Inside Out: Firms' Make-or-Buy Decisions and Aggregate Energy Intensity

Carole Marullaz *
University of Zurich

Giacomo Romanini†
Bank of Italy

November 19, 2025

Preliminary and incomplete - Please do not circulate

Abstract

We study how firms' make-or-buy decisions shape aggregate energy demand and emissions. The production of intermediate inputs requires energy, and firms are highly heterogeneous in their energy efficiency. Therefore, which firm is using energy matters for aggregate energy consumption, and whether firms decide to make or buy their intermediate inputs effectively allocates energy demand between them and their potential suppliers. We first theoretically establish that the lower the energy prices, the less consideration firms give to the energy component in their make-or-buy decision, prioritizing other factors. The resulting allocation increases aggregate energy relative to output. Because energy is not priced at its social cost due to the emission externality, the equilibrium allocation is inefficient and reduces welfare. Crucially, we show that a key moment, the joint distribution between firms' effective energy efficiency and their outsourcing capability, determines the importance of misallocation. However, firms' effective energy efficiency can only be identified when accounting for firm boundaries. We leverage detailed Italian firm-to-firm transaction data and develop and estimate a quantitative model of firms' make-or-buy decisions embedded in a production network to recover this novel moment and quantify the extent of energy misallocation in the Italian manufacturing sector and the welfare loss that it induces. The calibration is still work in progress, thus this draft does not contain any empirical results.

*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 the Bank of Italy.

1 Introduction

When energy prices fail to reflect the social cost of emissions, they distort firms' choices and generate misallocation. At the within-firm level, distorted prices tilt cost minimization toward excessive energy use relative to other inputs—even with homogeneous producers. However, recent empirical evidence has documented substantial heterogeneity in energy efficiency levels across otherwise comparable firms [Petrick et al., 2011, Lyubich et al., 2018, Capelle et al., 2023]. It is well understood that, with heterogeneous producers, distortions lead to an inefficient allocation of market share across firms producing the same product. In this paper, we argue that a third, overlooked margin amplifies the resulting inefficiencies: misallocation over the production of the same *input* also arises across downstream and upstream firms, that is, firms selling different *products*. Because make-or-buy decisions imply that firm boundaries effectively span multiple stages of production [Antràs, 2020, Chor et al., 2021, Boehm and Oberfield, 2023], heterogeneity in energy efficiency matters not only due to the allocation of energy across firms to produce a given product, but also due to the allocation of energy across firms to produce a given input. Since producing intermediates requires energy, when a firm outsources the production of an input, it effectively shifts its energy demand to its supplier. In other words, the make-or-buy decision allocates which firm uses energy to produce a given input. To evaluate the full misallocation induced by underpriced energy - and the corresponding potential gains from carbon pricing - we must account jointly for input allocation within firms, the cross-firm, within-product allocation among producers of the same good, and the cross-firm, within-input allocation created by idiosyncratic firm boundaries. Understanding these cross-firm allocative margins requires answering a fundamental question: which types of firms choose to produce inputs in-house, which choose to outsource, and to whom? Despite a rich literature on vertical integration and firm boundaries (e.g., [Grossman and Helpman, 2002, Antràs and Rossi-Hansberg, 2009, Boehm and Oberfield, 2020, Antràs, 2020, Fort, 2023]), we still lack empirical evidence on the joint distribution of energy efficiency and outsourcing capabilities[1] across firms, and on the relative importance of these factors in shaping the make-or-buy decision. This paper takes a first step toward filling this gap. First, it develops a simple theoretical framework to understand their respective contribution to the make-or-buy decision. Second, it leverages a comprehensive firm-to-firm transaction dataset covering the entire Italian manufacturing sector to recover their empirical distribution, and uses it to quantify the extent of misallocation induced by distorted energy prices and, in particular, by the make-or-buy decision of the firms.

Whether energy-intensive intermediate inputs are produced in-house or externally is a significant determinant of firm-level energy demand.[2] As illustrated in figure 1, the more upstream an intermediate input, the more energy intensive

[1] Make-or-buy decisions are based on multiple factors, such as specialization gains, distance to suppliers, necessity of the sourced inputs, or contract enforcement considerations, which tend to be highly idiosyncratic [Monteverde and Teece, 1982, Alfaro and Charlton, 2009, Lafontaine and Slade, 2013, Alfaro et al., 2019, Garg et al., 2019, Atalay et al., 2019, Antràs, 2020, Helper and Munasib, 2022, Boudreau et al., 2023, Grossi et al., 2023]. We group any factor unrelated to energy that influences the make-or-buy decision of the firms under the broad term outsourcing capability.

[2] In-house production is an important margin. For instance, [Garg et al., 2019] finds that Indian firms that produce at least one of their intermediate input in-house account for more than three-quarters of economic activity. [Rachapalli, 2021] also documents that multi-stage firms dominate Indian manufacturing production. Worldwide, within-firm trade represents more than a third of all global transactions [Antràs and Rossi-Hansberg, 2009].

its production [Shapiro, 2021]: this means that its production, if performed in-house, will account for a large share of the firm-level energy use.[3] We indeed find a robust and strong negative correlation between the level of outsourcing and the energy used within a firm's boundary in our data. This has two important implications: first, identifying the effective energy efficiency of a firm based on its energy demand requires knowing its boundary, that is, all the inputs it is producing in-house. Second, the make-or-buy decisions of the firms effectively allocate energy demand across firms. Consequently, when firms are heterogeneous in how efficient they are at processing energy, this allocation, that is, which firm produces its intermediate inputs, and which firm outsources them, and to which suppliers, becomes particularly relevant for aggregate energy use. We indeed observe that, within narrowly defined sectors, Italian firms exhibit significant heterogeneity in energy intensity, even after controlling for a large set of observable characteristics. The dispersion is larger than for other inputs' intensity, and this pattern has been robustly shown for different countries [Petrick et al., 2011, Lyubich et al., 2018, De Lyon and Dechezleprêtre, 2025], and often related to important differences in energy-biased technology [Hawkins-Pierot and Wagner, 2023, Carradori et al., 2023, Lanteri and Rampini, 2023, Capelle et al., 2023]. Therefore, the allocation of energy use across heterogeneously energy-efficient firms induced by their make-or-buy decisions is crucial for the aggregate energy used in production.

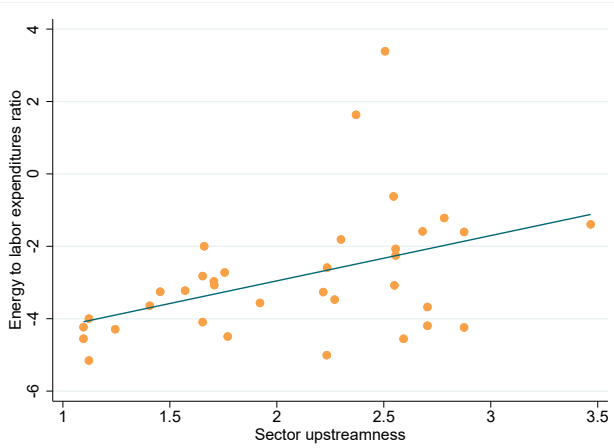


Figure 1: Correlation between the sectoral energy to labor expenditure ratio and the sector's own upstreamness

Notes: Source: Eurostat Italian SUT (2019). The graph plots the relation between the average sectoral energy expenditures to labor expenditures ratio (in log) and the sectoral upstreamness index of [Antràs et al., 2012].

Building on these observations, we begin with a simple model of make-or-buy decisions in which a downstream firm decides between outsourcing the production of its intermediate input or keeping the production in-house based on different firm-specific and relationship-specific factors, including its own and its supplier's energy efficiency. The model reveals that the lower energy prices, the less consideration the downstream firm gives to the energy component in its make-or-buy decision, relative to the other factors (that we summarize into the broad concept of outsourcing capabilities). Thus, when the price of energy is lower than its social cost, the equilibrium make-or-buy decision of the firm is inefficient, and leads

[3] In contrast, labor will primarily be used for the final task undertaken by the firm, meaning that intermediate input outsourcing is less relevant for the firm-level labor demand.

to misallocation across firms.[4] In turn, this misallocation increases aggregate energy intensity and decreases welfare, relative to the social planner's allocation. Finally, the model shows that the extent of misallocation depends on the joint distribution of firms' energy-biased technologies and outsourcing capabilities. To assess the empirical importance of this margin, we quantify the extent of misallocation in the Italian economy. We embed our model of make-or-buy decisions into a full production network, which guides our identification of the key firm-specific parameters: the firms' energy-biased technology and their outsourcing capability. Because this identification requires knowing the set of inputs that the firms produce in-house, which is unobserved, we develop a novel method to infer it based on firms' transactions. Equipped with these measures, we use our model to simulate a counterfactual allocation set under an efficient energy price, which gives us a benchmark to assess the magnitude and the impact of the distortions in the current network structure of the Italian manufacturing sector when accounting for energy. We then use our model to run additional counterfactual analysis that helps us characterize the effects of different carbon and industrial policy schedules, accounting for these different channels of adjustment. The calibration is still work in progress; therefore, this draft does not contain any quantitative results.

