j)\alpha^{el}_{\hat{\omega}}(1-\beta_{\hat{\omega}}) c_j^{el}(\hat{\omega})^{\frac{\rho_{\hat{\omega}}-\gamma}{(1-\rho_{\hat{\omega}})(1-\gamma)}} w^{\frac{-\rho_{\hat{\omega}}}{1-\rho_{\hat{\omega}}}}}
\end{aligned}
\tag{13}
$$

Firm-level input expenditure shares highlight again that, without accounting for firm boundaries ($\mathcal{I}_j$), input expenditure shares (residualized on factor prices) are not a sufficient statistic for effective input intensity. We are able to identify ν_j as the effective energy efficiency only if we control for the energy requirement of all the inputs that are *produced* by the firm, and not only the final product ω^c that is used to classify the firm into its sector.

The corresponding shares for outsourced intermediate inputs $\hat{\omega} \in \mathcal{O}_j$ are:

$$
\frac{p_j(\hat{\omega})x_{ij}(\hat{\omega})}{c_j(\omega^c)x_j(\omega^c)} = \tilde{\alpha}^{\hat{\omega}}_{\omega^c}(j)\psi_j^{\frac{\gamma}{1-\gamma}} \left(\frac{c_j(\omega^c)}{c_i(\hat{\omega})/q_{ij}}\right)^{\frac{\gamma}{1-\gamma}} \quad \text{and} \quad \frac{c_j(\hat{\omega})x_{ij}(\hat{\omega})}{p_e e_j + wl_j} = \frac{\tilde{\alpha}^{\hat{\omega}}_{\omega^c}(j)(c_i(\hat{\omega})/q_{ij})^{\frac{-\gamma}{1-\gamma}}}{\sum_{\omega\in\mathcal{I}_j\cup\{\omega^c\}} \tilde{\alpha}^{\omega}_{\omega^c}(j)\alpha^{el}_{\omega} c_j^{el}(\omega)^{\frac{-\gamma}{1-\gamma}}}
\tag{14}
$$

The effective energy efficiency ν_j and the distribution of outsourcing capabilities $\{q_{ij}\}$ for each firm are behind the variations in these moments, which we use in our calibration. Conditional on firm boundaries and factor prices, ν_j explains variation in the energy-to-labor expenditure ratio, while the distribution of $\{q_{ij}\}$ governs the across- and within-buyer dispersion in outsourced input expenditure shares. Jointly, they determine which inputs will be produced in-house and which will be outsourced.

3.3 Discussion

3.3.1 Two-stage results in the full model

The mechanism of the two-stage model carries over. Since ν_j lowers the in-house cost $c_j^{\mathcal{I}}(\hat{\omega})$, more energy-efficient firms are, ceteris paribus, less likely to outsource, especially at high energy prices (by Definition 1), whereas firms with high outsourcing capabilities are more likely to outsource. Suboptimal energy prices therefore distort make-or-buy decisions, and pricing energy at its social cost still decentralizes the efficient allocation and maximizes welfare. A common increase in the energy price affects producers' costs heterogeneously, shifting weight toward energy-efficient producers, either directly (higher-ν producers) or indirectly (producers with more energy-efficient suppliers). A switch still requires a trade-off, now between embedded energy efficiency and embedded Hicks-neutral productivity, the multi-stage counterparts of ν and q: as the energy price rises, embedded energy efficiency weighs more and eventually dominates.

What differs is that the weight of energy prices compounds along the supply chain. Outsourcing an

input assigns not only that input but also its entire upstream lineage to the supplier's supply chain, so embedded energy efficiency and productivity depend on the weighted primitives $\{q, \nu\}$ of all observed and potential pairs along the upstream supply chain, not only on those of the pair. For an energy-efficient firm, the energy-cost penalty of outsourcing accumulates across stages, and at high energy prices, such firms tend to integrate more inputs, even when they have strong outsourcing capabilities with some suppliers. The joint distribution of energy efficiency and outsourcing capabilities therefore still governs misallocation, but across entire supply chains, and the pairwise gaps $(\Delta \log q, \Delta\nu)$ are no longer sufficient statistics (Appendix A.3.3). Moreover, whereas in the two-stage model a higher energy price reduces firm-level embedded energy intensity through the extensive margin at every price level, in the full model it need not. The two-stage result still holds locally, at each switching point. However, switching thresholds differ across firm–input pairs, so a uniform price increase triggers upstream reallocations at different prices, which may overturn the local effect over some price intervals. Appendix A.3.2 clarifies this issue and shows that it occurs only in rare cases (1.4% of switches in our optimal counterfactual); while it could complicate solving numerically for the optimal allocation, it does not in our setting, where the optimal carbon price is pinned down by environmental damages.

3.3.2 Discussion of modeling assumptions

We close this section by discussing our main assumptions and their implications for our results.

General capabilities. We model outsourcing capabilities as a technological primitive that summarizes all non-energy determinants of make-or-buy decisions, from specialization gains and proximity to incompatibility ($q_{ij} \to 0$ prevents a relationship from forming). This lets the model accommodate the strong idiosyncrasies documented in firms' make-or-buy choices (Monteverde & Teece, 1982; Alfaro et al., 2019; Antràs, 2020), without taking a stand on their source. We are able to keep this generality by calibrating each of our model firms directly to an empirical firm, as explained in the next section. This is ideal to identify the joint distribution of outsourcing capabilities and energy efficiency in the data, but it implies that we cannot microfound it.

Distortions. To isolate the misallocation induced by the environmental externality, we abstract from other distortions: firms are perfectly competitive when selling to each other, and since outsourcing capabilities are technological, any frictions they contain, such as matching costs or markups, are taken as given. Our results should therefore be understood as the effects of a second-best environmental policy, which these frictions do not bias as long as they do not respond to energy prices. For comparison, Appendix B.13.1 reports an alternative calibration that treats the q_{ij} as purely distortive, measuring misallocation against a different benchmark.

Independent make-or-buy decisions. We abstract from fixed costs, economies of scale, and resource synergies across inputs, so that each make-or-buy decision is independent. This keeps in-house and outsourced production symmetric, so that producers can be compared on their technological primitives alone, and avoids a combinatorial discrete choice problem that existing tools (Antras et al., 2017;

Arkolakis et al., 2023) cannot solve without restricting the joint distribution of our two dimensions of heterogeneity. Reallocation driven by such interdependencies, for instance joint switches across inputs sharing a fixed cost, is therefore outside our quantification. The simplification is consistent with evidence that firm scale, span and position in the supply chain are largely orthogonal across European production networks (Magerman et al., 2026).

Firm-level energy efficiency. Each firm has a single effective energy efficiency ν_j, common to all the stages it performs, since energy-efficiency improvements in manufacturing typically apply across production processes rather than process by process (Wang et al., 2008; Ramírez-Santos et al., 2018; Kim et al., 2022). Energy use still differs across inputs through the input-specific parameters β_ω and ρ_ω. If, in practice, energy efficiency varies across the inputs a firm produces, ν_j captures its average, which biases our results if firms are systematically more or less efficient at the inputs they produce in-house than this average. In the data, however, firms show no systematic difference in energy efficiency relative to the average seller of the inputs they produce in-house, consistent with efficiency being a characteristic of the firm rather than of its core product (Figure B.8a).

No dynamics. The model is static and holds energy efficiency and outsourcing capabilities fixed. It therefore abstracts from relationship stickiness (Antràs & Chor, 2013; Martin et al., 2023), hence cannot speak to how long reallocation takes to unfold, and from the response of both primitives to energy prices, such as the adoption of energy-efficient technologies. It also ignores entry and exit. The least energy-efficient firms lose customers and household demand, so they shrink but do not exit; for their customers, who switch to another producer in either case, the two are equivalent. But these firms also adjust their own boundaries to mitigate their exposure. If this adjustment would not cover their fixed operating costs, they would exit instead, and we overstate the contribution of producer selection. However, firm turnover is low in Italian manufacturing, both relative to the rest of the Italian economy (Istat, 2014) and by European standards (Eurostat, 2024).

Identical recipes across firms. All firms producing a given product ω share the same recipe Ω_ω, as in Alfaro et al. (2019), Rachapalli (2021), and Alfaro et al. (2024): the technology is exogenous, and what is endogenous is which stages each firm performs in-house, which determines its production function. Our CES structure still lets expenditure shares vary across firms beyond the common input requirements. If firms producing the same product in fact use different recipes, we may misclassify some inputs as produced in-house; Appendix B.5 reports robustness exercises on our recipes.

Production network. We assume that the number of production stages is finite and the network is acyclic, which guarantees the uniqueness of the equilibrium (Dhyne et al., 2023; Alfaro-Ureña & Zacchia, 2025) and allows us to infer in-house sets and evaluate make-or-buy decisions recursively.[25] The model delivers that firms source each input from a single producer, even when several of their stages use it,

[25] Dhyne et al. (2023) show that the Belgian production network is consistent with an acyclic structure for more than 80% of the observed transactions, and that dropping inconsistent transactions has little impact on sectoral linkages and aggregate measures. Similarly, using firm-to-firm transaction data for Belgium, Estonia, Hungary, Italy, and Portugal, Magerman et al. (2026) find very little round-about production in firm-level production networks.

which our data support (Table B.6). We use this property to infer the in-house sets from observed input transactions.

4 Quantitative analysis

4.1 Data

To evaluate the aggregate energy intensity of the manufacturing sector relative to its socially efficient benchmark, accounting for the cross-producer allocation, we need to (i) identify which firm is *actually* producing each input, which requires identifying the set of inputs produced in-house by each firm and (ii) identify which firm should *optimally* be producing each input, which requires characterizing each potential producer by their efficiency. We fulfill both requirements leveraging Italian firm-to-firm transaction (B2B) data that uniquely identify both the buyer and the supplier (see Astuti and Piras (2026) for a complete description of the data). Our quantitative exercise is based on the 2019 cross-section, but we also use the 2021-2023 panel for robustness exercises and for part of our calibration. We combine it with customs and balance-sheet data, which assign each firm the ATECO 6-digit code of its main activity. At this level, manufacturing activities are defined by the good they produce (e.g., 23.64.00, the production of mortar), so each code identifies both an input and the activity that produces it, and each B2B transaction inherits the code of its seller (Appendix B.3.1 describes how we handle multi-product firms). All datasets are described in Appendix B.1. Appendix B.3 provides more details on how we prepare them, and Appendix B.4 reports descriptive statistics on our final sample, which contains 153,424 manufacturing firms - among which 145,498 report transactions with other manufacturing firms - and 2,928,934 transactions of intermediate inputs. We briefly discuss here the key variables that we use.

Factor expenditures. Firms' labor expenditures are directly observed in the balance sheet, while energy and intermediate input expenditures come from the B2B and customs data. Energy expenditures aggregate all purchases from firms classified as generating or selling energy (see Appendix B.2); in the baseline, we focus on total energy and do not differentiate firms by fuel mix. Intermediate input expenditures include only purchases of manufacturing inputs, where the allocation of production matters most for energy use. We exclude intangible and physical capital (such as long-lasting machinery and equipment), construction, buildings, and rents, which are not intermediate inputs, wholesale and retail trade, whose inclusion we test in Appendix B.13.3, as well as services, which could be produced in-house but have much lower energy requirements; all are counted on the seller's side as sales to the rest of the economy.[26]

Factor prices. Firms face heterogeneous factor prices in the data. When calibrating effective energy efficiency, we therefore use firm-specific factor prices, so that price differences are not loaded onto ν_j.

[26] The household in the model stands for the rest of the economy: final consumers, the public sector and the non-manufacturing sectors, whose purchases of manufacturing goods we treat as final demand. This assumes that these sectors are perfectly competitive and resell the manufacturing goods they buy as they are.

Our counterfactual exercises, by contrast, evaluate a common change in the energy price, so our baseline constructs equilibrium and counterfactual costs at identical factor prices, as in the theory, and a robustness exercise uses firm-specific prices instead. Under identical factor prices, we set the wage to the inverse markup $\frac{\varepsilon-1}{\varepsilon}$ in all scenarios and the equilibrium energy price to the average energy price faced by manufacturing firms, $p_e = 0.0830$ €/kWh. For firm-specific factor prices, we compute wages as labor expenditures divided by employment from the balance-sheet data, and energy prices as an average of fuel prices weighted by each firm's fuel quantity shares, using the fuel prices described in Appendix B.2.1.

4.2 Identification Strategy

Studying in-house production raises several challenges. Evaluating each firm's counterfactual make-or-buy decisions requires the system of intermediate input production costs, but we observe neither suppliers' prices nor in-house production costs, and both depend recursively on the upstream endogenous network. Moreover, while market transactions are observed for outsourced inputs, we do not observe the internal allocation of resources across in-house inputs. This asymmetry between the two alternatives prevents us from applying standard discrete choice methods, which require the alternatives to be characterized symmetrically (see Appendix B.7.1 for more details). Observed transactions, in turn, come with their own challenge because they are selected. But while we never observe the counterfactual demand for non-chosen alternatives, the selection itself is informative about their relative cost. Finally, because cross-producer misallocation operates at the extensive margin, affecting which producer is selected rather than how much the buyer demands from its suppliers, it cannot be captured by the wedges standard in the misallocation literature.

We tackle these challenges by exactly calibrating the model to firm-level data: each model firm corresponds to an empirical firm in the spirit of Choi et al. (2024) and Di Giovanni et al. (2024), characterized by its effective energy efficiency and a distribution of outsourcing capabilities, which enables us to recursively back out all unit costs and evaluate each make-or-buy decision in our counterfactual exercises. Exact calibration thus disciplines the model as tightly as the data allow. It reproduces the dispersion in energy intensity and firm boundaries across Italian firms without imposing a parametric family on the joint distribution of energy efficiency and outsourcing capabilities, precisely the object we set out to identify and about which little is known empirically. This is feasible because our setting delivers the full set of potential producers of each input: our data cover the universe of Italian manufacturing producers, hence all external suppliers, and we reconstruct the recipes of each product in Section 4.2.2, making each firm the potential in-house producer of its required inputs. The counterpart of this flexibility is that outsourcing capabilities absorb any unmodeled determinant of make-or-buy decisions, as well as any measurement error, so our counterfactual exercises are valid under the assumption that these residuals do not respond to energy price changes.

The calibration proceeds as follows. We first infer each firm's in-house input set by comparing the inputs it needs to those it purchases (Section 4.2.2). We then jointly recover firms' effective energy efficiency ν_j and input-specific energy requirements β_ω from variation in energy intensity and in-house portfolios

across firms and years (Section 4.2.3). Next, we calibrate outsourcing capabilities q_{ij} (Section 4.2.4), directly from intermediate input expenditure shares for observed pairs, and using revealed-preference bounds for non-selected ones. Finally, the TFP ψ_j and the demand shifters δ_ω match sales to the household (Appendices B.8 and B.9). Holding these parameters fixed across counterfactual scenarios, we build the counterfactual allocations in Section 4.3.

4.2.1 External calibration

We calibrate the elasticity of substitution between products and firms $\eta = \varepsilon = 2.5$ as in Boehm and Oberfield (2023). Following Baqaee and Farhi (2020), we set the elasticity of substitution between inputs to $\frac{1}{1-\gamma} = 0.2$ as a middle-ground between Boehm et al. (2014) and Atalay (2017)'s estimations, which is consistent with our data (Appendix B.13.4). We set the preference parameter $\alpha = 0.75$ to match the price elasticity of demand for manufacturing products (Kreinin, 1967).

Energy-to-labor elasticity of substitution: Bretschger and Jo (2021) find that energy and labor are typically strong complements across most NACE 2-digit manufacturing industries, though the degree of complementarity varies substantially. We use their estimates to calibrate our ρ_ω, aggregating our 6-digit product codes to their 2-digit industry classification. For the industries for which they don't report an estimate, we use their aggregate estimate of $\rho = -0.603$. These estimates, however, ignore in-house production. We carry a robustness exercise using an alternative strategy (described in Appendix B.6) that leverage the subsample of firms producing only their core product in-house.

4.2.2 Characterizing in-house production

To compare actual and potential producers of a given input, we need to identify *which* inputs are produced in-house by which firm, that is, the set $\mathcal{I}_j \ \forall j$. We compare the inputs that a firm needs to produce its core product to the inputs that it has access to in practice, that is, that we observe it purchases. To provide intuition, Figure 4 reports the intermediate input expenditure shares of firms producing mortar on two inputs that enter mortar production: cement and crushed minerals e.g., limestones. These inputs are observed as the most frequently purchased inputs by mortar producers. However, most firms are buying one or the other. From observing cement producers' transaction, we know that cement production requires crushed minerals. It therefore seems likely that a mortar producer buying crushed minerals, and not cement, makes the cement in-house. This is exactly the intuition behind our measurement approach.

We start by reconstructing the empirical analog of the recipes Ω_ω by identifying the complete set of inputs - both direct and indirect - that enter the production of each product ω. This set is inferred from our firm-to-firm transaction panel data structured as a directed acyclic graph (DAG), in which each node corresponds to an input and a directed edge from input $\hat{\omega}$ to input ω indicates that $\hat{\omega}$ is used in the production of ω, as in Boehm and Oberfield (2020) and Rachapalli (2021). An input used in several branches of a recipe is a single node reached through several paths, each corresponding to a stage. An input counts as a direct input of a product only if it does not already enter through one of the

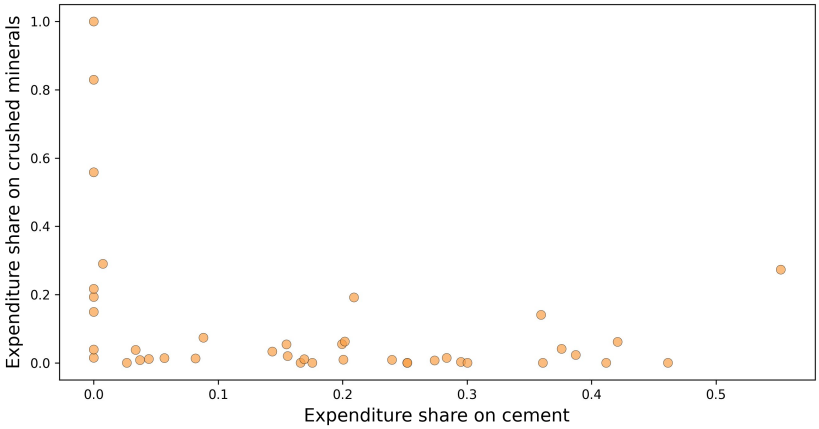


Figure 4: Input expenditure shares to produce mortar

Notes: Mortar producers' expenditure shares of total intermediate expenditures on cement and crushed minerals.

product's other inputs, so that recipes include the intermediate inputs that firms may produce in-house. For instance, although mortar producers sometimes buy limestone, other transactions show that it goes into cement, itself an input to mortar, so we attach it to cement rather than directly to mortar. Recipes are thus built recursively: each product's recipe contains the recipe of each of its inputs, which is the same wherever that input appears (see Appendix B.5). We also use the DAG to calibrate the direct input requirement $\alpha_\omega^{\hat{\omega}}$. For each input $\hat{\omega}$ directly connected to ω, we set $\alpha_\omega^{\hat{\omega}}$ proportional to the cross-firm, cross-year average expenditure on $\hat{\omega}$ as a share of costs.

Rather than exhaustively enumerating all inputs that could be used in niche or context-specific cases, we focus on the *crucial* inputs that systematically enter production and are plausible candidates for in-house production: we only consider transactions that compose more than 3% of the total input expenditures of a product in order to avoid picking up weak links. To also exclude one-off purchases, we compute these shares in each year and keep an input only if its share exceeds 3% on average across years and in at least two years.

Using the set of direct upstream inputs $\mathcal{U}_\omega^d$ and the recipe Ω_ω, combined with the set $\mathcal{O}_j$ of inputs purchased by each firm,[27] we can identify the set of inputs that they produce in-house. The key underlying assumption behind our strategy is that all firms producing the same product use the same recipe, as in e.g., Alfaro et al. (2019), Rachapalli (2021), and Alfaro et al. (2024).[28]

In-house input set. The set $\mathcal{I}_j$ of inputs produced in-house in 2019 by a firm producing ω consists of the inputs of its recipe Ω_ω that the firm reaches, moving upstream from ω, without passing through an input it purchases in that year. We exclude leaf inputs, that is, the most upstream inputs of each

[27] When a firm outsources a direct input $\hat{\omega}$, it purchases the fully assembled intermediate from a supplier, which is responsible for procuring all of $\hat{\omega}$'s own sub-inputs, so these upstream components are bundled into the transaction. Accordingly, the firm's outsourcing set $\mathcal{O}_j$ contains only direct inputs of either its core product ω_j^c or an input it produces in-house ($\omega \in \mathcal{I}_j$).

[28] Empirically, data limitations lead us to define a product by its six-digit code, which may still pool more granular products with different recipes; insofar as these products are similar, their recipes largely overlap, limiting the measurement error. Finer disaggregation would instead risk including perfectly substitutable varieties as distinct inputs in recipes, which, in our model, are sets of crucial inputs that are all necessary for production.

recipe, from in-house production. These inputs are conceptually closer to raw materials or extracted resources than to intermediate inputs that must be processed.[29] They are also likely to be more easily substitutable in the recipes. To capture all inputs produced in-house, we build an algorithm that moves upstream along the supply chains:

1. For each firm j, we start from the core product ω_j^c and consider all inputs $\hat{\omega}$ that are directly upstream of product ω_j^c ($\hat{\omega} \in \mathcal{U}^d_{\omega_j^c}$). These inputs are produced in-house (i.e., added to the set $\mathcal{I}_j$) if they are not purchased by the firm ($\hat{\omega} \notin \mathcal{O}_j$) and if they are not a leaf.

2. We move upstream and consider the in-house inputs $\hat{\omega} \in \mathcal{I}_j$ as the "products" and repeat Step 1. That is, we consider all inputs that are directly upstream of the in-house inputs $\hat{\omega}$ (inputs $\tilde{\omega} \in \mathcal{U}^d_{\hat{\omega}}$). Again, these inputs are produced in-house (i.e., added to the set $\mathcal{I}_j$) if they are not purchased by the firm ($\tilde{\omega} \notin \mathcal{O}_j$) and not a leaf.

3. We repeat Step 2 for each in-house input. In effect, we stop our algorithm every time we reach a purchased input or a leaf.

Our definition of in-house production corresponds to a narrowly defined form of vertical integration: inputs are classified as in-house only if their production is not observed, through a transaction, to take place in a separate firm. In practice, vertical integration takes many forms, and some subsidiaries might remain separate entities that would be identified as distinct firms in the data. Our empirical strategy does not count production carried out in such subsidiaries as in-house for the parent. We adopt this restrictive notion because it is most consistent with our assumption that all stages performed within a firm share a common energy-biased technology, an assumption that is less likely to characterize explicitly distinct subsidiaries. Since we do not consider financial frictions or scale economies, our aggregate results are not affected by this definition: affiliates classified as separate will be characterized by a similar effective energy efficiency if they are actually part of the same firm and our model is correctly specified. For our calibration, the relevant boundaries are thus those that delineate the limit of our empirical firms, as those are the ones that determine the firm-level variables that we match.

We find that firms produce on average 5 of their required inputs in-house (12% of their recipe), but none for the median firm, with substantial dispersion within and across sectors (Table B.7). Appendix B.5.1 reports further descriptive statistics and validation exercises.

4.2.3 Energy parameters

We now turn to identifying the key object of the paper: firms' effective energy efficiency ν_j. Firms differ in the energy they use both because they differ in efficiency and because they produce different sets of inputs in-house, and inputs themselves differ in how much energy they require, with more upstream

[29] We could not find evidence of any Italian manufacturing firm that has integrated the extraction of its primary material. The industry largely relies on buying and importing. Excluding them from the in-house set should therefore not bias our any of quantitative results.

products being systematically more energy intensive (Shapiro, 2021). Isolating ν_j therefore requires separating it from the energy requirement β_ω of each input, so that this heterogeneity is not loaded onto our firm-level measure. This is not straightforward: we observe no direct data on how firm-level energy demand splits across in-house inputs, and because in-house input sets are highly heterogeneous, we rarely observe firms producing the exact same bundle of inputs. Our in-house identification procedure nevertheless recovers which inputs each firm produces in-house, and partial overlaps between these sets suffice to separate input requirements from firm efficiency. Thus, we estimate β_ω and ν_j jointly from the energy-to-labor expenditure ratio:

$$\frac{p_e e_j}{w l_j} = \frac{\sum_{\omega \in \mathcal{I}_j \cup \{\omega^c\}} \tilde{\alpha}^{\omega}_{\omega^c}(j) \alpha^{el}_{\omega} \beta_\omega c_j^{el}(\omega)^{\frac{\rho_\omega - \gamma}{(1-\rho_\omega)(1-\gamma)}} \left(\frac{p_e}{\nu_j}\right)^{\frac{-\rho_\omega}{1-\rho_\omega}}}{\sum_{\omega \in \mathcal{I}_j \cup \{\omega^c\}} \tilde{\alpha}^{\omega}_{\omega^c}(j) \alpha^{el}_{\omega} (1-\beta_\omega) c_j^{el}(\omega)^{\frac{\rho_\omega - \gamma}{(1-\rho_\omega)(1-\gamma)}} w^{\frac{-\rho_\omega}{1-\rho_\omega}}} \tag{15}$$

Where $c_j^{el}(\omega) = \left(\beta_\omega \left(\frac{p_e}{\nu_j}\right)^{\frac{-\rho_\omega}{1-\rho_\omega}} + (1-\beta_\omega) w^{\frac{-\rho_\omega}{1-\rho_\omega}}\right)^{\frac{-(1-\rho_\omega)}{\rho_\omega}}$. The logic mirrors two-way fixed effects estimators à la Abowd et al. (1999): β_ω is a product-level parameter common across firms, while ν_j is a firm-level technology common to all the stages the firm performs. The two are separately identified because in-house input sets overlap across firms and firms face different factor prices. We estimate them by non-linear least squares, pooling the years 2019 and 2021-2023 and adding industry×year and technology change controls. We obtain a median $\hat{\beta}_\omega$ of 0.031 and a mean of 0.047, all estimates strictly interior to the parameter bounds, close to the 0.05-0.075 the literature typically calibrates (e.g., Casey et al. (2026)). Appendix B.6.2 describes the implementation, the identification requirements and a robustness exercise targeting measurement error in our recipes, and Appendix B.6.3 reports additional stylized facts.

Effective energy efficiency. At this stage, we can already document descriptive patterns in the cross-producer allocation of production by effective energy efficiency. To our knowledge, this is the first firm-level evidence on this margin. These results rely on the calibrated objects $\{\mathcal{I}_j, \nu_j, \{\nu_{k_j(\omega)}\}_\omega\}_j$, the identification of which requires only that labor and energy are combined through a CES aggregator, that recipes are sequential and common across producers, and, for producers with in-house intermediate stages, the normalization of input requirements described in Section B.5. Crucially, it does not depend on how we model outsourcing factors, nor on whether these factors reflect distortions, fixed costs, or other mechanisms.

First, we show the importance of controlling for firm boundaries when characterizing firms' effective energy efficiency. Figure 5 compares the effective energy efficiency calibrated with (ν) and without (ν^{naive}) controlling for firm boundaries. The bias ($\log(\nu) - \log(\nu^{naive})$) increases with ν, consistent with energy efficient firms integrating more energy-intensive inputs in-house. In other words, abstracting for firm boundaries shrinks the distribution of effective energy efficiency, as it makes energy efficient (inefficient) firms look less (more) efficient due to their relatively higher (lower) use of energy for in-house production (Figure B.5d).

Next, we show that make-or-buy decisions are far from solely based on effective energy efficiency con-

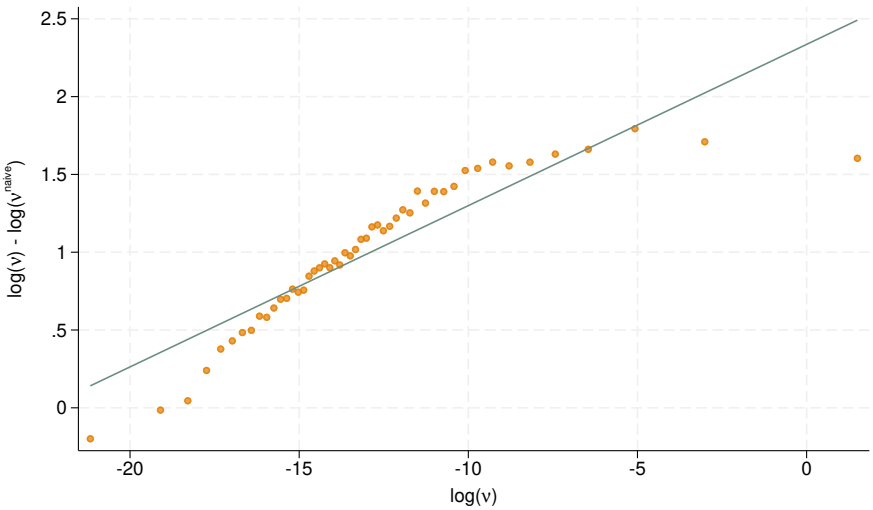


Figure 5: Difference in calibrated ν with and without make-or-buy, by level of ν

Notes: The graph plots the ratio (in logs) of firms' effective energy efficiency (ν) calibrated with and without accounting for firm boundaries, against the level of ν. The relation is measured for all single-product manufacturing firms in the sample.

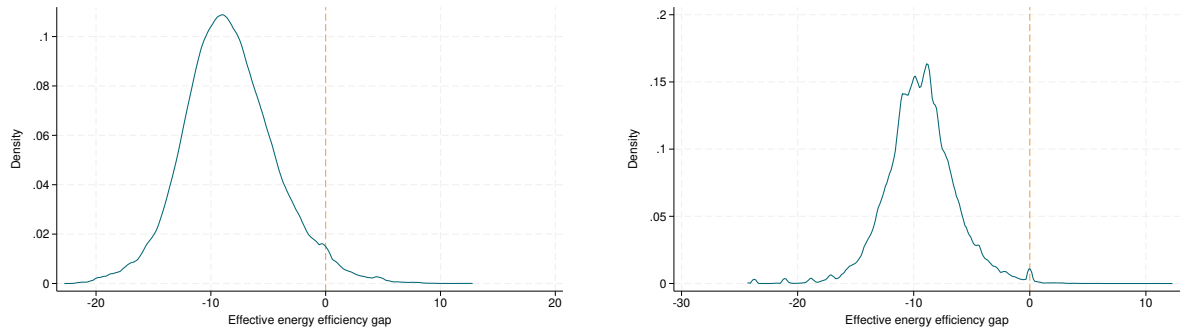
siderations, which might lead to substantial energy-use misallocation. Figure 6 reports the distribution, across inputs ω, of the effective energy efficiency of the actual producer $k = k_j(\omega)$ relative to that of the most energy efficient alternative producer available, $\log(\nu_k) - \log(\nu_{s^{eff}})$. Negative values indicate that the actual producer is the less energy-efficient of the two. For inputs produced in-house (Panel 6a), the actual producer is the using firm itself, and the alternative is the most energy-efficient external supplier: only approximately 3% of the inputs produced in-house do not have a more energy-efficient external option.[30] For outsourced inputs (Panel 6b), the actual producer is the selected supplier, and the alternatives include the in-house option: these inputs, too, are rarely produced by the most energy-efficient option, which is usually another external supplier (Figure B.7d). Compared with the in-house option alone, the vertical allocation is still not aligned with energy efficiency for more than half of them (Figure B.7c). Finally, in line with energy efficiency receiving a higher weight in the decision of less distorted inputs, Figure B.7b shows that the energy efficiency gaps of EU ETS-regulated inputs are smaller in magnitude.

Absent any other form of heterogeneity, these gaps would indicate misallocation. However, make-or-buy decisions also depend on other technological factors, captured by our outsourcing capabilities, which we identify next. These gaps are therefore necessary but not sufficient for misallocation: they may be justified, at least in part, by gains from outsourcing unrelated to energy.

They nonetheless already inform the direction of the potential misallocation. With many external suppliers but a single in-house option, and no systematic energy-efficiency advantage in producing in-house (Figure B.8a), most firms producing in-house have a more energy-efficient external alternative (Figure 6a).[31] Any extensive-margin reallocation induced by higher energy prices will therefore run mainly from

[30] This is robust to weighting inputs by importance using their total input requirements (Figure B.7a), and to comparing firms to the 95^{th} percentile supplier (Figure B.8a). In contrast, comparing in-house producers to their average supplier (Figure B.8a) suggests that there is no systematic gap between downstream firms and upstream suppliers, supporting our choice to calibrate ν_j as a firm-specific parameter, rather than a core-product component.

[31] The relevant comparison is the energy efficiency of the entire upstream supply chain, not firm-level ν alone. The



(a) In-house inputs (b) Outsourced inputs

Figure 6: Producer's effective energy efficiency relative to most energy efficient alternative

Notes: The figure reports the distribution of the selected producers' effective energy efficiency (ν_k) relative to that of their most energy-efficient alternative ($\nu_{s^{eff}}$), in log points. Panel (a) compares firms to their most energy-efficient external supplier ($s^{eff} \in \mathcal{S}_\omega$), for all their inputs produced in-house ($\omega \in \mathcal{I}_j$). Panel (b) compares the selected external supplier k to the most energy-efficient alternative $s^{eff} \in \{j\} \cup \mathcal{S}_\omega \setminus \{k\}$, for all outsourced inputs ($\omega \in \mathcal{O}_j$). The distributions are across all inputs produced in-house (panel (a)) or outsourced (panel (b)) by all single-product manufacturing firms in the sample. The vertical line marks a zero gap.

in-house to outsourcing or across external suppliers, a qualitative prediction confirmed in our quantitative analysis.[32]

4.2.4 Outsourcing capabilities

Outsourcing capabilities shape both margins of sourcing decisions: they determine whether production is carried out in-house or outsourced, and, conditional on outsourcing, they affect the value transacted. Estimating the entire vector of q_{ij} is complicated because the latter is only observed for a selected sample, but we use information from the former to account for this selection. Before describing our calibration strategy, it is useful to summarize the different types of relationships that we try to characterize through the outsourcing capabilities.

(a) **Observed transactions:** Firm j buys input ω from firm i and the transaction is directly observed in the data. For these relationships, the outsourcing capabilities q_{ij} are calibrated as the residual in intermediate input expenditure shares (Equation (14)), controlling for prices and iceberg trade costs δ_{ij}.[33] We build prices recursively using the previously identified effective energy-efficiency and in-house input set (see Appendix B.7.2).

(b) **Observed decisions:** Firm s sells input ω, but firm j buys input ω from firm $i \neq s$ or produces it in-house. Thus, no transaction between firm j and firm s is observed in the data. However, we know that firm s was not chosen by firm j as a supplier.

mechanism should carry over, though the set of firms with a more efficient alternative may differ.

[32] This does not mean that in-house production would be chosen at a *more* distorted energy price either: in-house production is not dominated by every outsourcing option, as Figure B.7c makes clear, but typically by at least one.

[33] In line with gravity models, we observe that buyers tend to select geographically closer producers. Because the first-order conditions only identify the effective price of observed transactions, the calibrated outsourcing capabilities of distant suppliers embed their trade costs. We therefore separate an iceberg trade cost $\delta_{ij} \geq 1$, increasing in buyer-supplier distance, from a distance-free capability q_{ij}, so that the effective price is $\delta_{ij} c_i(\omega)/q_{ij}$, and calibrate the distribution of q_{ij} jointly with the trade-cost elasticities so as to match how far observed suppliers are located within buyers' choice sets (Appendix B.7.3).

(c) **Unobserved decisions:** Firm s sells input $\hat{\omega} \in \Omega_\omega$, but firm j outsources input ω, so that $\hat{\omega}$ enters j's recipe only at stages left to its supplier. Firm j's internal make-or-buy decision on $\hat{\omega}$ still shapes its in-house cost of ω, and hence its decision to outsource it, but it is not implemented, and therefore not observed.

We thus work in two steps: we calibrate the q_{ij} for all observed pairs (types (a)), and then retrieve the q_{ij} for all unobserved ones (types (b) and (c)) by making a functional form assumption of their underlying distribution accounting for selection in a model-consistent way. We explain this second step below, while a more detailed description of our entire strategy can be found in Appendix B.7.

We start by assuming that outsourcing capabilities are lognormally distributed, $\log q_{ij} \sim \mathcal{N}(\mu, \sigma^2)$, and calibrate μ and σ by matching empirical moments to simulated moments of the observed sample (type (a)) through a Simulated Method of Moments (SMM) procedure. This sample is, however, selected: the q_{ij} of the non-chosen alternatives (type (b)) are unobserved, but, by a revealed preference argument, they are bounded from above by the cost ratio that drives the make-or-buy decisions. We simulate them below these upper bounds, use them to derive lower bounds on the capabilities of chosen relationships, and draw the simulated selected sample above these lower bounds, so that the simulated moments reflect the same selection as their empirical counterparts. Since we are interested in the relationship between firms' effective energy efficiency and their propensity to outsource, we also allow for a within-buyer component in outsourcing capabilities, modeled with a Student-t Copula $\mathcal{C}(q_{ij}, \dots q_{kj}; \rho_c, \tau_c)$ whose parameters we calibrate along with (μ, σ). Because simulating the unobserved capabilities requires the parameters $\hat{\theta} = \{\mu, \sigma, \rho_c, \tau_c\}$ that these capabilities then contribute to calibrate, we proceed iteratively. Given a guess of $\hat{\theta}$, each iteration proceeds as follows:

1. **Non-selected pairs.** For each buyer j and each non-selected external supplier $s \neq k_j(\omega)$ of input $\omega \in \mathcal{O}_j \cup \mathcal{I}_j$, we simulate q_{sj} from the lognormal truncated from above at $\bar{q}_{sj}$, with within-buyer dependence governed by the Copula. The truncation enforces consistency with the observed selection of producer $k_j(\omega)$:

   - If the input ω sold by s is produced in-house by firm j ($k_j(\omega) = j$): $\bar{q}_{sj} = \frac{\delta_{sj} c_s}{c_j^{\mathcal{I}}}$
   - If the input ω sold by s is outsourced to firm i by firm j ($k_j(\omega) = i$): $\bar{q}_{sj} = \frac{\delta_{sj} c_s}{\delta_{ij} c_i / q_{ij}}$

2. **Selected pairs.** Using the simulated alternatives (q_{sj}), we construct for each observed (i, j) pair a pseudo-lower bound

$$\tilde{q}_{ij}^m = \max_{s \neq i} \left\{ \frac{\delta_{ij} c_i \, q_{sj}}{\delta_{sj} c_s} \right\}$$

   Where s includes the in-house option, with $c_s = c_j^{\mathcal{I}}$, $\delta_{jj} = 1$ and $q_{jj} = 1$.[34] $\tilde{q}_{ij}^m$ represents the tightest lower bound on q_{ij} implied by the best, non-chosen alternative.[35] We draw q_{ij}^{sim} from our lognormal, conditional on $q_{ij}^{sim} \geq \tilde{q}_{ij}^m$, and compute simulated moments on this selected sample.

[34] This is a counterfactual, unobserved cost. We construct it recursively applying the same algorithm that we use to build any counterfactual allocation, but on q_{ij} drawn from the lognormal distribution evaluated at the parameter set $\hat{\theta}$ of the current iteration.

[35] The draws in Step 1 cannot be bypassed because the pseudo-lower bound is a function of observed suppliers costs,

The mean and variance of $\log q_{ij}$ anchor μ and σ. The Copula is identified by the average within-buyer variance of $\log q_{ij}$ and by the probabilities that at least two of a buyer's draws fall in the top 10% and top 2%, which separate the correlation ρ_c from the tail dependence τ_c.

We then update $\hat{\theta}$ to minimize the distance between empirical and simulated moments, and iterate until convergence. Once the parameters have been identified, we draw the unobserved q_{sj} from the same lognormal, independently across buyers and with within-buyer dependence governed by the Copula, truncating at the pair-specific bound for type (b) so that the draw is consistent with the observed selection, and leaving the draw untruncated for type (c), for which no decision is observed. Effective outsourcing costs still differ across pairs, through iceberg trade costs and through suppliers' own TFP, energy efficiency, and make-or-buy decisions.

Table 1: SMM estimates of the firm-pair outsourcing-capability distribution

Panel A: Estimated parameters			*Panel B: Moments*		
Parameter	Description	Estimate	Moment	Data	Simulation
μ	Mean	-5.21	$\mathbb{E}[\log q_{ij}]$	2.2154	2.2060
σ	Standard Deviation	2.24	$\mathrm{Var}(\log q_{ij})$	27.7465	23.3595
ρ_c	Within-buyer Copula correlation	0.108	$\mathbb{E}_j[\mathrm{Var}_\omega(\log q_{ij} \mid j)]$	15.2917	13.9374
τ_c	Copula degrees of freedom	2.05	Pr (≥ 2 q in top 10%)	0.0282	0.0219
			Pr (≥ 2 q in top 2%)	0.0015	0.0008

Notes. Parameters estimated by simulated method of moments on firm-to-firm transactions. Pairwise outsourcing capability follows $logq_{ij} \sim \mathcal{N}(\mu, \sigma^2)$; dependence across the inputs a given buyer sources is a one-factor Student-t Copula with correlation ρ_c and degrees of freedom τ_c. Panel B reports the empirical and simulated moments, both computed on the selected sample. The trade costs moments and results are reported in Appendix B.7.3.

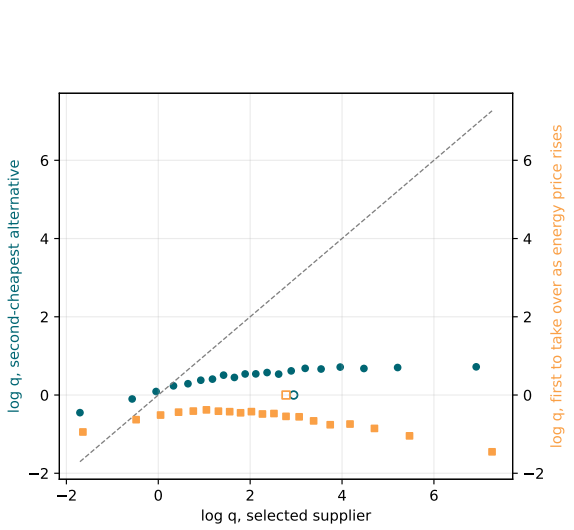
Table 1 reports the results of the calibration, also illustrated in Figure B.11a. The four parameters split into a marginal block and a dependence block. The marginal parameters, the mean μ and the variance σ^2, are the key driver of extensive margin reallocation: a large μ favors all outsourcing alternatives relative to in-house production, and can be expected to trigger in-house to outsourcing switches. Selected suppliers are characterized by outsourcing gains and are located nearby: once iceberg costs are counted, $E[\log(q_{ij}/\delta_{ij})] = 1.57$ for the selected suppliers, while the calibrated model implies $E[\log(q_{ij}/\delta_{ij})] \approx -6.73$ for a randomly drawn supplier located at mean distance. A large variance σ^2 reflects higher product differentiation across firms. Its effect on reallocation depends on where the untruncated mean μ sits relative to the pair-specific truncation bounds: if μ lies below them on average, which is the case in our calibration, a small σ^2 concentrates non-selected suppliers around that low value and makes them uncompetitive relative to the incumbent; if μ lies above them, a small σ^2 piles them up right under their respective bounds, making them close substitutes for the selected option. $\rho_c = 0.108$ is the within-buyer Copula correlation: a rather modest positive common component, meaning that a buyer that is an above-median outsourcer for one input tends to only be slightly so for its other inputs

counterfactual in-house costs, and unobserved outsourcing capabilities, thus an object for which the distribution is unknown. We therefore cannot use a closed-form result that would characterize that maximum.

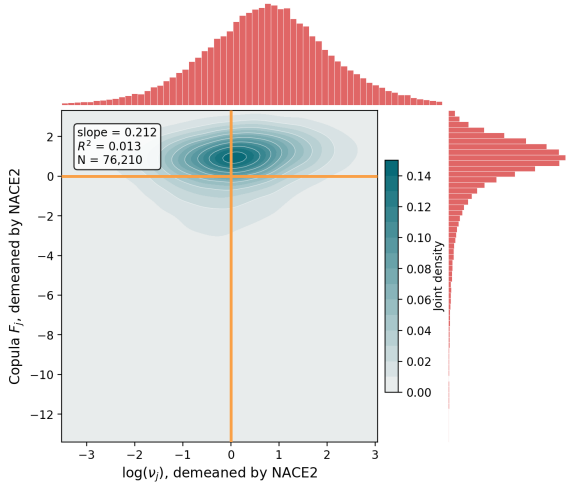
(implied Kendall's $\tau \approx 0.069$, Spearman ≈ 0.094). Finally, the Copula's degrees of freedom $\tau_c = 2.05$ indicates strong joint-tail dependence within buyer.

Figure 7a compares the outsourcing capability of each selected pair with those of two alternatives: the cheapest non-selected option, and the first option to take over as the energy price rises (Figure B.11b does the same for the alternative with the highest ν). Selected pairs mostly have higher capabilities, as expected, but points above the 45° line show that selection depends on unit costs, not on capabilities alone: a pair can be chosen over an alternative with a higher q. Moreover, the first option to take over is not necessarily the cheapest one at the equilibrium price, since higher prices shift weight from outsourcing capabilities toward energy efficiency.

The correlation between buyers' effective energy efficiency and their Copula within-buyer component is 0.0433^{***} (0.0241^{***} within sectors), meaning that energy efficient firms are, if anything, slightly more likely to outsource (precisely the trade-off of Lemma 2). This correlation, however, hides large heterogeneity. Figure 7b, the closest empirical counterpart of Figure 2, plots each buyer's energy efficiency ν_j against its within-buyer component F_j: a substantial mass of buyers lies in the first and third quadrants, where they face this trade-off, especially so in the first (energy-efficient firms with high outsourcing capabilities). Yet neither this correlation nor the joint distribution is a sufficient statistic for misallocation in the full production network (Appendix A.3.3): instead, we need to construct the counterfactual optimal allocation.



(a) Selected q relative to second cheapest



(b) Within-buyer component of outsourcing capabilities and effective energy efficiency

Figure 7: Calibrated distribution of outsourcing capabilities

Notes: Figure a) plots the selected pair's outsourcing capabilities against the outsourcing capabilities of the cheapest non-selected option (left y-axis) and against the first option that will be picked as energy price rises (right y-axis). White dots mark cases in which the alternative is in-house production, whose $q_{jj} = 1$. Figure b) plots the joint distribution of the effective energy efficiency (ν_j) and the Copula within-buyer component of the outsourcing capabilities (F_j) of each firm, both demeaned by NACE 2 sector. The orange lines split the graph into four quadrants.

4.2.5 Out-of-sample fit

The calibrated outsourcing capabilities $\{q_{ij}\}$ are, by construction, fitted to the observed network: they rationalize the make-or-buy and supplier choices we see but doing so, absorb any residual variation.

We therefore submit the calibration to two out-of-sample tests that its own moments do not target to clarify what our model can predict as well as its limitations. First, we use our parameters calibrated to 2019 moments to predict the allocation under the 2021, 2022 and 2023 energy prices (Table B.11). The model identifies well which in-house inputs are more likely to be outsourced: those in the top decile of predicted probability are outsourced about twice as often as those in the bottom decile. Because energy prices are the only change we feed the model, it cannot account for the technological shocks and supply chain disruptions that drove some of the reallocation observed between 2019 and 2023: about 4% of the inputs it expects to remain in-house are outsourced anyway, and it predicts almost none of the insourcing observed in the data. Since a more energy-efficient external alternative exists for most firms producing in-house, higher energy prices alone mostly trigger outsourcing, and insourcing in the data must come from other shocks.