To characterize the aggregate energy intensity of the Italian manufacturing sector relative to its socially efficient benchmark, we need to identify the unobserved factors that drive the make-or-buy decisions of each firm. We develop a theoretical static model of production organized in supply chains. Firms differ in terms of their energy-biased technology, but also in their outsourcing capabilities. Together, those characteristics drive the firm-level make-or-buy decision, which in turn endogenously shapes the production network. The allocation of energy use across firms matters insofar as firms are heterogeneous in their effective energy efficiency. Quantifying misallocation thus crucially relies on accurately identifying this parameter for each firm. However, as our model demonstrates, this is only feasible if we know the firm boundaries, that is, the set of inputs that is produced in-house. Indeed, energy consumption is observed at the firm-level, therefore energy intensity is determined both by the effective energy-biased technology of the firm, but also by the set of inputs that it produces in-house.[5] However, this set is generally unobserved. We develop a novel method to identify it based on firms' transactions. Assuming a sequential production process, we can infer that all the direct and indirect inputs required in the production of a firm's product (its "recipe") that are located downstream to the inputs that the firm purchases must have been produced in-house. We construct the "recipe" of a product by connecting every product in our sample to the inputs that enter its production directly or indirectly, using our firm-to-firm transaction data.[6] This in-house set allows us to accurately retrieve the energy-biased technology of a firm, based on its total energy demand. Finally, we identify its observed and unobserved relationship-specific outsourcing capabilities, based on a revealed preference argument. Intuitively, if we observe that a firm produces an input in-house, its pairwise outsourcing capabilities with all the potential

[4] In other words, when energy is under-priced, there is misallocation both within and across firms. Within firms, at the intensive margin, firms use too much energy relative to other inputs. Moreover, across firms, at the extensive margin, firms underweight energy in their make-or-buy decision. This misallocation across firms increases total energy use when firms are heterogeneous in their energy-biased technology.

[5] For a given energy-biased technology, a firm will appear relatively more energy intensive, in the sense that it uses more energy per unit of output, if it produces more inputs in-house, since those intermediate input production requires energy. Estimates that don't account for in-house production (e.g., [Carraro and De Cian, 2013, Hawkins-Pierot and Wagner, 2023]) confound the energy-biased technology with the degree of outsourcing.

[6] This is the firm-level equivalent to the total requirements coefficient in standard input-output tables (e.g., [Alfaro et al., 2019])

suppliers of this input must be relatively low. Those firm-level fundamentals are then used to characterize the efficient allocation of production that would arise under the social cost of energy, informing us on the extent of misallocation in the Italian economy. A key feature of our approach is that the model is estimated and solved directly using rich firm-level data,[7] where each model unit corresponds to an actual Italian firm. This allows us to fully match the empirical joint distribution of firm heterogeneity in effective energy use and make-or-buy choices, without imposing any parametric restrictions.

To summarize, the total energy used in aggregate production depends on the allocation of energy across firms that differ in their energy-biased technology. When energy is under-priced, firms tend to give a higher weight to other factors in their make-or-buy decision, leading to an inefficient allocation of intermediate input production, thus of energy use. The extent of misallocation and its implications for aggregate energy intensity and welfare depend on the joint distribution between the energy-biased technology and the outsourcing capabilities of each firm. These mechanisms highlight the importance of analyzing aggregate energy demand through the lens of organization theories and production networks. If the literature on distortions in production networks is extensive (e.g., [Baqaee and Farhi, 2019, Liu, 2019, La'O and Tahbaz-Salehi, 2022]), it has only recently started to include energy considerations, with a particular focus on the propagation of energy price shocks [King et al., 2019, Devulder and Lisack, 2020, Martin et al., 2024] or technology improvements [Dugoua and Dumas, 2021, Aghion et al., 2024]. Closer to our analysis is the work of [Klenow et al., 2024], who show that a unilateral carbon tax could reduce existing misallocation in a distorted production network by reallocating production toward high-quality, high markup firms. We contribute to this literature by extending supplier selection to consider firm boundary decisions in order to characterize the production network in terms of firm-level energy efficiency and outsourcing capabilities. This allows us to thoroughly assess the role of the make-or-buy decisions on energy allocation, and to quantify the total extent of misallocation induced by distorted energy prices.

More precisely, by specifically considering the role of firm-level make-or-buy decisions in shaping the production network, we join [Boehm and Oberfield, 2023] in bridging the gap between the endogenous production network literature [Huneeus, 2018, Lim, 2018, Acemoglu and Azar, 2020, Bernard et al., 2022, Dhyne et al., 2023, Acemoglu and Tahbaz-Salehi, 2025, Arkolakis et al., 2025], that usually restricts to firms producing only one input or stage in-house[8] and the literature on firm boundaries and outsourcing (e.g., [Grossman and Helpman, 2002, Antràs, 2020]) that usually abstracts from the production network, focusing on equilibrium make-or-buy decisions at only one stage of the supply chain, or on bilateral links formation. However, when firms are heterogeneous, supplier selection is crucial to quantify the extent of misallocation that is induced by inefficient make-or-buy decisions. At the same time, by explicitly allowing for the possibility to produce the inputs in-house, our model can account for the energy misallocation between firms *across* different stages, and not

[7] Because firms are more substitutable within rather than across industries, contributions of changes in allocative efficiency due to reallocation are usually much larger when measured using firm-level data [Baqaee and Farhi, 2020]. A similar exercise relying on industry-level input-output tables would risk to vastly underscoring the role of reallocation.

[8] In these models, one could theoretically include make-or-buy decisions if one thinks of in-house production as the outside option to trading with a new supplier. However, the asymmetry between potential external suppliers and in-house production generates additional challenges that make these already complex frameworks less suitable to specifically analyze firm boundary decisions.

only within stages.[9] Therefore, it is necessary to embed the make-or-buy decision into a production network.

This means that we also speak to the literature studying how make-or-buy decisions might lead to misallocation and their macroeconomic consequences on aggregate welfare.[10] [Fally and Hillberry, 2018, Chor and Ma,], [Boehm and Oberfield, 2020], and [Boehm, 2022] study the extent to which contracting frictions shape aggregate input sourcing patterns and welfare outcomes at the country level. In this paper, we focus on the misallocation arising from inefficient make-or-buy decisions that are induced by the distorted price of a particular input, namely energy, and quantify the welfare loss this induces. Importantly, energy is the misallocated input, and our welfare measure specifically accounts for the energy externality.[11]

Despite the importance of intermediate input production in the total energy demand of a firm, considering environmental distortions through the lens of firm boundaries is, to our knowledge, very novel. The trade-environment literature has improved our understanding of scope 3 emissions and of how intermediate input outsourcing can act as a substitute for using energy.[12] The theory predicts that firms facing high domestic energy prices will tend to outsource the production of their energy-intensive intermediate inputs abroad [Cole et al., 2014, Schenker et al., 2018, Bolz et al., 2024].[13] If some empirical studies seem to confirm those predictions, [Michel, 2013, Cole et al., 2014, Cherniwchan et al., 2017, Cole et al., 2021, Dai et al., 2021, Choi et al., 2023, Coster et al., 2024], many show that energy is only one of the many determinants of the outsourcing decision, and it effect seems relatively small [Li and Zhou, 2017, Antonietti et al., 2017, Caron, 2022, Cherniwchan and Taylor, 2022, Dechezleprêtre et al., 2022, Dussaux et al., 2023, Levinson, 2023]. Consistently, our model predicts that, when energy prices are distorted, little consideration is granted to the energy component of the decision. The focus of this literature has mainly been on supplier (firm or, more broadly, country) selection based on energy prices. We contribute to this literature by considering the firm itself as a non-trivial option of production. Indeed, as intermediate input production requires energy, the firm boundary effectively determines the total energy demanded by a firm. This leads us to consider jointly the different determinants behind the make-or-buy decision of the firms in a world in which firms differ in their energy-biased technology and energy prices are distorted. To our knowledge, we are

[9] The endogenous production network of [Eaton et al., 2022] also includes in-house production as an option in the potential supplier selection set. However, they model the former with a production process that differs in terms of input requirements, and not only in terms of firm-specific factor-biased technology, from outsourced production. This prevents a direct comparison of different energy allocations across downstream and upstream firms.

[10] The literature has made substantial progress in elucidating the drivers of firms' sourcing choices—both in terms of location and organizational form (see e.g., [Antràs, 2020] for a review). In contrast, the macroeconomic consequences of these firm-level decisions remain less well understood.

[11] [Le Moigne et al., 2024] make a related point at the country-level. They show that a carbon tax generates large environmental gains by inducing countries to specialize according to their environmental comparative advantage. This reallocation channel accounts for more than one-third of the emissions reductions brought about by a carbon tax. In our words, this means that, once the environmental externality is taken into account, there is potential for misallocation as environmental comparative advantage across agents (firms or countries) are not fully exploited when energy prices are inefficiently low.