Our next exercise examines one such shock directly: when a buyer's supplier exits, which sets the corresponding outsourcing capability to zero. In practice, we remove the selected supplier from a buyer's choice set and ask the model to predict what the buyer does instead, then confront that prediction with a choice observed in a *different* year in which that supplier was indeed unavailable (Table B.12). The predicted and realized probability of internalizing a given input align closely. Its failure to predict insourcing in the first exercise therefore reflects the shocks it is not fed, not an inability to generate insourcing. These external shocks could be matched by adding exogenous shifts to the outsourcing capabilities of particular groups of firms, but they are unrelated to the energy price change, so fitting them would improve the match without informing the mechanism we study. Keeping the energy price as the only shock is what allows us to attribute the predicted reallocation to the environmental externality.

4.3 Building counterfactual allocations

We evaluate counterfactual energy prices by re-solving the allocation, letting production reallocate across producers at both the intensive and the extensive margin, and decomposing the resulting change in aggregate outcomes into these margins. The key assumption behind all our counterfactual exercises is that the parameters $\{\nu_j, \{q_{ij}\}_{i \in \mathcal{S}_\omega,\, \omega \in \Omega_{\omega_j}}\}_j$, as well as the product recipes Ω_ω, are fundamentals of the domestic firms and their products, that is, they don't respond to the counterfactual energy price change. In our baseline, we always consider shocks that apply to both domestic and foreign prices (i.e., policies that include carbon-border adjustments) to abstract from offshoring and foreign carbon leakage effects.

Producer selection. We compute the unit cost of outsourcing and of in-house production for each firm and each input in its recipe, under the counterfactual price. We proceed recursively moving downstream, starting from inputs that only use leaf-inputs. For each input, we then compare the best outsourcing option to the in-house option to infer the make-or-buy decision of the firm. We characterize this decision for all inputs, since each one enters the in-house cost of the inputs downstream of it. Due to the sequentiality of production, however, decisions at stages located upstream of an outsourced input are not implemented: these stages are left to the supplier.

Sales shares. Aggregate energy intensity weighs firm-level energy intensity by sales shares (to manufacturing value-added).[36] We compute them recursively moving upstream. As explained in Appendix B.10.1, we start by computing the sales ratio of the firms who only sell to the "rest-of-the-economy" sector, which corresponds to our household. We then move upstream to firms that sell only to firms for which we have already computed sales, until we've covered all firms.

Counterfactual analysis. In all exercises, we can decompose the different effects:

1. **Intensive margin:** For a given selection of producers, firms reallocate resources between energy and labor for each of the inputs they produce in-house. They also reallocate resources across intermediate inputs. Moreover, the household shifts demand across firms and products. We measure this effect by keeping the equilibrium selection of producers fixed.

2. **Cross-seller reallocation:** Our producer selection includes both the choice of producing in-house or outsourcing, and the selection of the external supplier, conditional on outsourcing. We can measure the respective contributions of each of these sub-decisions by computing the gains of reallocating production across sellers keeping the equilibrium firm boundaries fixed.

3. **Vertical reallocation:** This channel captures the effect of energy misallocation across producers of the same input, specifically induced by the make-or-buy decision. We measure this margin as the aggregate effect net out of the intensive margin and the cross-seller reallocation.

5 Quantifying energy use misallocation

We are now in a position to quantify the environmental costs of make-or-buy decisions and the corresponding allocative gains from carbon pricing. To do so, we take the calibrated model and compute the change in aggregate energy intensity and welfare as a function of the price of energy. We measure the aggregate effect of misallocation as the gains from imposing a tax on energy of 0.06878 €/kWh, an 83% increase over the equilibrium price. This tax corresponds to a social cost of carbon of 190 €/tCO_2 (Ricke et al., 2018; Pindyck, 2019; Prest et al., 2023) and, under linear damages and a linear mapping from energy to emissions, achieves the socially optimal energy price.

Aggregate energy intensity falls by 47.6% and welfare rises by an equivalent variation of 2.10% of labor income relative to the observed equilibrium (Table 2).[37] Producer selection accounts for a fifth of the fall in energy intensity (19.1% of the log change) and 10.7% of the welfare gain (0.224 of the 2.10%): this is the cross-producer energy-use misallocation induced by underpriced energy. Including the firm itself as a potential producer matters particularly, since allowing firms to switch between in-house production

[36] In the equilibrium allocation, we observe total sales in the data. We match them using TFP and the demand shifter as described in Appendix B.8 and B.9.

[37] As discussed by Shapiro (2021), consumption plays a much larger role than energy in the welfare function, potentially due to the fact that the damage function is calibrated to modest historical warming, which may poorly reflect the cost of large future warming.

and outsourcing, on top of cross-seller reallocation, adds 10.1% of the log fall in energy intensity and 6.2% of the welfare gain, more than half of the welfare contribution of producer selection.

To put these results into context, we consider a 10% increase in energy prices. Since manufacturing value added is proportional to labor income in our model, and with the wage fixed by the numeraire, we can interpret the change in energy intensity as the aggregate elasticity of substitution between energy and labor in manufacturing, $-\Delta \ln(E/PY)/\Delta \ln p_e = 1.12$. Quantitative work finds a short-run aggregate elasticity of substitution between energy and non-energy inputs near zero and a long-run elasticity around one (Atkeson & Kehoe, 1999; Hassler et al., 2021). This gap is attributed to aggregate energy-saving technological change, which come from improvements in firms' technologies, on which the literature has mostly focused through the lens of directed technical change (Acemoglu et al., 2012; Fried, 2018; Casey et al., 2026), but also from the reallocation of production toward more energy-efficient firms. Our model isolates the latter, holding every firm's technology fixed, and separates reallocation through producer selection from reallocation at the intensive margin. This decomposition matters: without producer selection, the elasticity is 0.87: energy and labor are aggregate complements, and energy expenditure rises by 1.2% relative to labor expenditure when the energy price rises by 10%. Producer selection adds 0.25 and turns this rise into a fall (−1.1%): energy and labor become aggregate substitutes. As energy becomes more expensive, firms shift their sourcing toward energy-efficient producers, whose costs rise least, so that energy demand falls by 13.1% instead of 11.1%, while value added falls by 3.32% instead of 3.36%. Part of aggregate energy-saving technological change therefore comes from reallocation alone, without any firm improving its own technology, and a model with fixed producer selection would wrongly conclude that energy and labor are aggregate complements.

Firm-level estimates cannot be aggregated into this macro elasticity. Plant- and firm-level studies of French manufacturing find elasticities of energy use of about 0.5-0.6 for total energy and electricity, and 0.9 for gas (Marin & Vona, 2021; Fontagné et al., 2024); identified relative to producers in the same sector and year, they exclude entry, exit and reallocation across producers. In our model, a producer's own response, holding its efficiency and sourcing fixed, is close to these estimates, with an elasticity of about 0.47. The aggregate elasticity is larger because, as in Oberfield and Raval (2021), it combines producer-level substitution with reallocation across producers, which here also operates at the extensive margin, through producer selection.[38] Measuring the energy saved through this reallocation requires a framework that identifies both where production moves and the energy efficiency of the producers that gain it.

The magnitude of our intensive-margin macro-elasticity is driven by the elasticities of substitution within

[38] In their framework, the macro elasticity combines plant-level substitution with reallocation across plants and industries, which is driven by demand and by substitution between materials and primary factors. Producer selection spreads across all of these: supplier switching changes plants' sizes through demand, changes in how much plants outsource is reflected in materials substitution, and changes in which inputs they produce in-house alter their observed factor intensities, which would be read as plant-level substitution. Their materials margin, however, has a constant elasticity and a factor content common to all buyers, while our model microfounds it through make-or-buy decisions: it is not constant, and the energy embodied in purchased inputs depends on which producer makes them.

Table 2: Counterfactual energy intensity and welfare responses

	Energy price increase, $\Delta p_e/p_e$		
	+10%	+83%	+100%
Panel A. Energy intensity E/PY, % change			
Intensive margin only	−8.00	−40.76	−45.14
+ Cross-seller reallocation	−8.52	−44.12	−48.35
+ Vertical reallocation (all margins)	−10.13	−47.65	−51.72
Extensive margin contribution, %	**21.9**	**19.1**	**17.5**
Panel B. Welfare, equivalent variation in % of labor income			
Intensive margin only	0.608	1.873	1.843
+ Cross-seller reallocation	0.641	1.967	1.939
+ Vertical reallocation (all margins)	0.726	2.097	2.070
Extensive margin contribution, %	**16.2**	**10.7**	**11.0**

Notes: The table reports counterfactual changes in manufacturing energy intensity (Panel A) and Italian welfare (Panel B) for three increases in the energy price (+10%, +83% and +100%), starting from $p_e = 0.0830$€/kWh. Welfare changes are equivalent variations, exact because utility is quasilinear in the outside good, expressed in percent of labor income, which equals household income in the observed equilibrium. Energy intensity is the total energy used by the manufacturing firms in the sample (in quantity) relative to their total output sold to the domestic household. The damage term in Panel B uses consumption-based energy: the energy embodied in domestic manufacturing supply, including the energy used abroad to produce imported intermediates, which is also the base of the rebated carbon tax. Within each panel, margins of adjustment are added cumulatively: the first row holds both the cross-firm and the vertical (in-house vs. outsourcing) allocation fixed, so only the intensive margin adjusts; the second row also allows reallocation across sellers; the third row allows all margins. The extensive-margin share is the part of the total effect (third row) not delivered by the intensive margin alone (first row). For energy intensity, a ratio that responds multiplicatively, it is computed in logs: $\ln(1+\Delta_{\text{all}}) = \ln(1+\Delta_{\text{int}}) + R$, and the share is $R/\ln(1+\Delta_{\text{all}})$. For welfare, a difference in levels whose components (surplus, rebate and carbon terms) add by construction, it is computed arithmetically: $\Delta W_{\text{all}} = \Delta W_{\text{int}} + (\Delta W_{\text{all}} - \Delta W_{\text{int}})$, and the share is $(\Delta W_{\text{all}} - \Delta W_{\text{int}})/\Delta W_{\text{all}}$. Both decompositions are exact and additive across margins.

firms, between energy and labor ($\frac{1}{1-\rho_\omega} < 1\ \forall\omega$) and between inputs ($\frac{1}{1-\gamma} = 0.2$), and across firms and products in household demand ($\eta = \varepsilon = 2.5$), as shown in equation (25). The extensive margin contribution comes solely from reallocation across potential producers. They are perfect substitutes, but reallocation occurs discretely, and only if i) there exists at least one alternative producer whose cost rises less with the energy price, that is, whose supply chain is more energy efficient, and ii) this cost advantage is not offset by the associated loss in outsourcing capability. Table B.14 reports the share of pairs for which a more energy-efficient sourcing option exists (to a first order), satisfying (i), and the share of these potential switchers that actually switch when the energy price changes,[39] which reflects how often (ii) binds. Condition (i) is satisfied almost everywhere, consistent with the energy efficiency gaps (Figure 6): a more energy-efficient alternative exists for 97.8% of in-house inputs and 99.8% of outsourced inputs.

[39] Table B.13 confirms the force driving the switches: the higher ν_j, the more likely firms are to keep their inputs in-house, while inputs with high energy requirements (high β_ω) are the most responsive.

Condition (ii) is what binds: a 10% increase in the energy price moves fewer than 3% of potential switchers (2.4% of in-house inputs and 0.9% of outsourced inputs), and even a doubling moves fewer than a tenth. The cost advantage that the incumbent enjoys, of which the outsourcing-capability gap is the central component, is therefore large relative to the differential in energy exposure.[40] Higher carbon prices would move more, but fewer than half of the potential switchers would respond at all: the median pair requires an energy-price increase orders of magnitude beyond any plausible policy range. The pairs that do switch are those combining a small cost gap with a large differential in energy exposure, which is why their effect on aggregate outcomes is substantial: the extensive margin contribution comes mainly from the impact of the switches rather than their frequency.

Because the pairs closest to indifference switch first, the extensive margin is front-loaded, whereas the intensive margin responds almost proportionally to the log energy price. The contribution of producer selection is therefore larger for smaller price changes.[41] For a doubling of the energy price, the aggregate elasticity without producer selection is almost unchanged from the 10% case, at 0.87, but producer selection adds only 0.18, bringing the aggregate elasticity to 1.05. A 10% increase, seven times smaller than a doubling in log terms, thus already delivers almost a fifth of the producer-selection effect of a doubling, with a third as many sourcing switches.

5.1 Additional exercises

Environmental targets. Applying the requirements of the technology pathways for Italy consistent with an 80% reduction in CO_2 emissions by 2050 relative to 1990 (Gaeta et al., 2015), aggregate energy intensity should fall by approximately 53% relative to 2019. We use our model to compute the uniform carbon tax that delivers this reduction, first letting both margins adjust and then shutting down the extensive margin, as described in Appendix B.10.2. Ignoring extensive margin reallocation significantly overestimate the required carbon tax, which is 236.2 €/tCO_2 when firms can change producers, and 321.1 €/tCO_2 when only the intensive margin operates.

Non-linear damages. Static linear damages are a reasonable approximation for Italian manufacturing emissions, which are too small to affect marginal damages. However, to illustrate how the extensive margin would affect a global carbon tax, we change the welfare function to $U(X, Y, E) = X + \frac{Y^\alpha}{\alpha} - \Upsilon\left[(1-\phi_d)E + \frac{\phi_d E^2}{2E_0}\right]$, where E_0 is the baseline energy level and Υ is the social cost of carbon multiplied by the constant emission factor of energy, so that marginal damages at E_0 equal the social cost of carbon. In this exercise, we interpret proportional changes in E as proportional changes in global energy use. We set the elasticity of marginal damages with respect to energy at baseline to $\phi_d = 0.2$

[40] This cost advantage is also heterogeneous: reallocation spreads demand over more suppliers instead of concentrating it (Table B.15).

[41] Unlike its contribution to energy intensity, the contribution of producer selection to welfare rises again beyond the socially optimal price. There, energy is taxed above its social cost, so switching to more energy-efficient producers saves more energy than is socially desirable, but it still limits more the output loss by shifting production to producers that are cheaper at the higher price.

(Appendix B.10.3). The optimal energy price now depends on the counterfactual energy level E^*: $p_e^{SCC} = p_e + \Upsilon\left[(1-\phi_d) + \phi_d E^*/E_0\right]$, where $p_e = \chi$ is still fixed by the energy production side. Since E^* itself depends on p_e^{SCC}, we solve the counterfactual allocation iteratively until convergence. The optimal tax is 169€/tCO_2 with reallocation and 171€/tCO_2 when make-or-buy and supplier choices are held at their equilibrium values, both below the 190€/tCO_2 of the linear case: the tax roughly halves energy use, which lowers marginal damages. The extensive margin abates substantially more energy, but lowers the tax only slightly, since energy moves marginal damages only through their convex part, of weight ϕ_d.[42] Welfare gains are smaller than under linear damages for the same reason: as emissions fall, each additional unit abated avoids less damage. These patterns hold at higher values of ϕ_d, but the gap between the two taxes widens.

Tax on fossil fuels. We consider a tax that only applies to fossil energy sources, and accordingly attribute all damages to fossil energy, which accounts for 58% of baseline energy use. We model firms' energy mix as Leontief, hence shutting down within-firm fuel substitution, for comparability with our main results. We tax fossil energy at its marginal damage, 0.110€/kWh, which raises the price of fossil energy by 133% and the aggregate energy price by 83%, as much as the uniform Pigouvian tax on energy. The tax raises welfare by an equivalent variation of 2.21% of labor income, slightly more than the energy tax (2.10%), with the extensive margin contributing a similar share (9.6% against 10.7%). By contrast, both the decline in aggregate energy intensity (39.5%) and the extensive margin's share of it (15.3%) are smaller. Relative to the energy tax, fossil energy decreases by more (63.4% against 58.5%): even without fuel substitution, the tax concentrates the price increase on fossil-intensive firms, which shrink their energy demand the most. At the same time, since firms that rely mostly on non-fossil energy face a smaller price increase than under the energy tax, non-fossil energy decreases by much less (30.0% against 56.3%), so that the fossil share of energy drops from 58% to 42%, against 56% under the energy tax. The extensive margin accounts for 18% of this drop.

Sector-level policies. Between-firm misallocation arises even under a common distortion, and the optimal carbon tax is uniform across firms. However, in practice, Italian manufacturing firms face a set of sector-specific environmental policies. Two important examples are the EU ETS allocation rules based on product-specific energy intensity benchmarks and the regime for energy-intensive firms in specific sectors (Imprese a Forte Consumo di Energia), which grants substantial reductions in electricity costs and compensation for indirect carbon pricing. Firm boundaries play a particularly important role under sector-specific policies, which we formally study in Appendix A.2.4, summarizing here the main theoretical intuitions. Because eligibility is determined at the sector level and sectors are assigned based on the product sold, not on inputs produced in-house, these policies create systematic differences in the cost of energy and emissions across potential *producers* of the same input. Buyers can evade or capture

[42] This contrasts with the tax required for an environmental target, which must close the entire gap in energy response (reaching the energy level of the reallocating economy at its optimum requires 209€/tCO_2 without reallocation, against 169€/tCO_2 with it); here, energy enters the tax only through the convex part of damages.

the gains of the policy through their make-or-buy decision, generating additional extensive margin misallocation. We illustrate the issue with a counterfactual exercise in which the optimal carbon tax (the 83% increase) targets only EU ETS-covered sectors (3.4% of the firms in our manufacturing sample, accounting for $\sim 14.6\%$ of output).[43] As expected, the response is muted since fewer firms are targeted: holding firm boundaries at their equilibrium, aggregate energy intensity would decrease by 9.9%. The extensive margin, however, now works against the policy: relative to the frozen boundary allocation, endogenous make-or-buy decisions raise aggregate energy intensity by 1.8%, instead of lowering it as under a uniform price increase, so that the overall decrease is only 8.2%. This aggregate masks opposite movements across the two groups: endogenous make-or-buy decisions reduce covered firms' energy intensity by an additional 15.1%, but raise uncovered firms' by an additional 4.7%, whereas they lower it by 10.6% under a uniform price increase. The extensive margin is being used to relocate production toward producers facing the untaxed price, a form of domestic carbon leakage. Holding firm boundaries fixed, the induced contraction in output would lower uncovered firms' energy use by the equivalent of 30% of the energy covered firms save; with make-or-buy decisions, it instead rises by 9% of these savings (the domestic leakage rate). Relocated stages also use more energy: at the untaxed price, even an equally efficient producer performs them with roughly 27% more energy per unit of physical output, and the producers actually selected are on top of that often less energy efficient (18.5% of the switched inputs are produced with a worse supply-chain embodied energy efficiency relative to the equilibrium allocation, against 1.4% in the baseline). As a final note, these policies usually rely on measures of energy use per unit of output which, as we have shown, conflate true energy efficiency with organizational choices, resulting in a distorted targeting. These considerations underscore the importance of explicitly accounting for firm boundaries when designing and evaluating energy policies.

5.2 Sensitivity of the results

We have two separate blocks of results: the results on energy efficiency (Section 4.2.3), which rely on our calibration of the energy parameters $(\beta_\omega, \rho_\omega)$, our recipes and corresponding direct input requirements $\alpha_\omega^{\hat{\omega}}$, and our in-house set inference strategy, and the quantitative evaluation of energy-use misallocation, which additionally depends on the distribution of outsourcing capabilities. Both blocks are robust to calibrating our recipes using the 2019 transactions only, to including manufacturing-related services (e.g., plastic packaging services) as inputs, to setting a uniform elasticity of substitution $\rho_\omega = \rho = -0.67$, to externally calibrating the energy share β_ω from the Italian input-output table, to winsorizing ν, q, factor prices or energy intensity $(\frac{e_j}{c_j y_j})$ outliers, and to excluding foreign firms as potential suppliers of firms that don't already import in equilibrium. Moreover, our in-house set calibrated on comparison years $C = \{2021, 2022, 2023\}$ show a large overlap, and our main results carry over to a 2023 calibration.

In the second block, the extensive margin is, as expected, highly sensitive to the lognormal mean param-

[43] In practice, the EU ETS regulates installations, whose sector of activity might differ from their firm's industry classification, and only above a given threshold, so our firm coverage could be either too narrow or too broad. However, our regulated firms' sales shares closely align with studies identifying firms from the EUTL database directly (e.g., D'Arcangelo et al. (2022)).

eter μ, and somewhat to the variance parameter σ^2, which both drive how competitive the non-selected producers are (Table B.18). Moreover, high σ and low μ both make in-house production prohibitively uncompetitive. However, conditional on matching our targeted moments, the results are fairly robust to using different distributions for the outsourcing capabilities: we tried both a Pareto distribution with dependence modeled through a Copula as in the baseline, and a conditional nested logit with a buyer-specific scale parameter, but both were less successful at matching the data. Nonetheless, using simulated data, we could verify that, conditional on matching the moments, the results sensitivity to the scale (i.e., how good is in-house production relative to outsourcing) and dispersion (i.e., how concentrated the non-selected suppliers are) parameters is similar across the different specifications. The calibrated (μ, σ) also perform well to match the same moments evaluated over different sub-samples (excluding multi-product buyers, sellers, both, and on a randomly-selected subsample). Finally, the relative contribution of the extensive margin also increases with a lower input elasticity of substitution γ because costs become more sensitive as quantities adjust less.

Naive ν. We also re-run the counterfactual with the naive ν_j, calibrated from observed energy-to-labor expenditure ratios while ignoring in-house production, keeping the parameters of the outsourcing-capability distribution fixed. Since observed energy intensity rises mechanically with in-house production, the naive calibration makes in-house producers appear less energy efficient than they are. Outsourcing their inputs then seems to save more energy, and the contribution of producer selection to the fall in energy intensity at p_e^{SCC} doubles, from 19.1% to 39.7%, entirely through vertical reallocation whose contribution rises from 10.1% to 34.7%. The contribution of reallocation across sellers, by contrast, falls from 9.0% to 5.0%, because abstracting from firm boundaries compresses the distribution of energy efficiency across sellers of the same product (Figure B.5d), and with it the gains from switching suppliers.

6 Conclusion

Improvements in aggregate energy efficiency have been the main driver of the decoupling between GDP and energy consumption over the past decades (Hassler et al., 2021; Känzig & Williamson, 2024). They arise from firms improving their own technologies, but also from the reallocation of production toward more energy-efficient firms, and the large heterogeneity in energy efficiency across similar firms that we document suggests that substantial gains remain available through this second channel. In supply chains, this reallocation operates not only through how much of a given input each producer makes, but through which firm makes it at all, and this includes the buyer itself. We develop a unified framework to evaluate the allocation of energy across all potential producers. We show that when energy is underpriced, producer selection underweighs energy efficiency, misallocating energy use across producers and thereby raising aggregate energy intensity and lowering welfare. We quantify this channel by exactly calibrating a model of make-or-buy decisions embedded in a production network to Italian firm-to-firm transactions, from which we infer the inputs each firm produces in-house and recover producers' energy efficiency net of their boundaries.

We find large energy efficiency gaps: for most inputs, a more energy-efficient producer is available. However, whether the externality justifies switching to it is ultimately a quantitative question that depends on the outsourcing capabilities, which we recover from the ranking of options revealed by firms' make-or-buy choices. The answer is that most gaps are justified, but the few that are not carry large gains. Pricing emissions at the social cost of carbon lowers aggregate energy intensity by 47.6% and raises welfare by 2.10%, and reallocation through producer selection accounts for about a fifth of the former and a tenth of the latter. These gains come from the impact of switches rather than their frequency: only around 3.5% of the sourcing decisions change, yet they alone cut the energy needed per unit of output by 11.6%. This shows the importance of analyzing environmental distortions through the lens of firm boundaries: producer selection is a margin of adjustment available to policy that delivers large emission reductions even when no firm improves its own technology.

The magnitude of this channel depends on the joint distribution of energy efficiency and outsourcing capabilities, which our framework takes as given; explaining it is the natural next step, which we directly tackle in companion work in progress through the lens of firm investments. More broadly, our framework applies beyond energy. Any firm-level technology difference coupled with an outsourcing wedge makes the allocation across potential producers relevant for aggregate outcomes, especially in the presence of distortions. The outsourcing wedge itself does not have to be purely technological. In fact, multiple distortions prevent buyers to select the most efficient production mode (e.g., contractual frictions (Antràs & Chor, 2013; Boehm & Oberfield, 2020), matching and information frictions (Miyauchi et al., 2018; Bernard et al., 2022; Demir et al., 2024), trade frictions (Eaton et al., 2022; Panigrahi, 2022)). To isolate the environmental externality, this paper purposefully abstracts from any of those other frictions. A complementary next step is thus to use our framework to study the extent to which the environmental distortion interacts with those other frictions to shape the make-or-buy decisions.

References

Abowd, J. M., Kramarz, F., & Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, *67*(2), 251–333.

Acemoglu, D., Aghion, P., Bursztyn, L., & Hemous, D. (2012). The environment and directed technical change. *American economic review*, *102*(1), 131–166.

Acemoglu, D., & Azar, P. D. (2020). Endogenous production networks. *Econometrica*, *88*(1), 33–82.

Agostino, M., Giunta, A., Scalera, D., Trivieri, F., et al. (2016). Italian firms in global value chains: Updating our knowledge. *Rivista di politica economica*, *7*, 155–186.

Akcigit, U., Baslandze, S., & Lotti, F. (2023). Connecting to power: Political connections, innovation, and firm dynamics. *Econometrica*, *91*(2), 529–564.

Alfaro, L., Bloom, N., Conconi, P., Fadinger, H., Legros, P., Newman, A. F., Sadun, R., & Van Reenen, J. (2024). Come together: Firm boundaries and delegation. *Journal of the European Economic Association*, *22*(1), 34–72.

Alfaro, L., Brussevich, M., Minoiu, C., & Presbitero, A. F. (2025). *Bank financing of global supply chains* (tech. rep.). National Bureau of Economic Research.

Alfaro, L., Chor, D., Antras, P., & Conconi, P. (2019). Internalizing global value chains: A firm-level analysis. *Journal of Political Economy*, *127*(2), 508–559.

Alfaro-Ureña, A., & Zacchia, P. (2025). *Matching to suppliers in the production network: A quantitative framework* (tech. rep.).

Alpino, M., Citino, L., & Frigo, A. (2024). Firms in the energy crisis: Evidence from 2021-22. *Available at SSRN 5251730*.

Amador, J., & Cabral, S. (2015). Global value chains, labour markets and productivity. *The age of global value chains*, 107–20.

Andersson, J. J. (2019). Carbon Taxes and CO2 Emissions: Sweden as a Case Study. *American Economic Journal: Economic Policy*, *11*(4), 1–30. https://doi.org/10.1257/pol.20170144

Antras, P., Fort, T. C., & Tintelnot, F. (2017). The margins of global sourcing: Theory and evidence from us firms. *American Economic Review*, *107*(9), 2514–2564.

Antràs, P. (2020). Conceptual aspects of global value chains. *The World Bank Economic Review*, *34*(3), 551–574.

Antràs, P., & Chor, D. (2013). Organizing the global value chain. *Econometrica*, *81*(6), 2127–2204.

Antràs, P., Chor, D., Fally, T., & Hillberry, R. (2012). Measuring the upstreamness of production and trade flows. *American Economic Review*, *102*(3), 412–416.

Antràs, P., & Rossi-Hansberg, E. (2009). Organizations and trade. *Annu. Rev. Econ.*, *1*(1), 43–64.

Arkolakis, C., Eckert, F., & Shi, R. (2023). *Combinatorial discrete choice: A quantitative model of multinational location decisions* (tech. rep.). National Bureau of Economic Research.

Arkolakis, C., Huneeus, F., & Miyauchi, Y. (2025). *Production network formation, trade, and welfare* (tech. rep.). CESifo Working Paper.

Astuti, V., & Piras, D. (2026). Italian business-to-business invoicing data: A network analysis. *Available at SSRN 6239325*.

Atalay, E. (2017). How important are sectoral shocks? *American Economic Journal: Macroeconomics*, *9*(4), 254–280.

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

Baqaee, D. R., & Farhi, E. (2020). Productivity and misallocation in general equilibrium. *The quarterly journal of economics*, *135*(1), 105–163.

Bergquist, L., Lashkari, D., & Verhoogen, E. (2026). Wedges: A microeconomic perspective on misallocation.

Bernard, A. B., Dhyne, E., Magerman, G., Manova, K., & Moxnes, A. (2022). The origins of firm heterogeneity: A production network approach. *Journal of Political Economy*, *130*(7), 1765–1804.

Bernard, J.-T., & Kichian, M. (2021). The impact of a revenue-neutral carbon tax on gdp dynamics: The case of british columbia. *The Energy Journal*, *42*(3), 205–224.

Bianchi, P., Miller, L. M., & Bertini, S. (1997). The italian sme experience and possible lessons for emerging countries. *Bologna: Nomisma*.

Boehm, C., Flaaen, A., & Pandalai-Nayar, N. (2014). *Complementarities in multinational production and business cycle dynamics* (tech. rep.). Technical report, Working paper, University of Michigan.

Boehm, J. (2022). The impact of contract enforcement costs on value chains and aggregate productivity. *Review of Economics and Statistics*, *104*(1), 34–50.

Boehm, J., Dhingra, S., & Morrow, J. (2022). The comparative advantage of firms. *Journal of Political Economy*, *130*(12), 3025–3100.

Boehm, J., & Oberfield, E. (2020). Misallocation in the market for inputs: Enforcement and the organization of production. *The Quarterly Journal of Economics*, *135*(4), 2007–2058.

Boehm, J., & Oberfield, E. (2023). *Growth and the fragmentation of production* (tech. rep.). Tech. rep., mimeo. 7.

Bolz, S. J., Naumann, F., & Richter, P. M. (2024). Unilateral environmental policy and offshoring.

Borin, A., Conteduca, F., Leone, F., Mancini, M., & Zoi, P. (2025). How global are local value chains? *CEPR Discussion Paper*, (DP20910).

Boudreau, L., Cajal-Grossi, J., & Macchiavello, R. (2023). Global value chains in developing countries: A relational perspective from coffee and garments. *Journal of Economic Perspectives*, *37*(3), 59–86.

Bretschger, L., & Jo, A. (2021). Complementarity between labor and energy: A firm-level analysis. *Available at SSRN 4236136*.

Calligaris, S., Del Gatto, M., Hassan, F., Ottaviano, G. I., & Schivardi, F. (2018). The productivity puzzle and misallocation: An italian perspective. *Economic Policy*, *33*(96), 635–684.

Cappariello, R., & Felettigh, A. (2015). How does foreign demand activate domestic value added? a comparison among the largest euro-area economies. *A Comparison Among the Largest Euro-*

Area Economies (January 23, 2015). Bank of Italy Temi di Discussione (Working Paper) No, 1001.

Carvalho, V. M. (2014). From micro to macro via production networks. *Journal of Economic Perspectives*, *28*(4), 23–48.

Casey, G., Gao, Y., & Kruse-Andersen, P. (2026). Energy efficiency dynamics and climate policy. *Journal of the Association of Environmental and Resource Economists*.

Cherniwchan, J. M., & Taylor, M. S. (2022). *International trade and the environment: Three remaining empirical challenges* (tech. rep.). National Bureau of Economic Research.

Choi, J., Levchenko, A. A., Ruzic, D., & Shim, Y. (2024). Superstars or supervillains?: Large firms in the south korean growth miracle. *National Bureau of Economic Research*.

Chor, D., & Ma, L. (2020). Contracting frictions in global sourcing: Implications for welfare. *Manuscript). Singapore Management University*.

Cole, M. A., Elliott, R. J., & Okubo, T. (2014). International environmental outsourcing. *Review of World Economics*, *150*, 639–664.

Coster, P., Mejean, I., & di Giovanni, J. (2024). Firms' supply chain adaptation to carbon taxes. *FRB of New York Staff Report*, (1136).

D'Arcangelo, F. M., Pavan, G., & Calligaris, S. (2022). *The impact of the european carbon market on firm productivity: Evidence from italian manufacturing firms* (tech. rep.). Working Paper.

De Lyon, J., & Dechezleprêtre, A. (2025). *The great dispersion in energy productivity between firms* (tech. rep.). OECD.

Dechezleprêtre, A., Gennaioli, C., Martin, R., Muûls, M., & Stoerk, T. (2022). Searching for carbon leaks in multinational companies. *Journal of Environmental Economics and Management*, *112*, 102601.

Dechezleprêtre, A., Koźluk, T., Kruse, T., Nachtigall, D., & De Serres, A. (2019). Do environmental and economic performance go together? a review of micro-level empirical evidence from the past decade or so. *International Review of Environmental and Resource Economics*, *13*(1-2), 1–118.

Demir, B., Javorcik, B., & Panigrahi, P. (2024). *Breaking invisible barriers: Does fast internet improve access to input markets?* (Tech. rep.). CESifo Working Paper.

Dhyne, E., Kikkawa, A. K., Kong, X., Mogstad, M., & Tintelnot, F. (2023). Endogenous production networks with fixed costs. *Journal of International Economics*, *145*, 103841.

Dhyne, E., Kikkawa, A. K., Mogstad, M., & Tintelnot, F. (2021). Trade and domestic production networks. *The Review of Economic Studies*, *88*(2), 643–668.

Di Giovanni, J., Levchenko, A. A., & Mejean, I. (2024). Foreign shocks as granular fluctuations. *Journal of Political Economy*, *132*(2), 000–000.

Donald, E., Fukui, M., & Miyauchi, Y. (2025). *Unpacking aggregate welfare in a spatial economy* (tech. rep.). National Bureau of Economic Research.

Eaton, J., & Kortum, S. (2002). Technology, geography, and trade. *Econometrica*, *70*(5), 1741–1779.

Eaton, J., Kortum, S., & Kramarz, F. (2026). Firm-to-firm trade: Imports, exports, and the labor market. *Econometrica*, *94*(4), 1135–1170.

Eaton, J., Kortum, S. S., & Kramarz, F. (2022). *Firm-to-firm trade: Imports, exports, and the labor market* (tech. rep.). National Bureau of Economic Research.

Eurostat. (2024, January 10). *Business demography by size class and NACE Rev. 2 activity (2004–2020)* (bd_9bd_sz_cl_r2). https://doi.org/10.2908/BD_9BD_SZ_CL_R2

Fally, T., & Hillberry, R. (2018). A coasian model of international production chains. *Journal of International Economics*, *114*, 299–315.

Fetzer, T., Lambert, P. J., Feld, B., & Garg, P. (2024). *Ai-generated production networks: Measurement and applications to global trade* (tech. rep.).

Fontagné, L., Martin, P., & Orefice, G. (2024). The many channels of firm's adjustment to energy shocks: Evidence from france. *Economic Policy*, eiae011.

Fontaine, F., & Marullaz, C. (2024). Firms' response to energy price shocks. *Work In Progress*.

Forster, P. M., Smith, C., Walsh, T., Lamb, W. F., Lamboll, R., Cassou, C., Hauser, M., Hausfather, Z., Lee, J.-Y., Palmer, M. D., et al. (2025). Indicators of global climate change 2024: Annual update of key indicators of the state of the climate system and human influence. *Earth System Science Data Discussions*, *2025*, 1–72.

Fort, T. C. (2017). Technology and production fragmentation: Domestic versus foreign sourcing. *The Review of Economic Studies*, *84*(2), 650–687.

Fort, T. C. (2023). The changing firm and country boundaries of us manufacturers in global value chains. *Journal of Economic Perspectives*, *37*(3), 31–58.

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

Fullerton, D., & Heutel, G. (2017). The general equilibrium incidence of environmental taxes. *Distributional effects of environmental and energy policy*, 161–181.

Gaeta, M., Virdis, M. R., Alloisio, I., Borghesi, S., Cian, E. D., Parrado, R., Verdolini, E., et al. (2015). Deep decarbonization pathways, 16–36.

Golosov, M., Hassler, J., Krusell, P., & Tsyvinski, A. (2014). Optimal taxes on fossil fuel in general equilibrium. *Econometrica*, *82*(1), 41–88.

Grossman, G. M., & Helpman, E. (2002). Integration versus outsourcing in industry equilibrium. *The quarterly journal of economics*, *117*(1), 85–120.

Hassler, J., Krusell, P., & Olovsson, C. (2021). Directed technical change as a response to natural resource scarcity. *Journal of Political Economy*, *129*(11), 3039–3072.

Hsieh, C.-T., & Klenow, P. J. (2009). Misallocation and manufacturing tfp in china and india. *The Quarterly journal of economics*, *124*(4), 1403–1448.

Huneeus, F. (2018). Production network dynamics and the propagation of shocks. *Graduate thesis, Princeton University, Princeton, NJ*, *52*.

IEA, I. (2020). Global co2 emissions in 2019. *Paris.[Google Scholar]*.

Istat. (2014, August 7). *Demografia d'impresa. Anni 2007–2012* (Statistiche Report). Istituto Nazionale di Statistica. Rome. Retrieved October 8, 2026, from https://www.istat.it/wp-content/uploads/2014/08/Report_Demografia_dimpresa_anni_2007-2012.pdf

Känzig, D. R. (2023). *The unequal economic consequences of carbon pricing* (tech. rep.). National Bureau of Economic Research.

Känzig, D. R., & Williamson, C. (2024). Unraveling the drivers of energy-saving technical change.

Kim, J., Sovacool, B. K., Bazilian, M., Griffiths, S., Lee, J., Yang, M., & Lee, J. (2022). Decarbonizing the iron and steel industry: A systematic review of sociotechnical systems, technological innovations, and policy options. *Energy Research & Social Science*, *89*, 102565.

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

Klenow, P. J., Pastén, E., & Ruane, C. (2024). *Carbon taxes and misallocation in chile* (tech. rep.). Mimeo, Stanford.

Kreinin, M. E. (1967). Price elasticities in international trade. *The Review of Economics and Statistics*, 510–516.

Le Moigne, M., Lepot, S., Ossa, R., Ritel, M., & Simon, D. Z. (2024). *Greening ricardo: Environmental comparative advantage and the environmental gains from trade* (tech. rep.). WTO Staff Working Paper.

Le Quéré, C., Korsbakken, J. I., Wilson, C., Tosun, J., Andrew, R., Andres, R. J., Canadell, J. G., Jordan, A., Peters, G. P., & van Vuuren, D. P. (2019). Drivers of declining co2 emissions in 18 developed economies. *Nature Climate Change*, *9*(3), 213–217.

Levinson, A. (2023). Are developed countries outsourcing pollution? *Journal of Economic Perspectives*, *37*(3), 87–110.

Lyubich, E., Shapiro, J. S., & Walker, R. (2018). Regulating mismeasured pollution: Implications of firm heterogeneity for environmental policy. *AEA Papers and Proceedings*, *108*, 136–142.

Magerman, G., Palazzolo, A., Dhyne, E., Borsos, A., Kulikov, D., Linarello, A., Paulus, A., Romanini, G., & Saldias, M. (2026). Beyond firm size: Network position and shock transmission in firm-to-firm production networks across five economies. *Working Paper*.

Malik, S. (2026). The role of unreported intermediate inputs in firms' production structure. *Work In Progress*.

Manova, K., Moxnes, A., & Perelló, O. (2025). Productivity, matchability and intermediation in production networks.

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.

Martin, J., Mejean, I., & Parenti, M. (2023). Relationship stickiness, international trade, and economic uncertainty. *Review of Economics and Statistics*, 1–45.

Martin, R., De Preux, L. B., & Wagner, U. J. (2014). The impact of a carbon tax on manufacturing: Evidence from microdata. *Journal of Public Economics*, *117*, 1–14.

Marullaz, C., & Romanini, G. (2025). Energy efficiency within firm boundaries. *Work In Progress*.

Masson-Delmotte, V., Zhai, P., Pirani, A., Connors, S. L., Péan, C., Berger, S., Caud, N., Chen, Y., Goldfarb, L., Gomis, M. I., et al. (2023). Summary for policymakers. *Climate change 2021: The*

physical science basis. Contribution of working group I to the sixth assessment report of the intergovernmental panel on climate change, 3–32.

McFadden, D. (1974). The measurement of urban travel demand. *Journal of public economics*, *3*(4), 303–328.

Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *econometrica*, *71*(6), 1695–1725.

Metcalf, G. E., et al. (2003). Environmental levies and distortionary taxation:: Pigou, taxation and pollution. *Journal of Public Economics*, *87*(2), 313–322.

Metcalf, G. E., & Stock, J. H. (2020). Measuring the macroeconomic impact of carbon taxes. *AEA papers and Proceedings*, *110*, 101–106.

Miyauchi, Y., et al. (2018). *Matching and agglomeration: Theory and evidence from japanese firm-to-firm trade* (tech. rep.). Working paper.

Monarch, R. (2014). 'it's not you, it's me': Breakups in us-china trade relationships. *US Census Bureau Center for Economic Studies Paper No. CES-WP-14-08*.

Monteverde, K., & Teece, D. J. (1982). Supplier switching costs and vertical integration in the automobile industry. *The Bell Journal of Economics*, 206–213.

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

Ortubai, J. M. B., Piacentini, R. E. F., Serrador, A. G., & Rubio-Ramırez, J. F. (2025). Unraveling the impact of a carbon price shock on macroeconomic variables: A narrative sign restrictions approach.".

Panigrahi, P. (2022). *Triad trade and small worlds of large spatial production networks* (tech. rep.). Working Paper, Johns Hopkins University.

Peters, G. P., Andrew, R. M., Canadell, J. G., Fuss, S., Jackson, R. B., Korsbakken, J. I., Le Quéré, C., & Nakicenovic, N. (2017). Key indicators to track current progress and future ambition of the paris agreement. *Nature Climate Change*, *7*(2), 118–122.

Petrick, S., Rehdanz, K., & Wagner, U. J. (2011). Energy use patterns in german industry: Evidence from plant-level data. *Jahrbücher für Nationalökonomie und Statistik*, *231*(3), 379–414.

Pindyck, R. S. (2019). The social cost of carbon revisited. *Journal of Environmental Economics and Management*, *94*, 140–160.

Prest, B., Rennert, K., Newell, R., Pizer, W., & Anthoff, D. (2023). Updated estimates of the social cost of greenhouse gases for usage in regulatory analysis. *Resources for the Future, Washington, DC*.

Rachapalli, S. (2021). *Learning between buyers and sellers along the global value chain* (tech. rep.). Working paper.

Ramírez-Santos, Á. A., Castel, C., & Favre, E. (2018). A review of gas separation technologies within emission reduction programs in the iron and steel sector: Current application and development perspectives. *Separation and Purification Technology*, *194*, 425–442.

Rauch, J. E. (1999). Networks versus markets in international trade. *Journal of international Economics*, *48*(1), 7–35.

Restuccia, D., & Rogerson, R. (2008). Policy distortions and aggregate productivity with heterogeneous establishments. *Review of Economic dynamics*, *11*(4), 707–720.

Ricke, K., Drouet, L., Caldeira, K., & Tavoni, M. (2018). Country-level social cost of carbon. *Nature Climate Change*, *8*(10), 895–900.

Schenker, O., Koesler, S., & Löschel, A. (2018). On the effects of unilateral environmental policy on offshoring in multi-stage production processes. *Canadian Journal of Economics/Revue canadienne d'économique*, *51*(4), 1221–1256.

Shapiro, J. S. (2021). The environmental bias of trade policy. *The Quarterly Journal of Economics*, *136*(2), 831–886.

Stangl, J., Borsos, A., & Thurner, S. (2025). Using firm-level supply chain networks to measure the speed of the energy transition. *arXiv preprint arXiv:2503.01572*.

Syverson, C. (2004). Product substitutability and productivity dispersion. *Review of Economics and Statistics*, *86*(2), 534–550.

Wagner, U. J., Kassem, D., Gerster, A., Jaraite, J., Klemetsen, M. E., Laukkanen, M., Martin, R., Munch, J. R., Muûls, M., de Preux, L., et al. (2020). Carbon footprints of european manufacturing jobs: Stylized facts and implications for climate policy. *Available at SSRN 3577550*.

Wang, C., Larsson, M., Ryman, C., Grip, C.-E., Wikström, J.-O., Johnsson, A., & Engdahl, J. (2008). A model on co2 emission reduction in integrated steelmaking by optimization methods. *International journal of energy research*, *32*(12), 1092–1106.

A Theoretical Appendix

A.1 Cross-producer misallocation

This section formalizes cross-producer misallocation in a general environment, discusses the main measurement challenges, and explains how the framework applies to energy-use misallocation in our setting.

Producer selection hinges on the idea that firms select *one* producer, among a set of potential producers, to produce each of their intermediate inputs. For the theory, this is without loss of generality, since inputs can be defined at an arbitrarily granular level. In our data, however, this already holds at the six-digit ATECO product classification (Table B.6).[44] Because selection is discrete, when potential producers differ along two characteristics whose rankings do not exactly agree, producer selection involves a trade-off, and any distortion that tilts the weight placed on one characteristic in that decision can misallocate the production of intermediate inputs, and the resources they use, at the extensive margin. Moreover, because firms can produce part of their inputs in-house, this allocation is determined not only across the sellers of an input, which are typically observed in the data, but across all its potential producers, including the firms using that input in subsequent stages, whose in-house production is typically not observed. Overall, the characterization of cross-producer (mis)allocation therefore requires:

1. Identifying which firm *actually* produces the input for each firm
2. Identifying which firm *should* produce the input for each firm, under an optimal allocation.

If both coincide, there is no scope for cross-producer misallocation. Requirement 2 directly highlights that potential producers must be heterogeneous along some dimensions to generate misallocation. Moreover, if the sets in 1 and 2 don't match, there must be some sort of wedge in the make-or-buy decision that prevents the downstream firm to select the more efficient alternative. This implies two dimensions of heterogeneity: a firm-specific component that makes firms more or less efficient, and a relationship-specific wedge that distorts the make-or-buy decision.[45]

Formally, consider a downstream firm j and a mass of upstream firms $i \in \mathcal{S}_\omega$, characterized by an efficiency parameter ν_j, ν_i respectively. The downstream firm j needs an intermediate input ω in quantity $x_j(\omega)$, that it can either produce in-house using a primary factor $e_j(\omega)$, or outsource to one of the upstream supplier $i \in \mathcal{S}_\omega$. The firms' production function exhibit CRS and are identical up to their productivity ν, and they face the same factor price p_e that they take as given. However, the downstream

[44] This is confirmed more broadly in the literature, which has consistently shown that firms are largely undiversified, that is, they have very few suppliers, usually only one, for each of their inputs, and that diversification has even declined over the years (e.g., Antras et al. (2017) and Alfaro et al. (2025)). In particular, Alfaro et al. (2025) find that firms hit by a tariff shock drop their supplier before switching to a new one, instead of diversifying.

[45] In other words, cross-producer misallocation arises because a wedge distorts the make-or-buy decision. If this wedge is pecuniary, this extensive margin effect adds to the intensive margin (or direct effect in the terminology of Boehm and Oberfield (2020)) of the wedges, which drives misallocation within firms, across inputs, and across firms, as reflected by dispersion in TFPR (Hsieh & Klenow, 2009).

firm face a pairwise wedge $q_{ij} = \tilde{q}_{ij}\epsilon_{ij}$ (which can include firm i's markup) on the effective price paid to produce or outsource an input. Using Bergquist et al. (2026) terminology, the wedge q_{ij} can either be distortive, in the sense that it does not carry a value and does not constrain the Social Planner, or technological. In the latter case, it is part of the Social Planner's problem. The make-or-buy decision boils down to comparing the effective unit cost of producing ω in-house relative to outsourcing it to the cheapest supplier $i \in \mathcal{S}_\omega$:

$$\underbrace{\frac{c_\omega(p_e, \nu_j)}{\tilde{q}_{jj}\varepsilon_{jj}} \lessgtr \underbrace{\min_{i \in \mathcal{S}_\omega} \frac{c_\omega(p_e, \nu_i)}{\tilde{q}_{ij}\varepsilon_{ij}}}_{\text{Supplier selection}}}_{\text{cross-producer}}$$

The make-or-buy decision allocates production across heterogeneous firms. It includes, and extends, external supplier selection (that is, a Ricardian model with potential distortions (Eaton & Kortum, 2002)).If $\varepsilon_{ij} = 1 \; \forall i, j$, the make-or-buy decision is efficient. If at least one $\varepsilon_{ij} \neq 1$, the make-or-buy decision is distorted, and might result in misallocation if $\arg\min_{i \in \mathcal{S}_\omega \cup \{j\}} \frac{c_\omega(p_e, \nu_i)}{\tilde{q}_{ij}} \neq \arg\min_{i \in \mathcal{S}_\omega \cup \{j\}} \frac{c_\omega(p_e, \nu_i)}{\tilde{q}_{ij}\varepsilon_{ij}}$. Because the decision is discrete, the distortion does not automatically lead to misallocation: if $\tilde{q}_{ij}\varepsilon_{ij} = 1 \; \forall i, j$, the make-or-buy *allocation* remains efficient, and always aligns with the productivity (ν) ranking.

To show the impact of cross-producer misallocation on aggregate outcomes, we add some structure and consider a simple model in which downstream firms produce differentiated varieties by combining a continuum Ω of inputs $\omega \in \Omega$ using a CRS Cobb-Douglas technology with input shares β_ω summing up to 1. The firms can either produce their input in-house using a primary factor $e(\omega)$, or outsource its production to an external supplier. To simplify exposition, we consider only one supplier per input ω, and index the supplier by $i(\omega)$. All firms are again characterized by some productivity ν, and q_{kj} is the pairwise wedge. Under this specification, the unit cost of production of firm j is

$$c_j = \prod_{\omega \in \Omega} \left(\frac{c_\omega(p_e, \nu_k)}{q_{kj}} \right)^{\beta_\omega} \mathbf{1}\{k = k_j(\omega; p_e)\} = \prod_{\omega \in \mathcal{I}_j} \left(\frac{p_e}{q_{jj}\nu_j} \right)^{\beta_\omega} \prod_{\omega \in \mathcal{O}_j} \left(\frac{c_\omega(p_e, \nu_{i(\omega)})}{q_{i(\omega)j}} \right)^{\beta_\omega}$$

Where $\omega \in \mathcal{I}_j$ is the set of inputs produced in-house, $\omega \in \mathcal{O}_j$ is the set of inputs that are outsourced, $k_j(\omega; p_e)$ is the selected producer at price p_e, and $c_\omega(p_e, \nu_{i(\omega)}) = \frac{p_e}{\nu_{i(\omega)}}$ is supplier $i(\omega)$'s unit cost. Taking the factor price as given, firms first decide which inputs to produce in-house ($\omega \in \mathcal{I}_j$) and which inputs to outsource ($\omega \in \mathcal{O}_j$). Then, they choose their primary factor demand for in-house production. As in Boehm and Oberfield (2023), we suppose that $e_j(\omega)$ is chosen separately for each input $\omega \in \mathcal{I}_j$. The make-or-buy decision is independent across inputs, and boils down to comparing the effective unit cost of producing the input in-house or outsourcing it: $\frac{p_e}{q_{jj}\nu_j} \lessgtr \frac{p_e}{q_{i(\omega)j}\nu_{i(\omega)}}$.

The Cobb-Douglas aggregation of input shuts down the intensive margin misallocation across inputs by keeping input shares constant. To focus on the extensive margin misallocation induced by producer selection, we also shut down the between firm intensive margin channel by assuming a log-linear aggregate demand for downstream firms' varieties.[46]

[46] One might wonder how cross-producer misallocation would enter Hsieh and Klenow (2009)'s canonical framework. To

In that case, the aggregate price index is

$$logP = \int_{j\in\mathcal{N}} \log c_j dj = log\, p_e - \int_{j\in\mathcal{N}} log\, \nu_j dj - \int_{j\in\mathcal{N}} \log(q_{jj}) dj + \int_{j\in\mathcal{N}} \sum_{\omega\in\mathcal{O}_j} \beta_\omega log\left(\frac{\nu_j}{\nu_{i(\omega)}}\right) dj + \int_{j\in\mathcal{N}} \left(\sum_{\omega\in\mathcal{O}_j^*} \beta_\omega log\left(\frac{q_{jj}}{q_{i(\omega)j}}\right)\right) dj$$

We first consider the case of purely distortive wedges, which is the most standard case in the misallocation literature (e.g., Hsieh and Klenow (2009)). They are therefore set to 1 in the optimal allocation. Denoting variables in the optimal allocation by SP and in the equilibrium allocation by a $*$, we can measure the effect of misallocation on the aggregate price index[47] as:

$$logP^{SP} - logP^* = \int_{j\in\mathcal{N}} \underbrace{\sum_{\omega\in\mathcal{O}_j^{SP}\cap\mathcal{I}_j^*} \beta_\omega log\left(\frac{\nu_j}{\nu_{i(\omega)}}\right) - \sum_{\omega\in\mathcal{O}_j^*\cap\mathcal{I}_j^{SP}} \beta_\omega log\left(\frac{\nu_j}{\nu_{i(\omega)}}\right)}_{\text{Indirect (make-or-buy) effect}}$$
$$+ \underbrace{\left(\sum_{\omega\in\mathcal{O}_j^*} \beta_\omega log(q_{i(\omega)j}) + log(q_{jj}) \sum_{\omega\in\mathcal{I}_j^*} \beta_\omega\right)}_{\text{Direct effect}} dj$$

The distortion has two effects on the price index. First, the standard direct effect: wedges depress or inflate the aggregate price depending on whether they act as effective taxes or subsidies. Secondly, the novel effect of the make-or-buy decision, which arises at the *extensive* margin: the producing firm might not be the most productive firm. If $\frac{\nu_j}{\nu_i} >$ $(<)$ 1, the input ω produced by firm i should be produced in-house (outsourced). Quantifying the *make-or-buy* (indirect) effect on misallocation therefore requires identifying the productivity ν for all firms, such that it is possible to characterize the efficient allocation (i.e., which firm should actually produce) and quantify the welfare loss. But the distribution of ν is not enough: one needs to know the relevant comparison set that is, which firm could produce the input.

Before discussing the challenges that arise with quantification, we turn to the characterization of the effect of misallocation if the wedge is (partly) technological. Suppose that $q_{ij} = \tilde{q}_{ij}\varepsilon_{ij}$, where $\tilde{q}_{ij}$ is a

show this, we instead assume a CES aggregate demand such that:

$$logTFP_s = \frac{1}{1-\sigma} log\left(\int_{j\in\mathcal{S}_s} A_j(\mathcal{I},\mathcal{O})^{\sigma-1}\left(\frac{\overline{TFPR}}{TFPR_j}\right)^{\sigma-1} dj\right)$$

Where $A_j(\mathcal{I},\mathcal{O}) \equiv \nu_j^{-\sum_{\omega\in\mathcal{I}}\beta_\omega}\prod_{\omega\in\mathcal{O}}\nu_{i(\omega)}^{-\beta_\omega}$ is the firm's technology, $\overline{TFPR} = p_e\left(\int_{j\in\mathcal{S}_s}\frac{p_j y_j}{P_s Y_s}\left(q_{jj}\sum_{\omega\in\mathcal{I}_j}\beta_\omega + \sum_{\omega\in\mathcal{O}_j}\beta_\omega q_{i(\omega)j}\right)dj\right)^{-1}$, which is equal to $\frac{p_e}{q}$ if all wedges are equalized to q, while $TFPR_j \equiv \frac{p_j y_j}{\Gamma\prod e_j(\omega)^{\beta_\omega}} = p_e q_{jj}^{-\sum_{\omega\in\mathcal{I}}\beta_\omega}\prod_{\omega\in\mathcal{O}} q_{i(\omega)j}^{-\beta_\omega}$, which also simplifies to $\frac{p_e}{q}$ under equal wedges. We define $TFPR_j$ as the true technology residual by controlling for the normalizing constant $\Gamma = \prod_{\omega\in\Omega}\beta_\omega^{-\beta_\omega}$. The firm's technology $A_j(\mathcal{I},\mathcal{O})$ is endogenously determined by the make-or-buy decisions thus, we cannot assume a log-linear joint distribution for the wedges and A_j. Additionally to the intensive margin (or direct effect) of Hsieh and Klenow (2009), wedges decrease TFP through the extensive margin (an indirect effect) by distorting the make-or-buy allocation: if $\nu_j >$ $(<)$ ν_i and $\mathcal{O}_j$ $(\mathcal{I}_j)$ is not empty, A_j, hence TFP_s, would be larger. However, endogenous make-or-buy decisions also mitigate the direct effect by reducing the dispersion in wedges $q_{i(\omega)j}$, as long as the latter are correlated across j, within suppliers, and the $\mathcal{O}_j$ are not empty.

[47] With distortive wedges, the effect on aggregate TFP depends on whether those wedges are quasi-rent or if their revenue are still rebated to the household. In the latter case, $logTFP = log\frac{Y}{E} = log(p_e) - logP - log\left(\int_j(q_{jj}\sum_{\omega\in\mathcal{I}_j}\beta_\omega + \sum_{\omega\in\mathcal{O}_j} q_{i(\omega)j}\beta_\omega)dj\right)$ and $logTFP^{SP} - logTFP^* = logP^* - logP^{SP} + log\left(\int_j(q_{jj}\sum_{\omega\in\mathcal{I}_j^*}\beta_\omega + \sum_{\omega\in\mathcal{O}_j^*} q_{i(\omega)j}\beta_\omega)dj\right)$ because wedges distort the equilibrium revenues.

technological wedge. In that case,

$$logP^{SP} - logP^* = \int_{j\in\mathcal{N}} \underbrace{\sum_{\omega\in\mathcal{O}_j^{SP}\cap\mathcal{I}_j^*} \beta_\omega log\left(\frac{\nu_j\tilde{q}_{jj}}{\nu_{i(\omega)}\tilde{q}_{i(\omega)j}}\right) - \sum_{\omega\in\mathcal{O}_j^*\cap\mathcal{I}_j^{SP}} \beta_\omega log\left(\frac{\nu_j\tilde{q}_{jj}}{\nu_i\tilde{q}_{i(\omega)j}}\right)}_{\text{Indirect (make-or-buy) effect}}$$
$$+ \underbrace{\sum_{\omega\in\mathcal{O}_j^*} \beta_\omega log(\varepsilon_{ij}) + log(\varepsilon_{jj}) \sum_{\omega\in\mathcal{I}_j^*} \beta_\omega}_{\text{Direct effect}} dj$$

The direct effect is the same as above, but can be larger or smaller depending on how the pure distortions ε_{ij} and q_{ij} compare. The make-or-buy, or indirect, effect, now depends on the total technological relative advantage of the firms: if $\frac{\nu_j\tilde{q}_i}{\nu_i} > \ (<)1$, the input ω produced by firm i should be produced in-house (outsourced). Quantifying the *make-or-buy* (indirect) effect of misallocation therefore requires identifying the productivity ν *and* the technological $\tilde{q}_{ij}$ for all firms. However, if ν is Hicks-neutral, the two parameters cannot be separately identified.

Identification. With a sufficiently broad definition of productivity ν capturing all differences in firm-level technology, it can be argued that the residual outsourcing wedges are purely distortive ($q_{ij} = \varepsilon_{ij}$). The implication is that there is no relationship-specific technological wedge, arising, for instance, from customization. A direct route to identify the productivity parameters ν is to use variations in firms' size. Moving away from the simplifying example of a log-linear aggregate production function, and assuming that sales to household are not distorted,[48] productivity ν can be backed out from price-adjusted within-product variations in household expenditure shares on firms. However, the broad productivity ν encompasses multiple factors and might be challenging to credibly identify. In particular, one might fear that it also captures other distortions affecting household sales, such as markups, biasing the resulting misallocation. An alternative, which is the natural path in our application, is to narrow down the definition of productivity to improve its identification. We model firms' heterogeneity in productivity specifically as the energy-biased technology, and this refinement allows us to back out this productivity differences using energy-to-labor expenditure ratio. However, this narrower definition of productivity makes the assumption that q_{ij} is a purely distortive wedge implausible. There are many different (potentially technological) factors that drive the make-or-buy decision beyond energy considerations. In that case, q_{ij} becomes, at least partially, a technological wedge, which enters the social planner's problem (i.e., $\tilde{q}_{ij} \neq 1$). This has two implications: first, absent any other distortion, the allocation is efficient. Second, characterizing a counterfactual allocation requires identifying the wedge not only for the empirically observed relationships, but for *all potential* relationships. This setting thus confers more concreteness in the productivity term (ν) by adding another abstract parameter (q_{ij}) to absorb the remaining heterogeneity, but this parameter must also be empirically identified. In the context of energy, there is a clear and quantifiable distortion: the environmental externality.[49] Figure B.1 accordingly shows that, in our

[48] In the same spirit, Boehm and Oberfield (2020) assume that some inputs ("standardized inputs") are not subject to distortions to disentangle the technological from the distortive component of input transactions.

[49] This distortion enters the problem differently from ε_{ij}, since it doesn't affect the outsourcing wedge directly. However,

Italian data, inputs facing less distorted energy prices appear to be less misallocated at the extensive margin. To isolate the effect of the environmental distortion, we consider q_{ij} as a purely technological wedge: that is, $q_{ij} = \tilde{q}_{ij}$. This allows us to normalize the in-house wedge q_{jj} to one without loss of generality. To deal with the empirical challenge it raises, that is, the characterization of all the potential relationships, we rely on a revealed preference argument as explained in Section 4.2.4.

As a final comment, a natural concern is whether it is meaningful to compare producers when some inputs are made in-house by downstream firms while others are supplied by firms for whom the input is a core product. Our approach is to compare these actors precisely along the common object of interest: the production of the same input. Conditional on the input being produced, the underlying technological requirements - the "recipe" of the input - should plausibly be similar irrespective of organizational form: tyres require rubber whether they are produced by a specialized tyre manufacturer or internally by an integrated downstream firm. To the extent that in-house producers and specialized suppliers differ along observable characteristics (size, sales, location, factor prices, etc.), these differences can be controlled for in the empirical implementation. More fundamentally, in-house production may differ from specialized production because of scale economies or within-firm synergies; modeling these mechanisms explicitly would substantially complicate the framework and would introduce an additional wedge—beyond the energy externality - through which firms might inefficiently under-produce in-house. We abstract from these scale considerations not only to maintain a tractable production-network environment, but also to keep the cross-producer comparison transparent by disciplining it around a single firm-level dimension: effective energy efficiency. Moreover, by explicitly identifying which inputs are produced internally and accounting for their input-specific energy requirements, we partially absorb those differences in organization, scale, and synergies. We investigate these forces - and how they interact with energy mispricing - in a companion paper (Marullaz & Romanini, 2025). Finally, both prior work (Lyubich et al., 2018; De Lyon & Dechezleprêtre, 2025) and our data suggest that most dispersion in energy efficiency is explained by a firm-specific component rather than a (narrowly defined) product component. This supports modeling energy efficiency primarily as a firm characteristic, rather than a product-specific one.

A.2 Two-stage Model

This section provides additional results and derivations to complement the two-stage model results presented in the main text.

A.2.1 Efficient allocation

We derive the optimal allocation in the two-stage model. There are multiple products ω, a mass $\mathcal{S}_\omega$ of firms selling them, and each product requires combining multiple intermediate inputs $\hat{\omega} \in \Omega_\omega$. All inputs can be produced in-house with energy and labor by the downstream firms. Alternatively, except for the

even as a common distortion, it still results in inefficient make-or-buy decisions, as our model demonstrates.

core input $\hat{\omega}_c$, all intermediate inputs $\hat{\omega}$ can be outsourced to an external, upstream supplier from the set $\mathcal{S}_{\hat{\omega}}$ of suppliers of $\hat{\omega}$. Foreshadowing our quantitative model, welfare is quasi-linear in a final good X (produced exclusively with labor), manufacturing consumption and energy, and energy is produced solely using the final good X: $E = \frac{X^E}{\chi}$. The Social Planner maximize welfare subject to the economy resource constraint:

$$\max_{\{\{e_j(\omega),l_j(\omega),\{x_{ij}(\omega)\}_{i\in\mathcal{S}_\omega\cup\{j\}}\}_{\omega\in\Omega_{\omega_j}}\}_j, L^x} \mathrm{U} = X + U_Y(Y) + U_E(E)$$

Subject to

$$Y \leq F_Y(\{Y_\omega\}_\omega), \quad Y_\omega \leq G(\{y_j(\omega)\}_j), \quad X + X^E \leq L^x, \quad E \leq \frac{X^E}{\chi}$$

$$\sum_\omega \sum_{\hat{\omega}\in\Omega_\omega} \left(\int_{j\in\mathcal{S}_\omega} l_j(\hat{\omega})dj + \int_{i\in\mathcal{S}_{\hat{\omega}}} l_i di \right) \leq 1 - L^x, \quad \sum_\omega \sum_{\hat{\omega}\in\Omega_\omega} \left(\int_{j\in\mathcal{S}_\omega} e_j(\hat{\omega})dj + \int_{i\in\mathcal{S}_{\hat{\omega}}} e_i di \right) \leq E,$$

$$y_j(\omega) \leq \mathrm{F}\left(\{x_j(\hat{\omega})\}_{\hat{\omega}\in\Omega_\omega}\right)$$

$$x_j(\hat{\omega}) \leq q_{jj} F_{\hat{\omega}}\left(e_j(\hat{\omega}), l_j(\hat{\omega}); \nu_j\right) + \mathbf{1}\{\hat{\omega} \neq \hat{\omega}_c\} \int_{i\in\mathcal{S}_{\hat{\omega}}} q_{ij} x_{ij}(\hat{\omega}) di \; \forall \hat{\omega}$$

$$x_{i.}(\hat{\omega}) \leq F_{\hat{\omega}}\left(e_i, l_i; \nu_i\right)$$

$$\int_{j\in\mathcal{S}_\omega} x_{ij}(\hat{\omega}) \mathbf{1}\{\hat{\omega} \in \Omega_\omega\} dj \leq x_{.i}(\omega) \quad (\lambda_i(\omega))$$

$$x_{ij}(\hat{\omega}) \geq 0 \quad (\lambda^x_{ij}(\hat{\omega})), \quad l_j(\hat{\omega}) \geq 0 \quad (\lambda^l_j(\hat{\omega})), \quad e_j(\hat{\omega}) \geq 0 \quad (\lambda^e_j(\hat{\omega}))$$

Where all aggregators are separable in their arguments and λ are Lagrange multipliers associated with their respective resource constraint. Perfect substitutability of suppliers i will induce the discrete decision that results in $x_{ij}(\hat{\omega}) = 0 \; \forall i \neq k_j(\hat{\omega})$.

The optimal allocation satisfies the following conditions. First, equating the marginal value of labor in X-production and in downstream production pins the shadow wage: $w = 1$. Second, the shadow value of downstream firm j's output equals its marginal contribution to welfare:

$$\lambda_j(\omega) = \frac{\partial U_Y}{\partial Y} \cdot \frac{\partial log F_Y}{\partial log Y_\omega} \frac{\partial log G}{\partial y_j(\omega)} \quad \forall j \in \mathcal{S}_\omega.$$

For each downstream firm j and upstream stage $\hat{\omega} \in \Omega_\omega$, the shadow wage and energy price equal the marginal product of labor and energy in in-house production of $\hat{\omega}$, valued at the shadow cost of the intermediate bundle $x_j(\hat{\omega})$:

$$1 - \lambda^l_j(\hat{\omega}) = \mu_j(\hat{\omega}) \cdot q_{jj} \frac{\partial F_{\hat{\omega}}}{\partial l_j(\hat{\omega})} \quad \forall j, \hat{\omega}$$

$$\chi - \frac{\partial U_E}{\partial E} \frac{\partial E}{\partial e_j(\omega)} - \lambda^e_j(\hat{\omega}) = \mu_j(\hat{\omega}) \cdot q_{jj} \frac{\partial F_{\hat{\omega}}}{\partial e_j(\hat{\omega})} \quad \forall j, \hat{\omega}$$

where

$$\mu_j(\hat{\omega}) \equiv \lambda_j(\omega) \cdot \frac{\partial F}{\partial x_j(\hat{\omega})}$$

is the shadow value of a unit of the intermediate bundle $x_j(\hat{\omega})$ to downstream firm j, i.e. the product of the shadow value of j's final output and the marginal product of $\hat{\omega}$ in j's production.

For each upstream firm $i \in \mathcal{S}_{\hat{\omega}}$, the shadow wage and energy price equal the marginal product of labor and energy in upstream production, valued at the shadow cost $\lambda_i(\hat{\omega})$ of upstream output:

$$1 - \lambda_i^l = \lambda_i(\hat{\omega}) \cdot \frac{\partial F_{\hat{\omega}}}{\partial l_i} \quad \forall i \in \mathcal{S}_{\hat{\omega}}$$

$$\chi - \frac{\partial U_E}{\partial E} \frac{\partial E}{\partial e_j(\omega)} - \lambda_i^e = \lambda_i(\hat{\omega}) \cdot \frac{\partial F_{\hat{\omega}}}{\partial e_i} \quad \forall i \in \mathcal{S}_{\hat{\omega}}$$

For each downstream firm j, stage $\hat{\omega}$, and potential supplier $i \in \mathcal{S}_{\hat{\omega}}$, the KKT condition for $x_{ij}(\hat{\omega}) \geq 0$ is:

$$\lambda_i(\hat{\omega}) + \lambda_{ij}^x(\hat{\omega}) = \mu_j(\hat{\omega}) \cdot q_{ij} \quad \forall i \in \mathcal{S}_{\hat{\omega}},\ j \in \mathcal{S}_{\omega}$$

with $\lambda_{ij}^x(\hat{\omega}) \geq 0$ and $\lambda_{ij}^x(\hat{\omega}) \cdot x_{ij}(\hat{\omega}) = 0$. Since $\mu_j(\hat{\omega})$ is common across suppliers, the planner concentrates all sourcing on:

$$k_j(\hat{\omega}) = \underset{i \in \mathcal{S}_{\hat{\omega}} \cup \{j\}}{\arg\min} \frac{\lambda_i(\hat{\omega})}{q_{ij}}$$

where for the in-house option $i = j$, $\lambda_j(\hat{\omega})$ is replaced by the shadow cost of in-house production:

$$\lambda_j^{ih}(\hat{\omega}) \equiv \frac{w \cdot \partial F_{\hat{\omega}}/\partial l_j(\hat{\omega}) + p_e^{SCC} \cdot \partial F_{\hat{\omega}}/\partial e_j(\hat{\omega})}{[\partial F_{\hat{\omega}}/\partial l_j(\hat{\omega})]^2 + [\partial F_{\hat{\omega}}/\partial e_j(\hat{\omega})]^2} \cdot \frac{1}{q_{jj}}$$

with $p_e^{SCC} = \chi - \frac{\partial U_E}{\partial E}$ the social cost of energy. The firm outsources $\hat{\omega}$ if and only if $k_j(\hat{\omega}) \neq j$, i.e. if there exists $i \in \mathcal{S}_{\hat{\omega}}$ such that $q_{ij}/\lambda_i(\hat{\omega}) > q_{jj}/\lambda_j^{ih}(\hat{\omega})$.

We show that the efficient allocation can be decentralized as a competitive equilibrium with a Pigouvian tax τ^* on energy. In the decentralized equilibrium, downstream firms minimize cost taking as given the wage $w = 1$, the energy price p_e, and upstream input prices $\{p_i(\hat{\omega})\}_{i \in \mathcal{S}_{\hat{\omega}}}$. Upstream firms similarly minimize cost taking w and p_e as given. Equilibrium prices $\{p_i(\hat{\omega})\}$ satisfy the same conditions as the planner's shadow values $\{\lambda_i(\hat{\omega})\}$, with p_e in place of the social cost of energy p_e^{SCC}.

At an interior solution ($\lambda_j^l(\hat{\omega}) = \lambda_j^e(\hat{\omega}) = \lambda_i^l = \lambda_i^e = 0$), the planner's factor demand FOCs for both downstream and upstream firms imply the same cost-minimizing labor-energy ratio:

$$\frac{\partial F_{\hat{\omega}}/\partial l}{\partial F_{\hat{\omega}}/\partial e} = \frac{1}{p_e^{SCC}} \quad \forall j \in \mathcal{S}_{\omega},\ i \in \mathcal{S}_{\hat{\omega}},\ \hat{\omega}$$

where $p_e^{SCC} = \chi - \frac{\partial U_E}{\partial E}\big|_{Y^*, E^*}$ is the social cost of energy at the efficient allocation, reflecting both the direct utility cost of energy consumption and its effect on aggregate consumption through the resources used for its production. In the decentralized equilibrium without intervention, the private energy price is $p_e^{mkt} = \chi$, which omits the latter. The efficient Pigouvian tax τ^* is defined implicitly by:

$$\tau^* = p_e^{SCC} - p_e^{mkt} = \chi - \frac{\partial U_E}{\partial E}\big|_{Y^*, E^*} - \chi = -\frac{\partial U_E}{\partial E}\big|_{Y^*, E^*}$$

which characterizes τ^* implicitly since $\frac{\partial U_E}{\partial E}\big|_{Y^*, E^*}$ is endogenous. Under τ^*, the factor demand FOCs of all firms coincide with the planner's, and competitive equilibrium prices equal the planner's shadow values: $p_i(\hat{\omega}) = \lambda_i(\hat{\omega})$ for all i, $\hat{\omega}$.

In the decentralized equilibrium, downstream firm j sources $\hat{\omega}$ from:

$$k_j^{mkt}(\hat{\omega}, p_e) = \underset{i \in \mathcal{S}_{\hat{\omega}} \cup \{j\}}{\arg\min} \frac{p_i(\hat{\omega}, p_e)}{q_{ij}}$$

where for the in-house option the relevant cost is the firm's private unit cost of producing $\hat{\omega}$, evaluated at p_e^{mkt}. The efficient selection is:

$$k_j^*(\hat{\omega}) = \underset{i \in \mathcal{S}_{\hat{\omega}} \cup \{j\}}{\arg\min} \frac{\lambda_i(\hat{\omega})}{q_{ij}}$$

These coincide if and only if $p_i(\hat{\omega}) \propto \lambda_i(\hat{\omega})$ for all i, which holds under τ^* but fails otherwise.

These results imply the following corollary:

Corollary 1. *If $p_e^{mkt} = p_e^{SCC}$, given the allocation of production, the equilibrium allocation of resources within firms is efficient. Moreover, the allocation of production across firms is efficient.*

The first statement is simply a result of the first-welfare theorem. When firms face the correct energy price, the producer's cost-minimizing choices of energy and other primary factors - and, when outsourcing, the buyer's conditional demand for intermediates - coincide with the planner's intensive-margin allocation. The second statement results from a similar intuition: without any distortion, the make-or-buy decision is efficient because it minimizes the overall cost of production, taking into account the social cost of energy through p_e^{SCC}.

Trade-off. As discussed in the main text, the make-or-buy allocation trades off absolute advantage in energy use and in pairwise total factor use, represented by outsourcing capabilities. Misallocation arises because the private firms and the social planner weight those components differently: the social planner takes into account the environmental damage. To see this clearly, we specify the welfare function and aggregate production as:

$$\mathrm{U} = X + \sum_j \sum_\omega \beta_{j\omega} ln\left(\frac{y_j(\omega)}{\beta_{j\omega}}\right) - \Upsilon E$$

Where $\sum_j \sum_\omega \beta_{j\omega} = 1$. Re-arranging the FOCs, we obtain:

$$\lambda_j(\omega) y_j(\omega) = \beta_{j\omega} = (\chi + \Upsilon) \sum_{\hat{\omega} \in \Omega_\omega} e_i(\hat{\omega}) \mathbf{1}\{k_j(q, \nu, \Upsilon; \hat{\omega}) = i\} + \sum_{\hat{\omega} \in \Omega_\omega} l_i(\hat{\omega}) \mathbf{1}\{k_j(q, \nu, \Upsilon; \hat{\omega}) = i\}$$

Taking the sum over firms and products and using that $\sum_j \sum_\omega \beta_{j\omega} = 1$, we obtain:

$$(\Upsilon + \chi) E + L^p = 1$$

Since $X = 1 - L^p - \chi E$, welfare simplifies to

$$\mathrm{U} = \sum_j \sum_\omega \beta_{j\omega} ln\left(\frac{y_j(\omega)}{\beta_{j\omega}}\right)$$

Because sale shares are constant under Cobb-Douglas demand, and there is no income effect with quasi-linear utility, the optimal allocation maximizes the output of each firm separately. Recall that $\frac{y_j(\omega)}{\beta_{j\omega}} = \frac{1}{\lambda_j(\omega)}$. The energy and labor FOCs give:

$$\left(\frac{\partial F_{\hat{\omega}}}{\partial e_j(\hat{\omega})}\right)^{-1}(\nu_j) = \frac{1}{\chi + \Upsilon} \lambda_j(\omega) y_j(\omega) \frac{\partial log F}{\partial x_j(\hat{\omega})} q_{jj}$$

$$e_j(\hat{\omega}) = \tilde{E}\left(\frac{1}{\chi + \Upsilon} \lambda_j(\omega) y_j(\omega) \frac{\partial log F}{\partial x_j(\hat{\omega})} q_{jj}\ ; \nu_j\right)$$

Where $\tilde{E}$ is an increasing, homogeneous function of degree 0 that inverts $\left(\frac{\partial F_{\hat{\omega}}}{\partial e_j(\hat{\omega})}\right)^{-1}$. Similarly, we denote

$$l_j(\hat{\omega}) = \tilde{L}\left(\lambda_j(\omega) y_j(\omega) \frac{\partial logF}{\partial x_j(\hat{\omega})} q_{jj}\ ; \nu_j\right)$$

Plugging the optimal e_j and l_j in the production function:

$$\begin{aligned}
y_j(\omega) =& F\left[\left\{q_{ij}\, F_{\hat{\omega}}\left(\tilde{E}\left(\frac{1}{\chi+\Upsilon}\lambda_j(\omega) y_j(\omega)\frac{\partial logF}{\partial x_j(\hat{\omega})} q_{jj}\ ; \nu_j\right),\right.\right. \\
& \left.\left.\left.\tilde{L}\left(\lambda_j(\omega) y_j(\omega)\frac{\partial logF}{\partial x_j(\hat{\omega})} q_{jj}\ ; \nu_j\right)\right)\mathbf{1}\{k_j(q,\nu,\Upsilon;\hat{\omega}) = i\}\right\}_{\hat{\omega}\in\Omega_\omega}\right] \\
\Rightarrow \frac{y_j(\omega)}{\beta_{j\omega}} = \frac{1}{\lambda_j(\omega)} =& F\left[\left\{q_{ij}\, F_{\hat{\omega}}\left(\tilde{E}\left(\frac{1}{\chi+\Upsilon}\frac{\partial logF}{\partial x_j(\hat{\omega})} q_{jj}\ ; \nu_j\right), \tilde{L}\left(\frac{\partial logF}{\partial x_j(\hat{\omega})} q_{jj}\ ; \nu_j\right)\right)\mathbf{1}\{k_j(q,\nu,\Upsilon;\hat{\omega}) = i\}\right\}_{\hat{\omega}\in\Omega_\omega}\right] \\
=& F\left[\left\{q_{ij}\, F_{\hat{\omega}}\left(\tilde{E}\left(\frac{1}{\chi+\Upsilon}\ ; \nu_j\right), \tilde{L}\left(1\ ; \nu_j\right)\right)\mathbf{1}\{k_j(q,\nu,\Upsilon;\hat{\omega}) = i\}\right\}_{\hat{\omega}\in\Omega_\omega}\right] \\
=& F\left[\left\{q_{ij}\, \Lambda_{\hat{\omega}}\left(\frac{1}{\chi+\Upsilon}\ ; \nu_j\right)\mathbf{1}\{k_j(q,\nu,\Upsilon;\hat{\omega}) = i\}\right\}_{\hat{\omega}\in\Omega_\omega}\right]
\end{aligned} \tag{16}$$

Where we use that CRS functions are homogeneous of degree 1, and where $\Lambda_{\hat{\omega}} \equiv F_{\hat{\omega}}\left(\tilde{E}\left(\frac{1}{\chi+\Upsilon}\ ; \nu_j\right), \tilde{L}\left(1\ ; \nu_j\right)\right)$.

The optimal allocation $k_j(q,\nu,\Upsilon;\hat{\omega})\ \forall j, \hat{\omega} \in \Omega_{\omega_j}$ selects the producer with the highest combination of pairwise total factor efficiency q_{kj} and the energy efficiency ν_j, weighted by its shadow cost, inclusive of the environmental damage. Our definition of energy efficiency ν_j (definition 1) requires that $-log\Lambda_{\hat{\omega}}$ exhibits decreasing differences in $\chi+\Upsilon$ and ν_j. Therefore, as $\chi+\Upsilon$ increases, the optimal allocation favors a more energy efficient (higher ν_j) producer.

Due to the simplified Cobb-Douglas and quasi-linear utility setting, ν and q only affects welfare - hence the extensive margin allocation decision - through the shadow cost $\lambda_j(\omega)$. But due to the environmental damage ($\Upsilon > 0$), the shadow cost of energy is larger than its private production cost (χ), which pushes the social planner to select more energy efficient (high ν_j) producers. Absent any environmental externality ($\Upsilon = 0$), energy efficiency considerations are down-weighted, and might not even enter the decision if $\chi = 1$ (equal shadow costs for both factors). This is precisely what happens in a decentralized equilibrium, in which the firms don't internalize the environmental externality. $\chi < 1$ might even oppose the equilibrium and the optimal allocation, as the former will prioritize low ν_j (as long as $\chi+\Upsilon > 1$).

A.2.2 Firm-level energy intensity

This appendix formalizes the composition effect of firm-level energy intensity in terms of value added that is discussed in the main text. To simplify exposition, we consider a simplified setting with one supplier $i \in \mathcal{S}_\omega$ per input, and in which ν_i and q_{ij} are common to all inputs for a given buyer j. Define firm j's value added, i.e., the energy and labor expenditure of in-house production,[50] and its intermediate input expenditure as

$$VA_j \equiv \sum_{\omega\in\mathcal{I}_j} c_\omega(p_e;\nu_j)\, x_j(\omega), \qquad p_j^{\mathcal{O}} x_j^{\mathcal{O}} \equiv \sum_{\omega\in\Omega\setminus\mathcal{I}_j} \frac{c_\omega(p_e;\nu_i)}{q_{ij}}\, x_j(\omega). \tag{17}$$

[50] The intuition carries over to defining value added as labor expenditures only.

Under CRS, $c_j y_j = VA_j + p_j^{\mathcal{O}} x_j^{\mathcal{O}}$, so that Equation (7) can be written as

$$\frac{e_j}{c_j y_j} = \frac{e_j}{VA_j}\left(1 - \frac{p_j^{\mathcal{O}} x_j^{\mathcal{O}}}{c_j y_j}\right), \qquad \frac{e_j}{VA_j} = \sum_{\omega \in \mathcal{I}_j} \sigma_\omega(p_e; \nu_j)\, \frac{c_\omega(p_e; \nu_j)\, x_j(\omega)}{VA_j}. \tag{18}$$

Normalizing by value added removes the boundary effect, captured by the second factor, but e_j/VA_j remains a weighted average of input-level energy intensities over $\mathcal{I}_j$.

Lemma 3 (Sufficient statistic)**.** *Given p_e, e_j/VA_j is independent of the make-or-buy allocation if and only if $\sigma_\omega(p_e; \nu_j)$ is constant across $\omega \in \Omega$. In that case, it is strictly decreasing in ν_j, hence a sufficient statistic for ν_j.*

Proof sketch. If $\sigma_\omega(p_e; \nu_j) = \sigma(p_e; \nu_j)$ for all ω, the weights in (18) sum to one and $e_j/VA_j = \sigma(p_e; \nu_j)$, which is common to all producers with efficiency ν_j since F_ω is, and strictly decreasing in ν_j by Definition 1(i). Conversely, as q_{ij} falls from ∞ to 0, $\mathcal{I}_j$ expands from $\emptyset$ to Ω. Absent ties, inputs enter one at a time, and by Lemma 4 below, each entry leaves e_j/VA_j unchanged only if the entering input's $\sigma_\omega(p_e; \nu_j)$ equals e_j/VA_j before the switch. Hence, if e_j/VA_j is independent of the allocation, $\sigma_\omega(p_e; \nu_j)$ is constant across Ω. □

Outside this case, e_j/VA_j depends on which inputs are produced in-house. Because make-or-buy decisions operate at the extensive margin, the composition effect can be characterized input by input.

Lemma 4 (Composition at the margin)**.** *Suppose that $\mathcal{I}_j \neq \emptyset$ and that input ω enters it, i.e., $\Delta_{ij}(\omega; p_e)$ crosses zero from below. At the switch, the intermediate input expenditure share falls by $\frac{c_\omega(p_e;\nu_j)x_j(\omega)}{c_j y_j}$, and e_j/VA_j changes by*

$$\frac{c_\omega(p_e; \nu_j)\, x_j(\omega)}{VA_j + c_\omega(p_e; \nu_j)\, x_j(\omega)} \left(\sigma_\omega(p_e; \nu_j) - \frac{e_j}{VA_j}\right), \tag{19}$$

with VA_j and e_j/VA_j evaluated before the switch. The composition effect is therefore positive if and only if the entering input is more energy-intensive than the average input produced in-house.

Proof sketch. At the switch, $c_\omega(p_e; \nu_j) = c_\omega(p_e; \nu_i)/q_{ij}$: the effective input prices faced by j are continuous, and so are c_j, y_j and $x_j(\cdot)$. Only the assignment of ω's expenditure changes, from $p_j^{\mathcal{O}} x_j^{\mathcal{O}}$ to VA_j, together with the associated energy use $\sigma_\omega(p_e; \nu_j)\, c_\omega(p_e; \nu_j)\, x_j(\omega)$. □

By Definition 1(ii), $\Delta_{ij}(\omega; p_e)$ is strictly increasing in ν_j, so $\mathcal{I}_j$ only expands as ν_j rises. Between switches, e_j/VA_j varies at the intensive margin only: each $\sigma_\omega(p_e; \nu_j)$ falls by Definition 1(i), and the weights shift with relative in-house costs, unless F is Cobb-Douglas, in which case expenditure shares, and hence the intermediate input share, are constant between switches. The sign of the composition effect thus hinges on how inputs sort into in-house production. We illustrate with a CES technology.

CES illustration. Let $F_\omega(e, l; \nu) = \left(\beta_\omega^{1-\rho}(\nu e)^\rho + (1-\beta_\omega)^{1-\rho} l^\rho\right)^{1/\rho}$, with $\rho < 0$, the case in which this parametrization satisfies Definition 1(i). Since ν is energy-biased, $c_\omega(p_e; \nu)$ depends on p_e only through

p_e/ν, and Shephard's lemma gives

$$\frac{\partial \log c_\omega(p_e;\nu)}{\partial \log \nu} = -p_e\,\sigma_\omega(p_e;\nu), \qquad p_e\,\sigma_\omega(p_e;\nu) = \left[1 + \frac{1-\beta_\omega}{\beta_\omega}\left(\frac{\nu w}{p_e}\right)^{\frac{-\rho}{1-\rho}}\right]^{-1}. \tag{20}$$

An efficiency gain lowers unit cost in proportion to the energy cost share $p_e\sigma_\omega(p_e;\nu)$, which increases with β_ω: $\log c_\omega$ exhibits decreasing differences in (β_ω, ν). By (4), for given ν_i and q_{ij},

$$\frac{\partial \Delta_{ij}(\omega;p_e)}{\partial \beta_\omega} = \frac{\partial \log c_\omega(p_e;\nu_i)}{\partial \beta_\omega} - \frac{\partial \log c_\omega(p_e;\nu_j)}{\partial \beta_\omega} \gtrless 0 \quad \Longleftrightarrow \quad \nu_j \gtrless \nu_i, \tag{21}$$

i.e., a firm more energy-efficient than its supplier has a comparative advantage in producing energy-intensive inputs in-house.

Corollary 2. *In the CES case, suppose that ν_i and q_{ij} are common to all inputs and that β_ω differs across inputs.*

(i) ***Sorting.*** *If $\nu_j > \nu_i$, every input produced in-house is more energy-intensive than every outsourced input; if $\nu_j < \nu_i$, the reverse holds. By Lemma 2, $\mathcal{I}_j$ is a non-empty strict subset of Ω only if $q_{ij} > 1$ in the first case and $q_{ij} < 1$ in the second. Consequently, e_j/VA_j lies above the average of $\sigma_\omega(p_e;\nu_j)$ over Ω, weighted by expenditure shares in $c_j y_j$, in the first case, and below it in the second.*

(ii) ***Margin.*** *As ν_j rises, $\mathcal{I}_j$ expands. While $\nu_j > \nu_i$, each entering input is less energy-intensive than every input already produced in-house, so its entry lowers e_j/VA_j; while $\nu_j < \nu_i$, it is more energy-intensive than every such input, so its entry raises e_j/VA_j.*

(iii) ***Intermediate input share.*** *At the extensive margin, e_j/VA_j and $p_j^{\mathcal{O}} x_j^{\mathcal{O}}/c_j y_j$ move in the same direction if $\nu_j > \nu_i$ and in opposite directions if $\nu_j < \nu_i$, whatever shifts $\Delta_{ij}(\omega;p_e)$ (ν_j, q_{ij} or p_e).*

Proof sketch. By (21), $\Delta_{ij}(\omega;p_e)$ is strictly monotone in β_ω, so $\mathcal{I}_j$ is an upper (lower) set in β_ω if $\nu_j > (<)\ \nu_i$, and $\sigma_\omega(p_e;\nu_j)$ is strictly increasing in β_ω by (20). (i) follows since a weighted average of an increasing function over an upper (lower) set exceeds (falls short of) its average over Ω. As ν_j rises, $\Delta_{ij}(\omega;p_e)$ increases for every ω, since $\partial\Delta_{ij}(\omega;p_e)/\partial \log \nu_j = p_e\sigma_\omega(p_e;\nu_j) > 0$ by (20); $\mathcal{I}_j$ therefore expands while remaining an upper (lower) set, so each entrant lies below (above) every input in $\mathcal{I}_j$ in β_ω. (ii) and (iii) follow from Lemma 4, since each entry lowers the intermediate input share and exits reverse all signs. □

Corollary 2 shows why the correlation of the composition effect with ν_j is ambiguous in sign. Across firms, comparative advantage raises the measured intensity of firms more energy-efficient than their suppliers and lowers that of less efficient ones (i), generating a positive correlation. Among firms more efficient than their suppliers, however, higher efficiency extends in-house production to inputs less energy-intensive than those already made in-house (ii), generating a negative one; among less efficient firms, it extends in-house production to more energy-intensive inputs, reinforcing the positive one. Which force dominates depends on the joint distribution of efficiency gaps $\nu_j - \nu_i$ and outsourcing capabilities q_{ij}.

Since the intensive margin makes e_j/VA_j decrease with ν_j, a positive correlation attenuates measured efficiency differences, possibly to the point of reversing their ranking since $p_e\sigma_\omega(p_e;\nu)$ spans $(0,1)$ as β_ω does, while a negative one amplifies them.

For the same reason, e_j/VA_j need not fall with the intermediate input expenditure share. Within each case, the sign of their co-movement is set by (iii): positive for firms more efficient than their suppliers, negative otherwise. Across the two cases, firms more efficient than their suppliers have a more energy-intensive in-house composition by (i), but a lower intermediate input share only if the energy-intensive inputs they keep in-house account for a larger share of their costs than the energy-light inputs kept in-house by less efficient firms. This is a condition on expenditure shares, which depend on F, not on unit costs alone.

When ν_i and q_{ij} vary across inputs, $\mathcal{I}_j$ is no longer an upper or lower set in β_ω: Lemma 4 still applies input by input, but comparative advantage only tilts selection. The sign of both correlations is then a quantitative question, which our quantitative model addresses by identifying ν_j net of firm boundaries.

A.2.3 Aggregate Energy Intensity

Aggregating over firms, the change in aggregate energy intensity $\Delta\frac{E}{PY}$ following an increase in energy price $\Delta p_e = p_e^{SCC} - p_e$ can be decomposed into all three margins of Proposition 1. Denote the allocation at p_e as $\mathcal{A}(p_e) \equiv \{k_j(\hat{\omega};p_e)\}_{\hat{\omega}\in\Omega, j\in\mathcal{S}}$, the sales share of firm j as $s_j(p_e,\mathcal{A}(p_e)) \equiv \frac{c_j(p_e,\mathcal{A}(p_e))\, y_j(p_e,\mathcal{A}(p_e))}{P(p_e,\mathcal{A}(p_e))\, Y(p_e,\mathcal{A}(p_e))}$, and the expenditure share of firm j on input $\hat{\omega}$ as $s_{j\hat{\omega}}(p_e,\mathcal{A}(p_e)) \equiv \frac{\left(c_{\hat{\omega}}(p_e;\nu_{k_j(\hat{\omega};p_e)})/q_{k_j(\hat{\omega};p_e)\,j}\right) x_j(\hat{\omega})}{c_j(p_e,\mathcal{A}(p_e))\, y_j(p_e,\mathcal{A}(p_e))}$, where we make the dependence on the endogenous allocation $(k_j,\mathcal{A})$ explicit but suppress the wage to simplify notation. Aggregate energy intensity satisfies $\frac{E}{PY} = \int_{j\in\mathcal{S}} s_j(p_e,\mathcal{A}(p_e)) \sum_{\hat{\omega}\in\Omega} s_{j\hat{\omega}}(p_e,\mathcal{A}(p_e))\sigma_{\hat{\omega}}(p_e,\nu_{k_j(\hat{\omega};p_e)})\, dj$, so that, to first order in Δp_e for the intensive-margin terms,

$$
\begin{aligned}
\Delta\frac{E}{PY} \;=\; \int_{j\in\mathcal{S}} \Bigg\{ & s_j\left(p_e,\mathcal{A}(p_e)\right) \sum_{\hat{\omega}\in\Omega} s_{j\hat{\omega}}\left(p_e,\mathcal{A}(p_e)\right)\, \sigma_{\hat{\omega}}\left(p_e;\nu_{k_j(\hat{\omega};p_e)}\right) \\
& \times \Bigg(\underbrace{\frac{1}{\sigma_{\hat{\omega}}} \frac{\partial^2 \log c_{\hat{\omega}}\left(p_e;\nu_{k_j(\hat{\omega};p_e)}\right)}{\partial p_e^2}}_{\text{Int. margin: energy-labor within-firm } (\circ)} + \underbrace{\frac{\partial \log s_{j\hat{\omega}}\left(p_e,\mathcal{A}(p_e)\right)}{\partial p_e}\Big|_{\mathcal{A}\text{ fixed}} + \frac{\partial \log s_j\left(p_e,\mathcal{A}(p_e)\right)}{\partial p_e}\Big|_{\mathcal{A}\text{ fixed}}}_{\text{Int. margin: product demand at each stage } (\circ)} \Bigg) \Delta p_e \\
& + s_j\left(p_e^{SCC},\mathcal{A}(p_e)\right) \sum_{\hat{\omega}\in\Omega} \Bigg[\underbrace{s_{j\hat{\omega}}\left(p_e^{SCC},\mathcal{A}(p_e^{SCC})\right)\left[\sigma_{\hat{\omega}}\left(p_e^{SCC};\nu_{k_j(\hat{\omega};p_e^{SCC})}\right) - \sigma_{\hat{\omega}}\left(p_e^{SCC};\nu_{k_j(\hat{\omega};p_e)}\right)\right]}_{\text{Ext. margin: energy intensity}} \\
& \qquad + \underbrace{\sigma_{\hat{\omega}}\left(p_e^{SCC};\nu_{k_j(\hat{\omega};p_e)}\right)\left[s_{j\hat{\omega}}\left(p_e^{SCC},\mathcal{A}(p_e^{SCC})\right) - s_{j\hat{\omega}}\left(p_e^{SCC},\mathcal{A}(p_e)\right)\right]}_{\text{Ext. margin: expenditure share}} \Bigg] \\
& + \underbrace{\left[s_j\left(p_e^{SCC},\mathcal{A}(p_e^{SCC})\right) - s_j\left(p_e^{SCC},\mathcal{A}(p_e)\right)\right] \sum_{\hat{\omega}\in\Omega} s_{j\hat{\omega}}\left(p_e^{SCC},\mathcal{A}(p_e^{SCC})\right) \sigma_{\hat{\omega}}\left(p_e^{SCC};\nu_{k_j(\hat{\omega};p_e^{SCC})}\right)}_{\text{Indirect effect of ext. margin: across firms}} \Bigg\}\, dj
\end{aligned}
\tag{22}
$$

where terms marked $(\circ)$ are first-order approximations in Δp_e, while the extensive-margin terms are

exact and evaluated at p_e^{SCC} to isolate the reallocation effect from the price effect already captured by the intensive margin.