[12] [Köveker et al., 2025] shows that emissions embodied in intermediate inputs sourced outside of the firms compose roughly 70% of German firms' carbon footprint, with a large majority associated with domestically sourced inputs.

[13] Measuring carbon leakages to foreign countries is a key focus of this literature. However, the make-or-buy decision involves also, or even mostly, domestic partners [Antràs, 2005, Fort, 2017, Boehm and Oberfield, 2020, Dhyne et al., 2021, Carvalho et al., 2021, Fort, 2023].

the first to empirically identify the joint distribution between the firms' energy-biased technology and their outsourcing capabilities, which allows us to analyze the environmental implications of the make-or-buy decisions of firms.

Beyond quantifying misallocation, we also use our structural model to evaluate the aggregate effects of policy-induced energy price changes, accounting for the reallocation prompted by the make-or-buy decisions. A large empirical literature has estimated the impact of carbon policies on firm-level outcomes (e.g., [Martin et al., 2014, Dechezleprêtre et al., 2019, Cui et al., 2021, Fontagné et al., 2024]). However, micro-elasticity estimates based on idiosyncratic shocks or control groups cannot be aggregated to infer macroeconomic effects when there is reallocation across firms. Using difference-in-differences and synthetic control designs [Andersson, 2019, Metcalf, 2019], as well as specifications that control for macroeconomic variables influencing carbon price dynamics [Metcalf and Stock, 2020, Bernard and Kichian, 2021], a recent body of work has tried to estimate those macro elasticities directly. Consistent with firm-level evidence, this literature generally finds that carbon prices succeed in reducing emissions while exerting limited effects on aggregate outcomes. However, these results rely on the specific design of the policy under study, for instance, its sector coverage or the use of tax revenues. This might explain why, in contrast, more general macro-estimates [Känzig, 2023, Ortubai et al., 2025] point toward substantial environmental and macroeconomic effects when leveraging high-frequency variations in carbon prices. However, this approach is also not immune to confounding factors such as overlapping macroeconomic shocks that can bias the result. Moreover, estimates based on high-frequency shocks might fail to capture medium-run adjustments or adjustments that require a sufficiently important shock, while we know that firms' make-or-buy decision tends to be relatively sticky [Antràs and Chor, 2013, Monarch, 2014, Martin et al., 2023, Lafrogne-Joussier et al., 2023], such that firms might lever this channel only in response to relatively large and persistent changes. We complement these approaches by building a structural production network model that explicitly accounts for the allocative role of the make-or-buy decision of the firms as a margin of adjustments to energy price shocks. This allows us to fully measure the aggregate responses to a change in energy prices induced by firm-level adjustments and the reallocation of production both at the intensive at the extensive margin.

Outline. The rest of the paper is organized as follows. The coming section describes the data and the stylized facts that motivate the project. Section 3 studies the drivers and implications of the firms' make-or-buy decisions in a simple two-stage supply chain model to provide intuitions on the key mechanisms. We describe the full quantitative model in section 4, and, in section 5, develop a strategy to infer the set of inputs produced in-house, the energy-biased technology and the outsourcing capabilities of each firm from the Italian firm-to-firm data. This part is still work in progress, so this draft doesn't include the final results of the calibration. Section 7 lays out the next steps of the project. Additional details and proofs are provided in the appendix (A).

2 Data and Stylized facts

In this section, we present our data and some stylized facts that motivate the subsequent analysis. In particular, we document that within narrowly defined sectors, comparable firms are extremely heterogeneous in terms of their energy

intensity. We also demonstrate that intermediate input outsourcing is strongly negatively correlated with energy demand at the firm level.

Our calibration relies on characterizing the production network in terms of effective energy 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 provides the perfect setting for this analysis, because data on the universe of domestic firm-to-firm transactions, that can be combined with detailed custom data was recently made available. Moreover, the Italian manufacturing sector is characterized by its pronounced fragmentation:[Agostino et al., 2016][14] 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 and Cabral, 2015, Cappariello and Felettigh, 2015]. Finally, Italy is the second-largest manufacturing country and third CO_2 emitter in Europe.[15] Therefore, it offers a perfect setting to analyze the impact of firms' make-or-buy decisions and the resulting supply chain fragmentation on aggregate energy demand, while being both economically relevant and representative of similar western economies.

Data. We combine three datasets that describe the Italian production network in details: i) VAT data ii) custom data and iii) balance-sheet data. The first dataset contains the universe of business-to-business (B2B) domestic transactions in the Italian economy during the year 2019.[16] It is constructed based on firms' VAT declarations, and report the total yearly value of transactions, in euros, for each bilateral business relationship. Crucially, each business entity is uniquely characterized by an identifier common to all the other datasets.

We combine it with custom 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 firm-pair transactions at the yearly level. Products are classified at the CN 8−digit level.

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., [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 its main economic activity at the ATECO 6−digit level.[17] 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. Moreover, we exclude outliers in

[14]The success of Italian SMEs has notably been related to regional "clusters" of specialized firms that fragment the production process, and whose collaboration is facilitated by intermediate governance structures [Bianchi et al., 1997].

[15]Source: International Energy Agency (IEA)

[16]Years 2020 to 2023 are currently being made available, and we intend to use them to refine our calibration (see 5).

[17]The ATECO Italian classification code, or Classificazione delle Attività Economiche (Classification of Economic Activities), is the national standard for classifying businesses and self-employed individuals 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.

terms of turnover growth rates.

For our main stylized facts, 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 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 energy.(see appendix B.1 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 equipments, buildings and rent payments, construction, wholesale and retail, and services. Those transactions will be bundled together as non-manufacturing sales in our main analysis.

We only retain private, domestic manufacturing companies in the sample we analyze. Moreover, for our final analysis, we exclude firms classified as primary material-providers based on their ATECO 6−digit industry code, to focus on firms processing inputs instead of extracting them.[18] Our final sample contains 125801 manufacturing firms.

Stylized facts. Firstly, within narrowly defined economic activity, firms are extremely heterogeneous in terms of their energy efficiency [Petrick et al., 2011, Lyubich et al., 2018]. This feature is extremely robust across countries [Wagner et al., 2020, De Lyon and Dechezleprêtre, 2025]. We demonstrate that this regularity applies to Italian manufacturing firms as well. Figure 2 shows the cross-sectional distribution of the firms' energy expenditures to sales ratio (in logarithm). As a mean of comparison, we plot the same distribution for the labor expenditures and the capital expenditures to sales ratio. To ensure that we are comparing similar firms, we first residualize these variables on 6−digit product fixed effects. The observed dispersion in energy intensity can be driven by several factors, and the literature has pointed toward large differences in energy technologies.[19] We argue that it is also due to differences in firm boundaries.

[18] 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 from the final analysis.

[19] The dispersion subsists beyond observable characteristics such as investment, capital and labor intensity, and fuel mix.