The intensive-margin terms correspond to the first two channels of Proposition 1. The within-firm term is the response of $\sigma_{\hat{\omega}}$ for a given producer, negative by factor substitution. The product-demand terms capture how input and sales shares respond to p_e for a given allocation. By Definition 1(i), $c_{\hat{\omega}}$ increases proportionally more with p_e for low-ν producers; whether this raises or lowers $s_{j\hat{\omega}}$ and s_j depends on the elasticities of substitution between intermediate inputs and between final products, respectively.

The three extensive-margin terms are triggered by producer selection, the third channel. The two within-firm extensive-margin terms capture the two channels through which a switch affects aggregate energy intensity: the energy-intensity term reflects the change in how energy-intensively the input is produced, and the expenditure-share term the reallocation of economic weight across producers. The first is nonzero only for firm-input pairs whose producer differs between p_e and p_e^{SCC}; the second is nonzero for every input of a firm that switches at least one, since the switch changes c_j and reweighs all of j's inputs through the same substitution elasticities as the input-demand terms. Sales shares s_j, by contrast, have no extensive margin of their own: no firm's sales to the household jump to zero when a producer switches. They nonetheless respond discretely to re-sorting, through the cost jumps of j and of its competitors, and the across-firm indirect term collects this effect, which is nonzero even for firms whose own sourcing is unchanged. A switching firm's cost jumps down, since re-optimizing at p_e^{SCC} weakly lowers its unit cost relative to keeping $\mathcal{A}(p_e)$. How these jumps rank across firms, however, is not determined by ν alone: with a common set of external suppliers, high-ν downstream firms would secure a weakly lower cost at any price absent other heterogeneity, but pairwise-specific outsourcing capabilities break this ranking.

Both the planner and the private firm select producers by minimizing unit cost, and ν_k enters both decisions only through unit costs. The planner evaluates costs at p_e^{SCC}, which prices energy at its social cost; the private firm evaluates them at p_e and therefore underweights ν_k whenever $p_e < p_e^{SCC}$. A single distorted decision thus shows up in all three extensive-margin terms. The weights $s_j s_{j\hat{\omega}}$ and $\sigma_{\hat{\omega}}$ scale the aggregate effect of a switch: it matters more the larger the firm's expenditure on the input and the more energy-intensive its current production process.

The aggregate importance of the extensive margin therefore depends on both the *frequency* at which switches occur and their *impact* when they do, which we now formalize. Write $\Delta\sigma_{j\hat{\omega}}$ and $\Delta s_{j\hat{\omega}}$ for the bracketed differences in the second line of (22). Each extensive-margin term combines a *frequency*, the weighted share of pairs or firms for which it is nonzero, and an *impact*, its magnitude when it is. The three terms are nested in frequency: the energy-intensity term applies only to firm-input pairs that switch producer, the expenditure-share term to every input of a switching firm, and the across-firm indirect term potentially to all firms as soon as any firm switches. By Lemma 1 and Definition 1(i), $\Delta\sigma_{j\hat{\omega}} \leq 0$, and its magnitude increases with the ν-gap between the privately and socially optimal producers. The impact of the two share terms depends instead on how switches change unit costs and on the elasticities of substitution between inputs and between final products, respectively. Which pairs switch is governed by the joint distribution of outsourcing capabilities $\{q_{ij}\}$ and energy efficiencies $\{\nu_i, \nu_j\}$, which therefore

determines both the frequency of switches and the ν-gaps they carry.

The make-or-buy decision implies that even under Cobb-Douglas aggregation across both intermediate inputs and downstream products, a distortion affects aggregate energy intensity through between-firm reallocation. This is because the make-or-buy decision makes potential producers perfect substitutes regardless of the production function.[51] Moreover, when firm-level heterogeneity is not Hicks-neutral, even a common distortion induces such reallocation through the make-or-buy decision. Under Cobb-Douglas aggregation across firms and intermediate inputs,

$$\frac{E}{PY} = \int_{j \in \mathcal{S}} s_j \sum_{\hat{\omega} \in \Omega} s_{\hat{\omega}} \, \sigma_{\hat{\omega}}(p_e; \nu_{k_j(\hat{\omega}; p_e)}) \, dj,$$

where the sales and expenditure shares s_j and $s_{\hat{\omega}}$ are independent of prices. This shuts down the effect of product substitution on aggregate energy intensity: in (22), the product-demand terms, the extensive-margin expenditure-share term and the across-firm indirect term all vanish, leaving only the within-firm term and the extensive-margin energy-intensity term. Without make-or-buy, a common distortion p_e would thus affect aggregate energy intensity only through factor substitution, that is, through its direct effect on $\sigma_{\hat{\omega}}(p_e; \nu_j)$: relative quantities respond, but expenditure does not reallocate across firms, and the distortion operates only within producers. With endogenous make-or-buy decisions, by contrast, the identity of the producer $k_j(\hat{\omega}; p_e)$ depends on the distorted price, so that $\sigma_{\hat{\omega}}$ also varies through the selection of $\nu_{k_j(\hat{\omega}; p_e)}$. In other words, at a given energy price, E/PY is solely determined by the distribution of ν across selected producers, but this distribution is *endogenous* to the energy price.In the next section, we extend Cobb-Douglas production to energy and labor within firms, which means that the *impact* of a switch is solely determined by, and monotonically increasing in, the ν-gap ($\Delta\nu \equiv \nu_j - \nu_i$), while the *frequency* depends on its joint distribution with relative outsourcing capabilities ($\Delta \log q \equiv \log(q_{jj}) - \log(q_{ij})$). Under further distributional assumptions linking $\Delta \log q$ to the ν-gap, the switching frequency admits a closed form, which allows us to characterize extensive-margin misallocation as a function of the ν-gap distribution alone (Appendix A.2.4)

A.2.4 Cobb-Douglas Production Function

This section applies the results derived in the general two-stage model to Cobb-Douglas production. We focus on the case with only one upstream supplier and only one input, dropping the index ω.

Production. A mass 1 of downstream firms, indexed by j, need one intermediate input to produce. They can either produce it in-house, using energy and labor, or outsource its production to the upstream

[51] In the terminology of Baqaee and Farhi (2020), there should be no change in allocative efficiency in response to a shock in the purely Cobb-Douglas economy, because the resource allocation is fixed. This is the same intuition that arises from a purely vertical economy, in which each input only has one supplier. However, the make-or-buy decision breaks the Cobb-Douglas aggregation of potential producers - or adds a potential supplier to each input in the vertical economy - by making them perfect substitutes. This means that a shock can induce cross-firm reallocation i.e., changes in allocative efficiency.

supplier i:[52]

$$y_j = \begin{cases} q_{jj} \frac{1}{A_j} e_j^{1-\nu_j} l_j^{\nu_j} & \text{if in-house} \\ q_{ij} x_i & \text{if outsourced to } i \end{cases}$$

Where ν_j is the relative energy efficiency of firm j, q_{ij} its pairwise outsourcing capabilities and A_j a firm-specific normalizing constant equal to $\nu_j^{\nu_j}(1-\nu_j)^{1-\nu_j}$. We suppose that factor prices satisfy $w \leq p_e$, so that ν_i remains a pure technology parameter in the sense that higher ν_i reduces costs at any $p_e > w$.[53]

If the intermediate input is outsourced, it is produced by the upstream firm i that prices at marginal cost. Its production symmetrically requires labor and energy.

$$x_i = \frac{1}{A_i} e_j^{1-\nu_i} l_j^{\nu_i}$$

Importantly, firm i produces with its own effective energy efficiency ν_i.

Equilibrium. The downstream firms choose whether to produce the intermediate input in-house or to outsource it. Given the make-or-buy allocation, all firms chooses their factor and input demand to minimize costs. The Cobb-Douglas production function yields the usual unit cost for the production of the input for downstream firm j, when production occurs at firm $k \in \{i, j\}$:

$$\frac{c(p_e, w; \nu_k)}{q_{kj}} = \frac{p_e^{1-\nu_k} w^{\nu_k}}{q_{kj}}$$

Make-or-buy decision. Firm j decides on whether to produce in-house or to outsource by comparing the cost of of each option. With CRS and revenues decreasing in prices, this is equivalent to comparing the unit cost. Formally, firm j produces in-house if:

$$\frac{p_e^{1-\nu_j} w^{\nu_j}}{q_{jj}} < \frac{p_e^{1-\nu_i} w^{\nu_i}}{q_{ij}}$$

Which corresponds to Equation (3) in the main text. Higher $\frac{q_{ij}}{q_{jj}}$ will incentivize the firm to outsource, while a higher $\frac{\nu_j}{\nu_i}$ will lead the firm to produce in-house. Therefore, more energy efficient firms require a larger outsourcing capability to outsource.

Energy misallocation. For each firm, inefficient outsourcing arises if $\frac{c(p_e;\nu_j)}{q_{jj}} > \frac{c(p_e;\nu_i)}{q_{ij}}$ and $\frac{c(p_e^{SCC};\nu_j)}{q_{jj}} < \frac{c(p_e^{SCC};\nu_i)}{q_{ij}}$, that is:

$$\log\left(\frac{w}{p_e^{SCC}}\right) < \frac{\log(q_{jj}) - \log(q_{ij})}{\nu_j - \nu_i} \leq \log\left(\frac{w}{p_e}\right) < 0$$

[52] In this simplified framework, an outsourcing firm only acts as a retailer, producing no value-added. In the full model, firms need multiple intermediate inputs, and at least one of them must be produced in-house.

[53] This is WLOG because the factor prices are common across firms, so we can normalize production by a common shifter that normalizes the unit of labor and energy to ensure that $p_e > w$ always holds in the data. Moreover, all our intuitions hold under $p_e < w$ (so costs increasing in ν), but the sufficient condition for reallocation becomes a positive correlation ($\text{sgn}(\Delta \log q) = \text{sgn}(\Delta \nu)$).

With decreasing differences, this can only happen if $\nu_j > \nu_i$ (Lemma 1) i.e., distorted energy prices increase energy intensity by allocating production to the less efficient firm. Moreover, with $\nu_j > \nu_i$, observing outsourcing in the first place requires that $q_{ij} > q_{jj}$.

Likewise, inefficient in-house production arises if $\frac{c(p_e;\nu_j)}{q_{jj}} < \frac{c(p_e;\nu_i)}{q_{ij}}$ and $\frac{c(p_e^{SCC};\nu_j)}{q_{jj}} > \frac{c(p_e^{SCC};\nu_i)}{q_{ij}}$, that is:

$$\log\left(\frac{w}{p_e^{SCC}}\right) < \frac{\log(q_{ij}) - \log(q_{jj})}{\nu_i - \nu_j} \leq \log\left(\frac{w}{p_e}\right) < 0$$

This can only happen for firm-pair with $\nu_j < \nu_i$. Moreover, with $\nu_j < \nu_i$, observing in-house production in the first place requires that $q_{ij} < q_{jj}$. Energy misallocation therefore requires a trade-off between relative energy efficiency and relative outsourcing capabilities (Lemma 2), which means that the joint distribution of these primitives is the crucial moment determining the extent of misallocation.

Energy-use misallocation increases energy intensity. The energy used to produce one monetary unit of the product sold by firm j is: $\sigma(p_e, \nu_{k_j(p_e)}) = \frac{1-\nu_{k_j(p_e)}}{p_e}$. As shown in the main text (Equation (10)), it decreases with p_e both through the intensive and the extensive margins. This results carry over to aggregate energy intensity ($\mathcal{E}(p_e)$). We illustrate their respective contributions in Figure A.1. We simulate a large mass of firms and aggregate their production through Cobb-Douglas ($Y = \prod_0^1 y_j^{\alpha_j}$. where $\int_0^1 \alpha_j dj = 1$), such that $\frac{p_j y_j}{PY} = \alpha_j$ and $\mathcal{E}(p_e) = \frac{E}{PY} = \int_0^1 \frac{p_j y_j}{PY} \sigma(p_e, \nu_{k_j(p_e)}) dj = \int_0^1 \alpha_j \sigma(p_e, \nu_{k_j(p_e)}) dj$. We normalize $q_{jj} = 1\ \forall j$. As the energy price p_e rises, firms whose switching threshold $p_e^* = w \cdot q_{ij}^{1/(\nu_j - \nu_i)}$ is exceeded move from outsourced to in-house, reducing energy intensity through the extensive margin. The figure plots equilibrium energy intensity $\mathcal{E}(p_e)$ against p_e, together with the counterfactual in which the outsourcing composition $\nu_{k_j(p_e)}$ is held fixed at the make-or-buy allocation chosen under the initial price. The gap between the two curves isolates the contribution of the extensive margin: as more firms switch in-house, $\nu_{k_j(p_e)}$ rises, amplifying the decline in $\mathcal{E}(p_e)$ beyond what the intensive margin alone would deliver.

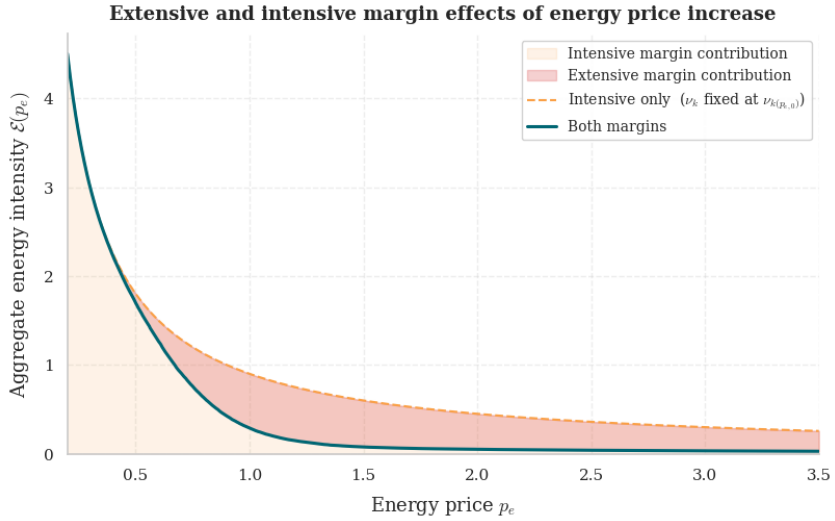


Figure A.1: Extensive and intensive margin effects of energy price changes

Notes: Simulated aggregate energy intensity at different prices of energy p_e. The dashed line keeps the make-or-buy allocation fixed at the initial allocation. The solid line allows for extensive-margin reallocation. The shaded areas correspond to the respective contributions of each margin. Parameters: $w = 1$, $\nu_i = 0.1$, $\nu_j = 0.9$; outsourcing wedges $\log(q_{ij})$ drawn i.i.d. from $\mathcal{N}(-0.18,\ 0.25^2)$; cost shares $\{\alpha_j\}$ drawn from a Dirichlet distribution (Dir($\mathbf{1}$)).

To fully characterize extensive margin misallocation, we specify the welfare function to be linear in

energy: $U = X + \frac{Y^\alpha}{\alpha} - \Upsilon E$, where X is a numeraire good produced competitively one-to-one with labor, which is supplied inelastically, meaning that $P_X = w = 1$. This specification implies that the optimal carbon tax, if rebated lump-sum to the household, is equal to the marginal damage Υ and independent of the underlying allocation. Moreover, we assume that energy is produced perfectly competitively using one unit of the numeraire good, hence $p_e = w = 1$ as well. In that case, the make-or-buy decision generates misallocation for the firms for which

1. $\frac{-\Delta \log(q)}{\Delta \nu} \geq 0$
2. $log(p_e + \Upsilon)\Delta\nu = log(1 + \Upsilon)\Delta\nu \geq (<) - \Delta logq \geq (<)\, 0 \quad \text{if} \quad \Delta\nu \geq\ (<)\, 0$

The first inequality is the trade-off discussed above. Firms facing such a trade-off will make an inefficient decision if their outsourcing capability gap $-\Delta \log(q) = \log(q_{jj}) - \log(q_{ij})$ (i.e., the relative (dis)advantage from in-house production) is smaller in magnitude than the relative effective energy efficiency gap $\nu_i - \nu_j$, weighted by the marginal damage Υ. In equilibrium, $p_e = w = 1$ implies that private firms don't consider the relative energy efficiency: their make-or-buy decision is solely driven by the outsourcing capability gap. However, the social planner weighs the latter against the energy efficiency gains, which arise by reducing the environmental damage.

Extensive-margin misallocation and the joint distribution of primitives The Cobb-Douglas case offers a tractable setting in which the *impact* of a switch is solely determined by, and increasing in, the ν-gap, while the *frequency* depends on its joint distribution with outsourcing capabilities. Parameterizing the latter, we can therefore characterize the extensive margin effect in terms of the ν-gap distribution only. Suppose that $\log(q_{jj}) - \log(q_{ij}) = \rho\, sgn(\nu_j - \nu_i)\, |\nu_j - \nu_i|^\theta\, \varepsilon_{ij}$, where $\rho \in \{-1, 1\}$ governs the sign of the correlation between relative outsourcing capabilities and relative energy efficiency (i.e., the ranking of absolute advantages), θ governs how the size of the q-gap scales with the ν-gap, and ε_{ij} is an idiosyncratic noise term.[54] We denote by H the distribution of $\nu_j - \nu_i$.

$$-\Delta \frac{E}{PY}\Big|_{\text{Ext}} = \int_{-1}^{1} \underbrace{\mathcal{P}\left(\log\left(\frac{p'_e}{w}\right) \geq -\rho\, |\nu_j - \nu_i|^{\theta-1} \varepsilon_{ij} \geq \log\left(\frac{p_e}{w}\right)\right)}_{\text{Frequency}} \times \underbrace{\frac{|\nu_j - \nu_i|}{p'_e}}_{\text{Impact}}\, dH(\nu_j - \nu_i) \tag{23}$$

The ν-gap matters through the impact, linearly, and through the switching frequency. It affects the latter through two channels, which combine in the exponent $\theta - 1$. First, the ν-gap enters the cost comparison directly: ceteris paribus, the more energy-efficient producer is relatively cheaper, and higher energy prices make this advantage more salient. Second, when the q-gap depends on the ν-gap ($\theta \neq 0$), the ν-gap has an indirect effect on the cost comparison through pairwise TFP. With $\rho = -1$, this indirect channel reinforces the direct one if $\theta < 0$ and offsets it if $\theta > 0$. Since selection depends on the ratio of the TFP disadvantage to the energy-efficiency advantage, $|\nu_j - \nu_i|^{\theta-1} \varepsilon_{ij}$, the direct channel dominates if and only if $\theta < 1$.

[54] Letting $\rho \in \mathbb{R}$ is isomorphic to normalizing the mean of ε_{ij} to 1, but this version cleanly isolates the sign of the correlation from any level shift.

We consider three cases. First, we impose $\varepsilon_{ij} = 1\ \forall i, j$, so that the q-gap is a deterministic function of the ν-gap and which pairs switch is known with certainty. Equation (23) then becomes

$$-\Delta \frac{E}{PY}\Big|_{\text{Ext}} = \begin{cases} \underbrace{\int_{l}^{u} \frac{\nu_j - \nu_i}{p'_e}\, dH(\nu_j - \nu_i)}_{\mathcal{O}\to\mathcal{I}} + \underbrace{\int_{-u}^{-l} \frac{\nu_i - \nu_j}{p'_e}\, dH(\nu_j - \nu_i)}_{\mathcal{I}\to\mathcal{O}} & \text{if } \rho = -1 \\ 0 & \text{if } \rho = 1 \end{cases}$$

where $u = (\log(p_e/w))^{\frac{1}{\theta-1}}$ and $l = (\log(p'_e/w))^{\frac{1}{\theta-1}}$ if $\theta < 1$, and $l = (\log(p_e/w))^{\frac{1}{\theta-1}}$ and $u = (\log(p'_e/w))^{\frac{1}{\theta-1}}$ if $\theta > 1$. With $\rho = 1$, energy efficiency and outsourcing capabilities favor the same producer, so there is no trade-off and no switching (Lemma 2).[55] With $\rho = -1$, the more energy-efficient producer always has a disadvantage in pairwise TFP, whose size varies with the ν-gap according to θ,[56] and the extensive margin is given by the mass of pairs between the bounds, weighted by their impact.

A larger price gap always expands the set of switching pairs. Consider a decrease in p_e, keeping p'_e fixed, which increases the distortion when $p'_e = p_e^{SCC}$. The upper (lower) bound $(\log(p_e/w))^{\frac{1}{\theta-1}}$ increases (decreases) if $\theta < (>)\ 1$, expanding the switching interval. Which pairs it draws in depends on how the ratio $|\nu_j - \nu_i|^{\theta-1}$ varies with the ν-gap. When $\theta < 1$, this ratio decreases with the ν-gap, so as p_e falls the interval expands from low toward high ν-gap (high-impact) pairs. When $\theta < 0$, the direct and indirect effects reinforce each other in this direction, since a larger ν-gap both raises the salience of energy efficiency *and* shrinks the TFP disadvantage. When $\theta > 1$, the ratio increases with the ν-gap, and the interval expands from high toward low ν-gap pairs. In the special case $\theta = 1$, the ratio is constant and all pairs cross simultaneously at $p_e = \exp(1)\, w$. In sum, ρ governs whether switching is possible at all, while θ determines which pairs are drawn in as the distortion grows.

We now introduce idiosyncratic noise through ε_{ij}, so that for a given ν-gap, the share of pairs that switch depends on the distribution of ε_{ij}. We impose $\rho = -1$ and $\varepsilon_{ij} > 0$ throughout to ensure negative correlation. Since the ν-gap affects both the width and the location of the switching interval, entering both bounds but with a larger weight in the upper one ($p'_e > p_e$), its effect on the switching frequency depends on this distribution. Under a uniform distribution, location is irrelevant as long as the interval lies within the support: only its width matters, which isolates the role of θ. Under a Pareto distribution, the right tail is thin, so a rightward shift of the interval reduces the mass of switchers even if the interval expands. This makes location operative and, as we show, reverses the effect of θ relative to the uniform case.

Second, suppose that $\varepsilon_{ij} \sim U[0, \bar{\varepsilon}]$, with $\bar{\varepsilon} \geq \log(p'_e/w)$. For $\theta \leq 1$, the switching interval then lies within

[55] θ determines the size of the q-gap, but not which producer it favors.

[56] The make-or-buy decision reduces the dispersion in costs - increasing aggregate production from concavity - by allowing inefficient firms to substitute toward more efficient production processes. However, a negative correlation between q and ν limits this channel, because it benefits firms differently: efficient firms can lower their costs even further through outsourcing, while inefficient firms have effectively fewer opportunities to do so. In a setting with two-dimensional heterogeneity, the make-or-buy decision can exacerbate the cost dispersion despite offering more flexibility to firms.

the support for all ν-gaps, and

$$-\Delta \frac{E}{PY}\Big|_{\text{Ext}} = \int_{-1}^{1} \frac{|\nu_j - \nu_i|}{p'_e} \mathcal{P}\left(\log \frac{p_e}{w} |\nu_j - \nu_i|^{1-\theta} \leq \varepsilon_{ij} \leq \log \frac{p'_e}{w} |\nu_j - \nu_i|^{1-\theta}\right) dH(\nu_j - \nu_i)$$
$$= \int_{-1}^{1} \log\left(\frac{p'_e}{p_e}\right) \frac{|\nu_j - \nu_i|^{2-\theta}}{p'_e \, \bar{\varepsilon}} dH(\nu_j - \nu_i)$$

For $\theta > 1$, the same expression holds for $|\nu_j - \nu_i| \geq (\log(p'_e/w)/\bar{\varepsilon})^{\frac{1}{\theta-1}}$, below which the interval exits the support. θ governs how the switching frequency varies with the ν-gap. When $\theta < 1$, the TFP disadvantage grows more slowly than the energy-efficiency advantage, the increased salience of energy efficiency dominates, and the frequency *increases* with the ν-gap. When $\theta < 0$,[57] the two effects reinforce each other, since a larger ν-gap also shrinks the TFP disadvantage: switching is then concentrated among high ν-gap pairs, which also have the largest impact.[58] When $\theta > 1$, energy-efficiency advantages come with disproportionately large TFP disadvantages and, over the range where the expression holds, the frequency *decreases* with the ν-gap; at $\theta = 2$, this exactly offsets the rise in impact, so every pair in that range contributes equally.

Finally, suppose that $\varepsilon_{ij} \sim Pareto(\varepsilon_m, \psi)$. As long as the lower edge of the switching interval exceeds the scale ε_m, that is, for all ν-gaps if $\theta \geq 1$ and $\log(p_e/w) \geq \varepsilon_m$, and for $|\nu_j - \nu_i| \geq (\varepsilon_m / \log(p_e/w))^{\frac{1}{1-\theta}}$ if $\theta < 1$,

$$-\Delta \frac{E}{PY}\Big|_{\text{Ext}} = \int_{-1}^{1} \varepsilon_m^{\psi} \frac{|\nu_j - \nu_i|^{1-\psi(1-\theta)}}{p'_e} \left(\left(\log\left(\frac{p_e}{w}\right)\right)^{-\psi} - \left(\log\left(\frac{p'_e}{w}\right)\right)^{-\psi}\right) dH(\nu_j - \nu_i)$$

The switching frequency now increases with the ν-gap if $\theta > 1$ and decreases if $\theta < 1$, reversing the uniform case. This is because an increase in the ν-gap not only expands ($\theta < 1$) or shrinks ($\theta > 1$) the switching interval, but also *shifts* it, as the upper bound responds more than the lower one. If $\theta < 1$, both bounds increase, the upper one more so: the interval expands but shifts to the right, where the Pareto tail is thin, which decreases the switching frequency. If $\theta > 1$, both bounds decrease with the ν-gap, because the TFP disadvantage increases more than the energy-efficiency advantage, the upper one especially so: the interval shrinks but shifts to the left, increasing the mass of switchers.

Together, these three cases establish that the extent of cross-producer misallocation depends not only on the sign of the correlation between energy efficiency and outsourcing capabilities (ρ), but also on its shape (θ) and on the distribution of idiosyncratic noise: the same ν-gap distribution can generate very different magnitudes of misallocation depending on how these primitives interact. This is precisely why quantifying cross-producer misallocation requires recovering the empirical joint distribution of these primitives, which we do for the Italian manufacturing sector in our quantitative exercise (Section 4).

[57] If $\theta = 0$, the size of the TFP disadvantage is independent of the ν-gap, which then affects selection only directly, through the salience of the price change $\log\left(\frac{p'_e}{p_e}\right)$: the frequency increases with the ν-gap. If $\theta = 1$, the direct and indirect effects exactly offset each other and the frequency is independent of the ν-gap; the joint distribution of primitives then matters only through ρ and the distribution of ε_{ij}.

[58] A steeper profile does not imply more misallocation in total: holding H and $\bar{\varepsilon}$ fixed, the integrand $|\nu_j - \nu_i|^{2-\theta}$ increases with θ on $\theta \leq 1$, since $|\nu_j - \nu_i| \leq 1$.

Extensive margin endogenizes aggregate technologies We show that, even with firm-level and pairwise heterogeneity, our model can aggregate to a sectoral or economy-wide representation under the Cobb-Douglas specification. Similar aggregate structures are widely used to study environmental policies (e.g., a static version of the canonical framework in Golosov et al. (2014) with no capital and a single polluting energy input). However, the make-or-buy decisions *endogenize* sectoral energy intensity and Hicks-neutral technology, which are typically treated as primitives in this canonical framework.

To simplify the exposition, we only report the most important derivations and assume that each firm uses a single input, that each sector has a unit mass of firms, and that aggregate sectoral production is log-linear, $\log Y_s = \int_{j\in\mathcal{S}_s} \log y_j \, dj$. We obtain, for each sector s:

$$\begin{aligned}
\log P_s &= \int_{j\in\mathcal{S}_s} \log(p_j)\, dj \\
&= \log(p_e) \int_{j\in\mathcal{S}_s} (1-\nu_{k_j(p_e)})\, dj + \log(w) \int_{j\in\mathcal{S}_s} \nu_{k_j(p_e)}\, dj - \int_{j\in\mathcal{S}_s} \log(q_{k_j(p_e)j})\, dj \\
&= (1-\nu_s)\log(p_e) + \nu_s \log(w) - \log(Q_s)
\end{aligned}$$

Underlying firm-level heterogeneity is therefore summarized by the sectoral energy efficiency $\nu_s \equiv \int_{j\in\mathcal{S}_s} \nu_{k_j(p_e)}\, dj$, which determines the sectoral energy share $1-\nu_s$, and the Hicks-neutral technology $Q_s \equiv \exp\Big(\int_{j\in\mathcal{S}_s} \log(q_{k_j(p_e)j})\, dj\Big)$. Abstracting from make-or-buy decisions, one would measure the aggregate effect of a counterfactual energy price by first calibrating ν_s and Q_s to match the observed distribution of energy intensities and outsourcing capabilities of each sector s. Our model shows, however, that these aggregate objects are not primitives of the economy. Because the identity of the producer $k_j(p_e)$ depends on energy prices through the make-or-buy decision, the distributions that discipline both ν_s and Q_s are endogenous to the allocation. As a result, calibrating them from the equilibrium allocation and holding them fixed when evaluating counterfactual policies would generally bias the measured effects.

Sector-specific policies. Cross-producer misallocation arises even under a common distortion. However, in practice, Italian manufacturing firms face a set of sector-specific environmental policies. Two important examples are the EU ETS free allocation rules based on product-specific energy intensity benchmarks and the regime for energy-intensive firms in specific sectors (Imprese a Forte Consumo di Energia), which grants substantial reductions in electricity costs and compensation for indirect carbon pricing. Endogenous make-or-buy decisions distort the effects of those policies through both the allocative and the measurement roles. We illustrate these issues in turn in the Cobb-Douglas context.

First, to focus on the allocative role, suppose that the policies can target effective energy intensity $1-\nu$. We take the EU ETS as an example of sector-specific policy. It grants free allowances based on product-specific energy intensity benchmarks, which we model as a firm-specific tax equalizing input energy intensity to the most efficient producer: $\frac{e_i}{C_i} = \frac{1-\nu_i}{(1+\tau_i)p_e} = \frac{e^m}{C^m} \Rightarrow 1+\tau_i = \frac{1-\nu_i}{1-\nu_m}$. This product-specific policy treats all suppliers equally, but not all *potential* producers: this can amplify, rather than reduce, cross-producer misallocation. Suppose that all upstream suppliers $i \in \mathcal{S}_{\omega^u}$ of a given input ω^u are subject

to the policy τ_i defined above, but that the buyers of that input, who can produce it in-house, are not subject to any policy because they are classified based on the downstream product $\omega^d \neq \omega^u$ they sell. We have that $\sigma_{\omega^u}(\nu_i, (1+\tau_i)p_e) = \sigma_{\omega^u}(\nu_m, p_e)\ \forall i \in \mathcal{S}_{\omega^u}$, where the second argument of σ is the energy price faced by the producer: energy intensity is correctly equalized across suppliers. However, suppose that initially, a buyer j was outsourcing to a more energy efficient supplier i ($\nu_j < \nu_i < \nu_m$):

$$\log\left(\frac{p_e}{w}\right)(\nu_j - \nu_i) \leq \log(q_{ij})$$

With the policy in place, we might observe a reversal:

$$\log\left(\frac{p_e}{w}\right)(\nu_j - \nu_i) + (1-\nu_i)\underbrace{\log\left(\frac{1-\nu_i}{1-\nu_m}\right)}_{\geq 0} \geq \log(q_{ij})$$

Therefore, the policy induces an energy inefficient switch from $k_j(\omega^u; p_e) = i$, with embodied energy intensity $\sigma_{\omega^u}(\nu_i, p_e) = \frac{1-\nu_i}{p_e}$, to $k_j(\omega^u; p_e, \{\tau_i\}_{i\in\mathcal{S}_{\omega^u}}) = j$, with embodied energy intensity $\sigma_{\omega^u}(\nu_j, p_e) = \frac{1-\nu_j}{p_e} \geq \frac{1-\nu_i}{p_e}$.

Second, the make-or-buy decision implies that the proxy used to design the policy - the observed energy intensity - conflates true energy efficiency with organizational choices, resulting in a distorted targeting of support. To illustrate, suppose that our downstream firms produce by combining multiple inputs $\omega \in \Omega$ through a Cobb-Douglas technology, and that, as before, all these inputs can be produced in-house or outsourced:

$$y_j = \prod_{\omega\in\Omega} x_j(\omega)^{\alpha^\omega}, \quad \sum_{\omega\in\Omega}\alpha^\omega = 1$$

The observed energy intensity of a downstream firm j producing ω^d is:

$$\frac{e_j}{C_j} = \frac{1-\nu_j}{p_e}\sum_{\omega\in\mathcal{I}_j}\alpha^\omega$$

where $\mathcal{I}_j$ is the set of inputs produced in-house. The product-specific benchmark m introduces a firm-specific wedge:

$$1+\tau_j = \frac{1-\nu_j}{1-\nu_m}\frac{\sum_{\omega\in\mathcal{I}_j}\alpha^\omega}{\sum_{\omega\in\mathcal{I}_m}\alpha^\omega}$$

where the benchmark firm m, selected on observed energy intensity, need not be the most energy efficient: $\nu_m \lesseqgtr \nu_j$. To simplify exposition, suppose that $\sum_{\omega\in\mathcal{I}_m}\alpha^\omega = \alpha^{\omega^c}$: the benchmark firm only produces its core input ω^c in-house. Even keeping the allocation fixed, the energy intensity of any input $\omega' \in \mathcal{I}_j$ produced in-house by j facing τ_j becomes:

$$\sigma_{\omega'}(\nu_j, (1+\tau_j)p_e) = \frac{1-\nu_j}{(1+\tau_j)p_e} = \frac{1-\nu_m}{p_e}\frac{\alpha^{\omega^c}}{\sum_{\omega\in\mathcal{I}_j}\alpha^\omega}$$

which is not aligned across producers and no longer depends on j's own energy efficiency ν_j: the policy sets it to a fraction $\alpha^{\omega^c}/\sum_{\omega\in\mathcal{I}_j}\alpha^\omega$ of the benchmark's input-level intensity $(1-\nu_m)/p_e$, determined by firm boundaries rather than energy efficiency. In effect, firms that are currently producing more inputs

in-house (large $\mathcal{I}_j$) will be relatively more affected, for a given effective energy efficiency ν_j, inducing them to outsource more.[59] Take input $\omega' \in \mathcal{I}_j$ and suppose that the uncovered external producer i of input ω' is energy inefficient ($\nu_i < \nu_j$). Before the policy, in-house production implies that:

$$\log\left(\frac{p_e}{w}\right)(\nu_j - \nu_i) \geq \log(q_{ij})$$

With the policy in place, j outsources ω' if:

$$\underbrace{\log\left(\frac{p_e}{w}\right)(\nu_j - \nu_i) - \log(q_{ij})}_{\geq 0 \text{ before the policy}} < (1-\nu_j)\log\left(\frac{1-\nu_j}{1-\nu_m}\sum_{\omega\in\mathcal{I}_j}\frac{\alpha^\omega}{\alpha^{\omega^c}}\right)$$

This requires $\tau_j > 0$, i.e. $\sum_{\omega\in\mathcal{I}_j}\alpha^\omega/\alpha^{\omega^c} > (1-\nu_m)/(1-\nu_j)$, which holds automatically if $\nu_j < \nu_m$ and, even if $\nu_j > \nu_m$, when j's in-house cost share is large enough. We might then observe an energy inefficient reversal for input ω' from $k_j(\omega'; p_e) = j$, with embodied energy intensity $\sigma_{\omega'}(\nu_j, p_e) = \frac{1-\nu_j}{p_e}$, to $k_j(\omega'; p_e, \tau_j) = i$, with embodied energy intensity $\sigma_{\omega'}(\nu_i, p_e) = \frac{1-\nu_i}{p_e} \geq \frac{1-\nu_j}{p_e}$. The same mechanism applies to the EU ETS product benchmark - the average emissions of the most efficient firms, which are identified by their total emissions relative to total quantity produced. Installations can then gain or lose transitional free allocations by approaching the benchmark through their make-or-buy decisions rather than their energy efficiency.

A.3 Multiple-stage Model

General Equilibrium. Let $\mathcal{D} = \bigcup_{\omega\in\Omega}\mathcal{S}_\omega$ denote the set of all firms in the economy, $L_m \equiv \int_{j\in\mathcal{D}}\sum_{\hat{\omega}\in\mathcal{I}_j\cup\{\omega_j^c\}} l_j(\hat{\omega})\,dj$ manufacturing employment and $E \equiv \int_{j\in\mathcal{D}}\sum_{\hat{\omega}\in\mathcal{I}_j\cup\{\omega_j^c\}} e_j(\hat{\omega})\,dj$ total energy use.

Definition 2 (Competitive equilibrium)**.** *Given a carbon tax $\tau \geq 0$, a competitive equilibrium consists of*

- *prices: a wage w, an energy price p_e, intermediate prices $\{p_i(\omega_i^c)\}_{i\in\mathcal{D}}$, household prices $\{p_j^h(\omega_j^c)\}_{j\in\mathcal{D}}$, and price indices $\{P_\omega\}_\omega$ and P, with the numeraire $P_x \equiv 1$;*
- *a set of decisions: for each firm $j \in \mathcal{D}$ and each input $\hat{\omega}$ in the recipe $\Omega_{\omega_j^c}$ of its core product ω_j^c, a producer $k_j(\hat{\omega}) \in \mathcal{S}_{\hat{\omega}} \cup \{j\}$, which, on the inputs that j reaches from ω_j^c without passing through an outsourced input, defines the in-house and outsourced sets $\mathcal{I}_j$ and $\mathcal{O}_j$; factor demands $\{e_j(\hat{\omega}), l_j(\hat{\omega})\}_{\hat{\omega}\in\mathcal{I}_j\cup\{\omega_j^c\}}$; purchases $\{x_{ij}(\hat{\omega})\}_{\hat{\omega}\in\mathcal{O}_j}$; and output $x_j(\omega_j^c)$;*
- *household choices $\{X_c, \{y_j\}_{j\in\mathcal{D}}\}$, energy production E, numeraire production $X = L_x$, and a transfer $\tilde{T}$,*

such that:

[59] Outside of the Cobb-Douglas framework, they would also lose sales share and shrink.

1. ***Household.*** *$\{X_c, \{y_j\}\}$ maximize $X_c + Y^\alpha/\alpha$, taking E as given, subject to $P_x X_c + \int_{j\in\mathcal{D}} p_j^h y_j\, dj = wL + \tilde{T} + \Pi$, where Y aggregates $\{y_j\}$ as in Section 3.*

2. ***Sourcing.*** *For every j and $\hat{\omega} \in \Omega_{\omega_j^c}$, $k_j(\hat{\omega}) \in \arg\min_{k\in\mathcal{S}_{\hat{\omega}}\cup\{j\}} c_k(\hat{\omega})/q_{kj}$, where, for $k = j$, $c_j(\hat{\omega})$ is the in-house cost $c_j^{\mathcal{I}}(\hat{\omega})$ given by (12) and $q_{jj} = 1$; $\mathcal{I}_j$ ($\mathcal{O}_j$) contains the inputs with $k_j(\hat{\omega}) = j$ ($k_j(\hat{\omega}) \neq j$) that j reaches from ω_j^c through in-house inputs only.*

3. ***Factor demands.*** *Given sourcing, $\{e_j, l_j, x_{ij}\}$ minimize the cost of producing $x_j(\omega_j^c)$ at prices $(p_e, w, \{p_i\})$.*

4. ***Pricing.*** *$p_i(\omega_i^c) = c_i(\omega_i^c)$ for sales to firms; $p_j^h(\omega_j^c) = \frac{\varepsilon}{\varepsilon-1} c_j(\omega_j^c)$ for sales to the household; $P_x = \frac{\varepsilon}{\varepsilon-1} w = 1$; energy is produced competitively from the numeraire bought at marginal cost, so $p_e = \chi w + \tau$.*

5. ***Market clearing.***

   (a) *Goods of each firm: $x_i(\omega_i^c) = y_i + \int_{j\in\mathcal{D}:\, k_j(\omega_i^c)=i} x_{ij}(\omega_i^c)\, dj$ for all $i \in \mathcal{D}$.*

   (b) *Labor: $L_x + L_m = L$.*

   (c) *Numeraire: $X_c + \chi E = L_x$.*

   (d) *Energy: the energy produced equals the energy used by all in-house stages, E.*

6. ***Government and profits.*** *$\tilde{T} = \tau E$ and $\Pi = (P_x - w)X_c + \int_{j\in\mathcal{D}} \left(p_j^h - c_j(\omega_j^c)\right) y_j\, dj$.*

A.3.1 Equilibrium conditions

We summarize the key equilibrium conditions related to energy at the firm and at the aggregate level. First, note that the production function of a firm producing $\hat{\omega}$ in-house can be rewritten as

$$x_j(\hat{\omega}) = \psi_j(\hat{\omega})\Big(x_j^{\mathcal{I}}(\hat{\omega})^\gamma + x_j^{\mathcal{O}}(\hat{\omega})^\gamma\Big)^{\frac{1}{\gamma}}$$

Where $\psi_j(\hat{\omega}) = \psi_j$ if $\hat{\omega}$ is the core product of the firm, and 1 otherwise, and where $x_j^{\mathcal{O}}(\hat{\omega}) = \Big(\sum_{\tilde{\omega}\in\mathcal{O}_j\cap\Omega_{\hat{\omega}}} (\tilde{\alpha}_{\hat{\omega}}^{\tilde{\omega}}(j))^{1-\gamma}(q_{ij}x_{ij}(\hat{\omega}$ with $i = k_j(\tilde{\omega})$ the supplier selected for $\tilde{\omega}$, is the quality-adjusted aggregate of all inputs that are outsourced by firm j, while

$$x_j^{\mathcal{I}}(\hat{\omega}) = \left(\sum_{\tilde{\omega}\in\{\hat{\omega}\}\cup(\mathcal{I}_j\cap\Omega_{\hat{\omega}})} (\tilde{\alpha}_{\hat{\omega}}^{\tilde{\omega}}(j)\alpha_{\tilde{\omega}}^{el})^{(1-\gamma)}(\beta_{\tilde{\omega}}^{1-\rho_{\tilde{\omega}}}(\nu_j e_j(\tilde{\omega}))^{\rho_{\tilde{\omega}}} + (1-\beta_{\tilde{\omega}})^{1-\rho_{\tilde{\omega}}} l_j(\tilde{\omega})^{\rho_{\tilde{\omega}}})^{\frac{\gamma}{\rho_{\tilde{\omega}}}}\right)^{\frac{1}{\gamma}}$$

Firm j's demands for energy and labor to produce in-house input or core product $\hat{\omega} \in \mathcal{I}_j \cup \{\omega^c\}$ are:

$$\frac{e_j(\hat{\omega})}{x_j(\omega^c)} = \tilde{\alpha}_{\omega^c}^{\hat{\omega}}(j)\alpha_{\hat{\omega}}^{el}\psi_j^{\frac{\gamma}{1-\gamma}}\beta_{\hat{\omega}}\left(\frac{c_j(\omega^c)}{c_j^{el}(\hat{\omega})}\right)^{\frac{1}{1-\gamma}}\left(\frac{c_j^{el}(\hat{\omega})}{p_e/\nu_j^{\rho_{\hat{\omega}}}}\right)^{\frac{1}{1-\rho_{\hat{\omega}}}}$$

$$\frac{l_j(\hat{\omega})}{x_j(\omega^c)} = \tilde{\alpha}_{\omega^c}^{\hat{\omega}}(j)\alpha_{\hat{\omega}}^{el}\psi_j^{\frac{\gamma}{1-\gamma}}(1-\beta_{\hat{\omega}})\left(\frac{c_j(\omega^c)}{c_j^{el}(\hat{\omega})}\right)^{\frac{1}{1-\gamma}}\left(\frac{c_j^{el}(\hat{\omega})}{w}\right)^{\frac{1}{1-\rho_{\hat{\omega}}}}$$

where ω^c is the core product sold by firm j. Therefore, firm-level demands for energy and labor are:

$$\frac{e_j}{x_j(\omega^c)} = \psi_j^{\frac{\gamma}{1-\gamma}} \sum_{\hat{\omega}\in\mathcal{I}_j\cup\{\omega^c\}} \tilde{\alpha}^{\hat{\omega}}_{\omega^c}(j)\alpha^{el}_{\hat{\omega}}\beta_{\hat{\omega}} \left(\frac{c_j(\omega^c)}{c_j^{el}(\hat{\omega})}\right)^{\frac{1}{1-\gamma}} c_j^{el}(\hat{\omega})^{\frac{1}{1-\rho_{\hat{\omega}}}} \left(\frac{p_e}{\nu_j^{\rho_{\hat{\omega}}}}\right)^{\frac{-1}{1-\rho_{\hat{\omega}}}}$$

$$\frac{l_j}{x_j(\omega^c)} = \psi_j^{\frac{\gamma}{1-\gamma}} \sum_{\hat{\omega}\in\mathcal{I}_j\cup\{\omega^c\}} \tilde{\alpha}^{\hat{\omega}}_{\omega^c}(j)\alpha^{el}_{\hat{\omega}}(1-\beta_{\hat{\omega}}) \left(\frac{c_j(\omega^c)}{c_j^{el}(\hat{\omega})}\right)^{\frac{1}{1-\gamma}} c_j^{el}(\hat{\omega})^{\frac{1}{1-\rho_{\hat{\omega}}}} w^{\frac{-1}{1-\rho_{\hat{\omega}}}}$$

which give the expenditure shares in (13). For outsourced inputs $\hat{\omega} \in \mathcal{O}_j$, bought from the selected supplier $i \in \mathcal{S}_{\hat{\omega}}$, the expenditure share is

$$\frac{c_i(\hat{\omega})x_{ij}(\hat{\omega})}{c_j(\omega^c)x_j(\omega^c)} = \tilde{\alpha}^{\hat{\omega}}_{\omega^c}(j)\left(\frac{\psi_j c_j(\omega^c)}{c_i(\hat{\omega})/q_{ij}}\right)^{\frac{\gamma}{1-\gamma}}.$$

Energy and intermediate inputs interact through two opposing channels. Conditional on outsourcing a given input, they are complements: intermediate inputs are typically complementary in production ($\gamma < 0$), so higher intermediate-input expenditure shares should be associated with higher energy expenditure shares. Yet they are also substitutes at the extensive margin, because a firm can replace in-house production, and the energy it uses, by outsourcing intermediate stages, which pushes energy demand down as intermediate-input purchases rise. Which force dominates is ultimately an empirical question. The aggregate energy required to produce one unit of GDP (energy intensity) is:

$$\begin{aligned}
\frac{E}{PY} &= \frac{1}{p_e}\sum_{\omega\in\Omega}\int_{j\in\mathcal{S}_\omega} \frac{c_j(\omega)x_j(\omega)}{PY}\frac{p_e e_j}{c_j(\omega)x_j(\omega)}dj \\
&= \frac{1}{PY}\sum_{\omega\in\Omega}\int_{j\in\mathcal{S}_\omega} c_j(\omega)x_j(\omega) \sum_{\omega'\in\mathcal{I}_j\cup\{\omega\}} \tilde{\alpha}^{\omega'}_{\omega}(j)\left(\frac{\psi_j c_j(\omega)}{c_j^{el}(\omega')}\right)^{\frac{\gamma}{1-\gamma}} \alpha^{el}_{\omega'}\beta_{\omega'}c_j^{el}(\omega')^{\frac{\rho_{\omega'}}{1-\rho_{\omega'}}}\left(\frac{p_e}{\nu_j^{\rho_{\omega'}}}\right)^{\frac{-1}{1-\rho_{\omega'}}} dj \\
&= \sum_{\omega\in\Omega}\int_{j\in\mathcal{S}_\omega} \underbrace{\left(\delta_\omega\left(\frac{P}{P_\omega}\right)^{\eta-1}\left(\frac{\varepsilon-1}{\varepsilon}\right)^{\varepsilon}\left(\frac{P_\omega}{c_j(\omega)}\right)^{\varepsilon-1} + \sum_{\tilde{\omega}\in\Omega}\int_{i\in\mathcal{S}_{\tilde{\omega}}} \mathbf{1}\{j\in\mathcal{S}_i^{\mathcal{O}}\}\tilde{\alpha}^{\omega}_{\tilde{\omega}}(i)\frac{c_i(\tilde{\omega})x_i(\tilde{\omega})}{PY}\left(\frac{\psi_i c_i(\tilde{\omega})}{c_j(\omega)/q_{ji}}\right)^{\frac{\gamma}{1-\gamma}} di\right)}_{\text{Demand to firm } j} \\
&\times \underbrace{\sum_{\omega'\in\mathcal{I}_j\cup\{\omega\}} \tilde{\alpha}^{\omega'}_{\omega}(j)\left(\frac{\psi_j c_j(\omega)}{c_j^{el}(\omega')}\right)^{\frac{\gamma}{1-\gamma}} \alpha^{el}_{\omega'}\beta_{\omega'}c_j^{el}(\omega')^{\frac{\rho_{\omega'}}{1-\rho_{\omega'}}}\left(\frac{p_e}{\nu_j^{\rho_{\omega'}}}\right)^{\frac{-1}{1-\rho_{\omega'}}} dj}_{\text{Energy used by firm } j \text{ per unit of production cost}} \\
&= \underbrace{\sum_{\omega\in\Omega}\int_{j\in\mathcal{S}_\omega} \delta_\omega\left(\frac{P}{P_\omega}\right)^{\eta-1}\left(\frac{\varepsilon-1}{\varepsilon}\right)^{\varepsilon}\left(\frac{P_\omega}{c_j(\omega)}\right)^{\varepsilon-1} \sum_{\omega'\in\mathcal{I}_j\cup\{\omega\}} \tilde{\alpha}^{\omega'}_{\omega}(j)\left(\frac{\psi_j c_j(\omega)}{c_j^{el}(\omega')}\right)^{\frac{\gamma}{1-\gamma}} \alpha^{el}_{\omega'}\beta_{\omega'}c_j^{el}(\omega')^{\frac{\rho_{\omega'}}{1-\rho_{\omega'}}}\left(\frac{p_e}{\nu_j^{\rho_{\omega'}}}\right)^{\frac{-1}{1-\rho_{\omega'}}} dj}_{\text{Energy used to directly produce for the household}} \\
&+ \underbrace{\sum_{\omega\in\Omega}\int_{j\in\mathcal{S}_\omega}\sum_{\tilde{\omega}\in\Omega}\int_{i\in\mathcal{S}_{\tilde{\omega}}} \mathbf{1}\{j\in\mathcal{S}_i^{\mathcal{O}}\}\tilde{\alpha}^{\omega}_{\tilde{\omega}}(i)\frac{c_i(\tilde{\omega})x_i(\tilde{\omega})}{PY}\left(\frac{\psi_i c_i(\tilde{\omega})}{c_j(\omega)/q_{ji}}\right)^{\frac{\gamma}{1-\gamma}} di \times \sum_{\omega'\in\mathcal{I}_j\cup\{\omega\}} \tilde{\alpha}^{\omega'}_{\omega}(j)\left(\frac{\psi_j c_j(\omega)}{c_j^{el}(\omega')}\right)^{\frac{\gamma}{1-\gamma}} \alpha^{el}_{\omega'}\beta_{\omega'}c_j^{el}(\omega')^{\frac{\rho_{\omega'}}{1-\rho_{\omega'}}}\left(\frac{p_e}{\nu_j^{\rho_{\omega'}}}\right)^{\frac{-1}{1-\rho_{\omega'}}} dj}_{\text{Energy used to produce for other manufacturing firms}}
\end{aligned} \tag{24}$$

where $\mathcal{S}_i^{\mathcal{O}}$ is the set of suppliers selected by firm i for the inputs it outsources, so that $j \in \mathcal{S}_i^{\mathcal{O}} \cap \mathcal{S}_\omega$ implies $\omega \in \mathcal{O}_i$. Aggregate energy intensity ($\frac{E}{GDP}$) increases if inefficient make-or-buy decisions allocate production to low $\nu-$firms. At the firm level, both the firm boundaries ($\mathcal{I}_j$) and the supplier selection distort the final cost $c_j(\omega_j^c)$ of the firm, hence its sales, if they are inefficient. But additionally, the make-or-buy decision also affects the firm's total energy use, which is visible in the term $\sum_{\omega\in\mathcal{I}_j\cup\{\omega_j^c\}} \tilde{\alpha}^{\omega}_{\omega_j^c}(j)\alpha^{el}_{\omega}\beta_{\omega}c_j^{el}(\omega)^{\frac{\rho_\omega-\gamma}{(1-\rho_\omega)(1-\gamma)}}\left(\frac{p_e}{\nu_j^{\rho_\omega}}\right)^{\frac{-1}{1-\rho_\omega}}$.

Aggregate energy intensity responds to a uniform change in the energy price at both the intensive and the extensive margin. Keeping the extensive margin fixed, let $s_{j\omega'}$ be the share of input $\omega' \in \mathcal{I}_j \cup \{\omega_j^c\} \cup \mathcal{O}_j$

in firm j's production cost, $\sigma^e_{el,j\omega'} \equiv \frac{p_e e_j(\omega')}{c^{el}_j(\omega')x^{el}_j(\omega')}$ the energy share of the energy-labor bundle of in-house input or core product $\omega' \in \mathcal{I}_j \cup \{\omega^c_j\}$, and $\sigma^e_j \equiv \sum_{\omega' \in \mathcal{I}_j \cup \{\omega^c_j\}} s_{j\omega'}\sigma^e_{el,j\omega'} + \sum_{\omega' \in \mathcal{O}_j} s_{j\omega'}\sigma^e_{k_j(\omega')}$ firm j's supply-chain energy share, where $k_j(\omega') \in \mathcal{S}^{\mathcal{O}}_j \cap \mathcal{S}_{\omega'}$ is the supplier of ω' selected by j. Let σ^j_ω be firm j's share of household expenditure on product ω, σ^ω_{PY} product ω's share of household expenditure, $\sigma^e_\omega \equiv \int_{j \in \mathcal{S}_\omega} \sigma^j_\omega \sigma^e_j dj$, and $\sigma^e_{PY} \equiv \sum_{\omega \in \Omega} \sigma^\omega_{PY}\sigma^e_\omega = \frac{\varepsilon}{\varepsilon - 1}\frac{p_e E}{PY}$. Then:

$$\begin{aligned}
\text{Firm level:}\quad \Sigma^e_j \equiv \frac{\partial \sigma^e_j}{\partial \log p_e} = &\underbrace{\sum_{\omega' \in \mathcal{I}_j \cup \{\omega^c_j\}} s_{j\omega'} \frac{-\rho_{\omega'}}{1-\rho_{\omega'}}\left(1 - \sigma^e_{el,j\omega'}\right)\sigma^e_{el,j\omega'}}_{\text{Energy-labor substitution}} + \underbrace{\sum_{\omega' \in \mathcal{O}_j} s_{j\omega'}\Sigma^e_{k_j(\omega')}}_{\text{Suppliers' responses}} \\
&+ \underbrace{\frac{-\gamma}{1-\gamma}\Big[\sum_{\omega' \in \mathcal{I}_j \cup \{\omega^c_j\}} s_{j\omega'}\sigma^e_{el,j\omega'}\left(\sigma^e_{el,j\omega'} - \tilde{\sigma}^e_j\right) + \sum_{\omega' \in \mathcal{O}_j} s_{j\omega'}\sigma^e_{k_j(\omega')}\left(\tilde{\sigma}^e_{k_j(\omega')} - \tilde{\sigma}^e_j\right)\Big]}_{\text{Input substitution}} \qquad (25)
\end{aligned}$$

$$\text{Product level:}\quad \Sigma^e_\omega = \int_{j \in \mathcal{S}_\omega} \left[\sigma^j_\omega \Sigma^e_j + (1-\varepsilon)\sigma^j_\omega \sigma^e_j\left(\tilde{\sigma}^e_j - \tilde{\sigma}^e_\omega\right)\right] dj$$

$$\text{Aggregate level:}\quad \Sigma^e_{PY} = \sum_{\omega \in \Omega}\left[\sigma^\omega_{PY}\Sigma^e_\omega + (1-\eta)\sigma^\omega_{PY}\sigma^e_\omega\left(\tilde{\sigma}^e_\omega - \tilde{\sigma}^e_{PY}\right)\right]$$

where $\tilde{\sigma}^e_j$, $\tilde{\sigma}^e_\omega$ and $\tilde{\sigma}^e_{PY}$ are the elasticities of $c_j(\omega)$, P_ω and P with respect to p_e, and $\partial \log(E/PY)/\partial \log p_e = \Sigma^e_{PY}/\sigma^e_{PY} - 1$. By Shephard's lemma and induction over the acyclic network, $\tilde{\sigma}^e = \sigma^e$ at every level when the extensive margin is fixed. Each substitution term is then one minus the elasticity of substitution, times the share-weighted variance of energy shares within the nest, e.g. $(1-\varepsilon)\int_{j \in \mathcal{S}_\omega} \sigma^j_\omega (\sigma^e_j - \sigma^e_\omega)^2 dj$ at the product level. The extensive margin modifies these responses directly, through $\mathcal{I}_j$ and $\mathcal{O}_j$, and indirectly, through $\tilde{\sigma}^e = \sigma^e + \text{Ext}$ at every level above the energy-labor nest, where Ext is the extensive-margin response.

A.3.2 Extensive margin in the multiple-stage model

As shown in Section 2, the extensive-margin response amplifies the intensive-margin effect of an increase in energy price. In a multi-stage production network, however, a higher energy price need not reduce energy intensity through the extensive margin at every price level.[60] While the two-stage results continue to hold locally at switching points, these thresholds differ across firm–input pairs, so a uniform price increase creates intervals over which upstream reallocations may temporarily overturn the local effect. The reason is that supply-chain embedded energy efficiency can no longer be ranked by the effective energy efficiency ν alone: it depends on the energy intensity of the entire supply chain, along which reallocations occur at different prices. These could create numerical issues to solve for the optimal

[60] What remains true, however, is that, since ν is not hicks-neutral but affects inputs asymmetrically, firms' costs - on which the make-or-buy decision is based - respond to an increase in energy prices heterogeneously, with more weights being given to selecting energy efficient producer (directly, through selecting higher $\nu-$producers, or indirectly, through selecting producers with more energy efficient suppliers). The heterogeneity in responses can reverse the cost comparison, leading to a make-or-buy switch, meaning that extensive margin allocation under any price $p_e \neq p^{SCC}_e$ is weakly inefficient. Analogously to the two-stage model, a reversal requires a negative correlation: the most energy efficient producer (inclusive of its upstream production, that is, embedded) has a relatively lower hicks-neutral embedded efficiency. The former becomes more important as energy prices rise, dominating the latter and leading to a reversal.

allocation, as the algorithm could get stuck in local optimum. However, we focus on settings in which the optimal carbon price is solely determined by the environmental damage, meaning that we only have to compare the aggregate energy intensity and welfare that an economy can achieve under different energy prices without having to solve for that price. Moreover, only 1.4% of all switches move production to a less energy efficient supply chain in our optimal counterfactual. We nonetheless formalize the issue below, illustrate it with an example, and show that monotonicity is restored under additional assumptions.

In the multiple-stage model, we define the supply chain embedded energy intensity of stage $\hat{\omega}$ as:

$$\mathcal{E}_j(\hat{\omega}; p_e) \equiv \frac{\partial \log C_{\hat{\omega}k}(p_e, \nu_k)}{\partial p_e} \mathbf{1}\{k_j(\hat{\omega}; p_e) = k\}$$

When j outsources to supplier k^*, $\mathcal{E}_j(\hat{\omega}; p_e) = \mathcal{E}_{k^*}(\hat{\omega}; p_e)$: the buyer inherits the supplier's embedded energy intensity. Therefore, this energy intensity measure corresponds to the total supply chain energy used to produce stage $\hat{\omega}$ for firm j by its chosen producer k, relative to the value of $\hat{\omega}$ produced. We now show that the extensive margin effect of a change in p_e does not necessarily reduce all supply chain embedded energy intensities individually.

Lemma 5. *Conditional on a switch, selected producers have weakly lower embedded energy intensity locally.*

Proof. Fix buyer j and stage ω. Let $p'_e < p''_e$ and suppose that the buyer switches from producer i at p'_e to producer s at p''_e, i.e.,

$$k_j(\omega; p'_e) = i, \qquad k_j(\omega; p''_e) = s, \qquad s \neq i,$$

and that $p_e \mapsto \log C_i(p_e)$ and $p_e \mapsto \log C_s(p_e)$ are differentiable at p''_e.[61] Define $\Delta_{si,j}(p_e) \equiv \log C_s(p_e) - \log C_i(p_e) - \log \frac{q_{sj}}{q_{ij}}$. Since q_{kj} is independent of p_e:

$$\frac{\partial \Delta_{si,j}}{\partial p_e} = \mathcal{E}_s(\hat{\omega}; p_e) - \mathcal{E}_i(\hat{\omega}; p_e) \tag{26}$$

By optimality at p''_e, $\Delta_{si,j}(p''_e) \leq 0$; by optimality at p'_e, $\Delta_{si,j}(p'_e) \geq 0$. Since $\Delta_{si,j}$ is continuous,[62] the intermediate value theorem gives at least one $p^*_e \in [p'_e, p''_e]$ with $\Delta_{si,j}(p^*_e) = 0$. Since $\Delta_{si,j}$ transitions from nonnegative to nonpositive at p^*_e, it is locally non-increasing there:[63]

$$\left.\frac{\partial \Delta_{si,j}}{\partial \log p_e}\right|_{p_e = p^*_e} \leq 0,$$

which by (26) gives $\mathcal{E}_s(\hat{\omega}; p_e) \leq \mathcal{E}_i(\hat{\omega}; p_e)$. □

Thus, at the switching point p^*_e, the newly selected producer s has weakly lower embedded energy intensity than the displaced producer i. We are however interested in the ordering at p''_e, not at p^*_e. The following condition is sufficient for the newly selected producer s to have a weakly lower embedded energy intensity than the displaced producer i at p''_e:

[61] If differentiability fails because C_i or C_s has a kink at p''_e, the conclusion holds with one-sided derivatives.

[62] $C_j(p_e)$ is continuous by induction on supply chain depth. At leaf nodes $C_j = c_j^{\mathcal{I}}$ is smooth. For a depth-d firm, $C_j(p_e) = \min_k\{C_k(p_e)/q_{kj}\}$ is the pointwise minimum of continuous functions, hence continuous. When j switches from k to k', it does so because $C_k/q_{kj} = C_{k'}/q_{k'j}$: the two options deliver equal cost at the switching point, so C_j transitions without a jump. Discontinuities arise only in $\sigma_j(p_e) = \partial C_j/\partial p_e$.

[63] With one-sided derivatives if $\Delta_{si,j}$ is not differentiable at p^*_e.

Condition 1 (Monotonic ranking of cost responses)**.** *Conditional on a switch from i to s,* $\log C_s(p_e)/C_i(p_e)$ *is weakly concave in* p_e *for all* $p_e \in [p_e^*, p_e'']$*. Equivalently,* $\mathcal{E}_s(\hat{\omega}; p_e^*) \leq \mathcal{E}_i(\hat{\omega}; p_e^*)$ *implies* $\mathcal{E}_s(\hat{\omega}; p_e) \leq \mathcal{E}_i(\hat{\omega}; p_e)$ *for all* $p_e \in [p_e^*, p_e'']$*.*

Condition 1 holds trivially if the price change is sufficiently small, such that $p_e'' \approx p_e^*$. However, for larger changes, it can be violated when upstream reallocation by the displaced producer i alters the relative cost slopes after the switch. Specifically, suppose that at some $\tilde{p}_e \in (p_e^*, p_e'']$, producer i switches its own input $\hat{\omega}$ from an energy-inefficient supplier to an energy-efficient one k with $\mathcal{E}_k < \mathcal{E}_{k'}$, reducing $\mathcal{E}_i$ below $\mathcal{E}_s$ while leaving the cost ranking unaffected. Such a violation is always transient. Define $R(p_e) \equiv C_k(\hat{\omega})/q_{ki}/C_{k'}(\hat{\omega})/q_{k'i}$; as long as the violation persists, $\sigma_k < \sigma_{k'}$, so

$$\frac{d \log R}{d \log p_e} = \mathcal{E}_k - \mathcal{E}_{k'} < 0,$$

and R is strictly decreasing. In a finite economy, the gap $\mathcal{E}_{k'} - \mathcal{E}_k$ is bounded away from zero, so R must reach 1 at some finite price, at which point either i switches back to k' or k itself reallocates, restoring $\mathcal{E}_s \leq \mathcal{E}_i$. However, the violation can persist over the entire interval $[p_e^*, p_e'']$ before self-correcting, because j's choice is determined by a *level* comparison - C_s/q_{sj} against C_i/q_{ij} - not a slope comparison. Since pairwise switching thresholds p_e^* differ across buyer–supplier pairs, no single price is guaranteed to lie outside all violation intervals. Additional assumptions on the joint distribution of the primitives are required to prevent these violations. In particular, condition 1 requires that the displaced producer i does not have systematically better outsourcing wedges with more energy efficient upstream suppliers. We discuss a couple of assumptions that, combined with decreasing differences in ν, are sufficient to rule this case out. Formal proofs are available upon request.

1. *Independence: q_{ij} is iid with mean* 1*, independent of ν.* Under iid match quality, ν_j governs only the in-house option, so two firms with different ν face the same distribution of outsourcing costs. The firm with higher ν is more selective about outsourcing and, since cost and energy intensity are both decreasing in ν, greater selectivity delivers lower expected energy intensity. The slope ordering $\mathcal{E}_s \leq \mathcal{E}_i$ holds in expectation.

2. *Log-supermodularity and level condition: for $\nu_k \geq \nu_{k'}$ and $\nu_j \geq \nu_{j'}$, $q_{kj}/q_{k'j} \geq q_{kj'}/q_{k'j'}$, and $q_{kj} \geq 1$ whenever $\nu_k \geq \nu_j$.* This rules out the case in which the less efficient firm i has sufficiently better bilateral match quality with a highly efficient upstream supplier that it accesses k while s does not: among suppliers with $\nu_k \geq \nu_s$, firm s always has a comparative advantage, so s selects an upstream option with ν at least as high as i's.

These distributional assumptions are sufficient[64] to guarantee that (on average) the reallocative effect of the make-or-buy decision decreases energy intensity. They are fairly restrictive, and don't apply in the data. However, we solely use them for illustrative purposes. The non-monotonicity that they exclude implies that there might be multiple local equilibria that prevent convergence in a numerical resolution of

[64] There exist other sufficient assumptions, such as common outsourcing wedges within an entire supply chain.

the optimal allocation. It also means that aggregate energy intensity might be highly non-linear in energy prices. As shown in the example below, however, those violations should remain rare under reasonable parametrizations, and have not affected our aggregate results in any of our quantitative exercises.

Example of a violation. Suppose firm j produces ω using one input $\hat{\omega}$, with the only external supplier of ω being i and the only external supplier of $\hat{\omega}$ being k. Production is Cobb-Douglas with energy efficiency ν:

$$x(\omega) = \frac{1}{A(\nu)} \left(e(\omega)^{1-\nu} l(\omega)^{\nu}\right)^{\alpha_{\omega}^{el}} x(\hat{\omega})^{\alpha_{\omega}^{\hat{\omega}}},$$

with $\nu_k < \nu_j < \nu_i$ (so k is most energy efficient). Let $q_{ij} > 1$ and $q_{ki} < 1$. At p_e^*, firm j switches from outsourcing ω to i to producing in-house; at this point i is producing $\hat{\omega}$ in-house. The local lemma is satisfied: $(1-\nu_j)/p_e^* \leq (1-\nu_i)/p_e^*$. At $\tilde{p}_e \in (p_e^*, p_e'']$, however, i switches $\hat{\omega}$ to k, giving i an embedded energy intensity of $(\alpha_{\omega}^{el}(1-\nu_i) + \alpha_{\omega}^{\hat{\omega}}(1-\nu_k))/\tilde{p}_e$, which can fall below j's in-house intensity $(1-\nu_j)/\tilde{p}_e$ if $\alpha_{\omega}^{el}(1-\nu_i) + \alpha_{\omega}^{\hat{\omega}}(1-\nu_k) < 1-\nu_j$. For the violation to persist through p_e'', j must not switch back to i or forward to k, which requires:

$$q_{ij} \leq \frac{1}{q_{ki}^{\alpha_{\omega}^{\hat{\omega}}}}.$$

Since $q_{ij} > 1 > q_{ki}$, a violation requires q_{ij} not too large and q_{ki} sufficiently small - a restrictive condition. In particular, if outsourcing wedges are positively correlated within a supply chain, this condition cannot hold and violations are ruled out.

A.3.3 Sufficient statistic for extensive margin misallocation

In the two-stage model with one supplier per input, the extent of misallocation induced by the make-or-buy decision is governed by the joint distribution of the q-gap $\Delta \log q \equiv \log(q_{sj}) - \log(q_{ij})$ and the effective energy efficiency of both suppliers (summarized by e.g., the ν-gap $\Delta\nu \equiv \nu_i - \nu_s$ in the Cobb-Douglas case or by the ratio $\frac{\nu_i}{\nu_s}$ in the CES case). The make-or-buy allocation is distorted for all pairs for which $\Delta \log q$ lies between the equilibrium cost cutoff and the optimal one. In a multi-stage network, this two-dimensional sufficient statistic no longer characterizes misallocation, because the relevant cost cutoff is now the *embedded* cost $\log(C_{s\omega}) - \log(C_{i\omega})$, which depends on the entire upstream allocation.

To simplify the comparison with the two-stage model, we again suppose that there is only one external supplier per input. The producer s is selected over the producer i if:

$$\Delta log q = log q_{sj} - \log(q_{ij}) \geq \log(C_{s\omega}(p_e)) - \log(C_{i\omega}(p_e))$$

Upstream make-or-buy adjustments imply that the cost cutoff responds to energy prices through both an intensive and an extensive margin:

$$\frac{\partial}{\partial p_e}\left[\log(C_{s\omega}(p_e)) - \log(C_{i\omega}(p_e))\right] = \mathcal{E}_s(\omega; p_e) - \mathcal{E}_i(\omega; p_e),$$

It is straightforward to show that a switch from producer s to producer i at p_e^* implies that $\mathcal{E}_s(\omega; p_e^*) - \mathcal{E}_i(\omega; p_e^*) > 0$.

We have that:

$$\mathcal{E}_s(\omega, p_e) \equiv \frac{\partial \log(C_{s\omega}(p_e))}{\partial p_e} \approx \underbrace{\sum_{\hat{\omega}\in\Omega_\omega} \frac{\partial \log(C_{s\omega}(p_e))}{\partial \log(C_{k\hat{\omega}}(p_e))} \frac{\partial \log(C_{k\hat{\omega}}(p_e))}{\partial p_e} \mathbf{1}\{k_s(\hat{\omega}; p'_e) = k_s(\hat{\omega}; p_e) = k\}}_{\text{Intensive margin}}$$
$$+ \underbrace{\sum_{\hat{\omega}\in\Omega_\omega} \frac{\partial \log(C_{s\hat{\omega}}(p_e))}{\partial \log(C_{m\hat{\omega}}(p_e))} (C_{k\hat{\omega}}(p'_e) - C_{m\hat{\omega}}(p_e)) \mathbf{1}\{k_s(\hat{\omega}; p'_e) = k \neq k_s(\hat{\omega}; p_e) = m\}}_{\text{Extensive margin}}$$

The change in total costs is a weighted sum of the change in each upstream input cost, with weights corresponding to the input share. It can be decomposed into intensive margin changes, holding the upstream allocation fixed, and extensive margin adjustments, reflecting changes in upstream make-or-buy decisions. The approximation comes from using the derivative as the weights for the extensive margin adjustments as well. The intensive margin is a weighted sum of the effective energy intensities of all active nodes along s's supply chain, with weights equal to the corresponding input requirements:

$$\frac{\partial \log(C_{s\omega}(p_e))}{\partial p_e}\Big|_{\text{Int. margin}} = \sum_{\hat{\omega}\in\mathcal{I}_s(p_e)\cap\mathcal{I}_s(p'_e)} \frac{\partial \log(C_{s\omega}(p_e))}{\partial \log(c^{el}_{s\hat{\omega}}(p_e))} \frac{\partial \log(c^{el}_{s\hat{\omega}}(p_e))}{\partial p_e}$$
$$+ \sum_{\hat{\omega}\in\mathcal{O}_s(p_e)\cap\mathcal{O}_s(p'_e)} \frac{\partial \log(C_{s\omega}(p_e))}{\partial \log(C_{k\hat{\omega}}(p_e))} \underbrace{\frac{\partial \log(C_{k\hat{\omega}}(p_e))}{\partial p_e}}_{=\mathcal{E}_k(\hat{\omega}, p_e)}$$

Where the intersections $\mathcal{I}_s(p_e) \cap \mathcal{I}_s(p'_e)$ and $\mathcal{O}_s(p_e) \cap \mathcal{O}_s(p'_e)$ restrict to stages whose sourcing decision is unchanged between p'_e and p_e - produced in-house at both prices, or outsourced to the same supplier at both prices, respectively. Without make-or-buy reallocation, this set remains fixed, and $\sum_{\hat{\omega}\in\mathcal{I}_j(p_e)} \frac{\partial \log(C_{s\omega}(p_e))}{\partial \log(c^{el}_{s\hat{\omega}}(p_e))} \frac{\partial \log(c^{el}_{s\hat{\omega}}(p_e))}{\partial p_e}$ is the empirical firm-level energy intensity. The change in cost then depends on both the firm-level energy intensity and on embedded outsourced input energy intensity, weighted by their input share, that is, the energy intensity of the entire supply chain up to that input, which we term as the embedded energy intensity. As an illustration, we consider Cobb-Douglas aggregators ($\rho = \gamma = 1$). The intensive margin simplifies to:

$$\mathcal{E}_s(\omega; p_e)\Big|_{\text{Int. margin}} = \sum_{\hat{\omega}\in\mathcal{I}_s(p_e)\cap\mathcal{I}_s(p'_e)} \tilde{\alpha}^{\hat{\omega}}_{\omega}(j) \frac{1-\nu_j}{p_e} + \sum_{\hat{\omega}\in\mathcal{O}_s(p_e)\cap\mathcal{O}_s(p'_e)} \tilde{\alpha}^{\hat{\omega}}_{\omega}(j) \mathcal{E}_{k_s(\hat{\omega};p_e)}(\hat{\omega}; p_e)$$

where in-house stages contribute $(1-\nu_j)/p_e$ and outsourced stages inherit the supplier's embedded intensity, recursively. The relevant primitive is therefore the joint distribution of the q-gap and the *embedded energy intensity gap* $\mathcal{E}_s - \mathcal{E}_i$, which aggregates all upstream ν's weighted by input requirements. The same firm-level ν can enter this expression multiple times when a firm produces several stages in-house, giving more weight to firms that are more deeply integrated.

Moving away from Cobb-Douglas, the intermediate input share $\frac{\partial \log(C_{s\omega}(p_e))}{\partial \log(C_{k\hat{\omega}}(p_e))}$ depends on the outsourcing capabilities of the outsourced inputs (e.g., in the CES case, $\frac{\partial \log(C_{s\omega}(p_e))}{\partial \log(C_{k\hat{\omega}}(p_e))} = \tilde{\alpha}^{\hat{\omega}}_{\omega}(j) \left(\frac{c_{k\hat{\omega}}}{c_{s\omega}}\right)^{\frac{-\gamma}{1-\gamma}} q_{ks}^{\frac{\gamma}{1-\gamma}}$). The relevant joint distribution extends to include the upstream outsourcing capabilities of all *observed* pairs.

Additionally, when allowing for upstream make-or-buy reallocation, even this joint distribution is no longer sufficient: $\mathcal{I}_j(p_e)$ and $\mathcal{O}_j(p_e)$ themselves respond to p_e endogenously as each producer j adjusts its own upstream sourcing. Thus, the effective energy efficiency ν and the outsourcing capabilities q_{ij} of all *potential* pairs become relevant for the changes in the cost cutoff.

B Empirical Appendix

B.1 Data

To evaluate the aggregate energy intensity of the manufacturing sector relative to its socially efficient benchmark, accounting for cross-producer misallocation, we need to (i) identify which firm is *actually* producing each input, which requires identifying the set of inputs produced in-house by each firm and (ii) identify which firm *should* optimally be producing each input, which requires characterizing each potential producer by their efficiency. To do so, we need firm-to-firm transaction data that uniquely identifies both the buyer and the supplier firm. This information is still difficult to obtain, as the number of countries with firm-to-firm transaction data is limited. The Italian economy is well suited to this analysis, because data on the universe of domestic firm-to-firm transactions, that can be combined with detailed customs data have recently been made available. Moreover, the Italian manufacturing sector is characterized by its pronounced fragmentation (Agostino et al., 2016):[65] more than 82% of Italian manufacturing companies employ fewer than 10 workers, and large manufacturing firms of more than 250 employees constitute only 0.3% of all firms. Italian manufacturing firms' integration in global value chains (GVC) has also recently surged, reaching levels similar to French or German firms (Amador & Cabral, 2015; Cappariello & Felettigh, 2015). Finally, Italy is the second-largest manufacturing country and third CO_2 emitter in Europe (IEA).

We combine three datasets that describe the Italian production network in detail: i) VAT data ii) customs data and iii) balance sheet data. The first dataset, fully described by Astuti and Piras (2026), contains the universe of domestic business-to-business (B2B) transactions in 2019 and 2021-2023.[66] It is constructed based on firms' VAT declarations. Since 1 January 2019, e-invoicing is mandatory for private entities resident or established in Italy; electronic invoices are routed through the Revenue Agency's Exchange System, yielding standardized digital records. We observe the total yearly value (euros) for each bilateral business relationship. Crucially, firms are uniquely identified by tax identifiers (VAT numbers/fiscal codes)[67] that are common to the other sources, enabling unambiguous merges (and VAT-

[65] The success of Italian SMEs has been related to regional "clusters" of specialized firms that fragment the production process (Bianchi et al., 1997).

[66] 2019 does not appear to be an outlying year: we rebuilt all our descriptive statistics using the 2023 cross-section and find extremely similar results.

[67] Firm identifiers are harmonized across sources, with targeted corrections for sole proprietorships and for entities belonging to VAT groups. A VAT group—comprising entities that are financially, economically, and organizationally linked—is treated as a single taxable person; intra-group transactions are non-VATable and the group files a single return. This is consistent with our model, and we therefore consolidate all group members and record transactions at the group

group consolidation where relevant). Each firm is assigned the 6-digit ATECO code of its main activity, which identifies the input it sells, after we refine this classification for multi-product sellers (Appendix B.3.1).[68]

We merge the above with customs data prepared by the Italian National Institute of Statistics (ISTAT) warehouse for international trade statistics, based on the microdata collected by the Italian Customs and Monopolies Agency (CMA). It reports information on transactions involving both extra-EU and intra-EU suppliers. While the former are recorded at the single-transaction level, the latter are aggregated by months or quarters. We aggregate all buyer-input transactions to the annual level. Products are classified at CN 8-digit and mapped to CPA 6-digit using the Bank of Italy concordance. We then construct a mapping of the 6-digit CPA code into ATECO 6-digit codes.

Finally, we retrieve firm-level information from balance sheet data, which covers the universe of Italian limited liability companies. The dataset is provided by CERVED through the registry of the Italian Chamber of Commerce, and has been used extensively to study Italian companies (e.g., Calligaris et al. (2018) and Akcigit et al. (2023)). It contains information on the firms' location, value-added, turnover, total and fixed assets, employment, and intermediate inputs expenditures, entry date, and their main economic activity at the ATECO 6−digit level. Employment is harmonized across sources and, where missing, imputed from labor-cost data. For the subsample that we use to calibrate the model, we drop observations with non-positive or missing labor costs, turnover and costs of goods sold.

We need a measure of firms' total energy demand, which is not directly reported in the balance sheet data. However, we can directly observe firms' energy purchases in the B2B and customs data. Therefore, following the literature (e.g, Stangl et al. (2025)), we define a firm's total energy expenditures as the sum of all expenditures to firms classified as generating or selling energy (see Appendix B.2 for details). In line with our model, we focus on intermediate inputs, and remove transactions on intangible capital, physical capital such as long-lasting machines or equipment, buildings and rent payments, construction, wholesale and retail, and services. We count these transactions, on the seller's side, as sales to the rest of the economy.

We only retain private, domestic manufacturing companies in the sample we analyze. Moreover, for our final analysis, we exclude firms classified as leaf inputs-providers, to focus on firms processing inputs.[69] Appendix B.3 provides more details on how we prepare the data, and Appendix B.4 reports descriptive statistics on our final sample. Our final sample contains 153424 manufacturing firms - among which 145498 report transactions with other manufacturing firms - and 2928934 transactions. Table B.1 reports descriptive statistics on the structure of the production network of the Italian manufacturing sector. Akin

level.

[68] The ATECO Italian classification code, or Classificazione delle Attività Economiche (Classification of Economic Activities), is the national standard for classifying businesses by their economic activity in Italy. Developed by ISTAT, it is directly derived from the European Union's NACE system for its first four digit, and further refined with two additional digits to reflect national specificity. We also use the Firm Registry by Infocamere to verify our firms' classification and fill missing identifiers.

[69] We still keep those upstream firms, as well as foreign firms, to build the supply chains and the costs faced by our manufacturing firms, but we exclude them in the final analysis.

to other production networks studied using similar datasets (Dhyne et al., 2021; Borin et al., 2025), the distribution of the number of buyers per seller and the number of suppliers per buyer both exhibit substantial heterogeneity and point to a highly skewed network structure, with a minority of firms concentrating a large share of transactions. The last two rows further indicate that buyers typically rely on only one supplier per input.

Table B.1: Network statistics: suppliers and buyers (2019)

Statistic	Mean	Std. dev.	Min	P25	P50	P75	Max
Number of domestic firms per product	386	780	1	34	40	451	10393
Number of buyers per domestic seller	32	105	1	3	9	28	11289
Number of suppliers per domestic buyer	30	50	1	5	15	37	2270
Number of suppliers per buyer × input	1.84	2.48	1	1	1	2	346
Number of domestic suppliers per buyer × input	1.78	2.33	1	1	1	2	346

B.2 Energy expenditures

Following the literature (e.g., Stangl et al. (2025)), we define a firm's total energy expenditures as the sum of all its purchases from firms classified as generating or selling energy. The corresponding ATECO codes are listed in Table B.2. We also manually classify any transaction with the main Italian energy retailers, Eni, Edison, Enel and Shell, as energy transactions, as some of those were identified as retailers in the dataset. For our manufacturing firms, which exclude firms that process and resell energy, these purchases correspond to the actual fuel used to power their production. We validate our measure of energy expenditures using the Indagine sulle imprese industriali e dei servizi (Invind) representative survey ad-hoc energy sections included in 2022 and 2023, and the Cassa per i Servizi Energetici e Ambientali (CSEA) data.

This measure ignores in-house energy generation. To evaluate its prevalence, we look at purchases from firms that install or maintain energy-generating equipment (Table B.3): fewer than 4% of firms in our final sample make such purchases.[70]

B.2.1 Energy prices

We build total energy prices as the harmonic mean of fuel prices, weighted by the different fuel expenditure shares in the firms' energy mix. We retrieve fuel prices from different datasets. First, for firms importing a fuel, we directly use the importing price computed from imported values and quantities reported in the customs data. Second, for a subsample of firms, we observe electricity and gas prices directly in the Bank of Italy's Invind survey (Alpino et al., 2024). For the remaining firms, we follow the methodology of De Lyon and Dechezleprêtre (2025) and model fuel prices as a log-linear function

[70] This is an upper bound, because 42.22.00 includes telecommunications projects unrelated to energy generation.

Table B.2: List of ATECO 2007 codes of activities whose output is treated as an energy input

ATECO Code	Activity
05.10.00	Mining of hard coal
05.20.00	Mining of lignite
06.10.00	Extraction of crude petroleum
06.20.00	Extraction of natural gas
08.92.00	Peat extraction
19.10.01	Manufacture of pitch and pitch coke
19.10.09	Manufacture of other coke oven products
19.20.10	Petroleum refining
19.20.20	Manufacture of petroleum derivatives
19.20.30	Blending of LPG (Liquefied Petroleum Gases) and bottling
19.20.90	Manufacture of other refined petroleum products and fossil fuel products
24.46.00	Processing of nuclear fuel
35.11.00	Production of electricity
35.12.00	Transmission of electricity
35.13.00	Distribution of electricity
35.14.00	Trade of electricity
35.21.00	Manufacture of gas
35.22.00	Distribution of gaseous fuels through mains
35.23.00	Trade of gas through mains
35.30.00	Steam and air conditioning supply
46.12.01	Activities of agents involved in the wholesale of liquid and gaseous fuels
46.12.02	Activities of agents involved in the wholesale of solid fuels
46.71.00	Wholesale trade of petroleum products and lubricants for motor vehicles, heating fuels
47.30.00	Retail sale of automotive fuel
47.78.40	Retail sale of fuel for domestic use and heating
49.50.10	Transport of gases via pipeline

Table B.3: ATECO 2007 codes of activities installing or maintaining energy-generating equipment

Code	Activity
33.12.53	Repair and maintenance of distilling or rectifying plant for petroleum refineries and chemical industries
33.20.01	Installation of electric motors, generators, transformers and electricity distribution and control apparatus
33.20.05	Installation of steam or other generating boilers
42.22.00	Construction of utility projects for electricity and telecommunications

of quantity consumed to account for the fact that large energy consumers tend to negotiate preferential rates. This allows us to infer fuel prices from observed expenditures and the elasticity of fuel prices to quantity, that we estimate using a detailed energy survey covering a representative sample of French manufacturing firms (EACEI).

B.3 Mapping the data to the model

We take the model to the manufacturing data by identifying each input ω with a six-digit ATECO code (2007 version). A finer level of disaggregation would risk treating perfectly substitutable varieties as distinct inputs in the recipes, which we model as sets of crucial inputs that are necessary for production. The cost is that all varieties within a six-digit code are assumed to share the same recipe and the same energy requirement. We assume that each firm's activity code corresponds to the *single* product it sells (see Appendix B.3.1 for a discussion on how we deal with multi-product firms).[71] For approximately 10% of our manufacturing firms, our data only report a more aggregated code. This seems to be random and due to misreporting. For most of them, this broader category still corresponds to only one 6−digit code, that we assign to those firms. For the remaining, we assign them randomly to a six-digit code within their category, with frequency weights based on the transaction shares of those six-digit codes within these broader categories.

Input space. We focus on intermediate inputs, and remove transactions on intangible capital, physical capital such as long-lasting machines or equipments, buildings and rent payments, construction, wholesale and retail, and services. Wholesale and retail firms could act as intermediaries (Manova et al., 2025). Therefore, we verify that our recipes are robust to their inclusion in a robustness exercise described in Appendix B.13.3.

We define the set $\mathcal{R}$ as leaf nodes, i.e., products for which we don't model the production explicitly and that cannot be produced in-house. These inputs are conceptually closer to raw materials or extracted resources than to intermediate inputs that must be processed.[72] They are also likely to be more easily substitutable in the recipes.[73] We treat their cost as exogenous and fix it at the average cost across suppliers of that input in all counterfactuals, and set their outsourcing capabilities to one ($q_{ij} = 1\ \forall i, j$). To eliminate discrepancies arising from heterogeneous sourcing of leaf inputs - which are not the focus

[71] The ATECO classification is based on the NACE classification, and similarly organizes economic activity into 615 detailed four-digit industries, further subdivided into 1224 six-digit subcategories of product. To give a sense of the ATECO classification's granularity, consider the NACE sector 25.93 ("Manufacture of wire products, chain and springs"). The ATECO classification further splits it into the three six-digit categories 25.93.10 ("Manufacture of wire products"), 25.93.20 ("Manufacture of chain") and 25.93.30 ("Manufacture of springs"). Similarly, the ATECO classification refines several NACE sectors classified as "other" or "residuals". For instance, take the NACE sector 20.59 ("Manufacture of other chemical products n.e.c."). It comprises products 20.59.10 ("Manufacture of chemicals for photographic use"), 20.59.20 ("Manufacture of organic chemicals obtained from basic products derived from fermentation processes or vegetable raw materials "), 20.59.30 ("Chemical treatment of fatty acids "), 20.59.40 ("Manufacture of various chemicals for industrial use"), 20.59.50 ("Manufacture of chemicals for office use and non-industrial consumption"), 20.59.60 ("Manufacture of auxiliary products for the textile and leather industries"), 20.59.70 ("Manufacture of electrochemical products (excluding the production of chlorine, soda, and potash) and electrothermal products") and finally 20.59.90 ("Manufacture of all other various chemical products n.e.c.").

[72] We could not find evidence of any Italian manufacturing firm that has integrated the extraction of its primary material. The industry largely relies on buying and importing. Excluding primary materials from the in-house set should therefore not bias any of our quantitative results.

[73] For instance, aluminum can be produced with either bauxite or aluminum scrap. If a firm is purchasing the former but not the latter, it would be wrong to infer that the latter is produced in-house.

of our analysis - we assume that all firms requiring a given leaf input purchase it at this fixed price, regardless of whether a corresponding transaction is observed in the data. Leaf inputs are also never sold to the household, so that household expenditure PY only covers manufacturing products. This set includes primary material excluding all energy inputs, since energy expenditures are already accounted for separately when we build firms' cost. Following Boehm and Oberfield (2023), we also include in $\mathcal{R}$ products classified as "residual" ("not elsewhere classified", "others" etc), or whose definition is very broad.[74] Firms producing these products likely have different input mixes because they produce different goods, rather than because they adopted different organizational forms.

Foreign firms. We don't have a firm identifier for foreign firms. However, they remain potential suppliers. We aggregate foreign firms into a single seller per input×country. We treat foreign firms and the firms for which we don't observe energy or labor expenditures alike as a sort of exogenous competitive fringe. Under each price scenario, they remain potential producers, but their costs are exogenously set at the average cost of their product among our sample of domestic firms. This means that our counterfactual increase in energy price should be understood as either driven by domestic policies that also implement a corresponding carbon-border adjustment mechanism (CBAM) targeting imports, or an exogenous increase in energy price coming from foreign markets, that affects domestic and foreign firms alike.

Single supplier. The model assumes that each firm has only one supplier per input. Most buyer-input pairs have a single supplier, and when they have several, the top supplier accounts on average for 72% of their purchases of that input (Table B.6).[75] For buyer-input pairs with multiple suppliers, we keep the supplier with the highest transaction value.

Non-manufacturing sector. We measure PY, the manufacturing firms' sales to the non-manufacturing sector, as firms' total sales net of their exports and of their sales of intermediate inputs to manufacturing buyers. PY therefore includes sales to the public sector and to the non-manufacturing sectors, which we do not explicitly model in the production network, as well as the transactions excluded from the input space.

[74] In particular, these include: Manufacture of other various food products n.e.c., Manufacture of felt and other various textile articles n.e.c., Manufacture of other products of wood n.e.c., Manufacture of all other various chemical products n.e.c., Manufacture of other various plastic products n.e.c., Manufacture of other non-metallic mineral products n.e.c., Manufacture of other various fabricated metal products n.e.c., Manufacture of other electrical equipment n.e.c., Manufacture of other various general-purpose machinery n.e.c., Manufacture of other agricultural and forestry machinery n.e.c., Manufacture of other machine tools (including parts and accessories) n.e.c, Manufacture of machinery for the food, beverage, and tobacco industries (including parts and accessories), Manufacture of all other various special-purpose machinery n.e.c. (including parts and accessories) , Manufacture of other parts and accessories for motor vehicles, and Manufacture of other various articles n.e.c.

[75] The literature has also consistently shown that firms are largely undiversified, that is, they have very few suppliers - usually only one - for each of their inputs (e.g., Antras et al. (2017) and Alfaro et al. (2025)). Diversification has even declined over the years. In particular, Alfaro et al. (2025) find that firms drop their supplier before switching to a new one, instead of diversifying, when hit by a tariff shock.

B.3.1 Multi-product firms

At each step of our calibration, we must deal with multi-product firms. We exclude multi-product firms when building our recipes, to avoid including inputs that actually enter the composition of other products produced by those firms. We also evaluate how much the recipes of the different products of multi-product firms overlap. Moreover, we build the in-house set of multi-product firms separately for each of their products, considering each recipe separately in the comparison with the firms' set of purchased inputs.[76] We then combine these sets to create the firms' total in-house set $\mathcal{I}_j^{full} = \bigcup_{\omega \in \mathcal{P}_j} \mathcal{I}_j(\omega)$, where $\mathcal{P}_j$ is the set of products sold by firm j, which allows us to identify their effective energy efficiency ν_j based on the firm-level energy and labor expenditures.[77] Indeed, to be consistent with our model, we assume that the effective energy efficiency is a fixed characteristic of the firms, common across all inputs and products. In the B2B data, we treat the transactions with multi-product suppliers as follows. We assign them as applying to an input that enters the recipe of the buyer's product. If this corresponds to more than one input among the multiple inputs sold by the supplier,[78] we pick the input that has the highest within-supplier product sale share.[79] Finally, in the counterfactual we define multi-product firms as potential suppliers of all their products (for products with a within-firm product sale share s_ω larger than 5%). We construct the unit cost of each product of a multi-product firm recursively based on the outsourced and in-house inputs that enter the recipe of that product, exactly as we do for the unit cost of single-product firms.[80] Consistent with our model, the effective energy efficiency ν, the outsourcing capabilities[81] and the make-or-buy decision of each input remain the same within a buyer firm, regardless of the product for which this input is used. Since costs are built recursively, and only depend on upstream inputs, the make-or-buy decision and our calibration strategy for the q_{ij}[82] is indeed common across products, within

[76] The assumption that multi-product and single-product firms require the same product recipe for a given product is standard in the literature (e.g., Boehm and Oberfield (2020)), mainly due to the lack of data on which input is used to produce which product in the multi-product firms portfolio.