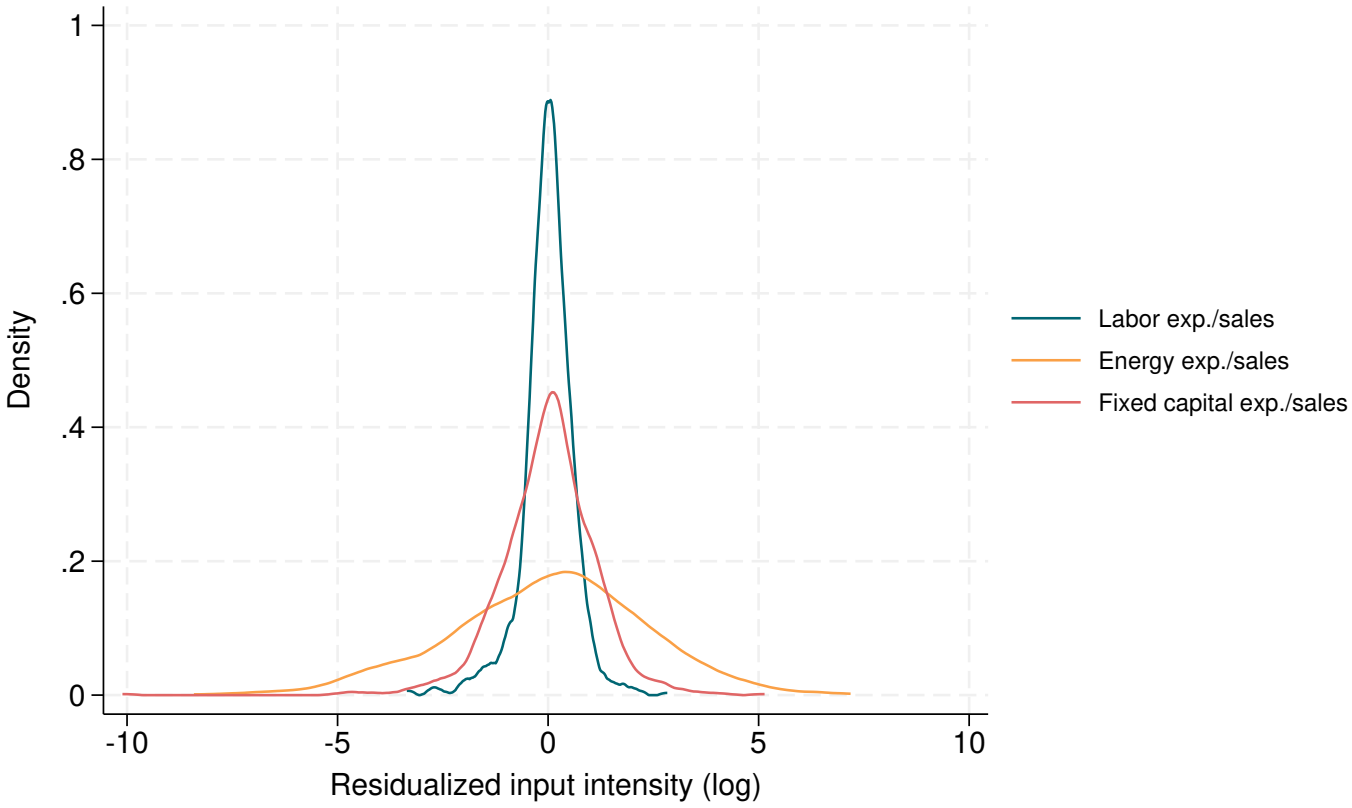


Figure 2: Dispersion of input intensity

Notes: The figure plots the distribution of the residuals of the ratio of energy, capital and labor expenditures to sales (in logarithm). The residuals are obtained from a cross-sectional (2019), firm-level regression on sectors (4-digit).

Indeed, firms that outsource the production of their intermediate input consume less energy in their own production, for a given level of sales. Figure 3 illustrates the negative correlation between intermediate input expenditures and energy expenditures per unit sold. This is robust to considering the intermediate input expenditures to sales ratio. Intuitively, ceteris paribus, whether a firm is producing its intermediate input in-house or is outsourcing their production to an external supplier determines the total amount of energy that is used within the boundary of the firm. Because the production of intermediate inputs are relatively energy intensive, as illustrated above in figure 1, the outsourcing behavior of the firm is therefore a key determinant of firm-level energy usage. This has two important implications: first, identifying the effective energy efficiency of a firm requires knowing its boundaries that is, all the inputs it is producing in-house. Second, the make-or-buy decisions of the firms effectively allocate energy demand across heterogeneous firms. In turn, this allocation determines the aggregate energy intensity of the economy.

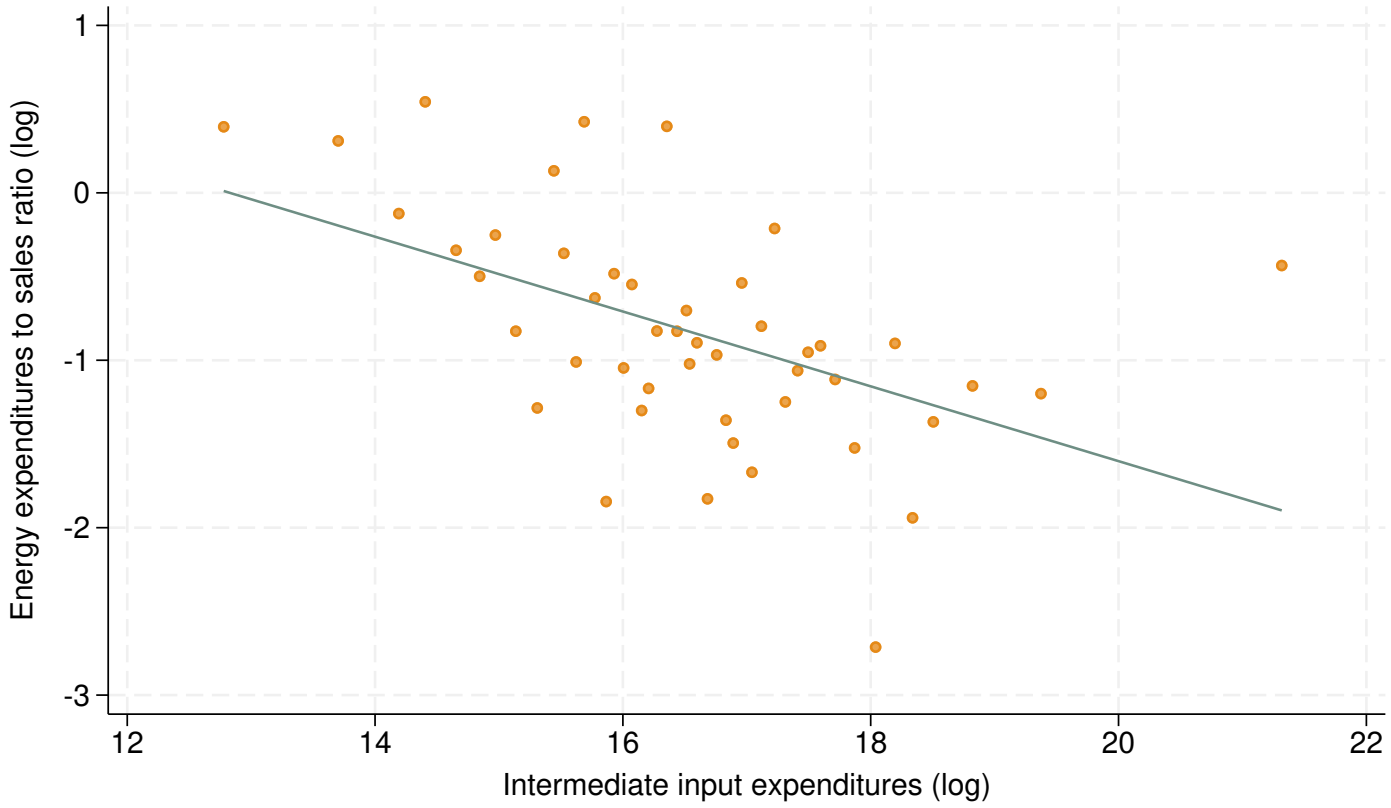


Figure 3: Correlation between firms' energy per unit sold and intermediate input expenditures

Notes: The graph plots the relation between energy intensity, defined as energy expenditures per sales (in logarithm) and intermediate input expenditures (in logarithm). The relationship is measured controlling for economic activity at the 6-digit level.

Building on these facts, we turn to a simple model that provides intuitions on the key mechanisms through which firm-level make-or-buy decisions affects the aggregate energy demand of production, and establish that those decisions might entail energy misallocation. In section 4, we develop a quantitative model of the entire manufacturing sector to evaluate the misallocation present in the production network when accounting for energy.

3 A simple two-stage model of fragmentation

To highlight the key mechanisms at play, we describe a simple static partial equilibrium model of make-or-buy decision in which a downstream firm decides whether to produce its intermediate inputs in-house or to outsource their production to an upstream firm. In short, intermediate input production requires energy and firms differ in their energy-biased technology, as well as their outsourcing capabilities. Those two parameters drive the make-or-buy decision of the downstream firm, which in turn determines the allocation of energy across the two firms, hence the total amount of energy used used to produce the final good. The model reveals that, when energy prices are distorted relative to the social cost of energy, the downstream firm under-considers the energy-biased technology in its make-or-buy decision, leading to misallocation and higher level of energy used per unit of final output. More details and complementary proofs are provided in appendix A.1.

Emission externality and welfare. A representative household supplies one unit of labor inelastically, receives utility from consumption and is harmed by energy. Formally, $U(y, e)$, where $\frac{\partial U}{\partial y} > 0$, $\frac{\partial U}{\partial e} < 0$, $\frac{\partial^2 U}{\partial^2 y} < 0$ and the Inada conditions hold. The welfare function is informative of two elements: first, the energy externality creates a distortion. Absent any policy intervention, energy won't be priced sufficiently.[20] Secondly, similar to spatial equilibrium models (e.g., [Donald et al., 2025]), but different from the standard misallocation literature (e.g., [Hsieh and Klenow, 2009, 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 in the model, output will be maximized.

Downstream firm. A downstream firm, indexed by j, produces using an intermediate input $\hat{\omega}$. It can either produce it in-house, using energy and labor, or outsource its production to an upstream supplier s:

$$y_j = \begin{cases} q_{jj} \left(\beta^{1-\rho}(\nu_j e)^\rho + (1-\beta)^{1-\rho} l^\rho\right)^{\frac{1}{\rho}} & \text{if in-house} \\ q_{ij} x_i(\hat{\omega}) & \text{if outsourced to } i \end{cases}$$

Where ν_j is the energy-biased technologyof firm j and $\rho < 1$ parametrizes the elasticity of substitution between labor and energy. q_{ij} reflects any relationship-specific factor that makes outsourced inputs differ from in-house inputs.[21] For instance, if outsourcing only generates specialization gains, $q_{ij} > q_{jj}$. In contrast, if outsourcing implies a loss of control over input quality for the downstream firm, $q_{ij} < q_{jj}$. Following the empirical literature [Bretschger and Jo, 2021a], we also impose that $\rho < 0$ for illustrative purposes, but this assumption is innocuous as the main insights remain if labor and energy are substitutes.