[77] This aggregation assumes that the different products are aggregated through a CES with elasticity of substitution $\frac{1}{1-\gamma}$ exactly like the input-nest.

[78] This applies to fewer than 40% of transactions involving a multi-product firm: the average (median) number of candidate inputs is 1.60 (1.00), while the 75^{th} percentile is 2.

[79] The implication is that if we wrongly assign input A instead of input B to a transaction, we will infer that B is produced in-house. This biases the estimation of the effective energy efficiency of the buyer, since we control for the in-house input set, which is built by weighting each in-house input by its energy requirement β_ω. The bias thus depends on how strongly A and B differ in their energy requirements.

[80] This treatment allows for some of the multi-product firms' products to also enter the recipes of their other products, meaning that they will face make-or-buy decisions and produce them or outsource them in the equilibrium and counterfactual scenarios, just as any other input. This case remains rare in our data.

[81] In contrast, we make the outsourcing capabilities buyer×supplier×input traded-specific, so they can vary within a multi-product supplier. They are described as buyer×supplier specific in the model. However, whenever we infer or refine the unobserved q_{ij} from observed decisions (type (b)), we use cutoffs that are buyer×input-specific. We therefore key q by $(j, s, \hat{\omega})$ triples rather than (j, s) pairs. This is equivalent for single-product suppliers, but not for multi-product suppliers. Concretely, for observed transaction triples $(j, s, \hat{\omega}) \in \mathcal{T}$, we use the FOC-calibrated $q_{sj,\hat{\omega}}$ directly. For multi-product firm s's triples $(j, s, \hat{\omega})$ that fall into type (b), we change the outsourcing capability to align with the input-specific cutoff, which is determined by the buyer×input-specific option that was selected by buyer j for input $\hat{\omega}$.

[82] To calibrate multi-product firms' outsourcing capabilities q_{ij} for each input $\hat{\omega} \in \Omega_{\omega_j^c}$ they purchase for their core product ω_j^c, we do account for the fact that the recorded purchase of $\hat{\omega}$ might be used for other secondary products as well.

firms.

We compute the household sales of a multi-product firm product by product: each of its products is a separate variety in the market of that product, whose household demand depends on its own unit cost $c_j(\omega)$ relative to the price index P_ω, exactly as for single-product firms. The product sale shares s_ω only serve to split the firm's observed household sales across its products when we calibrate the demand shifters (Appendix B.8).[83] We compute energy use in the same way. For each product $\omega \in \mathcal{P}_j$, we compute the energy used per unit of production cost as in Equation (13), replacing the core product by ω, with in-house total requirements $\tilde{\alpha}_\omega^{\hat{\omega}}(j)$ computed over the firm's total in-house set $\mathcal{I}_j^{full}$, and the same ν_j and ψ_j for all products. The firm's energy use is the sum, over its products, of this energy intensity times the model's sales of the product, to households and to other firms, in the same scenario. It therefore covers the energy used at all in-house stages, including those serving secondary products, while the energy embodied in outsourced inputs is accounted for at their suppliers.[84]

Since the balance sheet data assigns only one ATECO per firm, we take several measures to identify multi-product firms. In our baseline, we use the customs data and define multi-product firms as firms with product export shares (out of their total exports) of several distinct 6−digit products above 5%.[85] We identify 21233 multi-product firms in our final sample that contains 153424 manufacturing firms. Table B.4 reports some descriptive statistics on those multi-product firms.

This implicitly assumes that all firms that only sell domestically are single-product firms.[86] We mitigate this assumption with a robustness exercise (still work-in-progress) that flags firms based on their *likelihood* to produce multiple products. In particular, we infer the likelihood that a firm produces multiple products using the recipes built with Fetzer et al. (2024)'s product×input dataset. In practice, we evaluate the share of distinct inputs (weighted by their transaction shares) purchased by a firm that enter the recipe of its reported product. A high share indicates that the firm's observed input purchases are consistent with the technological requirements of producing the single product by which it is classified. Conversely, a low share suggests that the firm purchases many inputs unrelated to that product's recipe, consistent with the interpretation that the firm produces multiple outputs. We then identify firms that are most likely multi-product firms as firms with a low share. Within that exercise, we consider two versions. In the first version, we treat the likely-multi-product firms as multi-product only in the recipe step, that

We weight the unit cost of $\hat{\omega}$ (which is independent of its use) by the total requirement weights $\tilde{\alpha}_j^{\hat{\omega},\nu} = \sum_{\omega \in \mathcal{P}_j} s_\omega \tilde{\alpha}_\omega^{\hat{\omega}}(j)$ across all its core and secondary products ω. We use firm-level energy-labor cost (accounting for all in-house inputs regardless of their use) and energy and labor expenditures in the denominator to create the input expenditure shares.

[83] Since ψ_j is calibrated on the core product (Appendix B.8), the model matches the household sales of the core product, while the split of the remaining household sales across secondary products follows from their relative costs.

[84] Leaf products sold by multi-product firms are not modeled and have no product line of their own. The energy used at in-house stages that no product line accounts for, such as stages serving a leaf product, is weighted by the observed product sale shares s_ω.

[85] We use export shares as the product sales share s_ω. For the 11 firms for which the core product is not exported, we build product sales share based on the total non-manufacturing sales, assigning the residual domestic sales to the core product.

[86] Standard trade models could rationalize this assumption through a Melitz (2003)-type argument: multi-product firms are larger, and large firms are exporters. Thus, domestic firms should primarily be single-product firms.

Table B.4: Descriptive Statistics for Multi-Product Firms

	Panel A: Shares of (in %)		
	Manufacturing firms	Sellers	Buyers
mp. firms	13.84	14.45	12.61
Transactions involving mp. firms	-	35.91	42.09

	Panel B: Within Multi-Product Firms						
	Mean	Std. Dev.	Min	P25	Median	P75	Max
Nb. of products per mp. firm	2.78	1.11	2.00	2.00	2.00	3.00	11.00
Product sales share	0.34	0.27	0.05	0.11	0.24	0.55	0.95

Notes: Panel A reports the share of multi-product (mp.) firms and of the transactions that involve mp. firms for the final sample of manufacturing firms, and for the sample of manufacturing sellers and buyers in the firm-to-firm transaction data. Panel B reports descriptive statistics for multi-product firms. The first row shows the distribution of the number of products produced by firms operating in more than one product category. The second row reports the distribution of product-level sales shares within multi-product firms.

is, we exclude them when building the recipes. We don't treat those firms as multi-product firms in any other step of our quantitative exercise, because we don't actually observe the other products they might be producing. In the second version, we infer the probability that a multi-product firm produces a given set of products based on the technological similarity between its observed input purchases and the recipes of each product. Formally, let $\mathcal{O}_f$ denote the set of distinct inputs purchased by firm f, and let Ω^U_ω denote the technological recipe of product ω, defined as the set of *distinct* inputs required to produce ω either directly or indirectly. Intuitively, a firm's input bundle $\mathcal{O}_f$ should closely overlap with the recipe Ω^U_ω of its actual product(s). To quantify this overlap, we compute for every firm–product pair (f, ω) a similarity index $s_{f\omega}$ capturing the degree of correspondence between $\mathcal{O}_f$ and Ω^U_ω. We compute a weighted Jaccard index that accounts for the intensity of each input in the firm's purchases:

$$s^w_{f\omega} = \frac{\sum_{\hat{\omega} \in \mathcal{O}_f \cup \Omega^U_\omega} \min(v_{\hat{\omega} f}, v_{\hat{\omega}\omega})}{\sum_{\hat{\omega} \in \mathcal{O}_f \cup \Omega^U_\omega} \max(v_{\hat{\omega} f}, v_{\hat{\omega}\omega})}$$

where $v_{\hat{\omega} f}$ denotes the share of total transaction value that firm f allocates to input $\hat{\omega}$, normalized so that $\sum_{\hat{\omega} \in \mathcal{O}_f} v_{\hat{\omega} f} = 1$, and $v_{\hat{\omega} p}$ represents the relative importance of input $\hat{\omega}$ in the recipe of product p (set to $1/|\Omega^U_\omega|$ if unknown). This measures the overlap between the two input bundles relative to their union, while the normalization ensures that the weighted similarity is independent of firm size and reflects only the composition of inputs. The resulting similarity scores are transformed into probabilities by normalizing across all candidate products:

$$P(\omega \mid f) = \frac{s^\omega_{f\omega}}{\sum_{\omega'} s^{\omega'}_{f\omega'}}$$

For each firm, additional to its reported product, we retain the top 3 (with robustness using up to 5) products with the highest normalized similarity $P(\omega \mid f)$, which represent the most plausible outputs that could rationalize its observed input structure. This procedure thus recovers an empirical distribution of

potential product portfolios consistent with firms' observed input choices and the technological structure of production. It assumes that the same inputs can be used for multiple products, consistent with the multi-product firm literature (Boehm et al., 2022), and assesses the likelihood for each product independently of the rest of the portfolio.

B.4 Descriptive statistics of the sample

Table B.5 reports descriptive statistics on our final samples. The upper panel reports descriptive statistics on the final sample of manufacturing firms that we use for our analysis, which are all manufacturing firms for which we have energy and labor expenditures information. The lower panel reports descriptive statistics on the final sample of firm-to-firm transactions that we use for our calibration, which contain transactions between all manufacturing firms and their upstream suppliers. 90% of the transactions are between two firms that are part of our final sample, and 81% of our domestic suppliers are part of our final sample.

Panel A of Table B.6 reports descriptive statistics on the number of distinct suppliers from which a buyer purchases the same input. Most inputs are sourced from a unique supplier. The second column confirms this pattern using the number of distinct countries a firm sources a given input from. Panel B additionally shows that, when firms source an input from multiple suppliers, the lion's share of their sourcing still comes from one supplier.

B.4.1 Additional stylized facts

Analogously to Figure 1, Figure B.1 plots firms' intermediate expenditure shares, a proxy for outsourcing, against their energy intensity relative to their suppliers', splitting the sample by whether the purchased input's production is covered by the EU ETS cap-and-trade program.[87] Panel B.1a measures energy intensity relative to sales, Panel B.1b per unit of labor expenditure, which corresponds to value added net of energy costs in our model. Since the only distortion in our model is the unpriced environmental externality, inputs covered by the ETS face less distorted energy prices and should therefore be less misallocated.[88] Panel B.1a is consistent with this prediction: the negative slope is flatter for regulated inputs, meaning the tendency of relatively energy-intensive buyers to produce in-house is weaker where energy prices are less distorted. This interpretation remains subject to caution: energy intensity conflates firm boundaries and effective energy efficiency (the confound our structural approach directly addresses), and ETS coverage is itself assigned to activities with high energy requirements, usually more upstream, so the two groups differ in composition, not only in regulation. Panel B.1b partially addresses the first of these concerns. Because bringing inputs in-house raises energy use and labor use together, this denominator removes the mechanical span effect of firm boundaries, though not their composition effect, as we pointed in our discussion of Equation (7). The slope indeed flattens for both groups, and turns

[87] We map the EU ETS list of regulated activities into ATECO-classified inputs following Coster et al. (2024).
[88] The EU ETS carbon cost contributes to roughly $9-12\%$ to the final energy price of the covered firms.

Table B.5: Descriptive statistics of the clean datasets (2019)

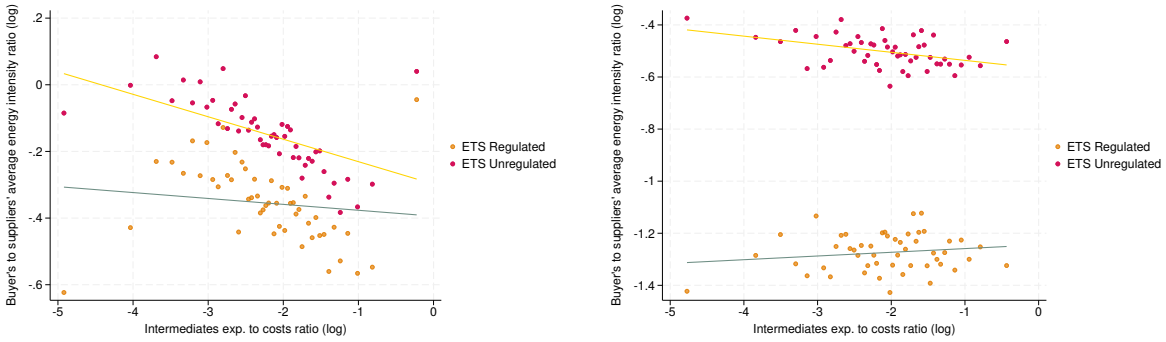
	Manufacturing firms - balance-sheet
Number of firms	153424
Total energy expenditures (bio €)	54.3
Energy price (€/kWh)	0.083 (0.07)
Total labor expenditures (bio €)	174.8
Total sales (bio €)	1074.0
Total employment (headcount, mio)	2.72
Avg. (sd) employment per firm (headcount)	17.8 (183)
Share of firms that export	42%
Share of firms that import	23%
Italian Manufacturing Employment (headcount, mio)	3.9
Italian Manufacturing Output (bio €)	1249
Italian Manufacturing VA (bio €)	351
Italian GDP (bio €)	2130
	Manufacturing firm-to-firm transactions dataset
Total firm-to-firm manufacturing transactions (bio €)	259.8
Number of firms	179691
Number of suppliers	141344
Number of buyers	175088
Number of buyers with balance-sheet information	117908
Number of manufacturing inputs	473
Number of manufacturing products	474
Number of transactions	2928934
Share of firms that export	28%
Share of firms that import	15%
Share of domestic transactions	89%
Share of domestic suppliers	95%
Average (sd) number of buyers per domestic supplier	32.2 (105.9)

Notes: The upper panel reports descriptive statistics on the final sample of manufacturing firms that we use for our analysis, which are all manufacturing firms for which we have energy and labor expenditures information. The middle panel reports aggregate variables from Istat for the Italian manufacturing sector. The lower panel reports descriptive statistics on the final sample of firm-to-firm transactions that we use for our calibration, which contain transactions between all manufacturing firms and their upstream suppliers. Transactions are limited to one supplier per input-buyer pair. When averages are reported, standard deviations are in parenthesis.

Table B.6: Descriptive Statistics on Suppliers per Buyer-Input

	A: Number of Suppliers per Buyer-Input		**B: Supplier Transaction Share**	
Statistic	Domestic suppliers	Foreign countries	All suppliers	Top supplier
Mean	1.78	1.72	0.26	0.72
Standard deviation	2.33	1.78	0.31	0.20
Minimum	1	1	0.000000003	0.036
25th percentile	1	1	0.018	0.57
Median	1	1	0.11	0.74
75th percentile	2	2	0.45	0.90
Maximum	346	78	0.9999	0.9999

slightly positive for regulated inputs, again pointing to less misallocation where energy is priced closer to its social cost.



(a) Energy exp. relative to sales

(b) Energy exp. relative to labor exp.

Figure B.1: Buyer's energy intensity relative to its suppliers' and intermediate input expenditure share, by EU ETS coverage

Notes: The figures plot the correlation between the intermediate input expenditure share of the buyer and the energy intensity of the buyer, relative to the average energy intensity of its suppliers, by inputs that are and are not regulated by the EU ETS. Energy intensity is defined as energy expenditure relative to sales in panel (a) and per labor expenditures in panel (b), controlling for products (6-digit). The relation is built for the sample of manufacturing single-product firms active in 2019.

B.5 Characterizing in-house production

Product recipes. As in Boehm and Oberfield (2023), the production of a product takes the form of a production tree: it combines inputs both sequentially (vertically) and horizontally, and each of these inputs is itself produced according to a production tree. A product is not used as an input in its own recipe, but the same input can enter a recipe at several stages if it is required in different branches of the tree. To characterize the recipes, we use our firm-to-firm transaction data and construct a directed acyclic graph (DAG) that links all products to their direct and indirect inputs. Starting from the observed input transactions, we construct a forward graph mapping each product to its inputs as in Boehm and Oberfield (2020) and Rachapalli (2021). We then recursively traverse this graph to recover the entire transitive closure of inputs for each product, that is, its recipe.

We clean the data to eliminate cycles and reduce noise. Following Carvalho (2014), Boehm and Oberfield (2020), and Rachapalli (2021), we only consider transactions that compose more than 3% of the total input expenditures of a product in order to avoid picking up weak links. To also exclude one-off purchases, we compute these shares in each year and keep an input only if its share exceeds 3% on average across years and in at least two years; the resulting recipes are common to all years. This deliberately constructs conservative (compact) recipes. Rather than exhaustively enumerating all inputs that could arise in niche or context-specific cases, we focus on the crucial inputs that systematically enter production, which are the most plausible candidates for in-house production. Outside the food industry, we manually drop beverage purchases, which are amenities for workers rather than inputs firms would plausibly produce in-house.[89] We also exclude transactions with multi-product firms from the recipe construction. Finally,

[89] We keep food purchases, as some firms run in-house cafeterias whose energy use enters their energy bills. However,

we exclude the 13 products that have only one producer to reduce noise, and characterize them as leaf inputs. Importantly, we ensure acyclicity both ex ante, by cleaning our data, and ex post, by implementing runtime safeguards in our recursive algorithm to prevent the propagation of cycles. This guarantees that every product's recipe is well-defined and finite.[90]

A standard DAG treats a direct transaction between two firms as the input of the supplier being a direct input of the product of the buyer. However, with in-house production, this is not always true: firm j producing ω_j^c might buy input $\hat{\omega}$ to produce $\tilde{\omega}$ in-house, and use $\tilde{\omega}$ to produce ω_j^c. We ensure that the DAG takes this feature into account by applying a transitive reduction to our acyclic graph. A transitive reduction is the minimal sub-graph that has exactly the same reachability as the original graph, but strips out every "redundant" edge that is implied by a longer path. Intuitively, whenever we observe a firm that produces a given product purchasing a particular input, we include that input in the product's production tree (recipe). We classify an input as a *direct* input of the product only if there is no other firm producing the same product that purchases some *other* input for which the former input already appears in that other input's production tree (i.e., we do not "skip" an intermediate stage).[91] A useful implication of this construction is invariance: each product is assigned a single production tree that is independent of where that product appears within the production trees of other inputs.

Using this in-house-consistent DAG, we recover three sets for each product ω. Defining $tr_{\hat{\omega},\omega} = 1$ if there is a path from input $\hat{\omega}$ to product ω in the DAG, we first build the recipe $\Omega_\omega = \{\hat{\omega} : tr_{\hat{\omega},\omega} = 1\}$ as the set of inputs that are required, directly or indirectly, to produce ω (upstream set). Then, we build $\mathcal{V}_\omega = \{\tilde{\omega} : tr_{\omega,\tilde{\omega}} = 1\}$ as the set of products that can be produced, directly or indirectly, using input ω (downstream set). Finally, we build $\mathcal{U}_\omega^d \subseteq \Omega_\omega$ as the set of direct inputs of ω, that is, those linked to ω by an edge of the DAG. With these different elements, combined with the set $\mathcal{O}_j$ of inputs purchased by each firm, we can identify the set of inputs that each firm produces in-house. Before that, to evaluate the robustness of the recipes, we draw random subsamples containing 60% of transactions, rebuild the product recipes, and — product by product — compare them with the baseline: the average and median overlaps exceed 70% of inputs.

In-house input set. A natural characterization of the set of inputs produced in-house by a firm producing product ω in 2019 is the intersection between the set of inputs upstream of product ω (i.e., its recipe Ω_ω) and the set of inputs downstream of the inputs that the firm purchases in that year ($\cup_{\hat{\omega}\in\mathcal{O}_j}\mathcal{V}_{\hat{\omega}}$), or in a branch that is entirely produced in-house. Formally, this would be $\mathcal{I}_j = \Omega_\omega \cap \{\cup_{\hat{\omega}\in\mathcal{O}_j}\mathcal{V}_{\hat{\omega}}\}$. Alternatively, we could also use $\mathcal{I}_j = \Omega_\omega \setminus \mathcal{O}_j^e$, where $\mathcal{O}_j^e = \mathcal{O}_j \cup \left(\cup_{\hat{\omega}\in\mathcal{O}_j}\Omega_{\hat{\omega}}\right)$ is the set of all inputs outsourced by the firm either directly or indirectly. However, applied as is, both strategies fail to capture the in-house set defined in our model, because inputs are used in multiple supply chains. Figure B.2 illustrates the issues with a production process for some product ω. Stages labeled h are performed

since prepared meals are arguably not a crucial input in most recipes, we also run robustness checks excluding them.

[90] Our cycles breaking algorithm excludes fewer than 4% of our total transactions.

[91] This rule can still wrongly create a direct link between a product and an input that, in reality, are separated by an intermediate for which we never observe a transaction. This does not affect our analysis, because any such missing intermediate is unobserved for *all* firms and is therefore excluded uniformly.

in-house, stages labeled o produce inputs that the firm purchases, that is, *observed in the data*, and stages labeled oe lie upstream of a purchased input and are therefore left to its supplier. $\mathcal{I}_j = \Omega_\omega \setminus \mathcal{O}_j^e$ wrongly excludes $\omega_e(h)$ from $\mathcal{I}_j$ due to $\omega_e(oe)$, since input ω_e enters $\mathcal{O}_j^e$. $\mathcal{I}_j = \Omega_\omega \cap \{\cup_{\hat{\omega}\in\mathcal{O}_j} \mathcal{V}_{\hat{\omega}}\}$ wrongly includes $\omega_g(o)$ and $\omega_l(oe)$ in $\mathcal{I}_j$ due to $\omega_j(o)$. Therefore, we need to move upstream recursively, as in the algorithm of the main text, to accurately identify all inputs produced in-house.

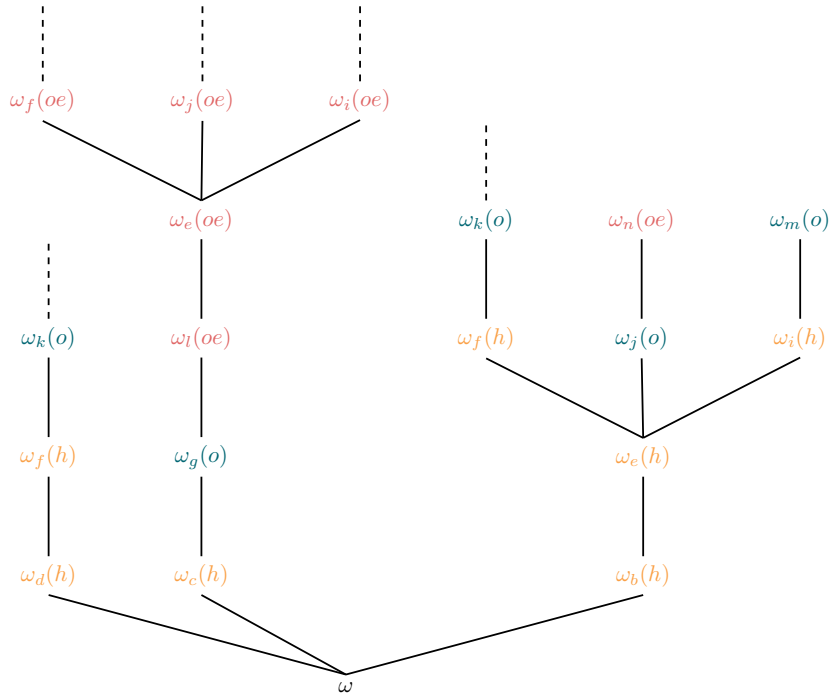


Figure B.2: Illustration of the production process of product ω

Notes: Each node is a stage. Stages labeled (h) are performed in-house, stages labeled (o) produce inputs that the firm purchases, and stages labeled (oe) lie upstream of a purchased input and are left to its supplier. Dashed lines indicate upstream stages not shown.

The same input can be produced in-house at several stages if it is required in different branches of the production tree (e.g., $\omega_f(h)$ in Figure B.2). As discussed above, a given input always has the same recipe, regardless of its location in a supply chain, and therefore the same effective cost for a given firm. The firm thus makes the same make-or-buy decision at every stage of that input that it decides on.[92] An input that we observe being purchased therefore supplies all the in-house stages that use it, and our algorithm never classifies it as produced in-house. Due to the sequentiality of production, however, a stage producing a given input may be performed in-house while another stage producing the same input lies upstream of a purchased input, and is therefore left to the supplier (e.g., $\omega_e(oe)$ in Figure B.2). Accurately measuring the resources a firm uses to produce an input in-house thus requires tracking every stage at which it produces that input in-house. This is consistent with our CES structure: with a common substitution elasticity $\frac{1}{1-\gamma}$ across stages and no supply-chain-specific requirement shifters, expenditure shares are pinned down solely by effective unit costs, so a given input enters in-house costs additively, once per stage at which the firm produces it in-house, as captured by the in-house total requirement $\tilde{\alpha}(j)$.

[92] This would differ in a probabilistic setting in which make-or-buy decisions account for the importance of each supply chain.

Input requirements. We calibrate the direct input requirement $\alpha^{\hat{\omega}}_{\omega}$ when constructing the recipes. For a direct input $\hat{\omega} \in \mathcal{U}^d_{\omega_j}$ purchased from supplier i, firm j's cost share is

$$\frac{p_{ij}(\hat{\omega})x_{ij}(\hat{\omega})}{c_j(\omega^c_j)x_j(\omega^c_j)} = \psi_j^{\frac{\gamma}{1-\gamma}} \tilde{\alpha}^{\hat{\omega}}_{\omega_j}(j) \left(\frac{c^{\mathcal{I}}_i(\hat{\omega})/\tilde{q}_{ij}}{c_j(\omega^c_j)} \right)^{\frac{-\gamma}{1-\gamma}}. \tag{27}$$

Only the product $\tilde{\alpha}^{\hat{\omega}}_{\omega^c_j}(j)\, \tilde{q}_{ij}^{\gamma/(1-\gamma)}$ enters (27), and Equation (29) backs out $\tilde{q}_{ij}$ so that (27) holds exactly for any requirements. Expenditure data therefore do not separately identify requirements from the level of outsourcing capabilities, and requirements are pinned by a normalization. Within each product ω, we set $\alpha^{\hat{\omega}}_{\omega}$ proportional to the cross-firm, cross-year average of the left-hand side of (27), among firms with core product ω that purchase $\hat{\omega}$. Since the transitive reduction rules out that a direct input $\hat{\omega} \in \mathcal{U}^d_{\omega}$ also enters the recipe of another input of ω, these firms use $\hat{\omega}$ only at the core-product stage, so that $\tilde{\alpha}^{\hat{\omega}}_{\omega}(j) = \alpha^{\hat{\omega}}_{\omega}$. This normalizes the average relative-price term in (27) to be equal across the direct inputs of ω. Its common level is set by α^{el}_{ω}, calibrated from labor and energy shares in the Italian Input-Output table, through the normalization $\alpha^{el}_{\omega} + \sum_{\hat{\omega} \in \mathcal{U}^d_{\omega}} \alpha^{\hat{\omega}}_{\omega} = 1$.

This normalization is not innocuous: it sets the level of outsourcing capabilities relative to in-house production ($q_{jj} = 1$), and hence each firm's distance to its make-or-buy threshold. We test it against revealed preference. Firm j outsources $\hat{\omega}$ only if $c_i(\hat{\omega})/q_{ij} \leq c^{\mathcal{I}}_j(\hat{\omega})$, and hence only if $c_i(\hat{\omega})/q_{ij}$ does not exceed the cost at which j could produce $\hat{\omega}$ performing all non-leaf stages upstream of it in-house. Since $\tilde{q}_{ij}$ is increasing in $\tilde{\alpha}^{\hat{\omega}}_{\omega^c_j}(j)$ when $\gamma < 0$, a violation is consistent with a requirement set too low for the corresponding pair. Under our calibration, only 4% of observed transactions violate this necessary condition, which seems to indicate that the bias is not too large.[93] Nonetheless, we are currently implementing an alternative that identifies requirements from the make-or-buy margin rather than normalizing them: for each product ω, the share of requirements going to direct inputs is calibrated to match the observed frequency with which producers of ω outsource them, iterating between this calibration, the unit-cost recursion of Appendix B.7.2, and the distribution of outsourcing capabilities until convergence.

B.5.1 Descriptive statistics on in-house production

Table B.7 reports the mean and percentiles of the number and share of distinct inputs required in the recipes (Panel A), produced in-house (Panel B) and outsourced (Panel C) for our sample of manufacturing firms. Manufacturing product recipes sequentially combine 38 distinct inputs across 106 stages, on average, but this mechanically hides a large dispersion depending on where the stages lie in the supply chain: some manufacturing inputs use a single leaf input (e.g., cutting and processing of marble, timber planing, wheat milling) that is, a recipe of size 1. In contrast, the manufacturing of vehicles, or of precision instruments require almost a hundred upstream inputs. For instance, television cameras and diodes come after the processing of 89 and 65 crucial manufacturing inputs, respectively.[94]

[93] Violations may also reflect measurement error in reconstructed costs, which Equation (29) absorbs into $\tilde{q}_{ij}$.

[94] As a mean of comparison, Rachapalli (2021) reports for India that, on average, a product is directly vertically linked (upstream or downstream) to 12 other products out of 5754. Her measure of direct linkages is constructed from observed

On average, 5 intermediate inputs that a single-product firm needs directly or indirectly to produce its product are produced in-house, but the median firm produces none of its inputs in-house. Yet, this hides a lot of dispersion both within and across sectors: 25% of the single-product firms produce more than 17% of all their required inputs. At the 2-digit level, apparel, leather and transport equipment are the sectors with the highest average in-house share, whereas tobacco producers produce the least in-house. In-house inputs are also very diverse; on average, more than 97% (86%) of the in-house inputs of a given firm have different 4−digit (3−digit) codes. Moreover, simple OLS regressions show that the number of inputs produced in-house significantly increases with the price of energy, both in raw value and relative to the wage.

We find that firms do not internalize inputs at random: in-house production is systematically tilted toward inputs that are more central in the technological structure of their production process. We measure the technological importance of in-house inputs using network-implied input requirements. For each firm j, the weighted in-house count $\mathcal{I}_j^w$ sums the requirement coefficients of inputs produced internally. The weighted share of in-house inputs is much larger than the unweighted share. Firms, on average, produce 9% of their required stages in-house (12% in terms of distinct inputs), but 31% of their total input requirements. To assess whether firms select important inputs disproportionately, we compare observed weighted importance to a random benchmark that preserves the number of in-house inputs. The importance ratio (observed over expected weighted count) has a mean of 4.78 and a median of 1.82. 87% of firms exhibit a ratio above one, 68.2% exceed 1.5, and only 7.0% fall below 0.5. Rank-based evidence yields the same conclusion: in-house inputs have an average rank of 10.0 (median 6.2), compared to an expected rank of 38.7 under random selection. The mean rank ratio (observed over expected) is 0.40 (median 0.35), and 99% of firms select inputs that are more important than random (rank ratio < 1). Finally, concentration among the most important inputs is pronounced. On average, 34% of top-quartile inputs are produced in-house. The resulting concentration ratio (top-quartile share over overall share) has a mean of 3.86 (median 2.50), with 54% of firms over-representing top-quartile inputs. This result is partly mechanical from our assumption on sequential production, which clusters in-house inputs close to the final product. Nonetheless, it still supports our choice to consider make-or-buy decisions on crucial inputs only. At the input level, we find that inputs that are more specialized (based on (Rauch, 1999)'s specification) are less likely to be produced in-house.

Figure B.3 compares our product recipes with those derived from the Fetzer et al. (2024) dataset; despite substantial differences - especially in product-classification granularity and filtering - the sectoral rankings are strongly correlated, supporting the validity of our algorithm. Our baseline recipes pool all years. To assess their stability, we also construct them separately for each year. Although only 35% of inputs are shared across years by count, these shared inputs typically account for 80-90% of the requirements (i.e., the Leontief mass): year-specific recipes agree on the economically important inputs. Figure B.4 confirms that the weights of shared inputs are also stable. Points on the 45 line indicate inputs whose importance

input transactions, and therefore omits production stages carried out in-house. Together with her finer product classification and broader sector coverage, this difference in measurement can explain the gap relative to our average number of direct links.

Table B.7: Recipes and firm boundaries

	Mean	Std. Dev.	Min	P25	Median	P75	Max
Panel A. Recipes (non-empty)							
Number of distinct inputs in recipes	38.43	24.01	1	13	40	59	89
Number of direct manufacturing inputs in recipes	2.06	0.98	1	1	2	3	5
Number of direct downstream manufacturing products	1.65	3.11	0	0	1	2.3	34
Stage depth	12.73	6.21	1	7	14	18	23
Panel B. in-house inputs							
Number of in-house inputs (single-product firms)	4.72	9.25	0	0	0	4	58
Share of in-house inputs to recipe (single-product firms)	0.121	0.185	0	0	0	0.173	0.733
Number of in-house inputs (multi-product firms)	2.94	6.74	0	0	0	2	56
Share of in-house inputs to recipe (multi-product firms)	0.071	0.131	0	0	0	0.102	0.667
Panel C. outsourced inputs							
Number of outsourced inputs in recipe (single-product firms)	1.53	0.92	1	1	1	2	11
Share of outsourced inputs to recipe (single-product firms)	0.060	0.057	0.0018	0.0133	0.0366	0.1000	0.40
Number of outsourced inputs in recipe (multi-product firms)	2.51	1.64	1	1	2	3	14
Share of outsourced inputs to recipe (multi-product firms)	0.020	0.027	0.0005	0.0055	0.0113	0.022	0.33

Notes: This table reports descriptive statistics on the number of distinct inputs in the recipes (panel A), produced in-house (panel B) and outsourced (panel C) by the manufacturing firms in the sample. Stage depth (fourth row of panel A) is the longest path from any upstream input to the product. In-house inputs are inferred by the authors from the products' recipes and the firms' input expenditures, and do not include core products. Outsourced inputs are manufacturing inputs purchased by the firms that enter the recipe of their product. Inputs and products are classified at the 6-digit ATECO level. The sample is restricted to manufacturing products with non-empty recipes.

for the product is unchanged between 2019 and 2023, and points above (below) the line inputs that became more (less) central to the recipe in 2023. Important inputs (with large $\tilde{\alpha}$) concentrate close to the 45 line.

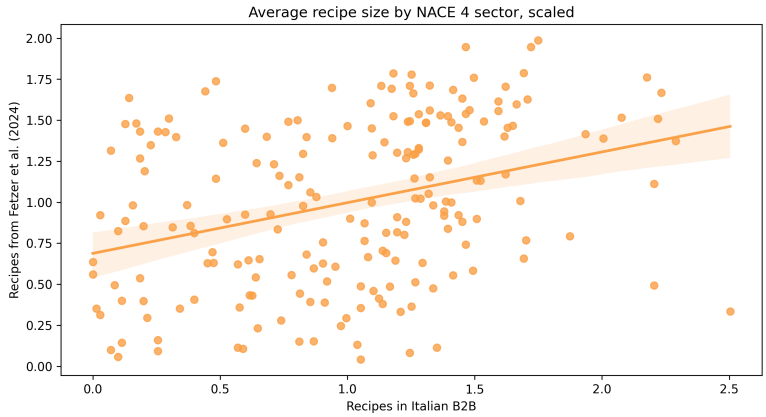


Figure B.3: Sectoral relative mean size of recipe built from Italian B2B vs Fetzer et al. (2024) data

Notes: The plot reports the correlation between the sectoral average number of inputs in product recipes, scaled by the total average number of inputs, for recipes built using the Italian firm-to-firm transaction data versus Fetzer et al. (2024) input-product dataset.

Finally, we validate our in-house inference method through two exercises. First, Table B.8 reports how often an input that a firm made in-house in 2019 is outsourced in a later year. These churns can be caused by i) measurement: a firm may in fact source an input from outside but not buy it in every single year and so that in 2019 we wrongly record it as made in-house ii) genuine change: firms do revise what they make and what they buy over time. The small magnitudes are consistent with either story, but the pattern across years point toward the second explanation: the share rises steadily from 2021 to 2023, and the fraction of firms with zero change fall from 61.7% to 58.0%. Occasional gaps in purchasing would

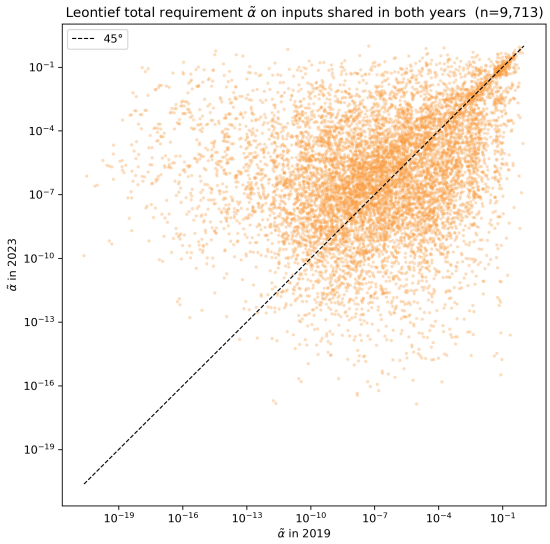


Figure B.4: Stability of Leontief total input requirements ($\tilde{\alpha}$) between 2019 and 2023

Notes: Each point is a (product, input) pair such that the input appears in the product's input requirements vector in both years i.e. we restrict to the intersection of inputs per product, conditional on the product itself being observed in both years (419 of 421/424 products with non-empty recipes). The x-axis plots $\tilde{\alpha}^{2019}_{\omega,\hat{\omega}}$ and the y-axis plots $\tilde{\alpha}^{2023}_{\omega,\hat{\omega}}$, where $\tilde{\alpha}_{\omega,\hat{\omega}}$ is the Leontief total requirement of an input $\hat{\omega}$ per unit of product ω, computed by recursing through the production DAG. Both axes are log-scaled; the dashed line is the 45 identity. Points on the 45 line indicate inputs whose importance for the product is unchanged across years. Points above (below) the line indicate inputs that became more (less) central to the recipe in 2023.

not grow with distance from 2019, whereas genuine changes accumulate. Second, Table B.9 uses these switches to test the sequential structure itself. In our production function, an input is needed only at the stages that use it, so a firm that starts buying ω should also stop buying whatever it was purchasing only in order to produce ω. Its remaining purchases, which feed stages it still performs in-house, should be unaffected. Comparing the two groups within the same firm nets out any reorganization of purchasing unrelated to the switch. The prediction holds in every comparison year: inputs left unnecessary by a switch are dropped 4.4 to 5.2 percentage points more often than the firm's other purchased inputs, and the gap is unchanged once we absorb input fixed effects, so it does not simply reflect those inputs being of a different kind. Two further patterns support the structure. Inputs that remain needed are dropped at the same rate as the purchases of firms that did not switch at all, and the gap is stable across years even though overall churn in purchasing rises with distance from 2019.

B.6 Energy parameters

This section provides additional details on the calibration of the product-specific parameters of the energy-labor nest, the elasticity of substitution governed by ρ_ω and the energy requirement β_ω, along with the firm-specific effective energy efficiency ν_j.

B.6.1 Energy-to-labor elasticity of substitution

$\frac{1}{1-\rho_\omega}$ corresponds to the energy-to-labor elasticity of substitution of a given product ω. Bretschger and Jo (2021) find that energy and labor are typically strong complements across most industries, though the degree of complementarity varies substantially. To account for this heterogeneity, we employ alternative

Table B.8: Inputs made in-house in 2019 that are purchased in a later year

	2021	2022	2023
Share of the 2019 in-house input set purchased in year C, across firms			
Mean	0.092	0.097	0.104
Median	0.000	0.000	0.000
75th percentile	0.071	0.077	0.091
Firms with a share of 0	61.7%	60.1%	58.0%
Pooled over all inputs			
Share of inputs purchased	0.0418	0.0455	0.0485
Firms compared	45,989	44,154	42,986

Notes: The sample is the 56,463 firms with a non-empty in-house input set in 2019; a firm enters a column only if it is also present in that year's data (10,474, 12,309 and 13,477 firms drop out in 2021, 2022 and 2023 respectively). For each firm we compute the fraction of its 2019 in-house inputs that it buys from another firm in year C; the last panel pools all firm-input pairs instead of averaging across firms.

Table B.9: Purchases of upstream inputs after a switch from making to buying

	2021	2022	2023
Share of 2019 purchased inputs no longer purchased in year C			
Inputs no longer needed after the switch	0.324	0.347	0.346
Inputs still needed	0.281	0.312	0.310
All inputs, firms that did not switch	0.286	0.313	0.323
Difference between the first two rows, within firm			
Mean across firms	0.052	0.045	0.044
	(0.007)	(0.007)	(0.007)
Regression, firm fixed effects	0.054	0.049	0.043
	(0.006)	(0.006)	(0.006)
Regression, firm and input fixed effects	0.044	0.048	0.049
	(0.006)	(0.006)	(0.006)
Firms compared	39,034	37,567	36,473
of which switch at least one input	15,542	15,499	15,818
of which have inputs no longer needed	10,072	9,966	10,128

Notes: The sample is firms present in both 2019 and year C that sell the same product in each and have a non-empty in-house and purchased input set in 2019. A firm switches when at least one of its 2019 in-house inputs is purchased in year C. A purchased input is no longer needed if the firm reaches it only through a switched input, and still needed otherwise. All of the structure (the in-house set, the purchased set and the recipe) is measured in 2019; year C contributes only purchases. The first two rows are computed over switching firms with at least one input of the first type, the third over firms that did not switch. The mean across firms is taken over the firms that have inputs of both types (6,453, 6,375 and 6,529). The regressions are linear probability models of an indicator for no longer purchasing the input on an indicator for the input being no longer needed, over all firm-input pairs (103,776, 100,099 and 97,486), with standard errors clustered by firm.

strategies to identify the elasticity parameter ρ_ω.

Strategy I: Using Bretschger and Jo (2021) estimates. Bretschger and Jo (2021) report the elasticity of substitution for 19 NACE 2−digit manufacturing industries. In our baseline, we use their estimates to calibrate our ρ_ω, aggregating our 6−digit product codes to their 2−digit industry classification. For the industries for which they don't report an estimate, we use their aggregate estimate of $\rho = -0.603$. This strategy highly restricts product heterogeneity. Moreover, Bretschger and Jo (2021) ignore in-house production. Formally, they estimate, for each industry separately

$$\log \frac{e_{jt}}{l_{jt}} = a_0 + \underbrace{\frac{1}{1-\rho_\omega}}_{\equiv a_1} \log\left(\frac{w_{jt}}{p_{e_{jt}}}\right) + a_j + a_t + \xi_{jt}$$

However, recall that in our model, the ratio of energy to labor demand takes the form:

$$\log \frac{e_j}{l_j} = \log \frac{\sum_{\omega \in \mathcal{I}_j \cup \{\omega_j^c\}} \tilde{\alpha}^\omega_{\omega_j^c}(j)\, \alpha^{el}_\omega \beta_\omega\, c_j^{el}(\omega)^{\frac{-1}{1-\gamma}} \left(\frac{c_j^{el}(\omega)}{p_e}\right)^{\frac{1}{1-\rho_\omega}} \nu_j^{\frac{\rho_\omega}{1-\rho_\omega}}}{\sum_{\omega \in \mathcal{I}_j \cup \{\omega_j^c\}} \tilde{\alpha}^\omega_{\omega_j^c}(j)\, \alpha^{el}_\omega (1-\beta_\omega)\, c_j^{el}(\omega)^{\frac{-1}{1-\gamma}} \left(\frac{c_j^{el}(\omega)}{w}\right)^{\frac{1}{1-\rho_\omega}}}$$

The price elasticity of the energy to labor demand is a non-linear weighted average of the product-specific elasticity of substitution that we are after, which depends on the in-house input set of the firms. In other words, their estimating equation identifies the true elasticity of substitution only in special cases, for instance if all the energy-labor bundles of a firm share the same ρ_ω and either the same β_ω, or $\rho_\omega = \gamma$. Due to the large variations in the types of inputs that are produced in-house even by similar firms, this assumption might seem too restrictive. Therefore, we use insights from our model to develop, as a robustness, a strategy that estimates model-consistent elasticity of substitution ρ_ω for each of our products, accounting for in-house production. This exercise is still work-in-progress, hence results are not included in this section, but we briefly describe it below.

Strategy II: Estimating ρ_ω on outsourcing firms. For the subsample of single-product firms that only produce their core product ω_j^c in-house ($\mathcal{I}_j = \emptyset$), the energy to labor expenditure ratio simplifies to

$$\log \frac{p_e e_j}{w l_j} = \log\left(\frac{\beta_{\omega_j^c}}{1-\beta_{\omega_j^c}}\right) + \frac{\rho_{\omega_j^c}}{1-\rho_{\omega_j^c}} log\Big(\frac{w}{p_e}\Big) + \frac{\rho_{\omega_j^c}}{1-\rho_{\omega_j^c}} \log \nu_j \tag{28}$$

We can estimate the product-level elasticity of substitution if we observe sufficiently many of these "single-product" firms producing the same product. To increase variations, we leverage the time series of the data. However, the time series currently available remains pretty short (2019 to 2023, excluding 2020 to avoid the COVID-pandemic shock), so we also carry a robustness exercise using French data, as explained below. We instrument energy prices using national prices weighted by the fuel mix of the firms. Effectively, there is also selection in the sense that, to estimate $\frac{\rho_\omega}{1-\rho_\omega}$, we are interacting factor prices with $\mathbf{1}\{\mathcal{I}_j = \emptyset\}$, and the selection process is correlated with ν_j. Instrumenting alone is not sufficient to deal with selection. However, our model shows that selection is problematic in this equation solely due to its correlation with the unobserved energy-biased technology ν_j. We leverage the panel structure

of our data and take the first-difference to partial out ν_j for each firm that maintains $\mathbf{1}\{\mathcal{I}_j = \emptyset\}$ over the two periods, under the assumption that the energy-biased technology remains fixed over time (as strongly supported by Figure B.6). In effect, we estimate

$$\Delta \log \frac{p_{e_{jt}} e_{jt}}{w_{jt} l_{jt}} = \frac{\rho_\omega}{1-\rho_\omega} \Delta \log \left(\frac{w_{jt}}{p_{e_{jt}}} \right) + \lambda_t + u_{jt}$$

Where λ_t is a time fixed effect, and u_{jt} is an iid error term which is not correlated with any component driving the make-or-buy decision.

For robustness, we run the same regression as in Strategy II on the same subsample of French firms. The longer time period considered offers more variations, and a detailed energy survey allows for a cleaner instrument. We use the Annual production survey (EAP) over the years 2010 to 2019 (to avoid the COVID years), a confidential representative survey at the firm×product level that contains information on the values and quantity of all inputs purchased by a firm in a given year, and all products produced by that firm. The survey is cross-sectional, but it also reports the same information for the previous year, meaning that we observe each firm for at least two consecutive years. We again focus on single-product and $\mathcal{I}_j = \emptyset$ firms. Single-product firms are directly identifiable in the survey. We apply the recipes that we built in our Italian data to infer the set of inputs produced in-house, exactly as described in Section 4.2.2. To run our regression, we complement this dataset with two additional firm-level data: the survey on energy consumption in industries (EACEI), which contains detailed annual information on costs and quantities consumed by energy type for a representative sample of firms. This allows us to refine the firms' fuel mix and to instrument energy prices with a shift-share based on initial fuel shares and national prices, as in Marin and Vona (2021) and Fontaine and Marullaz (2024). We also use the balance sheet data (FARE-FICUS), from which we retrieve wage and labor expenditures.