Intermediate suppliers. If intermediate input $\hat{\omega}$ is outsourced, it is produced by an upstream firm i that prices at marginal cost. Its production again necessitates labor and energy.

$$x_i(\hat{\omega}) = \left(\beta_{\hat{\omega}}^{1-\rho}(\nu_i e_i(\hat{\omega})^\rho + (1-\beta_{\hat{\omega}})^{1-\rho} l_i(\hat{\omega})^\rho\right)^{\frac{1}{\rho}}$$

Importantly, firm i produces with its own energy-biased technology ν_i.

Equilibrium. The downstream firm chooses whether to produce input $\hat{\omega}$ in-house or to outsource it. Conditional on producing in-house, it chooses its labor and energy demand to minimize costs. The upstream firm chooses its labor and energy demand to minimize costs, conditional on producing for the downstream firm. We characterize the model's equilibrium in two steps. First, we derive firm-level optimal behavior conditional on a given make-or-buy strategy. Second, we endogenize the make-or-buy decision. We conclude by discussing the model's implications for aggregate energy intensity.

[20] We could define the social cost of carbon as $p_e^{SCC} = -\frac{\partial U/\partial e}{\partial U/\partial y}$, but it is endogenous to the equilibrium. Therefore, it would make our benchmark evaluation very sensitive to our welfare functional form, potential measurement errors and factors or other distortions omitted from the model. Therefore, as in [Hsieh and Klenow, 2009], we prefer to consider an exogenous benchmark that we see as a feasible objective for Italy and take it as exogenous.

[21] Despite their prevalence in firm boundaries theory [Antràs, 2020], we abstract from incomplete contracts because our quantitative exercise mainly focuses on domestic outsourcing.

Given the equilibrium make-or-buy strategy, the producing firm k of input $\hat{\omega}$ chooses labor and energy to minimize its costs. The CES production function yields the usual unit cost of production for downstream firm j:

$$\frac{c_j(\hat{\omega})}{q_{kj}} = q_{kj}^{-1}\left(\beta\left(\frac{p_e}{\nu_k}\right)^{\frac{-\rho}{1-\rho}} + (1-\beta)w^{\frac{-\rho}{1-\rho}}\right)^{\frac{-(1-\rho)}{\rho}}$$

Where p_e and w are respectively the price of energy and the wage. We define p_e^{SCC} as the energy price that would match the social cost of energy once emissions are taken into account.

Make-or-buy decision. Firm j decides on whether to produce it in-house or to outsource it by comparing the unit cost of of each option. Formally, firm j produces input $\hat{\omega}$ in-house if $c_{ij}^{\mathcal{I}}(\hat{\omega}) < \frac{c_{ij}^{\mathcal{O}}(\hat{\omega})}{q_{ij}}$:

$$q_{jj}^{-1}\left(\beta\left(\frac{p_e}{\nu_j}\right)^{\frac{-\rho}{1-\rho}} + (1-\beta)w^{\frac{-\rho}{1-\rho}}\right)^{\frac{-(1-\rho)}{\rho}} < q_{ij}^{-1}\left(\beta\left(\frac{p_e}{\nu_i}\right)^{\frac{-\rho}{1-\rho}} + (1-\beta)w^{\frac{-\rho}{1-\rho}}\right)^{\frac{-(1-\rho)}{\rho}}$$

The make-or-buy decision depends on the energy-biased technology of both firms $\{\nu_j, \nu_i\}$ as well as on the relative outsourcing capabilities $\frac{q_{ij}}{q_{jj}}$. In particular, 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. Importantly, the weight given to each component differs, because the energy-biased technology only economizes on energy, while the outsourcing capabilities economize on both labor and energy. We can now turn to deriving our main results.

Lemma 1 *When $\rho < 0$, the lower energy prices, the less weight is given to the relative energy-biased technology in the make-or-buy decision, and the more weight is given to the relative outsourcing capabilities. The resulting between-firm allocation weakly increases energy per unit of final output.*

Intuitively, at high energy prices, it becomes important for the downstream firm to economize on energy. Therefore, it gives more weight to the relative energy-biased technology in its decision.[22] The between-firm allocation resulting from the make-or-buy decision yields weakly lower energy per unit of final output at high energy prices: the firm that produces the input, hence the effective energy user, is relatively weakly more energy efficient. Because in-house production and outsourcing are mutually exclusive alternative, this is an extensive margin effect, hence the *weak* monotonicity. Therefore, the extensive margin amplifies the well-known intensive margin of higher energy prices reducing within-firm energy demand.

We call p_e^{SCC} the price of energy inclusive of the (exogenously determined) social cost of carbon. The next proposition naturally ensues from lemma 1.

Proposition 1 *If $p_e < p_e^{SCC}$, the weight given to the energy component in the make-or-buy decision is inefficiently low. Therefore, the allocation of production across firms is weakly inefficient.*

Formal proofs are derived in appendix A.1. In other words, when energy is under-priced, too little consideration is given to the energy-biased technology in the make-or-buy decision.[23] Intuitively, both the private firms and the social planner

[22] The opposite arises if $\rho > 0$: less weight is given to the relative energy-biased technology. However, this also results in weakly lower energy intensity. In other words, for all ρ, the firm pays more attention to reducing its energy demand at higher energy prices.

[23] If $\rho > 0$, *too much* weight is given to the energy-biased technology relative to the outsourcing capabilities. However, this misallocation also results in a higher energy intensity, because a higher technology ν_j is no longer an energy-saving technology but an energy-augmenting technology: it increases the energy used by the firm.

care about both the relative hicks-neutral outsourcing capabilities $\frac{q_{ij}}{q_{jj}}$, and the relative energy-biased technology $\frac{\nu_j}{\nu_i}$. However, they assign different weights, driven by the price of energy p_e, to these different factors in their make-or-buy decision. This is because firms optimize on costs, while the social planner minimizes both costs and energy demand. With complementary inputs, firms benefit relatively more from a high energy-biased technology when facing higher prices, as it contributes to absorbing those high prices (formally, the cost function is supermodular in factor prices and factor-biased technology). In contrast, with substitutable inputs, firms benefit relatively less from a high energy-biased technology when facing higher prices, because those high prices induce them to substitute away from energy (formally, the cost function is submodular in factor prices and factor-biased technology).
These results imply the following corollary:

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

The first statement is simply a result of the first-welfare theorem. Since everything is efficient in the model, the equilibrium allocation of both labor and energy chosen by the producing firm at the intensive margin corresponds to the efficient allocation. The second statement results from a similar intuition: without any distortion, the outsourcing decision is efficient because it minimizes the overall cost of production, taking into account the social cost of energy through p_e^{SCC}. We want to highlight that, because it minimizes the effective unit cost, the make-or-buy decision maximizes output, for a given energy price. That is, the extensive margin (between firms) allocation is efficient in terms of maximizing output.[24] The inefficiency arises once we take into account that energy enters welfare. The rationale is mirroring spatial equilibrium models in which agents make their location choices based on utility considerations. Therefore, the allocation does not necessarily maximize aggregate real GDP. In our case, the allocation maximizes output but not welfare, because firms *don't* consider the energy externality in their make-or-buy decision.

We now show that this inefficient decision induces misallocation of production across firms, and higher level of total energy used per unit of output produced (energy intensity).[25] For each input, the make-or-buy decision leads to misallocation if[26]

$$\frac{1-\beta}{\beta}\left(\frac{w}{p_e}\right)^{\frac{-\rho}{1-\rho}} > (<) \frac{\left(\left(\frac{q_{ij}\nu_i}{q_{jj}\nu_j}\right)^{\frac{\rho}{1-\rho}} - 1\right)\nu_j^{\frac{\rho}{1-\rho}}}{1-\left(\frac{q_{ij}}{q_{jj}}\right)^{\frac{\rho}{1-\rho}}} > (<) \frac{1-\beta}{\beta}\left(\frac{w}{p_e^{SCC}}\right)^{\frac{-\rho}{1-\rho}} \text{ if } \rho < (>)\, 0 \tag{1}$$

This only arises if there is a trade-off between the input-specific outsourcing capabilities $\frac{q_{ij}}{q_{jj}}$, and the energy-biased technology, that is if $\frac{\nu_j}{\nu_i} > \frac{q_{ij}}{q_{jj}} > 1$ or if $1 > \frac{q_{ij}}{q_{jj}} > \frac{\nu_j}{\nu_i}$.[27] Therefore, the joint distribution of the outsourcing capabilities and the energy-biased technology of the firm is the crucial moment that will characterize the extend of misallocation. We will

[24] The intensive margin (within-firm) allocation is also efficient for a given energy price, that is, to maximize output.

[25] In standard model of misallocation (e.g., [Hsieh and Klenow, 2009]), firms are subject to heterogeneous distortions. In our model, they are subject to a common distortion ($p_e < p_e^{SCC}$, but this distortion affects them differently depending on their energy-biased technology.

[26] Even if this inequality does not hold, the make-or-buy decision is still distorted, because too little consideration is given to the energy component. However, it does not necessarily lead to misallocated production of the input, since the decision is discrete.

[27] Such trade-off can easily be micro-founded by considering that firms need to pay a fixed cost to improve their energy-biased technology, a

identify it in our data.

As an illustration, we suppose that $q_{jj} = 1$, $\nu_i = 1$. If $q_{ij} > (<)\ 1$, a trade-off arises when $\nu_j > (<)\ 1$, that is, if the downstream firm is more (less) energy efficient than its upstream supplier i. Figure 4 plots the optimal and equilibrium allocation, set respectively under p_e^{SCC} and $p_e = 0.47 p_e^{SCC}$,[28], for calibrated values of $\frac{p_e^{SCC}}{w} = 15.5$, $\beta = 0.2$ and $\rho = -0.7$. If the firm $\{q_{ij}, \nu_j\}$ falls above the line, outsourcing is chosen. Conversely, if the firm $\{q_{ij}, \nu_j\}$ falls in the area below the line, in-house production is chosen. The shaded areas represent the inefficient allocations, where the decision of the firm differs from the decision of the social planner. As ν_j moves away from $\nu_i = 1$ (i.e., as the heterogeneity in energy-biased technology increases), inefficient outsourcing is more and more likely to arise.

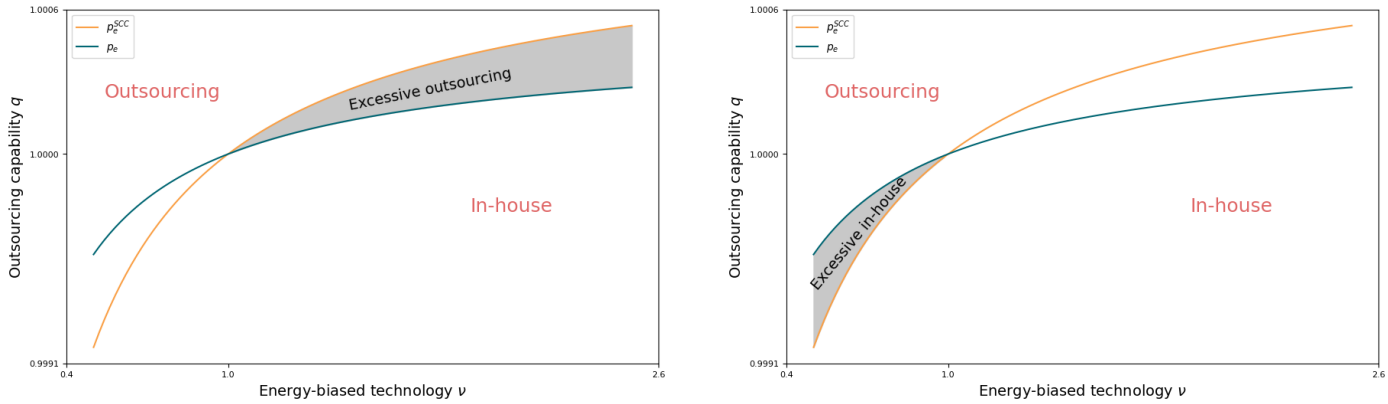


Figure 4: Specialization gains versus energy-biased technology

Probability of Misallocation For $\rho < 0$, the probability that a firm makes an inefficient decision is characterized by

$$\mathcal{P}\left(\frac{1-\beta}{\beta}\left(\frac{w}{p_e}\right)^{\frac{-\rho}{1-\rho}} > \frac{\left(\left(\frac{q_{ij}\nu_i}{q_{jj}\nu_j}\right)^{\frac{\rho}{1-\rho}} - 1\right)\nu_j^{\frac{\rho}{1-\rho}}}{1 - \left(\frac{q_{ij}}{q_{jj}}\right)^{\frac{\rho}{1-\rho}}} > \frac{1-\beta}{\beta}\left(\frac{w}{p_e^{SCC}}\right)^{\frac{-\rho}{1-\rho}} \right) \tag{2}$$

question that we explore in a companion paper [Marullaz, 2025]. For instance, suppliers might be highly specialized and produce a customized input (high q_{ij}, low $q_{ik}\ \forall k \neq j$, that is, a negative correlation between the q_{ik} of supplier i). This customization reduces their customer base, meaning that they are too small to afford paying the fixed cost to improve their energy-biased technology. However, the downstream firm j will outsource most of its inputs to gain from specialization, especially at low energy prices which lowers the importance of the energy component. The downstream firm is however large enough to cover the fixed cost of adoption, resulting in a trade-off between the relative energy-biased technology and the outsourcing capabilities. The mirroring case of general suppliers also results in a trade-off: those suppliers have a positive correlation between all their $q_{ik} \forall k$, and can serve all downstream firms. This gives them the scale to invest in their energy-biased technology. However, the downstream firms lose quality from this generality (small q_{ik}), and they choose to produce the input in-house. With investment costs increasing convexly with the number of inputs, the downstream firms won't invest in energy-biased technology improvements, resulting, again, in a trade-off between the relative energy-biased technology and the outsourcing capabilities. Nonetheless, the goal of this project is not to explain this joint distribution, but to match it accurately for all firms to measure total misallocation. The lack of empirical knowledge on the joint distribution of the relative energy-biased technology and outsourcing capabilities, and the large dispersion in make-or-buy practices, have led us to keep a very general and flexible characterization of the outsourcing capabilities, and to remain agnostic on those microfoundations.

[28] For a SCC of 170€ per metric tonne of CO_2 [Pindyck, 2019, EPA, 2022], the average energy price in Italy is approximately 47% of the SCC.

Define $\Upsilon_{ij} = \frac{\left(\left(\frac{q_{ij}\nu_i}{q_{jj}\nu_j}\right)^{\frac{\rho}{1-\rho}} - 1\right)\nu_j^{\frac{\rho}{1-\rho}}}{1-\left(\frac{q_{ij}}{q_{jj}}\right)^{\frac{\rho}{1-\rho}}}$, then, the probability is characterized by the distribution of firms:

$$F_\Upsilon\left(\frac{1-\beta_\omega}{\beta_\omega}\left(\frac{w}{p_e}\right)^{\frac{-\rho}{1-\rho}}\right) - F_\Upsilon\left(\frac{1-\beta_\omega}{\beta_\omega}\left(\frac{w}{p_e^{SCC}}\right)^{\frac{-\rho}{1-\rho}}\right)$$

It is clear that the distribution F_Υ, that is, the joint distribution of the buyer-supplier level energy-biased technologies and other-outsourcing-determinants bundle $\{\nu_j, \nu_i, q_{ij}, q_{jj}\}$, will determine the degree of misallocation. However, to our knowledge, these objects have never been empirically related. This is a key contribution of our quantitative analysis.

The specification of Υ_{ij} prevents us to express F_Υ analytically. This means that we will have to exactly identify ν_j and q_{ij} for all firms in our data. In our quantitative section (5), we tackle this challenge, guided by our model, and develop a novel strategy that makes use of our rich firm-to-firm transaction data. However, we can discuss some of its characteristics. In particular, misallocation can only arise if $\Upsilon_{ij} > 0$[29], that is if the numerator and the denominator have the same sign: $sgn\left(\left(\frac{q_{ij}\nu_i}{q_{jj}\nu_j}\right)^{\frac{\rho}{1-\rho}} - 1\right) = sgn\left(1 - \left(\frac{q_{ij}}{q_{jj}}\right)^{\frac{\rho}{1-\rho}}\right)$. This is satisfied if $\frac{\nu_j}{\nu_i} > (<) \frac{q_{ij}}{q_{jj}} > (<) 1$. This implies a negative correlation between $\frac{q_{ij}}{q_{jj}}$ and $\frac{\nu_i}{\nu_j}$, which is intuitive: misallocation arises if firms with a relatively high level of energy-biased technology also tend to have high level of outsourcing capabilities, which incentivize them to outsource despite the environmental loss, or if firms tend to outsource to suppliers with a relatively low energy-biased technology, because their strong outsourcing capabilities dominates this negative component. A typical example is when luxury clothing firms outsource to specialized, family-owned entity to benefit from the specialized skills and craftsmanship of their suppliers, but the latter are financially constrained and cannot afford to reduce their environmental footprint, or to control the environmental impact of their own supply chain.