B.6.2 Energy parameters

We estimate the input-level energy requirements β_ω and the firm-level effective energy efficiency ν_j by non-linear least squares on the energy-to-labor expenditure ratio (Equation (15)), pooling the years 2019 and $2021 - 2023$:

$$\left(\hat{\beta}, \hat{\nu}\right) = \arg\min \sum_{j,t} \left(\log \frac{p_{e_{jt}} e_{jt}}{w_{jt} l_{jt}} - m_{jt}\left(\beta, \nu_j; \mathcal{I}_j, \tilde{\alpha}_j, \rho, \gamma\right) \right)^2$$

where m_{jt} is the model-implied log expenditure ratio, evaluated at the firm's in-house input set $\mathcal{I}_j$, its requirement weights $\tilde{\alpha}_j$, and its factor prices in year t. We parametrize β_ω through its log-odds, which guarantees $\beta_\omega \in (0, 1)$.

Identification. The structure is that of a two-way fixed effects model à la Abowd et al. (1999): β_ω is a product effect common across firms, ν_j a firm effect common to all the stages it performs. For single-product firms with $\mathcal{I}_j = \emptyset$, the ratio reduces to the log-linear Equation (28), in which $\log\left(\frac{\beta_\omega}{1-\beta_\omega}\right)$ and $\frac{\rho_\omega}{1-\rho_\omega} \log \nu_j$ enter additively; the multi-input case mixes them non-linearly. Relative β_ω's are identified as long as the firm×product network is connected - the analog of the connected set in Abowd et al. (1999), where firms producing several inputs in-house play the role of the movers - and our diagnostics show

that a single connected component covers virtually all products and firms. Firm-level variations in factor prices provide an additional source of identification, which does not rely on the make-or-buy margin. One direction, however, is not identified: lowering all $\log \nu_j$ by a common κ and raising each $\log\left(\frac{\beta_\omega}{1-\beta_\omega}\right)$ by $\kappa \frac{\rho_\omega}{1-\rho_\omega}$ leaves Equation (28) unchanged for every firm, the analog of the location normalization in two-way fixed effects models (in the multi-input case, this direction is only approximately flat). Since ρ_ω varies across sectors, this normalization sets not only the level of the β_ω's but also their relative level across sectors with different ρ_ω: the data identify $\frac{1-\rho_\omega}{\rho_\omega} \log\left(\frac{\beta_\omega}{1-\beta_\omega}\right)$ up to a common constant. We pin it with the energy expenditure shares of the Italian supply and use tables (SUT). These shares are industry aggregates: an industry's share includes the energy its firms use to produce inputs in-house, and therefore overstates the β of the industry's own product when inputs with high energy requirements are made in-house. This reallocation approximately cancels out in the aggregate, where the industry shares weighted by the in-house exposure of each industry's firms equal the β_ω's weighted by each product's in-house exposure - the exposure of an input being its requirement weight times its labor-and-energy cost share, i.e. its weight in the firm's energy-labor nest. We therefore anchor this exposure-weighted average, rather than each β_ω, to the tables.

Implementation. Since ν_j is a technology, we hold it fixed within firm over the sample period, but we allow it to break when the firm's fuel mix changes markedly (a change in the renewable share above 40 percentage points), which signals an actual technological change. We also include NACE $2-$digit $\times$ year shifters on $\log \nu_{jt}$, which absorb industry-level trends - in particular the adjustment to the $2022-2023$ energy crisis - so that neither β_ω nor the firm effects load aggregate shocks. We minimize the objective by block coordinate descent, alternating exact one-dimensional updates of the firm effects and of the industry shifters with quasi-Newton updates of the β_ω's, each weakly decreasing the objective, and re-imposing the normalization after each round; standard errors are obtained by bootstrap clustered at the firm level. As a validation, we confirm on simulated data that the procedure recovers the true parameters. We also recalibrate the ν_j separately, year by year, and verify that they line up with our estimated firm effects. This separate calibration allows us to confirm that ν_j is persistent over time, as a true technology should be.

Results. Table B.10 summarizes the calibration: using 460,478 firm-year observations over 2019 and $2021-2023$ from 119,522 distinct firms, we obtain a median $\hat{\beta}_\omega$ of 0.032 and a mean of 0.049 with all estimates strictly interior to the parameter bounds. The estimated parameters remain highly correlated with the external Italian-SUT benchmarks (0.93 on the logit scale) while reducing the sum of squared residuals by 21% relative to the SUT-based warm start, confirming that the internal calibration refines rather than overturns the external energy-share information.

B.6.3 Effective energy efficiency

A well-documented empirical regularity is that firms exhibit substantial heterogeneity in energy intensity even within narrowly defined industries (Petrick et al., 2011; Lyubich et al., 2018). This dispersion is

Table B.10: Internal calibration of the energy-share parameters β_ω (panel NLS)

Estimated $\hat{\beta}_\omega$	
p5 / p25	0.007 / 0.017
Median	0.031
Mean	0.047
p75 / p95	0.065 / 0.089
Min / Max	0.004 / 0.65
Corr. with external (SUT) logit-β	0.934

Notes: Nonlinear least squares on the pooled firm panel, with 99 NACE2×year shifters and fuel mix change controls. Cells in components with fewer than 5 firms are pinned down to their external value. All free $\hat{\beta}_\omega$ lie strictly inside the parameter bounds. The correlation with the external SUT benchmark is computed on the logit scale over free cells.

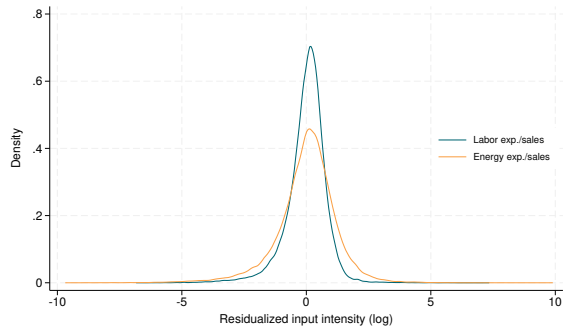
persistent across countries (Wagner et al., 2020; De Lyon & Dechezleprêtre, 2025), and the starting point for energy-efficient reallocative gains. We show that the same pattern holds among Italian manufacturing firms. Figure B.5a plots the cross-sectional distribution of energy expenditures to sales (in logs), residualized on 6-digit product fixed effects. For comparison, we overlay the corresponding distribution for labor expenditure shares, which is markedly tighter. The observed dispersion in energy intensity can be driven by several factors, and the literature has pointed toward large differences in energy technologies.[95] This comparison is across sellers of the same product. The relevant benchmark, however, is across all producers of the same *input*, including firms producing it in-house. Figure B.5b shows that the heterogeneity persists there.

As our model demonstrates, however, those measures of energy intensity conflate firm boundaries and effective energy efficiency. Figure B.5c shows that this substantial dispersion across comparable producers remains once we use the effective energy efficiency calibrated from our model, and in fact increases (Figure B.5d). The latter is expected: the measurement role of firm boundaries compresses the cross-seller dispersion in observed firm-level energy intensity. High-ν firms, which, all else equal, have lower observed intensity, should produce relatively more in-house and so raise their observed intensity, while low-ν firms outsource more and lower theirs. This compression should be even more salient at higher energy prices as they strengthen those forces (for a fixed set of producers).[96] Formally, a higher price raises the covariance between effective energy efficiency and the in-house share. Because observed intensity falls with efficiency but rises with the in-house share, this compresses observed firm-level dispersion, unless the price also widens the dispersion in that share enough to offset it. Thus, reduction in the dispersion

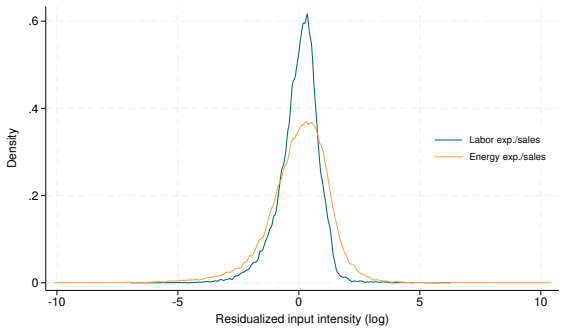
[95] The dispersion remains robust to controlling for investment, factor intensities, factor prices, and fuel mix. Excluding multi-product firms does not alter the results.

[96] Higher energy prices do not necessarily reduce the dispersion in *embedded* energy intensity, despite shifting production toward more energy efficient firms. The role of distortions in sustaining productivity dispersion in competitive environments is well-established: absent sufficiently strong countervailing wedges, production is concentrated in the most efficient producers (Syverson, 2004). With make-or-buy, however, this logic breaks down: if higher energy prices tend to systematically favor in-house production for energy-efficient buyers, the more energy-efficient option is exclusively accessible to the buyer itself. With large heterogeneity across buyers, embedded dispersion across supply chains can increase even as average embedded energy intensity falls.

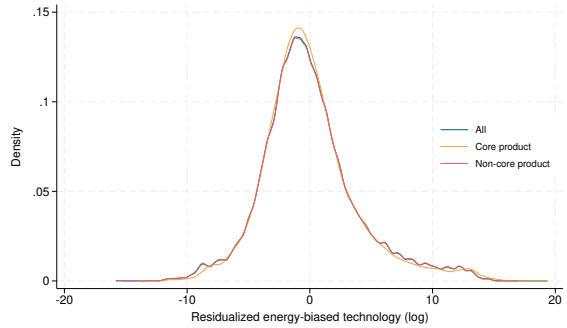
of observed proxies following a change in energy price does not necessarily indicate a reduction in the dispersion of effective energy efficiency. In our data, the standard deviation of effective energy efficiency naively implied by observed intensity (controlling for prices) in 2019 was 1.06 times its 2023 level, having compressed as energy prices rose. In contrast, the dispersion in effective energy efficiency calibrated year by year barely changed.



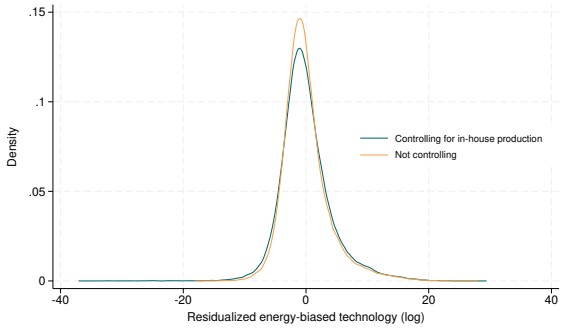
(a) Across sellers



(b) Across producers



(c) Effective energy efficiency ν



(d) With vs without controlling for $\mathcal{I}_j$

Figure B.5: Dispersion of input intensity and energy efficiency

Notes: Figure a) and b) plot the distribution of the residuals of the ratio of energy and labor expenditures to sales (in logarithm), while figure c) and d) plot the residuals of the calibrated effective energy efficiency ν. Figure c) uses the calibrated ν as described in our main calibration, while Figure d) adds a ν calibrated without controlling for firm boundaries. The residuals are obtained from a cross-sectional (2019), firm-level regression on input (6-digit) fixed effects. Product corresponds to the core product of each seller in figure a), and to any input produced by the firm in figure b). Figure c) splits the two categories into a distribution controlling for core product and one for non-core inputs fixed effects. Figure b) additionally controls for the number of inputs produced in-house.

Persistence. Figure B.6 shows the effective energy efficiency (ν_j) calibrated using the 2019 and 2023 cross-sections separately. The strong positive correlation confirms that our calibration does not pick up noise, and truly informs on the underlying effective technology of the firm. It also suggests that effective energy efficiency is a fairly persistent firm characteristic.

Effective energy efficiency gap. Figure B.7 reports, across inputs, the gap between the effective energy efficiency of each input's actual producer and that of an alternative producer, so that negative values indicate a more energy-efficient alternative. Panel B.7a weights the in-house comparison of the main text by each input's total requirement in producing the firm's product. Panel B.7b splits the same comparison, unweighted, by EU ETS coverage: regulated inputs show smaller gaps, in line with energy efficiency receiving a higher weight in the decision of less distorted inputs.[97] Panel B.7c compares

[97] We classify inputs as regulated by mapping the EU ETS list of activities into ATECO input codes, following Coster

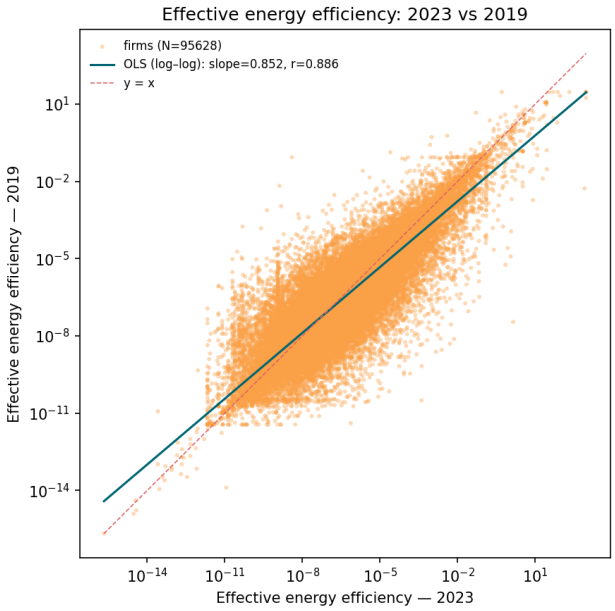


Figure B.6: Correlation between the calibrated ν_j in 2019 and 2023

Notes: The correlation is for the sample of Italian manufacturing firms operating both in 2019 and 2023.

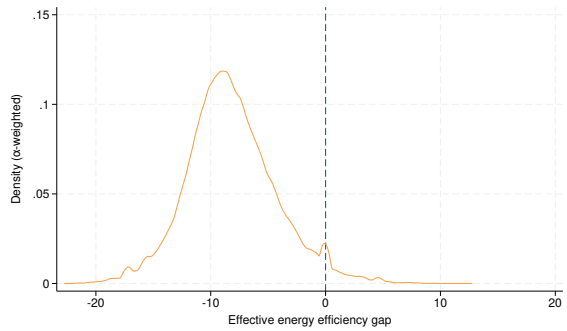
the actual supplier of each outsourced input to the buyer itself, isolating the make-or-buy decision from supplier selection. Panel B.7d compares the actual supplier to the most energy-efficient alternative, either among all potential producers, including the buyer, or among other sellers only. The two distributions barely differ, so the most energy-efficient alternative is usually another external supplier.

Similarly, Figure B.8a reports the distribution, across all inputs produced in-house, of the energy efficiency gap between the in-house producer and the supplier at the 95^{th} percentile and the average and median suppliers. The former confirms that the negative mass in energy efficiency gap of in-house inputs is not driven by outliers. The latter shows that there does not seem to be any systematic advantage of upstream suppliers, relative to downstream producers in terms of effective energy efficiency, supporting our calibration of ν as a firm-specific parameter rather than a core-product-specific one. If anything, the slightly larger positive tail suggests that our specification succeeds in capturing a technology taken into account by the firms in their producer selection - albeit not sufficiently when energy prices are distorted - and not random noise.

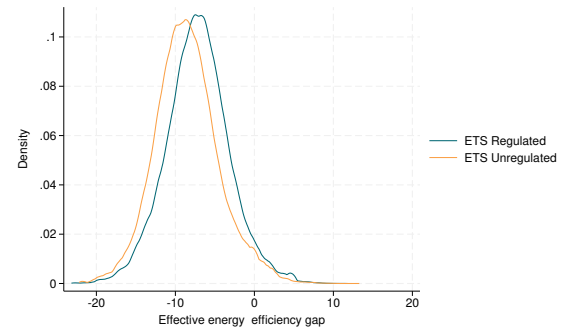
Figure B.8b splits inputs for which a decision is observed ($\omega \in \mathcal{I}_j \cup \mathcal{O}_j$) by their share of more energy-efficient non-chosen alternatives. Inputs with high shares are more likely to have at least one alternative which the environmental externality could justify switching to.

Finally, Figure B.9 reports the correlation between the effective energy efficiency of the firms and the number of inputs they produce in-house, relative to the total number of inputs in their recipes. The U-shape shows that the most energy-efficient firms tend to produce more inputs in-house, but that it is also

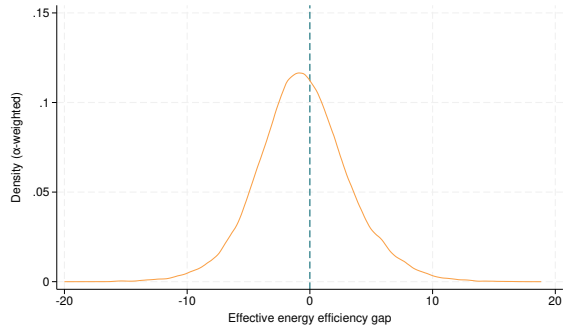
et al. (2024). In practice, the EU ETS covers installations based on their main activity, so an input produced in-house is correctly classified only if it corresponds to the main activity of the installation producing it, or if that activity is also covered.



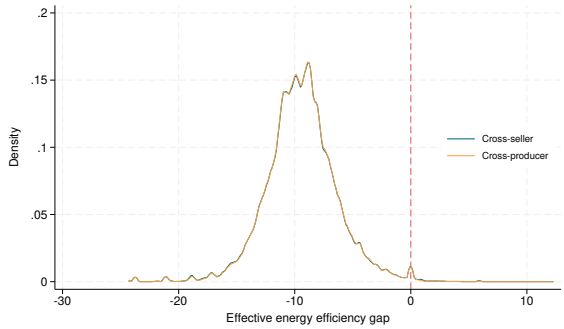
(a) In-house inputs



(b) In-house inputs by EU ETS



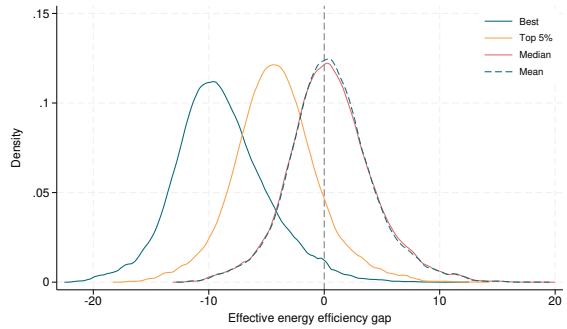
(c) Outsourced inputs (buyer only)



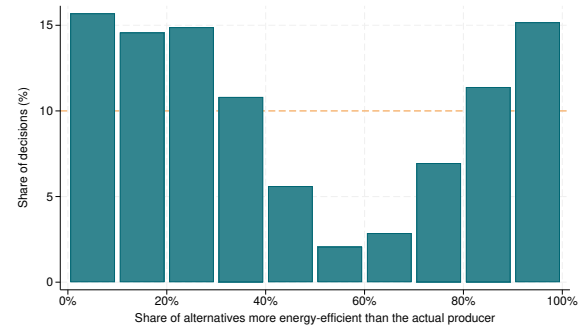
(d) Outsourced inputs (best producer vs supplier)

Figure B.7: Actual producer's effective energy efficiency relative to alternative producers

Notes: The figures report the distribution, across inputs, of the actual producer k's effective energy efficiency relative to that of an alternative producer s, in log points: $\log \nu_k - \log \nu_s$. Panel (a) covers inputs produced in-house ($\omega \in \mathcal{I}_j$, so that $k = j$) and compares the firm to its most energy-efficient external supplier, $s = s^{eff} \in \mathcal{S}_\omega$. Panel (b) splits the same comparison by EU ETS coverage. Panel (c) covers outsourced inputs ($\omega \in \mathcal{O}_j$) and compares the selected supplier $k \in \mathcal{S}_j^{\mathcal{O}}$ to the buyer, $s = j$. Panel (d) compares the selected supplier to the most energy-efficient alternative among all potential producers, $s = s^{eff} \in \{j\} \cup \mathcal{S}_\omega \backslash \{k\}$, and among external suppliers only, $s = s^{eff} \in \mathcal{S}_\omega \backslash \{k\}$. Panels (a) and (c) weight each input by its total input requirement $\tilde{\alpha}^{\omega}_{\omega^c_j}$ in producing the firm's product ω^c_j; panel (b) and (d) are unweighted. The sample covers all single-product manufacturing firms. The vertical line marks a zero gap.



(a) Energy efficiency gaps



(b) Share of more efficient alternatives

Notes: Figure (a) reports the distribution of actual producers' effective energy efficiency (ν_k), relative to the effective energy efficiency (ν_s) of their alternative producers, in percentage. It compares actual producers to the most energy efficient external producer ($s \in \mathcal{S}_\omega$), the one at the 95^{th} percentile of energy efficiency, and the average and median ones, for all in-house inputs. Figure (b) splits inputs with an observed make-or-buy decision ($\omega \in \mathcal{I}_j \cup \mathcal{O}_j$) by the share of a more energy-efficient non-chosen alternatives. The dotted line at 10% would represent a uniform split.

the case for energy inefficient firms, suggestive of energy-use misallocation induced by the make-or-buy decisions.

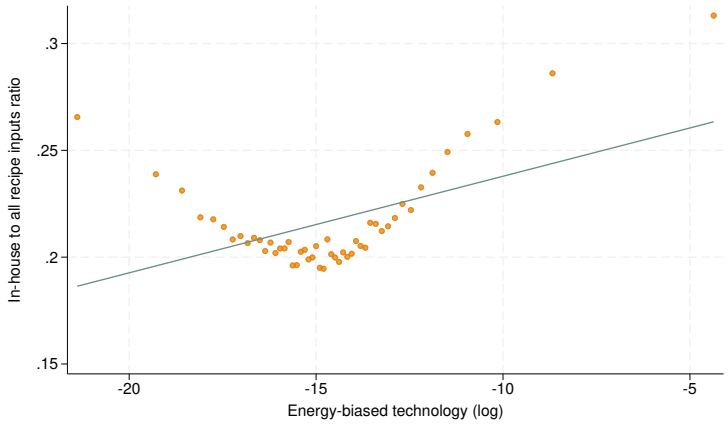


Figure B.9: Firms' effective energy efficiency and in-house share

Notes: The graph plots the relation between firms' effective energy efficiency ν and the number of inputs produced in-house relative to the total number of inputs in recipes. The relationship is measured controlling for product (6-digit).

B.7 Outsourcing capabilities

B.7.1 Inapplicability of standard methods.

Make-or-buy decisions, inclusive of supplier selection, are a discrete choice model. However, the standard type I extreme value distributional assumptions (e.g., McFadden (1974) and Eaton and Kortum (2002)) used to solve these models are not applicable in our context. Selection informs on the ranking of effective costs c/q, but with two dimensions of heterogeneity $\{\nu, q\}$, c/q does not follow a known distribution except under very restrictive and inopportune assumptions on the joint distribution $F(\nu, q)$. Sequential production and our exact calibration allow us to control for the costs, making q the only stochastic variable at each stage. Assuming a standard type I extreme value distribution for q would make the supplier selection probability tractable. However, in-house production is also a competing alternative that enters the selection problem, and its cost is not observed. Imposing $q_{jj} = 1$, we can make progress and compute the cost of each input produced in-house. This normalization breaks the requirement of all alternatives being type I extreme value distributed, but we could accommodate this asymmetry by considering a two-layer decision (in-house vs outsourcing, then conditional on outsourcing, supplier selection). However, the counterfactual in-house costs of outsourced inputs (i.e., the losing alternatives of the upper layer decision) remain unobserved and, with sequential production, dependent on upstream make-or-buy decisions hence on the distribution we precisely try to calibrate. The two-layer structure would deliver a closed form for the bottom layer, but not for the upper one, which still requires the same recursive computation over upstream make-or-buy decisions that our baseline strategy performs. The gain in tractability is therefore confined to the layer that is not binding, while the cost is a distributional assumption our data do not support: the empirical distribution of outsourcing capabilities is much better approximated by a lognormal. We therefore retain the lognormal and solve the selection problem numerically, as described below.

B.7.2 Unit costs computation

Our strategy relies on measuring $c_j^{\mathcal{I}}(\omega)$ and $c_i(\omega)$ for all inputs,[98] even the ones for which we don't observe a transaction. For this, we need the unit price of each input, which, from perfect competition on the input market, is the unit cost of each seller. We calibrate them recursively.[99]

1. For each non-leaf input ω and domestic firm j, we compute $c_j^{el}(\omega)$ using $\{\nu_j, \beta_\omega, \rho_\omega\}$.

2. For each non-leaf input ω and seller for which we do not observe labor or energy expenditures (which includes all foreign sellers), we set their costs exogenously at the average cost of ω among our sample of domestic firms. Those firms don't make any make-or-buy decision. We proceed in the same way under each counterfactual price scenario.

3. For each leaf-input ω^{leaf}, we use $c_i(\omega^{leaf}) = \bar{c}(\omega^{leaf}) = \bar{c}^{el}(\omega^{leaf})$, the average cost across producers of that leaf input. This cost remains fixed across counterfactual exercises.

4. We define the set $\mathcal{K}$ as the set of inputs for which we have identified all suppliers' c_i. We then compute the unit costs of each in-house inputs and core products $(c_j^{\mathcal{I}}(\omega))$,[100] for each firm, recursively. The core product is always produced in-house. Therefore, $c_j(\omega_j^c) = \frac{c_j^{\mathcal{I}}(\omega_j^c)}{\psi_j}$.

   (a) Using all the transactions of input $\omega \in \mathcal{K}$, we identify $\tilde{q}_{ij}$ for all suppliers i of $\omega \in \mathcal{K}$ and buyer j of $\omega \in \mathcal{K}$ as

$$\log(\tilde{q}_{ij}) = \frac{1-\gamma}{-\gamma}\left(\log\frac{p_e e_j + w l_j}{p_{ij}(\omega)x_{ij}(\omega)} + \frac{\gamma}{1-\gamma}\log\left(\frac{c_j^{el,tot}}{c_i^{\mathcal{I}}(\omega)}\right) + \log\tilde{\alpha}^{\omega}_{\omega_j^c}(j)\right) \tag{29}$$

   Where $x_{ij}(\omega)$ is j's total purchase of ω from i across all its uses, $c_j^{el,tot} = \left(\sum_{\omega' \in \mathcal{I}_j \cup \{\omega_j^c\}} \tilde{\alpha}^{\omega'}_{\omega_j^c}(j)\, \alpha^{el}_{\omega'} c_j^{el}(\omega')^{\frac{-\gamma}{1-\gamma}}\right)^{\frac{-(1}{}}$ and $\tilde{q}_{ij} = \psi_i q_{ij}/\delta_{ij}$ includes TFP ψ_i and iceberg trade costs δ_{ij}.[101] Moreover, defining A as

[98] For observed transactions (type (a)), the central identification challenge is to separate the outsourcing primitive embedded in $\tilde{q}_{ij}$ from firm-level endogenous costs. For two firms j, k with the same core product and in-house structure, each sourcing ω from i and $\hat{\omega}$ from s, the expenditure ratio $\frac{p_{ij}(\omega)x_{ij}(\omega)/p_{sj}(\hat{\omega})x_{sj}(\hat{\omega})}{p_{ik}(\omega)x_{ik}(\omega)/p_{sk}(\hat{\omega})x_{sk}(\hat{\omega})} = \left(\frac{q_{ij}/q_{sj}}{q_{ik}/q_{sk}}\right)^{\frac{\gamma}{1-\gamma}}$ depends neither on costs nor on requirements. Precisely because these cancel, however, it carries no information on the level of q_{ij}. The ratio of energy and labor to intermediate expenditure (Equation (29)) identifies $\tilde{q}_{ij}$ in level, relative to in-house production ($q_{jj} = 1$), but only given the in-house total requirement $\tilde{\alpha}^{\omega}_{\omega_j^c}(j)$, and it relies on a measure of costs. Appendix B.5 describes the normalization that pins requirements.

[99] Computing unit costs directly from the data means we need not assume a joint distribution over both primitives. We only need to take a stand on the marginal distribution of outsourcing capabilities, leaving their correlation with effective energy efficiency unrestricted. Since this correlation is the object of interest in our counterfactual exercise, this is a significant advantage. However, any misspecification in the cost reconstruction is mechanically absorbed by q_{ij}. This is innocuous provided that such measurement error does not respond to energy price changes.

[100] We need $c_j^{\mathcal{I}}(\omega)$ for all in-house inputs, and not only for the core product, to calibrate the outsourcing capabilities of the transactions that include inputs that are produced in-house based on the revealed preferences.

[101] With a Hicks-neutral productivity ψ_i, we cannot disentangle q_{ij} from ψ_i using expenditure shares. However, we can still build the suppliers' costs for each buyer firm, as q_{ij} and ψ_i enter in the same way (that is, inside $\tilde{q}_{ij}$). This allows us to build the cost of the core product of each firm, which we use to identify ψ_i using i's household sales. Once we've identified the ψ_i, we can retrieve q_{ij} net of the ψ_i. We do so before using them for our SMM calibration, to mitigate the within-supplier independence assumption.

the matrix of direct requirements with $A_{\omega\hat{\omega}} = \alpha_{\omega}^{\hat{\omega}}$, and D_j the diagonal matrix with $(D_j)_{\omega\omega} = \mathbb{1}\{\omega \in \mathcal{I}_j\}$, we have that $\tilde{\alpha}_{\omega}^{\hat{\omega}}(j) \equiv \sum_{K\geq 0} \sum_{\omega_1,\ldots,\omega_K \in \mathcal{I}_j} \alpha_{\omega}^{\omega_1}\alpha_{\omega_1}^{\omega_2}\cdots\alpha_{\omega_K}^{\hat{\omega}} = \left[(I - AD_j)^{-1}A\right]_{\omega\hat{\omega}}$ (and $\tilde{\alpha}_{\omega}^{\omega}(j) \equiv 1$). Acyclicity makes AD_j nilpotent, so the sum is finite.

(b) The unit cost of producing $\omega \in \mathcal{I}_j \cup \{\omega_j^c\}$ in-house is:

$$c_j^{\mathcal{I}}(\omega) = \left(\alpha_{\omega}^{el} c_j^{el}(\omega)^{\frac{-\gamma}{(1-\gamma)}} + \sum_{\hat{\omega}\in\Omega_{\omega}\cap\mathcal{I}_j} \tilde{\alpha}_{\omega}^{\hat{\omega}}(j)\,\alpha_{\hat{\omega}}^{el} c_j^{el}(\hat{\omega})^{\frac{-\gamma}{(1-\gamma)}} + \sum_{\hat{\omega}\in\Omega_{\omega}\cap\mathcal{O}_j\subseteq\mathcal{K}} \tilde{\alpha}_{\omega}^{\hat{\omega}}(j)\left(\frac{c_i^{\mathcal{I}}(\hat{\omega})}{\tilde{q}_{ij}}\right)^{\frac{-\gamma}{1-\gamma}}\right)^{\frac{-(1-\gamma)}{\gamma}} \tag{30}$$

Where i is the observed supplier of $\hat{\omega}$ to buyer j. Since costs depend on the input but not on the stage, all the stages producing a given input within a firm have the same unit cost and the same make-or-buy decision, so that an input reached through several in-house chains enters once, with $\tilde{\alpha}_{\omega}^{\hat{\omega}}(j)$ summing the requirements of all the stages at which it enters.

(c) We do this for all firms that only purchase inputs $\omega \in \mathcal{K}$. If this includes all the sellers of ω ($\mathcal{S}_{\omega}$), we add ω to $\mathcal{K}$.

(d) We repeat steps 4a, 4b and 4c until we have identified all unit costs for all suppliers.

5. Once we have recovered all empirical unit costs, we can proceed to characterize the distribution of the unobserved outsourcing capabilities, and draw them for our counterfactual analysis from these distributions. Before that, we externally validate the calibrated unit costs by assessing their ability to predict observed prices for the subsample of exporting firms. In particular, we compare calibrated costs to firm-product–level export prices (Figure B.10). The strong correspondence between calibrated costs and observed export prices suggests that the calibration successfully recovers latent marginal costs that are relevant for firms' pricing decisions. We also note that the outsourcing capabilities that characterize specialized inputs, based on Rauch (1999) tend to be twice as large, on average and at the median, as those of non-specialized inputs, reflecting that they seem to capture specialization gains, among other outsourcing factors.

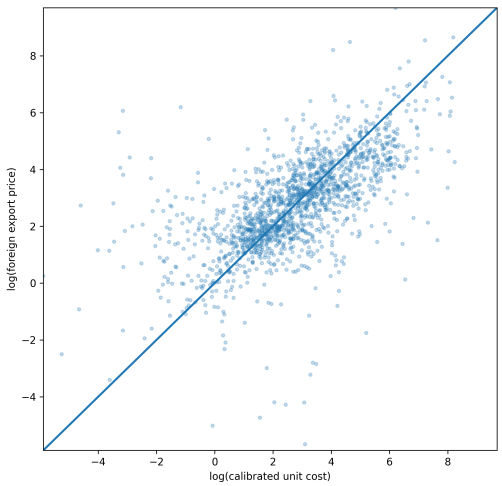


Figure B.10: Calibrated unit cost and observed price

Notes: Correlation between the calibrated unit cost and the observed unit export price for the subsample of exporting firms.

B.7.3 Estimating the parameters

We exploit the fact that the outsourcing capabilities not only determine the intermediate expenditure shares, but also the make-or-buy decision itself. As discussed, the former is observed only for a selected sample. Therefore, we use information from the latter to account for this selection. Intuitively , we can think of the q_{ij} as decomposed into a firm-specific component that determines selection, and a pairwise component that explains the residual variations in input expenditure share within firm, which are observed conditional on the input being outsourced. Our method allows for a general combination of these two components: the dependence between buyer and idiosyncratic components is non-linear and supports tail dependence.

To build the unobserved outsourcing capabilities, we calibrate their distribution to match empirical moments on the observed outsourcing capabilities through a Simulated Method of Moments (SMM) procedure. First, we assume that $\log(q_{ij}) \sim \mathcal{N}\left(\mu, \sigma^2\right)$, with within-buyer correlation modeled through a Student-t Copula $\mathcal{C}(q_{ij}, \dots q_{kj}; \rho_c, \tau_c)$ with parameters (ρ_c, τ_c). However, observed outsourcing capabilities are selected. We therefore construct the simulated *conditional* moments to account for this selection, using pseudo-bounds from the distribution of unobserved outsourcing capabilities. The distribution of unobserved capabilities q_{sj} is identified using information contained in observed make-or-buy decisions through a revealed-preference argument. Intuitively, when a supplier is not chosen, the make-or-buy decision implies that its effective cost must exceed that of the chosen alternative, which in turn delivers a natural upper bound on its outsourcing capability relative to the observed option. A key challenge is that the construction of unobserved q_{sj} depends on the parameters that we seek to calibrate, while the calibration of these parameters itself requires imputations of q_{sj}. We therefore adopt an iterative procedure that jointly imputes unobserved outsourcing capabilities and updates the parameters until convergence. Monte-Carlo exercises consistently validate the finite-sample properties of our Copula-based SMM estimator.

Algorithm. The key difficulty is that the capabilities q_{sj} for non-selected suppliers are unobserved, yet selection implies that they must satisfy inequality constraints relative to the observed q_{ij}. The algorithm exploits this structure by simulating unobserved alternatives to construct *pseudo-lower bounds* for observed winners, and by matching observed moments of q_{ij} to model-implied conditional moments.

1. **Measurement error** We model each observed $\log q_{ij}$ as the true capability, on which firms select suppliers, plus a transitory measurement error ϵ, and measure the error variance σ_ϵ^2 before running the algorithm, as the part of the variance of $\log q_{ij}$ that is not shared across years by the same buyer-supplier-input link, and hold it fixed. In our baseline calibration, we estimate $\sigma_\epsilon = 1.421$, about 7% of the observed variance.

2. **Initialization.** We initialize the parameter vector

$$\hat{\theta} = (\mu, \sigma, \rho_c, \tau_c)$$

3. **Simulation of unobserved alternatives.** For each buyer j and each non-selected supplier $s \neq k_j(\omega)$ of input $\omega \in \mathcal{O}_j \cup \mathcal{I}_j$, we simulate q_{sj} from a lognormal distribution truncated from above,

$$\log q_{sj} \sim \mathcal{N}(\mu, \sigma^2) \quad \text{truncated from above at } \bar{q}_{sj}$$

where the truncation enforces consistency with the observed selection of supplier i:

- If the input sold by s is produced in-house by firm j: $\bar{q}_{sj} = \frac{\delta_{sj} c_s}{c_j^{\mathcal{I}}}$
- If the input sold by s is outsourced to firm i by firm j: $\bar{q}_{sj} = \frac{\delta_{sj} c_s}{\delta_{ij} c_i / q_{ij}}$

Within-buyer dependence is generated by a one-factor Student-t Copula:[102] for each buyer j there is a latent factor F_j (shared across all of j's draws), and each q_{sj} is a function of F_j and a supplier-specific shock. Conditional on F_j, the $\{q_{sj}\}_s$ are independent across s. We can only model one dimension of correlation. We choose the within-buyer, reflecting an unobserved outsourcing capability parameter, because the buyer is the one making the make-or-buy decision. This assumes that the suppliers' propensity to sell is fully captured by dispersion in cost, distance and TFP.

4. **Posterior update of the buyer factor.** Given the candidate parameters and the observed selection patterns, we update the posterior of each F_j via sampling–importance resampling (SIR).[103] The buyer-specific log-likelihood aggregates one term per observation of j:

   - Outsourced observations contribute the winner's log-density plus the sum of log-cdf terms across the simulated losers, minus a selection correction that conditions on the buyer outsourcing rather than in-housing.
   - In-housed observations contribute the log-probability of the joint event "every external alternative s in ω's pool failed to beat in-house," equal to

$$\sum_{s=1}^{K} \log F_{\tau_c}\big(\xi_{sj}^{\mathcal{I}}\big), \qquad \xi_{sj}^{\mathcal{I}} \;=\; \frac{T\big(\delta_{sj} c_s / c_j^{\mathcal{I}}\big) - \sqrt{\rho_c}\, F_j}{\sqrt{1-\rho_c}},$$

   where $T(x) = F_{\tau_c}^{-1}\Big(\Phi\Big(\frac{\log x - \mu}{\sigma}\Big)\Big)$ maps the lognormal marginal to the t_{τ_c} quantile and F_{τ_c} is the t_{τ_c} cdf. In q-space a supplier loses when its capability falls *below* the threshold, hence the cdf rather than the survival.

   We resample one F_j per buyer with weights proportional to the exponentiated log-likelihood, and use the resampled F_j in subsequent steps so that within-buyer joint moments reflect the conditional distribution of F_j given the data.

[102] Monte-Carlo experiments with the Gaussian limit ($\tau_c \to \infty$) showed that it produces wider CIs on ρ_c and weaker identification signals.

[103] This is the standard "selection-corrected SIR" idea: candidate draws of the latent factor are reweighted by the model-implied probability of the observed selection events, then resampled. The conditioning on outsourcing vs. in-housing is what makes the simulator's within-buyer dependence consistent with the observed selection

5. **Construction of pseudo-lower bounds.** Using the simulated alternatives, construct for each observed (i, j) pair a pseudo-lower bound $\tilde{q}_{ij}^m = \max_{s \neq i} \left\{ \frac{\delta_{ij} c_i \, q_{sj}}{\delta_{sj} c_s} \right\}$. The pseudo-lower bound represents the tightest lower bound on q_{ij} implied by the best alternative supplier.[104] The set of alternatives includes the in-house option with $c_s = c_j^{\mathcal{I}}$, $\delta_{jj} = 1$ and $q_{jj} = 1$. However, this is a counterfactual, unobserved cost, which depends on upstream costs and unobserved make-or-buy decisions.[105] In each iteration, we construct it recursively, applying the same recursive make-or-buy algorithm that we use to build any counterfactual allocation. This requires knowing the upstream (unobserved) q_{ij}, which we draw from the lognormal distribution evaluated at the parameter set $\hat{\theta}$ of the current iteration.

6. **Simulation of model-implied moments.** Let $g(\cdot)$ denote a moment function. For each observed winner (i, j), we compute the model-implied conditional expectation of $g(q_{ij})$ given that supplier i was selected, that is, the expectation under the lognormal marginal truncated from below at the pseudo-lower bound of Step 5:

$$m_{\text{sim},ij}(\theta) = \mathbb{E}_\theta \left[g(Q \, e^{\sigma_\epsilon \epsilon}) \;\middle|\; Q \geq \tilde{q}_{ij}^m(\theta; \mathcal{Q}_j) \right], \qquad \epsilon \sim \mathcal{N}(0,1) \text{ independent of } Q,$$

where $\mathcal{Q}_j = \{q_{sj}\}_{s \neq i}$ denotes the collection of simulated alternatives for buyer j. We evaluate this expectation by drawing directly from the truncated distribution, rather than by drawing from the unconditional lognormal and discarding the draws that violate the bound. For each observed pair we draw R replicates $u_{ij}^{(r)}$, map them through the buyer's Copula, and invert the truncated lognormal cdf:

$$Q_{ij}^{(r)} = \exp\left\{ \mu + \sigma \, \Phi^{-1}\left(F_{ij}^{lb} + u_{ij}^{(r)} \left(1 - F_{ij}^{lb}\right) \right) \right\}, \qquad F_{ij}^{lb} = \Phi\left(\frac{\log \tilde{q}_{ij}^m - \mu}{\sigma} \right),$$

so that $Q_{ij}^{(r)} \geq \tilde{q}_{ij}^m$ holds by construction and every draw is retained. Because the empirical moments are computed on observed capabilities, we then add the measurement error to each draw, after the truncation since selection operates on true capabilities: $Q_{ij}^{\text{obs},(r)} = Q_{ij}^{(r)} \exp\{\sigma_\epsilon \, \epsilon_{ij}^{(r)}\}$ with $\epsilon_{ij}^{(r)} \sim \mathcal{N}(0,1)$. The simulated moment is then

$$m_{\text{sim},ij}(\theta) = \frac{1}{R} \sum_{r=1}^{R} g\left(Q_{ij}^{\text{obs},(r)} \right)$$

The pseudo-lower bound is itself a random variable, since it depends on the alternatives $\mathcal{Q}_j$ simulated in step 3. We do not integrate over it pair by pair: within an iteration we hold one draw of $\mathcal{Q}_j$ per observed pair, and the averaging over the distribution of the bounds comes from pooling across the many observed pairs in step 6. The underlying uniforms are frozen across parameter evaluations (common random numbers), so that the objective is smooth in $\hat{\theta}$; the bounds themselves still move with $\hat{\theta}$, and are refreshed at each iteration together with the posterior of F_j.

[104] The draws in step 3 cannot be bypassed because the pseudo-lower bound is a function of observed suppliers costs, counterfactual in-house costs, and unobserved outsourcing capabilities, thus an object for which the distribution is unknown. We therefore cannot use a closed-form result that would characterize that maximum.

[105] The recursive equilibrium cost algorithm delivers i) the cost $c_s(\omega)$ of all suppliers $s \in \mathcal{S}_\omega$ of all inputs ω, and ii) the in-house cost $c_j^{\mathcal{I}}(\omega)$ of every firm $j \in \mathcal{D}$ for every input ω in its recipe, including those it outsources.

7. **Moment matching.** Aggregate simulated moments across observations and match them to the corresponding empirical moments computed from observed winners:

$$\hat{\theta} = \arg\min_{\theta} \|m_{\text{data}} - m_{\text{sim}}(\theta)\|^2 .$$

Once the parameters have been identified, we draw the unobserved q_{sj} from the truncated lognormal for type (b), ensuring consistency with the observed producer selected, and from the general lognormal distribution for type (c). To put observed suppliers on the same footing as unobserved ones, whose capabilities are drawn from the distribution of true capabilities free of measurement errors, we first replace each observed $\log q_{ij}$ with its posterior mean under the revealed-preference restriction that the chosen supplier is cheaper than in-house production (while the restriction on other external suppliers is ensured by their own truncated draws).

Moment Conditions. We match seven moments. Two anchor the marginal lognormal distribution and three exploit the panel structure of buyers' transactions across inputs to identify the Copula correlation ρ_c and the tail-dependence parameter τ_c. Two further moments, described at the end of this section, discipline the trade-cost parameters (ζ, ζ_f).

Block 1: Marginal distribution of winners.

1. **Moment 1: Mean log capability.**

$$m_1 = \mathbb{E}[log(q_{ij})] .$$

Captures the average level of outsourcing capability among selected suppliers. Conditional on selection, it primarily identifies the location parameter μ of the lognormal distribution: μ governs the level from which observed and unobserved alternatives are drawn, and therefore the mean of $\log(q_{ij})$ after selection.

2. **Moment 2: Variance of log capability.**

$$m_2 = \text{Var}(log(q_{ij})) .$$

Captures the dispersion of capabilities among winners and identifies the lognormal scale parameter σ, which is the standard deviation of $\log(q_{ij})$ before selection. A higher σ implies greater heterogeneity in high-capability draws, which translates into higher variance of $\log(q_{ij})$ even after conditioning on selection.

Block 2: Within-buyer dependence moments.

The two marginal moments anchor (μ, σ) but cannot identify the Copula on their own. Each buyer is observed in multiple transactions across different inputs, and the joint distribution of a buyer's q_{ij} across those transactions carries direct information about within-buyer dependence. We exploit this with three buyer-level moments.

3. **Moment 3: Within-buyer variance of log capability.**

$$m_3 = \mathbb{E}_j[\mathrm{Var}_\omega(\log(q_{ij}) \mid j)]\,,$$

where the inner variance is taken across a buyer j's observed transactions and the outer expectation is over buyers with at least two transactions ($n_j \geq 2$). Under the one-factor Copula structure, $\mathrm{Var}_\omega(\log(q_{ij}) \mid j)$ decreases monotonically in the Copula correlation ρ_c: higher ρ_c implies that a buyer's transactions cluster tightly around their personal mean, shrinking the within-buyer variance. This moment provides supplementary signal for ρ_c.

4. **Moment 4: Joint-tail probability at the 90^{th} percentile.**

$$m_4 = \Pr\big(\big|\{i : q_{ij} \geq Q_{0.90}\}\big| \geq 2 \,\big|\, n_j \geq 2\big)\,,$$

where $Q_{0.90}$ is the 90^{th} percentile of the pooled marginal $\log q$ distribution, i.e. the top decile of capabilities. Under independence of capabilities across a buyer's inputs, this probability equals the binomial baseline; positive within-buyer dependence inflates it. The 90^{th}-percentile threshold captures *body* dependence and is the principal anchor for ρ_c.

5. **Moment 5: Joint-tail probability at the 98^{th} percentile.**

$$m_5 = \Pr\big(\big|\{i : q_{ij} \geq Q_{0.98}\}\big| \geq 2 \,\big|\, n_j \geq 2\big)\,.$$

The same construction at a deeper threshold ($Q_{0.98}$, the top 2% of marginal $\log q$). The Gaussian Copula has zero asymptotic tail dependence: as the threshold becomes extreme, the joint-tail probability decays toward the binomial baseline. The Student-t Copula, in contrast, has positive tail dependence, controlled by the degrees-of-freedom parameter τ_c: small τ_c inflates m_5 even far into the tail, while $\tau_c \to \infty$ recovers the Gaussian limit. The pair (m_4, m_5) therefore separates body correlation from tail dependence, identifying τ_c orthogonally to ρ_c.