Energy intensity. Misallocation increases aggregate energy intensity.[30] To illustrate this result, suppose that $\frac{q_{ij}}{q_{jj}} > 1$. Moreover, suppose that there is misallocation. For instance, $\hat{\omega}$ is outsourced to supplier i if $p_e < p_e^{SCC}$, but $\hat{\omega}$ is produced in-house under $p_e = p_e^{SCC}$ We denote variables in a $p_e < p_e^{SCC}$ allocation with a star, and variables in a $p_e = p_e^{SCC}$ with SCC. We can decompose the effect of a different energy price through both the extensive and the intensive margin adjustments:

$$\frac{(e/y)^*}{(e/y)^{SCC}} = \underbrace{\left(q_{ij}\frac{\nu_i}{\nu_j}\right)^{\frac{\rho}{1-\rho}} \times \left(\frac{c_i^{\mathcal{O}}(p_e^{SCC})/q_{ij}}{c_j^{\mathcal{I}}(p_e^{SCC})}\right)^{\frac{1}{1-\rho}}}_{\text{Extensive margin: Allocative efficiency}} \times \underbrace{\left(\frac{c_i^{\mathcal{O}}(p_e)/p_e}{c_i^{\mathcal{O}}(p_e^{SCC})/p_e^{SCC}}\right)^{\frac{1}{1-\rho}}}_{\text{Intensive margin}} > 1$$

Energy intensity is larger under the equilibrium $p_e < p_e^{SCC}$ allocation. The intensive margin effect is standard. The increase in the relative price of energy reduces within-firm energy demand through the well-known substitution effect: $\frac{c_i^{\mathcal{O}}(p_e)/p_e}{c_i^{\mathcal{O}}(p_e^{SCC})/p_e^{SCC}} = \left(\frac{1+\frac{1-\beta}{\beta}\left(\frac{w}{p_e/\nu_i}\right)^{\frac{-\rho}{1-\rho}}}{1+\frac{1-\beta}{\beta}\left(\frac{w}{p_e^{SCC}/\nu_i}\right)^{\frac{-\rho}{1-\rho}}}\right)^{\frac{-(1-\rho)}{\rho}} > 1.$

[29] The misallocation probability in equation (2) also highlights that if $\beta \to 1$ (no labor) or $\rho \to 0$ (Cobb-Douglas), there is no potential for misallocation because ν_j turns into a hicks-neutral technology and is no longer distinctly affected by the price of energy.

[30] In Appendix A.1.1, we relate our concept of misallocation to the concept more standardly used in the literature by using wedges.

Additionally, the make-or-buy decision induces a novel extensive margin effect. Misallocation increases total energy usage by allocating production to a firm that is less energy efficient.[31] Formally, $\left(\frac{q_{ij}\nu_i}{q_{jj}\nu_j}\right)^{\frac{\rho}{1-\rho}} > 1$, as this guarantees equation 1 to hold true. Intuitively, since $q_{ij} > 1$, the firm has decided to outsource to its supplier under the low energy price-allocation. The gains from outsourcing capabilities outweigh the loss from outsourcing to a relatively less energy efficient firm, because energy prices are low. This results in a too high level of energy used, relative to a world where the social costs of energy are fully taken into consideration.[32] Moreover, $p_e^{SCC} > p_e$ and allocation choices imply that $\frac{c_j^{\mathcal{O}}(p_e^{SCC})}{q_{ij}} > c_j^{\mathcal{I}}(p_e^{SCC}) > c_j^{\mathcal{I}}(p_e) > \frac{c_i^{\mathcal{O}}(p_e)}{q_{ij}}$: because the make-or-buy decision minimizes costs, the costs reduction from reallocation (at the new price) also contributes to reducing energy intensity as overall, a more productive firm produces.

This is because the technologies ν and q affect firm-level energy and output differently. Through demand, output is determined by the marginal cost function, which is linear in q^{-1} but convex in ν_j^{-1}. In contrast, $e_j = \beta\left(\frac{p_e/\nu_j^{\rho}}{c_j}\right)^{\frac{-1}{1-\rho}} y_j$: additional to the effects on output, which determines the scale, energy is also directly affected by ν_j and q_j through the (inverse) relative effective price of energy $\left(\frac{p_e/\nu_j^{\rho}}{c_j}\right)^{\frac{-1}{1-\rho}} = \left[\frac{1}{q_{jj}\nu_j}\left(1+\left(\frac{w\nu_j}{p_e}\right)^{\frac{-\rho}{1-\rho}}\right)^{\frac{-(1-\rho)}{\rho}}\right]^{\frac{-1}{1-\rho}}$. While $q_{jj}\nu_j$ have the same effect in the first term, the additional (factor-biased) part of the energy technology enters the second term in a convex way. Since the make-or-buy decision is solely based on comparing the effective unit cost, it maximizes output. However, by affecting the energy use, it decreases welfare.

Finally, this is a firm-level result. In appendix A.1, we show that our results aggregate to multiple firms. Moreover, in the aggregate, (intensive margin) reallocation of production across firms through changes in sales share will also play a role. This corresponds to the allocative efficiency effect of [Baqaee and Farhi, 2020], which, in our case, affects energy intensity by reallocating resources towards or away from less productive producers.

We illustrate this result with a simple simulation in figure 5. To observe different decisions, we augment the model with a downstream firms producing using multiple inputs, combined in a CES with elasticity γ.[33] For each input, the make-or-buy decision remains the same and separate.[34] We can now consider the make-or-buy decision for all inputs required by the downstream firm. We set $\rho = \gamma = -0.7$ following [Bretschger and Jo, 2021a]'s estimate of the elasticity

[31] Without firm heterogeneity, the make-or-buy decision of the firm is irrelevant. Appendix A.1.2 makes this argument explicit.

[32] Notice that with $\rho > 0$, the firm would have given more weight to $\frac{\nu_j}{\nu_i}$ in its make-or-buy decision, leading to in-house production in equilibrium. Misallocation would require outsourcing under the efficient energy price, which implies $\frac{1-\beta}{\beta}\left(\frac{w}{p_e}\right)^{\frac{-\rho}{1-\rho}} < \frac{\left(\left(\frac{q_{jj}\nu_j}{q_{ij}\nu_i}\right)^{\frac{\rho}{1-\rho}}-1\right)\nu_j^{\frac{\rho}{1-\rho}}}{1-\left(\frac{q_{jj}}{q_{ij}}\right)^{\frac{\rho}{1-\rho}}} < \frac{1-\beta}{\beta}\left(\frac{w}{p_e^{SCC}}\right)^{\frac{-\rho}{1-\rho}}$ when $\frac{q_{ij}}{q_{jj}} > 1$. Thus, we must have that $\left(\frac{q_{jj}\nu_j}{q_{ij}\nu_i}\right)^{\frac{\rho}{1-\rho}} > 1$, meaning that again, $\frac{(e/y)^*}{(e/y)^{SCC}} \propto \left(\frac{q_{jj}\nu_j}{q_{ij}\nu_i}\right)^{\frac{\rho}{1-\rho}} > 1$.

[33] In particular, $y_j(\omega) = |\Omega_\omega|^{-(1-\gamma)}\left(\int_{\hat{\omega}\in\Omega_\omega} x_{ij}^{\gamma}(\hat{\omega})d\hat{\omega}\right)^{\frac{1}{\gamma}}$, and a continuum of upstream suppliers can each produce one input $\hat{\omega}$ and price at marginal costs.

[34] Because there is no fixed cost in the make-or-buy decision, the scale of the firm is irrelevant. Therefore, the decision can be taken for each input separately. This assumption is similar to [Eaton et al., 2022] and avoids falling into the realm of a combinatorial discrete choice problems [Antras et al., 2017, Arkolakis et al., 2023]. The need to explicitly characterize suppliers to evaluate the energy-biased technology related to both the outsourcing and the in-house option would render such a combinatorial discrete choice problem intractable.

of substitution between energy and labor, and $\beta = 0.2$ based on the average energy expenditure share in manufacturing sectors recovered from the OECD Italian input-output table. We normalize $q_{jj} = 1$ and consider upstream firms with an identical energy-biased technology ν_i that we normalize to 1, and draw the outsourcing capabilities randomly from a uniform distribution: $q_{ij} \sim U[1, 3]$. Finally, we consider a downstream firm with $\nu_j = 2.15$, which corresponds to a firm at a one standard deviation above the mean in the distribution of energy efficiency in figure 2. We consider a range of energy price to wage ratio from 1 to 15.5, based on recent estimates of the social cost of carbon of 170€ per metric tonne of CO_2 [Pindyck, 2019, EPA, 2022] and an average manufacturing wage of = 11€/hour in Italy, which would give us $\frac{p_e^{SCC}}{w} = 15.5$