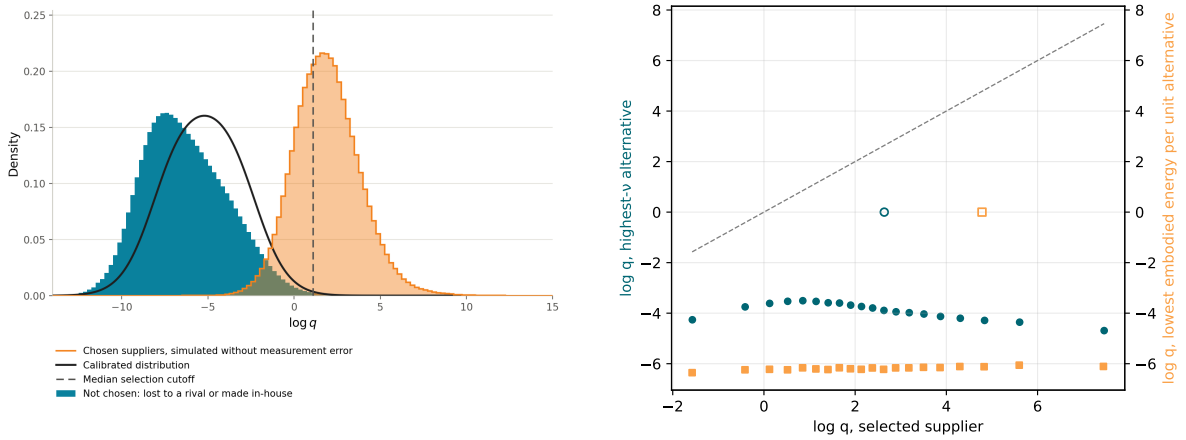
Iceberg trade costs For an observed transaction, the first-order conditions pin down $\tilde{q}_{ij} = \psi_i q_{ij}/\delta_{ij}$, which already includes the iceberg trade cost δ_{ij}. We fit the parametric marginal distribution of capabilities to the distance-free q_{ij}, so that the calibrated outsourcing capability in Section 4.2.4 corresponds to q_{ij}/δ_{ij}. We parametrize

$$\delta_{ij} = \begin{cases} \big(1 + d_{ij}/\bar{d}\big)^{\zeta} & \text{if } i \text{ is a domestic supplier,} \\ \exp(\zeta_f) & \text{if } i \text{ is a foreign supplier,} \end{cases}$$

where d_{ij} is the distance in kilometers between the two firms, $\bar{d} = 150\,\text{km}$ is a fixed reference distance, ζ is the domestic distance elasticity and ζ_f a flat foreign premium. We discipline (ζ, ζ_f) with two additional moments. The first is the average distance percentile of the selected supplier among the domestic members of its choice set. In the model it is computed from exact win probabilities conditional on outsourcing. The second is the share of transactions sourced from foreign suppliers. We find $\hat{\zeta} = 2.27$ and $\hat{\zeta}_f = 1.26$ (with data moments 0.232 and 0.033, and model moments 0.248 and 0.029).

B.7.4 Distribution of outsourcing capabilities

Figure B.11 shows the result of our calibration for the distribution of outsourcing capabilities. Figure B.11a compares the empirical distribution of the outsourcing capabilities of the selected pairs to the calibrated distribution. The observed distribution sits to the right of the calibrated distribution, in line with those pairs being selected. The two distributions overlap, partly because observed q carries measurement error and partly because suppliers are selected on unit cost rather than on q alone: a supplier can lose to a rival with a lower q when its own cost is higher. The non-chosen pairs sit to the left of the calibrated distribution. Most of them are candidates rejected in favor of in-house production, and they are concentrated at low q: the in-house threshold does not depend on a winner's capability, so a make decision reveals that all candidates were weak. Candidates that lost to a rival are more dispersed, and their upper tail reaches into the distribution of the selected pairs, since the threshold is loose when the winner is strong. Figure B.11b compares the selected pair's outsourcing capabilities to those of the alternative with the highest ν and of the alternative with the most energy-efficient supply chain.



(a) Distribution of selected and all q

(b) Selected q relative to most energy efficient

Figure B.11: Calibrated distribution of outsourcing capabilities

Notes: Figure a) plots the empirical distribution of the outsourcing capabilities for selected pairs, along with the calibrated distribution for all pairs. Figure b) plots the selected pair's outsourcing capabilities against the outsourcing capabilities of the option with the highest ν (left y-axis) and against the alternative whose supply chain embodies the lowest energy per unit produced (right y-axis). White dots mark cases in which the alternative is in-house production, whose $q_{jj} = 1$.

B.8 Hicks-neutral productivity

Once we have built the costs of all firms, as described in Appendix B.7, we identify the Hicks-neutral productivity ψ_j[106] using sales share to household. In particular, the household sales ratio of two firms

[106] In addition to the outsourcing capabilities q_{ij} and the energy-biased technology ν_j, we can identify either the firm-specific Hicks-neutral technology (TFP) or a technology specific to the firm's in-house inputs, q_{jj}, but not both, because in-house transactions are never observed. We choose to identify TFP and normalize q_{jj} to 1, as it makes the interpretation of the q_{ij} as gains or losses from outsourcing more intuitive.

producing the same product[107] and selling it to the household is

$$\frac{p_j^h(\omega)\, y_j}{p_s^h(\omega)\, y_s} = \left(\frac{p_j^h(\omega)}{p_s^h(\omega)}\right)^{1-\varepsilon} = \left(\frac{c_j(\omega)}{c_s(\omega)}\right)^{1-\varepsilon}$$

Therefore, we calibrate $\frac{\psi_j}{\psi_s}$ using

$$\frac{\psi_j}{\psi_s} = \left(\frac{p_j^h(\omega)\, y_j}{p_s^h(\omega)\, y_s}\right)^{\frac{1}{\varepsilon-1}} \frac{c_j^{\mathcal{I}}(\omega)}{c_s^{\mathcal{I}}(\omega)},$$

where $c_j^{\mathcal{I}}(\omega)$ is the in-house unit cost of the core product, before TFP, given by Equation (30). Recall that Equation (29) identifies $\tilde{q}_{ij} = \psi_i q_{ij}/\delta_{ij}$ for the observed transactions between buyer j and supplier i. We use our measure of ψ_i to filter it out of $\tilde{q}_{ij}$, while the iceberg trade cost δ_{ij} is separated in the SMM (Appendix B.7.3). In effect, this means that we isolate the firm-specific productivity from the outsourcing capability, such that q_{ij} solely summarizes factors that are specifically related to outsourcing. Therefore, after having identified the observed $\tilde{q}_{ij}$ but before running the SMM algorithm, we identify the ψ_i using our measure of costs, and clear the outsourcing capabilities from their ψ_i. This reduces the importance of any supplier-specific component in the q_{ij} for aggregate objects, hence mitigates the impact of our within-supplier independence assumption on the q_{ij} calibration.

B.9 Demand shifter

To calibrate the product-specific demand shifter $\delta(\omega)$, we use the aggregate sales of product ω to the rest of the economy, the household in our model. In particular, firm j's domestic sales sales_j^R are equal to total sales minus exports and minus sales of intermediate inputs to other manufacturing firms. They are then aggregated at the product level across firms selling ω. For multi-product firms, we first split sales_j^R across their products using their product sale shares $s_{\hat{\omega}}$. In the model, we have that:

$$\frac{\text{sales}^R(\omega)}{PY} = \delta(\omega)\left(\frac{P}{P_\omega}\right)^{\eta-1}$$

Where PY aggregates $\text{sales}^R(\omega)$ across all manufacturing products and $P_\omega = \frac{\varepsilon}{\varepsilon-1}\left(\int_{j\in\mathcal{S}_\omega} c_j(\omega)^{1-\varepsilon} dj\right)^{\frac{1}{1-\varepsilon}}$.

$$\frac{\text{sales}^R(\omega')}{\text{sales}^R(\omega)} = \frac{\delta(\omega')}{\delta(\omega)}\left(\frac{P_\omega}{P_{\omega'}}\right)^{\eta-1} \quad \forall \omega \neq \omega'$$

In equilibrium, the household's demand implies $PY = Y^\alpha$, hence $P = (PY)^{\frac{\alpha-1}{\alpha}}$. Therefore, we retrieve $\delta(\omega)$ from

$$\frac{P^{1-\eta}}{\delta(\omega)} = \sum_{\omega'} \frac{\delta_{\omega'}}{\delta_\omega} P_{\omega'}^{1-\eta} \Rightarrow \delta(\omega) = (PY)^{\frac{(\alpha-1)(1-\eta)}{\alpha}} \left[\sum_{\omega'\neq\omega} \frac{\delta_{\omega'}}{\delta_\omega} P_{\omega'}^{1-\eta} + P_\omega^{1-\eta}\right]^{-1}$$

[107] For multi-product firms, we use their core product cost and weigh their total sales by the core product share.

B.10 Counterfactual analysis

B.10.1 Counterfactual sales

We need each firm's production cost relative to GDP, $c_j(\omega)x_j(\omega)/PY$, in the efficient allocation to compute the efficient aggregate energy intensity in (24). Since firms sell to each other at cost, this equals their sales to other firms plus their sales to the rest-of-the-economy net of the markup. We compute these ratios recursively, moving upstream.

1. Given unit costs and prices $p_i^h(\omega) = \frac{\varepsilon}{\varepsilon-1}c_i(\omega)$, we compute the product price indices $P_\omega = \left(\int_{i\in\mathcal{S}_\omega} p_i^h(\omega)^{1-\varepsilon}di\right)^{\frac{1}{1-\varepsilon}}$ and the aggregate price index $P = \left(\sum_{\omega\in\Omega}\delta_\omega P_\omega^{1-\eta}\right)^{\frac{1}{1-\eta}}$.

2. We initialize $\mathcal{H}$, the set of firms whose production costs we have computed, with the firms that sell only to the rest-of-the-economy sector. Their relative production costs are

$$\frac{c_i(\omega)x_i(\omega)}{PY} = \delta_\omega\left(\frac{\varepsilon-1}{\varepsilon}\right)^\varepsilon c_i(\omega)^{1-\varepsilon}P^{\varepsilon-1}\left(\frac{P}{P_\omega}\right)^{\eta-\varepsilon}$$

3. We move upstream to firms outside $\mathcal{H}$ whose buyers all belong to $\mathcal{H}$. Their relative production costs are

$$\begin{aligned}\frac{c_j(\omega)x_j(\omega)}{PY} &= \delta_\omega\left(\frac{\varepsilon-1}{\varepsilon}\right)^\varepsilon c_j(\omega)^{1-\varepsilon}P^{\varepsilon-1}\left(\frac{P}{P_\omega}\right)^{\eta-\varepsilon}\\ &\quad+\sum_{\hat\omega\in\Omega}\int_{i\in\mathcal{H}\cap\mathcal{S}_{\hat\omega}}\mathbf{1}\{j\in\mathcal{S}_i^{\mathcal{O}}\}\tilde\alpha_{\hat\omega}^{\omega}(i)\frac{c_i(\hat\omega)x_i(\hat\omega)}{PY}\left(\frac{\psi_i c_i(\hat\omega)}{c_j(\omega)/q_{ij}}\right)^{\frac{\gamma}{1-\gamma}}di\end{aligned}$$

   Where $\hat\omega$ is the core product of buyer $i\in\mathcal{H}$, $\frac{c_i(\hat\omega)x_i(\hat\omega)}{PY}$ its relative production cost, and $\mathcal{S}_i^{\mathcal{O}}$ its set of selected suppliers. The relative sales of j to buyer i thus add up, with their requirement weights $\tilde\alpha_{\hat\omega}^{\omega}(i)$, all the stages of i at which the input ω produced by j enters.

4. We add those firms to $\mathcal{H}$.

5. We start again at step 3 and repeat until $\mathcal{H}$ contains all our firms.

B.10.2 Energy Targets

The Deep Decarbonization Pathways Project constructs technology pathways for Italy consistent with an 80% reduction in CO_2 emissions by 2050 relative to 1990 (Gaeta et al., 2015). Its Energy Efficiency scenario requires aggregate energy intensity to fall from 3.89 to 1.60 MJ per dollar between 2010 and 2050. We rebase this target to the 2019 baseline of our model: Italian energy intensity fell by roughly 13% over 2010-2019, so reaching the 2050 level implies a further reduction of about 53% relative to 2019. The target is set for economy-wide energy intensity, not manufacturing alone. However, since our object of interest is the gap between the two tax levels, the calibration of the target mainly fixes the scale of the exercise, and the relative contribution of the extensive margin is largely invariant to it. If anything, applying the economy-wide target to manufacturing overstates the effort required of the sector:

in the DDPP's own scenarios industrial emissions fall by $50-55\%$ relative to 2010, against $90-95\%$ in residential and services, so our estimate should be read as an upper bound on the tax needed to meet manufacturing's share of the aggregate target.

We then compute the uniform carbon tax that delivers this reduction, first letting both margins adjust and then shutting down the extensive margin. Because the energy price is only set by the supply side in our model (energy demand does not feed back into it, and the wage is fixed), aggregate energy intensity is a function of the energy price, $EI(p_e)$, and each of its values requires solving one counterfactual equilibrium. Finding the tax that delivers the target is therefore a one-dimensional inversion rather than a fixed-point problem, so we solve the counterfactual once at each of 11 energy prices between the baseline pre-tax price $\bar{p}_e$ and $4\bar{p}_e$, including the Pigouvian price $\bar{p}_e + \Upsilon$. For each scenario, we interpolate $\ln EI$ linearly in $\ln p_e$ between adjacent grid points and read off the price p_e^* at which EI reaches $(1-X)$ times its baseline equilibrium value, where X is the target reduction. Energy intensity falls monotonically along the grid in every scenario, so each target is reached at a unique price. The tax is $p_e^* - \bar{p}_e$ per kWh, which we convert to a tax per tonne of CO_2 using the 2019 average emission factor of Italian electricity, 0.362 kg CO_2/kWh.

As with all our counterfactual results, this procedure assumes that the tax passes fully into a single energy price that is the same for all firms and that the price has no general-equilibrium feedback, thanks to our welfare specification. Moreover, it relies on energy intensity being log-linear between grid points. This is close to exact when the extensive margin is shut down, since the price elasticity of energy intensity then stays between -0.83 and -0.79 across the grid. With both margins active, discrete switches in firm boundaries and suppliers make the elasticity vary between -1.19 and -0.79, so taxes that fall between grid points are approximations. The resulting error is small: interpolating in levels instead of logs, or with a monotone cubic spline, changes every tax by less than 3.1 €/tCO_2 and the ratio of the two taxes by less than 0.01.

B.10.3 Non-linear damages

To illustrate how the extensive margin would affect a global carbon tax, we change the welfare function to

$$U(X, Y, E) = X + \frac{Y^\alpha}{\alpha} - \Upsilon \left[(1-\phi_d)E + \frac{\phi_d E^2}{2E_0} \right]$$

where E_0 is the baseline energy level and Υ is the social cost of carbon multiplied by the constant emission factor of energy, so that marginal damages at E_0 equal the social cost of carbon. In this exercise, we interpret proportional changes in E as proportional changes in global energy use. With damages quadratic in end-of-horizon temperature $T = T_0 + dE$, the elasticity of marginal damages with respect to energy at baseline is $\phi_d = dE_0/(T_0 + dE_0)$, the share of end-of-horizon warming caused by the energy use affected by the policy. The parameter d combines the emission factor with the transient climate response to cumulative emissions over a 25-year horizon, which corresponds to the medium-term horizon our static model represents. Using a best-estimate response of 0.45°C per 1000 $GtCO_2$ (Masson-

Delmotte et al., 2023) and global energy-related emissions of 33.3 $GtCO_2$ in 2019 (IEA, 2020), baseline energy use held constant over 2019–2044 adds $dE_0 \approx 0.37°C$. Extrapolating human-induced warming, 1.36°C in 2024 and rising by 0.27°C per decade (Forster et al., 2025), gives $T_0 + dE_0 \approx 1.9°C$ in 2044, hence $T_0 \approx 1.5°C$ and $\phi_d \approx 0.2$. Over the likely range of the climate response ($0.27 - 0.63°C$ per 1000 $GtCO_2$), ϕ_d lies between 0.12 and 0.28. The optimal energy price now depends on the counterfactual energy level E^*: $p_e^{SCC} = p_e + \Upsilon\left[(1-\phi_d) + \phi_d E^*/E_0\right]$, where $p_e = \chi w$ is still fixed by the energy production side. Since E^* itself depends on p_e^{SCC}, we solve the counterfactual allocation iteratively until convergence.

B.11 Out-of-sample fit

We run two out-of-sample tests using untargeted moments. First, we use our parameters calibrated to 2019 moments to predict the allocation under the 2021, 2022 and 2023 energy prices. Second, we remove the selected supplier from a buyer's choice set and ask the model to predict what the buyer does instead, then confront that prediction with a choice observed in a *different* year in which that supplier was indeed unavailable.

Predicting make-or-buy switches. We move each 2019 firm's energy price by its own observed change to 2021, 2022 or 2023 and compute, for each make-or-buy decision observed in 2019, the model's probability that it switches, integrating over the unobserved outsourcing capabilities of candidate suppliers (Table B.11). The model identifies which in-house inputs are more likely to be outsourced: those in the top decile of predicted probability are outsourced about twice as often as those in the bottom decile. Because energy prices are the only change we feed the model, it cannot anticipate switches driven by other shocks: even in the bottom decile of predicted probability, about 4% of the inputs it expects to remain in-house are outsourced anyway, and it predicts almost none of the insourcing observed in the data (Panel B). Since a more energy-efficient external alternative exists for most firms producing in-house, higher energy prices alone mostly trigger outsourcing (Section 4.2.3), and insourcing in the data must come from other shocks, which shift outsourcing capabilities. Our next exercise examines one such shock directly, the exit of a buyer's supplier, which sets the corresponding outsourcing capability to zero: there, the model closely matches the realized in-house share. Its failure to predict insourcing in the first exercise therefore reflects the shocks it is not fed, not an inability to generate insourcing.

Unavailable suppliers. We take, in our calibration year, every decision node (j, ω, i) that is outsourced at the equilibrium price, and we retain those whose incumbent domestic supplier i is absent from *all* of the raw datasets of our comparison year $C \in \{2021, 2022, 2023\}$. We interpret a domestic incumbent absent in year C as one that has exited: the test thus asks what the buyer did once the relationship was no longer available to it. For each such event we re-solve j's problem holding the baseline calibration fixed, with i deleted from the choice set: the next-best option is computed with the counterfactual's own decision machinery, with the same CES cost recursion, the same draws for unobserved

Table B.11: Predicted switching probability vs. observed switches

	2021	2022	2023
Panel A: Inputs made in-house in 2019, switching to outsourcing			
Realized switching rate (%)	5.43	5.67	5.59
Mean predicted probability (%)	4.65	7.05	9.07
Realized rate, bottom decile of predicted probability (%)	3.87	4.38	3.85
Realized rate, top decile of predicted probability (%)	9.42	8.13	8.62
Decisions	35,525	35,049	34,869
Panel B: Inputs outsourced in 2019, switching to in-house			
Realized switching rate (%)	13.94	15.83	17.77
Mean predicted probability (%)	0.91	0.67	1.07
Decisions	13,705	13,604	13,370

Notes: Each 2019 firm's energy price moves by its own observed change between 2019 and year Y. For a stratified random sample of about 60,000 direct make-or-buy decisions of 2019 per year (strata: 2019 status $\times$ buyer NACE-2 sector, proportional allocation), among decisions whose firm and input are still present in year Y, the predicted probability is the model's probability that the decision switches, integrating over the unobserved outsourcing capabilities of candidate suppliers. Rates are computed over the decisions that are still a direct make-or-buy decision in year Y (*Decisions*); year-Y status is constructed from the year-Y data as for 2019. Deciles are deciles of the positive predicted probabilities (fewer than 1.5% of in-house decisions have a zero probability); the 95% firm-cluster bootstrap confidence intervals of the bottom- and top-decile rates do not overlap in any year. For each decile, the realized rate is the share of its decisions that switch. Panel B reports no deciles because the model assigns a positive insourcing probability to only about 1% of outsourced inputs (126 to 237 decisions per year).

pairs, the same iceberg trade costs, and delivers a probability of producing ω in-house and a probability distribution over alternative suppliers. We then score these predictions against j's realized choice over years C and $C+1$ (year C only for $C = 2023$, the last year in our data), so that a buyer that runs down its inventory or replaces its supplier with a delay is not mistaken for one that produces in-house, and separately validate (i) the *decision* (make versus buy), for which the relevant choice set is the baseline pool minus the incumbent, and (ii) the *identity* of the replacement supplier, for which we restrict the pool to firms that exist in the year of the replacement.[108] Because the buyer is forced to re-select for a reason unrelated to energy prices, this exercise isolates the other determinants of supplier choice: comparing the model's predicted replacement to the realized one shows which characteristics drive selection outside our mechanism. Table B.12 contains the results. We report the 2021 results for completeness, but because it contains very few events, our discussion is based on 2022 and 2023.

The model performs well in levels the predicted and realized probability of internalizing ω match within one to five percentage point, and mispredict fewer than 40% of the events. The identity of the replacement supplier shows both the strength and the limit of the model. Outsourcing capabilities of non-selected pairs are drawn independently from the calibrated distribution, disciplined only by the within-buyer copula and the truncation that rationalizes the observed choice. Thus, in an exercise in which real

[108] We report as the headline the events for which the buyer accounts for less than half of the supplier's sales, because high-dependence pairs might reflect a specific vertical disintegration margin. Moreover, we similarly score a control group of links whose supplier survived: their realized persistence is the ceiling any predictor could reach, and their voluntary switching rate is the noise floor below which no hit rate is informative.

switching is unrelated to energy, the model often misses the exact identity of the supplier, conditional on outsourcing: the realized 2023 supplier is ranked first in 0.7% of cases, still 10 times better than under a uniform null but only 1.45 times better than the market-share null that predicts the largest seller of the input. The model does not pin down which firm the buyer turns to, but the distribution it places over candidates reveals where the information it carries comes from. Comparing the realized supplier to the model's own predictive distribution, it is neither systematically more nor less energy-efficient than the model implies (mean PIT of 0.50 in 2022, 0.47 in 2023, against 0.5 under correct calibration), and does not have systematically higher or lower outsourcing capabilities (both 0.49). It is, despite our inclusion of trade costs, still slighlty nearer (mean PIT of distance of 0.37 and 0.38).

B.12 Make-or-buy switching margins

Table B.13 reports switching and retention shares following a doubling of the energy price from $p_e = 0.083$ to $2 \times p_e = 0.166$, by quartile of the buyer firm's effective energy efficiency (ν) or the input energy requirement (β_ω). More energy efficient firms are less likely to switch any input (column 2) and switch a lower share of their inputs (column 3).[109] This is because most switches are outward: more energy efficient firms actually tend to bring more initially outsourced inputs in-house (column 6) but to outsource fewer initially in-house inputs (column 5), and the latter dominates the switching rate gradient. Accordingly, they retain more of their production in-house (column 4). Finally, inputs that require more energy (high β_ω) are more likely to switch (column 7), in both directions.

Table B.14 reports, for each initial sourcing mode, the share of pairs for which a cleaner sourcing option exists (that is, one whose cost rises less with the energy price) and the share of these potential switchers that actually switch when the energy price increases. Inputs produced in-house can reallocate only by outsourcing; outsourced inputs can reallocate along two routes, to a different supplier or back to in-house production. The first two rows therefore report the aggregate response of each initial mode (for outsourced inputs, a change of sourcing through either route, each pair counted once) and the last two rows decompose the outsourced margin by route. We also note that the price increase required for at least 50% of potential switchers to reallocate exceeds the highest plausible price we consider.

The shares in Table B.14 are predicted switches rather than counterfactual outcomes, and the denominator is the set of sourcing decisions each firm actually makes in equilibrium. We take each of these decisions separately and ask, holding everything else at its equilibrium value, how large an energy price increase it would take for a cleaner alternative to become cheaper than the option currently in use; the decision counts as a switch if that threshold lies below the counterfactual price. This differs from the full counterfactual in two respects. First, it is a first-order comparison that holds the rest of the network

[109] The larger gradient at input-level shares, relative to switching firms, is partly mechanical from the direction of the switches: when a firm starts to outsource an input, the stages upstream of it are left to the supplier. If they were initially performed in-house, the corresponding inputs become indirectly outsourced, increasing the share of switched inputs.

Table B.12: Predicted versus realized choice when the incumbent is unavailable

	Comparison year C		
Baseline $B = 2019$	2021	2022	2023
Panel A. Sample			
Outsourced links at the equilibrium price	166,027	166,027	166,027
of which incumbent absent in C (events)	165	594	1,116
of which classifiable	122	462	824
Realized in-house share	0.31	0.27	0.29
Panel B. Make-or-buy margin			
Mean predicted P(in-house)	0.23	0.26	0.25
Share of mispredicted events	0.30	0.33	0.38
Panel C. Identity of the replacement supplier			
Buy events	84	339	579
Top-1 hit rate	0.000	0.003	0.007
95% CI (buyer-clustered)	[0.000, 0.000]	[0.000, 0.009]	[0.002, 0.014]
uniform null	0.0007	0.0006	0.0007
market-share null	0.0033	0.0044	0.0048
Top-decile hit rate	0.32	0.28	0.30
uniform null	0.07	0.08	0.08
market-share null	0.29	0.24	0.28
Panel D. Predictive calibration (PIT) of the realized supplier			
Mean PIT, q_{ij}	0.49	0.49	0.49
Mean PIT, distance to buyer	0.44	0.37	0.38
Mean PIT, ν_i	0.51	0.50	0.47

Notes: Events are equilibrium-outsourced decision nodes of the baseline year whose incumbent supplier is absent from every raw dataset of year C. Reported predictions are the resulting choice frequencies across 100 simulations, evaluated with the counterfactual's exact decision rule. Panel A: classifiable events are those for which the buyer is in C and ω is in its C recipe. For $C = 2021$ and 2022, outcomes follow a two-year rule over C and $C + 1$: an event is classifiable if the buyer sourced ω in C or $C + 1$, or bought it in neither year while observed in both, with ω inferred in-house in at least one of them; replacements first bought in $C + 1$ are scored against the candidates present in $C + 1$. 2023 uses the one-year rule, as 2024 is not yet available. Panel B scores the make-or-buy margin on the pool from which only the incumbent is removed. The mean predicted P(in-house) is compared with the realized in-house share of Panel A and checks the aggregate level of insourcing. An event is mispredicted when the model puts the probability of in-house above one half and the buyer outsourced, or below one half and the buyer made it in-house. Panel C conditions on outsourcing: it scores the identity of the replacement on the classifiable events whose buyer sourced ω from a supplier in C or, for 2021 and 2022, in $C + 1$ (84, 339 and 579 events). The top-decile rate asks whether the realized supplier is among the $\lceil 0.1|\text{pool}| \rceil$ candidates with the highest choice probability; its nulls are the uniform rate $k/|\text{pool}|$ and the share of events whose realized supplier is among the k largest sellers of ω in the pool by transaction value. Panel D is the probability integral transform (PIT) of the realized choice. The draws put a probability π_i on every candidate i in the pool, which induces a predictive distribution over any supplier characteristic x, $\hat{F}(v) = \sum_i \pi_i \mathbf{1}\{x_i \leq v\}$; the PIT of an event is $\hat{F}$ evaluated at the realized supplier. Under a correctly calibrated predictive distribution the PIT values are uniform on $[0, 1]$: mean 0.5 and a flat decile histogram, so the maximum absolute deviation of a decile share from 0.1 is a shape statistic, and a mean below 0.5 says the realized supplier is systematically lower in that characteristic than the model's weight implies.

at its equilibrium allocation,[110] so it misses the switches that become attractive only once other firms along the chain have themselves reorganized, and counts some that cease to be attractive once they have. Second, it scores every decision as if the firm still faced it at the counterfactual price, whereas in the

[110] The first-order approximation still tracks the full counterfactual. When the energy price doubles, it predicts that 7.5% of initially in-house nodes and 1.0% of initially outsourced nodes reverse their own make-or-buy decision, and that 5.0% of outsourced nodes switch supplier; the full counterfactual produces 7.5%, 0.8% and 4.4%.

Table B.13: Share of switching firms and in-house retention (%), by ν_j and β_ω quartile

	By ν_j quartile					By β_ω quartile
Quartile	Switching firms (%)	Switching pairs (%)	Kept in-house (%)	In → Out (%)	Out → In (%)	Switching pairs (%)
Q1	34.41	44.24	47.31	5.93	0.06	11.27
Q2	24.37	16.66	80.53	2.96	0.13	18.08
Q3	20.14	7.20	92.13	1.31	0.17	19.20
Q4	14.33	1.97	98.43	0.34	0.27	22.90

Notes: The table compares the equilibrium at $p_e = 0.083$ with the counterfactual at $2 \times p_e = 0.166$, by quartile of the buyer's effective energy efficiency ν_j (within its core product) or of the input energy requirement β_ω. Pairs are the buyer×input combinations that the buyer produces in-house or buys directly at either price; inputs embedded in a purchased input are excluded. The second column reports the share of firms, among those outsourcing at least one input at the initial price, whose set of suppliers of directly outsourced inputs changes – because an input is brought in-house, newly outsourced, or bought from a different supplier. The third, fifth, sixth and seventh columns are shares of all pairs in the quartile: the fifth and sixth report pairs moving from in-house to outsourced and from outsourced to in-house, and the third and seventh report all pairs whose make-or-buy status changes, which also includes pairs that leave or enter the set because an input they feed into, directly or indirectly, changes status. Supplier changes on inputs that stay outsourced are not counted. The fourth column is instead a share of the pairs produced in-house at the initial price; inputs not kept in-house are either outsourced or left to the supplier of a newly outsourced input downstream.

model a firm that starts buying an input also stops choosing how to source everything upstream of it, which now comes embodied in what it buys.

Table B.14: Make-or-buy switching margins: switchable population and response to a carbon price

	Cleaner alternative		Flips at $1.1 \times p_e$		Flips at $1.83 \times p_e$		Flips at $2 \times p_e$	
Margin flips at	exists (1)	Second best	of all nodes	of (1)	of all nodes	of (1)	of all nodes	of (1)
Initially in-house	0.978	0.303	2.35%	2.40%	6.94%	7.09%	7.46%	7.63%
Initially outsourced	0.998	0.666	0.89%	0.89%	5.01%	5.02%	5.73%	5.74%
via supplier switch only	0.998	0.445	0.89%	0.89%	4.43%	4.44%	4.95%	4.96%
via re-insourcing only	0.519	0.519	0.00%	0.00%	0.71%	1.36%	0.96%	1.85%

Notes: First-order break-even $\Delta \ln p_e^\star = (\ln c_{alt} - \ln c_{chosen})/(s_{chosen} - s_{alt})$ per decision node, where s_k is the embodied energy cost share of option k and the alternative is the cleaner candidate that overtakes the option in use first. "Cleaner alternative exists" is the share of the margin's nodes for which at least one candidate producer of the input embodies a lower energy cost share than the option the buyer currently uses. Next-best is the case in which the cheapest non-selected option available at current prices is cleaner; on the re-insourcing margin the only alternative is the buyer's own production, so the two columns coincide by construction. The six "flips at" columns give the predicted flip share at three energy-price increases of the price grid, 1.1×, 1.83× and 2× the equilibrium energy price $p_e = 0.083$, first as a share of all the margin's nodes and then as a fraction of those satisfying (1). The "Initially outsourced" row is the pair-level union of the two routes below it: condition (1) holds if a cleaner alternative of *any* kind exists, and a pair flips if either route overtakes the option in use, counted once; the route rows evaluate each alternative separately against the same population, so their entries need not sum to the aggregate. On every margin, the energy price at which the median node satisfying (1) would flip lies above $4 \times p_e$, the highest price of our grid: half of the potential switchers would not respond within any plausible policy range.

B.12.1 Destination of reallocation

Table B.15 checks whether reallocation concentrates demand on a few suppliers, the case in which capacity constraints, absent from the model, could bind. When increasing the energy price to p_e^{SCC} (+83%), reallocation spreads demand over more suppliers instead of concentrating it: the number of active suppliers rises along with the number of outsourced inputs, and concentration falls within inputs and is unchanged across suppliers (Panel A). Buyers that change supplier do not converge on the same producer: only a small share of an input's movers choose its most popular destination, and about a third

go to suppliers that had fewer than two buyers before (Panel B). Large gains are rare and accrue mostly to large firms. Concentration rises for only a handful of inputs, and a higher concentration does not imply that a supplier absorbs much more demand: for instance, for pulp, concentration rises as some buyers stop outsourcing it, but its main supplier keeps the same 20 buyers. A genuine pile-up is the exception: one agrochemicals producer goes from 20 to 108 buyers, becoming the input's main supplier.

Table B.15: Supplier concentration with and without reallocation

	No reallocation	Reallocation
Panel A: Allocation of outsourced inputs to suppliers		
Outsourced buyer×input pairs	166,027	175,109
Active suppliers	36,232	40,270
Buyers per supplier, maximum	585	582
Buyers per supplier, Gini	0.60	0.59
Within-input HHI, pair-weighted mean	155	133
Inputs with HHI $>$ 2,500 (%)	8.9	5.3
Panel B: Reallocation relative to no reallocation		
Movers (buyer×input pairs)		27,429
Movers to the input's most-chosen supplier (%)		6.0
Effective number of destinations per input		245
Movers to suppliers with $<$ 2 buyers before (%)		34.8
Suppliers losing all buyers (%)		4.3
Continuing suppliers with $\geq 5\times$ more buyers (%)		1.0
Net buyer inflow captured by the top 1% gainers (%)		15.6
Sales percentile of suppliers gaining $\geq$ 10 buyers, median		93

Notes: Allocation of outsourced buyer×input pairs to suppliers following an increase in energy price from $p_e = 0.083$ to 0.152. *No reallocation* holds firm boundaries and suppliers at their equilibrium, which reproduces the observed allocation; *Reallocation* lets make-or-buy and supplier choices adjust. Counts are of buyer×input pairs, not quantities. HHI is computed over suppliers within each input (0–10,000). Movers are pairs whose supplier differs from the no-reallocation allocation (supplier switches and newly outsourced inputs); the most-chosen-supplier share and the effective number of destinations (inverse HHI of movers' suppliers) are pair-weighted means across inputs. Continuing suppliers have buyers in both allocations. Sales percentiles are among sample firms, at equilibrium.

B.13 Extensions and robustness exercises

B.13.1 ν-only counterfactual

Our baseline treats the outsourcing capabilities as a technological wedge, so that the equilibrium allocation maximizes output and only welfare is distorted. In this appendix we take the opposite view and shut down every source of producer heterogeneity except effective energy efficiency: we set $q_{ij} = 1$ for all pairs *and* $\psi_j = 1$ for all firms, and re-solve the entire allocation with the recursive algorithm of Section 4.3, at different energy prices. Because ν is then the only primitive that differs across producers, every potential producer faces the same input prices, and a producer with a higher ν can replicate any sourcing

plan of a less efficient one at a lower cost, since both its own stage and every input it could produce in-house are cheaper. The highest-ν potential producer of each input is therefore selected at every energy price, and its embedded supply chain is the most energy-efficient.

The ν-only allocation. Almost nothing of the observed allocation survives. We compare the two allocations at each of the 807,756 equilibrium make-or-buy decisions (buyer–input pairs): at 98.3% of them, the input is produced by a different firm in the ν-only economy. Of these changes, 79.2% are inputs made in-house in equilibrium that the ν-only economy outsources, 20.7% are inputs bought in both allocations but from a different supplier, and only 708 go from bought to made. This is the model's logic taken to its limit: downstream firms face a large pool of external producers, so absent pairwise heterogeneity they can rarely do better internally. Node counts overstate how many decisions firms revise, because buying an input also hands over every input below it in the production tree: the ν-only economy has only 185,558 make-or-buy decisions left, 6,579 of them in-house. Counting instead the inputs that enter or leave a firm's set of directly purchased inputs (supplier changes excluded), we observe that the adjustment is diffuse rather than concentrated: 62.6% of firms revise their boundary, by 1.73 inputs on average, and the top decile of firms accounts for only 37.6% of these revisions.

The ν-only allocation is, moreover, exactly identical at the energy prices we consider: since the highest-ν producer is the cheapest at every price, no price change can induce a switch, the multi-stage counterpart of Lemma 2. All the ν-only energy intensity price response that remains in Table B.16 is intensive-margin.

Table B.16: Aggregate energy intensity under the ν-only allocation

p_e (€/kWh)	$EI_{\nu\text{-only}}$	$EI_{\nu\text{-only}}/EI_{eq}$
0.1660	0.4611	0.6199
0.0830 (p_e^{eq})	0.8439	0.6223
0.0415	1.5544	0.6197

Notes: Aggregate energy intensity in kWh per euro of manufacturing output, in the ν-only scenario and in the baseline. The benchmark is the equilibrium allocation evaluated at the corresponding energy price (EI_{eq}).

Allocating each input to its most energy-efficient embedded producer would reduce aggregate energy intensity by 41 to 42 percent, at every price. We read this as an upper bound on what producer selection can deliver: it is the allocation a planner would choose if energy efficiency were the only thing that mattered, and the gap to it is the room that the outsourcing capabilities and TFP occupy.

The implied wedge. For each equilibrium node (j, ω) we define

$$wd_{j\omega} \;=\; \frac{c_i}{c_s} \;\geq\; 1,$$

where i is the equilibrium producer, s the ν-only producer, and both costs are evaluated in the ν-only economy. wd is the production-cost penalty that the relationship capability q_{ij} had to clear for j to select i over the more energy efficient s; absent any technological difference between producers, it is a pure

wedge. The median advantage is modest - 1.062 for inputs internalized in equilibrium, 1.074 for supplier switches, and 1.073 for the few inputs the ν-only economy brings in-house - but the distribution has a very fat right tail. It is larger for inputs covered by the EU ETS (with a geometric mean of 1.42, against 1.15 for unregulated inputs), consistent with differences in energy efficiency weighing more heavily for inputs with high energy requirements.

Table B.17 relates $\log(wd)$ to the difference in observables of both suppliers $\Delta x \equiv x_i - x_s$ (meaning that the regression sample only contains flips in which both the equilibrium producer and the ν-only producer are external suppliers). The wedge is systematically larger when the equilibrium producer is foreign rather than domestic, when it has larger sales than the ν-only producer and, surprisingly, when it is located farther from the buyer, while it is smaller when the buyer has a longer common history with the equilibrium producer than with the ν-only producer. These are the observables that the firm-boundaries literature associates with make-or-buy frictions, so the wedge is not orthogonal noise. At the same time, the specification leaves more than 95% of the variation unexplained, which we read as consistent with our treatment of q_{ij} as a largely idiosyncratic, relationship-specific technological wedge rather than a systematic function of observables.

Table B.17: The ν-only wedge and relationship observables

	$\log(wd)$
Δ foreign	0.3109*** (0.0111)
Δ log size	0.0713*** (0.0019)
Δ overlap years	-0.0047*** (0.0002)
Δ geo km	1.5×10^{-4}*** (4.35×10^{-6})
N / R^2	164,573 / 0.045

Notes: HC1 standard errors in parentheses; *** $p < 0.01$. The dependent variable is $\log(wd) = \log(c_i/c_s)$, the log cost penalty of the equilibrium producer i relative to the ν-only producer s, both evaluated in the ν-only economy. The sample is restricted to flips between two external producers, so that every characteristic is observed for both and enters in difference, $\Delta x \equiv x_i - x_s$: *geo km* is the haversine buyer-producer distance, *overlap years* the number of overlapping years of activity between buyer and producer, *foreign* an indicator for a foreign producer, and log size the producer's log sales. The specification includes buyer NACE2 fixed effects. Coefficients are descriptive correlations, not causal effects.

B.13.2 Sensitivity Analysis

Table B.18 reports the counterfactual changes in manufacturing energy intensity following a doubling in energy prices at different values of the lognormal parameters (μ, σ). All other parameters, including the Copula parameters (ρ_c, τ_c) are kept fixed at their baseline calibration value (Table 1).[111] In the last column, we report the share of inputs whose 2019 supplier is dropped in 2023 (the validation exercise described in Section B.11) and that the model assigns to be produced in-house. The corresponding data share is 0.29, to which our baseline calibration is close (Table B.12). This share highlights the mechanism: at high (low) μ, in-house production never (always) outperforms external suppliers, and is

[111] Intensive margin moves across exercises - even though the allocation is held fixed and unobserved q are thus never selected - because the posterior adjustment of observed $\log q_{ij}$ depends on the calibrated distribution.

therefore never (always) the next-best alternative.

Table B.18: Results sensitivity to different parameters

	100% increase in p_e				Supplier dropped in 2023
	$\%\frac{\Delta(E/PY)}{E/PY}$				
	Both margins	Intensive margin only	Extensive margin contribution (%)	Share of out→in flips (%)	**P**(in-house) (%)
$\mu = -5.2,\ \sigma = 2.24$ (base)	−49.57	−42.37	19.5	5.64	25.4
$\mu = 0,\ \sigma = 2.24$	−53.87	−43.45	26.30	0.62	0.65
$\mu = -13,\ \sigma = 2.24$	−44.86	−41.51	9.90	95.07	99.95
$\mu = -5.2,\ \sigma = 5$	−52.07	−43.52	22.32	1.08	1.15
$\mu = -5.2,\ \sigma = 0.5$	−48.77	−40.48	22.45	84.06	96.72

Notes: The table reports the counterfactual results at different lognormal parameters (μ, σ). All values are in percentages. Columns two to five report the counterfactual changes following a doubling in energy price. Column two and three report the counterfactual changes in manufacturing energy intensity, which is measured as the total energy used by the manufacturing firms in the sample (in quantity) relative to their total output sold to the domestic household. The second column allows for all margins of reallocation. The third column keeps the cross-firm allocation fixed. The fourth column shows the contribution of the extensive margin. The fifth column reports the share of pairs that flip from outsourcing to in-house production, out of all vertically switching pairs. The last column reports the predicted share of inputs whose 2019 supplier is dropped in 2023 that the model assigns to be produced in-house by the model. The initial energy price is $p_e = 0.0830$€/kWh.

The aggregate results respond only mildly to the Copula within-buyer dependence ρ_c. Because the within-buyer component F_j is negative on average, a higher ρ_c makes external suppliers less competitive relative to in-house production, dampening the outward reallocation through which the extensive margin mostly operates. This is amplified among the firms most prone to switch, those producing the most in-house, since the calibration assigns them a lower F_j to rationalize their large in-house share. The pairs that do switch, however, overcome a larger penalty and thus exhibit higher energy efficiency gains. The switching frequency falls while its average impact rises, leaving aggregate energy intensity and welfare mostly unresponsive. In any case, all our calibration exercises agree on a relatively low ρ_c.

B.13.3 Wholesalers as intermediaries

In our baseline analysis, we exclude transactions where the seller is a wholesaler, as these do not represent direct production relationships. Wholesale firms can act as intermediaries, effectively buying inputs from upstream manufacturing firms and selling them to downstream manufacturing firms (Manova et al., 2025). Borin et al. (2025) highlight their importance in connecting the domestic Italian production network to foreign trade. Therefore, we will miss some potentially relevant input-product linkages and firm-to-firm transactions when excluding them from our sample. As a robustness check, we incorporate wholesaler transactions by tracing them back to their manufacturing suppliers. This robustness exercise is still work-in-progress, hence results are not included in this section.

We identify wholesalers using the ATECO classification codes corresponding to wholesale trade sectors (codes 46.XX). For each wholesale code, we manually map it to the closest corresponding manufacturing product category. We choose the most disaggregated category and drop the wholesale code if we cannot find a sufficiently good match.[112] We then verify that a firm classified as a wholesaler purchases inputs

[112] For instance, we match the code 46.34.20, "Wholesale of non-alcoholic beverages", to the manufacturing code 11.07.00, "Manufacture of soft drinks and mineral waters", or the code 46.44.40, "Wholesale of cutlery, flatware and cookware" to the manufacturing code 25.71.00, "Manufacture of cutlery and flatware". In contrast, we drop the residual code 46.90.00, "Non-specialized wholesale trade".

from the matched manufacturing sector, ensuring that the wholesaler genuinely intermediates in that product market. Specifically, we retain a wholesaler w if:

(i) The firm's product code corresponds to a wholesale trade category; and

(ii) The wholesaler purchases at least one input whose code exactly matches the corresponding manufacturing code (for codes reported at 6-digit level, we allow for 5-digit prefix matches).

For each retained wholesaler, we identify its main supplier as the manufacturing firm from which it sources the highest transaction value. When a wholesaler w appears as a seller to buyer j for input ω, we replace the transaction as follows: the seller identity is substituted with w's main manufacturing supplier i, and the input code is updated to reflect the corresponding manufacturing product. This procedure effectively "sees through" the intermediary, attributing the transaction to the underlying producer.

We exclude wholesaler-intermediated transactions when constructing product recipes and when calibrating the parameters governing the distribution of outsourcing capabilities q_{ij}. This exclusion reduces noise, as wholesaler transactions do not directly reveal firm-level outsourcing capabilities, but also capture wholesaler-specific characteristics. Since we use the main supplier's identity as the effective seller, we employ that supplier's price $c_i(\omega)$ when computing the buyer's input cost recursively. Consequently, the outsourcing capability q_{ij} associated with a wholesaler-intermediated transaction absorbs all wholesaler-specific costs and markups. This approach is valid for our counterfactual analysis insofar as these intermediation costs and markups are not affected by the change in energy prices (conditional on the wholesaler's suppliers' price), since the q_{ij} are held fixed in the counterfactual. Because wholesalers' internal operating costs are less dependent on energy, relative to manufacturing firms, this assumption is not too restrictive.

B.13.4 Estimating the elasticity of substitution between intermediate inputs

We verify our choice for the externally calibrated value for the elasticity of substitution between intermediate inputs $\frac{1}{1-\gamma}$, by estimating it using our detailed customs data. Customs data contains both expenditures and quantities of inputs purchased at the quarterly level, from which we back out unit values. Therefore, assuming that the firm boundaries do not change at such frequency, we have that

$$\Delta \log \frac{p_{ij}(\omega)\, x_{ij}(\omega)}{p_{sj}(\hat{\omega})\, x_{sj}(\hat{\omega})} = \frac{-\gamma}{1-\gamma} \Delta \log \frac{p_{ij}(\omega)}{p_{sj}(\hat{\omega})} + \frac{\gamma}{1-\gamma} \underbrace{\Delta \log \frac{q_{ij}}{q_{sj}}}_{=0}$$

Where $\frac{p_{ij}(\omega)}{p_{sj}(\hat{\omega})}$ is the relative price of imported inputs ω and $\hat{\omega}$, and $\frac{q_{ij}}{q_{sj}}$ is the ratio of buyer j's outsourcing capability with the supplier i of input ω to its capability with the supplier s of input $\hat{\omega}$. This regression correctly identifies γ as long as there is no within-firm variation in relative outsourcing capabilities $\frac{q_{ij}}{q_{sj}}$ at this frequency. We run it on all imported manufacturing intermediate inputs, excluding energy imports, pooling 2019 and 2021-2023, taking year-on-year differences and absorbing the reference input $\hat{\omega}$ with buyer-quarter fixed effects. Because prices are measured as unit values, classical measurement error in

quantities attenuates the OLS estimate of the elasticity, which is 0.81; we therefore instrument a firm's own price change with the contemporaneous change in the price of the same input paid by all other importers. The instrument is strong (first-stage $F = 195$) and yields an estimated elasticity of 0.14 (s.e. 0.10), statistically indistinguishable from our calibrated value of $1/(1-\gamma) = 0.2$, which we retain. The estimate is stable across specifications, ranging from 0.12 to 0.20 whether inputs are defined at the CN 8-digit or at the ATECO 4-digit level, whether or not we pool source countries, and whether or not we trim extreme unit-value changes.