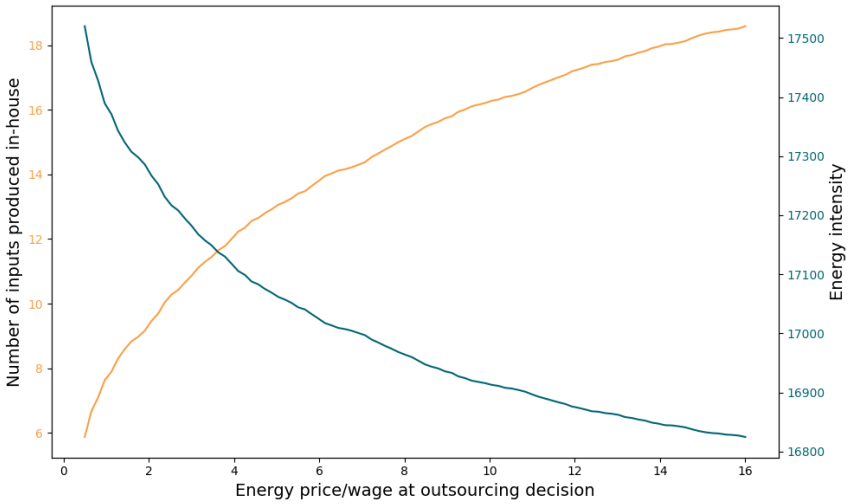


Figure 5: In-house input set, energy intensity and energy prices

Figure 5 shows that, as the price of energy increases, more weight is given to the energy component in the outsourcing decision of the firm. Because the downstream firm is more energy efficient, this tends to increase the number of inputs that it produces in-house. In turn, energy intensity (evaluated at $p_e = p_e^{SCC}$ to shut down the intensive margin) decreases because energy misallocation is reduced.

The main take-away of this section is that lower energy intensity requires relatively high ν_j firms to produce in-house, but low energy prices bias the decision away from ν_j. We next turn to quantify the extent of this misallocation.

4 A quantitative model of fragmentation in a production network

Through the lens of the simple model, we have highlighted the fact that energy is under-considered in the make-or-buy decision of the firms, resulting in misallocation. We now turn to quantify the importance of this misallocation in the Italian manufacturing economy. This is an ex-post analysis, in the sense that it requires knowledge of the both equilibrium, observed Italian manufacturing sector allocation, and the counterfactual allocation that we use as our efficient benchmark. However, we do not have an empirical "efficient" benchmark to which we could compare the Italian economy,

as opposed to e.g., [Hsieh and Klenow, 2009]'s quantitative exercise. We therefore need to measure our counterfactual allocation. This forces us to impose some parametric structure to our framework, such that we can calibrate it using our rich firm-level data, and build the counterfactual allocations of interest. Nonetheless, we apprehend the data as the unique representation of the model: each empirical firm corresponds to a model unit, that we will characterize in terms of an energy-biased technology and a distribution of outsourcing capabilities. This implies that we fully capture the joint dispersion across Italian firms in energy intensity and firm boundaries, with as little functional form assumptions as possible on the underlying distributions of these parameters, on which little is known. To do so, we embed a model of firm boundaries into a production network[35] that guides our identification of the key firm-level parameters, and our simulation of the efficient benchmark economy. In particular, we extend the baseline model so that a manufacturing supply chain can be composed of multiple stages (vertical), and each stage might use multiple inputs (horizontal). Each manufacturing firm, regardless of the stage of its final good, can produce its intermediate inputs in-house or outsource them to other firms, meaning that the make-or-buy decision of each firm jointly shapes the production network. Below, we lay out the main components of the model, while additional details can be found in appendix A.2. Moreover, in Appendix A.2.4, we discuss some of our assumptions and in appendix A.3, we relate our model to the existing production network and make-or-buy strands of literature.

Household. A representative household supplies L unit of labor inelastically, receives utility from consumption and is harmed by energy. The household consumes an aggregate bundle $C = (X, Y)$ of a numeraire good X and a manufacturing good Y, with quasi-linear preferences: $U(C, E) = X + \frac{Y^\alpha}{\alpha} + \tilde{U}(E)$, where $\frac{\partial \tilde{U}}{\partial E} < 0$. The numeraire good is perfectly competitive and solely produced with labor, $X = L^X$.[36] The manufacturing good Y aggregates the different products produced by the manufacturing firms.[37] In particular, $Y = \left(\sum_\omega \delta_\omega^{\frac{1}{\eta}} Y_\omega^{\frac{\eta-1}{\eta}}\right)^{\frac{\eta}{\eta-1}}$, where $\delta(\omega)$ is a product demand shifter, and where the product bundle for product ω aggregates the different varieties of that product produced by the N_ω firms that sell product ω: $Y_\omega = \left(\int_{j\in N_\omega} y_j(\omega)^{\frac{\varepsilon-1}{\varepsilon}} dj\right)^{\frac{\varepsilon}{\varepsilon-1}}$. 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 $P_\omega = \left(\int_{j\in N_\omega} p_j(\omega)^{1-\varepsilon} dj\right)^{\frac{1}{1-\varepsilon}}$ and the aggregate price index is $P = \left(\sum_\omega \delta_\omega P_\omega^{1-\eta}\right)^{\frac{1}{1-\eta}}$.

We suppose that, in equilibrium, energy prices p_e are set on international markets and taken as given by manufacturing firms and domestic energy retailers. The efficient benchmark will consider the fragmentation of production under the *feasibly* efficient energy price p_e^{SCC} (see below).

[35] Misallocation arising from make-or-buy decisions directly depends on the outside option, that is, which supplier would have been chosen in the outsourcing case. However, to correctly measure the energy-biased technology of this supplier, one needs to know its organizational structure as well. Therefore, quantifying misallocation requires building the entire production network, with heterogeneous firms at each stage choosing their boundary.

[36] This leaves $L^m = L^{tot} - L^X = 1 - L^X$ units of labor in the manufacturing sector.

[37] In our calibration, manufacturing sales to the household includes all public and non-manufacturing sectors that we don't explicitly model in the production network. This relies on the implicit assumption that these sectors are perfectly competitive, and that the household can either consume the manufacturing good directly or buy it as it is from these sectors. We could also explicitly model a retailer that works as an intermediary, combining the manufacturing products into the consumption good Y, as in [Eaton et al., 2022].

Labor markets are perfectly competitive, thus the wage is equalized across all sectors and we normalize it to 1. Manufacturing firms also behave competitively, hence make zero profit. The household owns all firms in the economy. Therefore, additional to labor income, it receives the revenue from energy retailers. Its budget constraint satisfies

$$P_x X + PY = wL^{tot} + p_e E + D_f$$

Since some intermediate inputs are purchased to foreign firms, we include the wedge D_f to eliminate the empirical trade imbalances, as in [Di Giovanni et al., 2024].[38] These functional form assumptions imply that

$$X = L^x = L^{tot} - L^m \quad P_x = w = 1 \quad PY = Y^\alpha$$

Therefore, manufacturing production is determined by the supply side, in which different sector will respond to cost changes differently depending on their demand shifts, and the numeraire good production absorbs the labor adjustments.

Our specification of welfare implies that $p_e^{SCC} = U'(E)$, which is endogenous to the equilibrium. It makes our benchmark evaluation very sensitive to potential measurement errors and factors or other distortions omitted from the model. Therefore, as in [Hsieh and Klenow, 2009], we prefer to consider an exogenous benchmark that we see as a feasible objective for Italy. In particular, we set this *feasible* efficient energy price at 170€ based on estimations of the SCC by [Nordhaus, 2017, Pindyck, 2019] and predictions for the ETS energy price in 2040.

Modules production. Supply chains can be interpreted as a tree of production modules. Each module, indexed by ω, is produced by combining energy and labor with a set Ω_ω of intermediate inputs produced at previous stages. In other words, Ω_ω corresponds to the recipe of input ω. To acquire an input $\hat{\omega} \in \Omega_\omega$ required in a module ω, a firm can either produce it in-house, following $\hat{\omega}$'s own production module, or outsource its production to a supplier. Figure 6 illustrates the production process. Formally, the production of input ω by firm j takes the form:

$$x_j(\omega) = \left(\left(\beta_\omega^{1-\rho(\omega)} (\nu_j e_j(\omega))^{\rho(\omega)} + (1-\beta_\omega)^{1-\rho(\omega)} l_j(\omega)^{\rho(\omega)} \right)^{\frac{\gamma}{\rho(\omega)}} + \int_{\hat{\omega} \in \Omega_\omega} (q_{ij} x_{ij}(\hat{\omega}))^\gamma d\hat{\omega} \right)^{\frac{1}{\gamma}} \tag{3}$$

Where $\rho(\omega)$ represents the degree of substitutability between energy and labor in the production of ω, and β_ω represents the share of energy needed in the production of input ω. As in the simple model, ν_j is firm j's energy-biased technology[39]

[38] As is standard in the literature, we keep the trade imbalances D_f invariant in the counterfactual exercises (e.g., [Allen and Arkolakis, 2014]).

[39] Importantly, energy-biased technologies are firms' characteristics. Numerous energy studies have shown that energy-efficiency improvements in manufacturing processes should be considered at the highest level of integration, due to the synergies across stages and the cross-cutting nature of the different options [Wang et al., 2008, Ramírez-Santos et al., 2018, Kim et al., 2022]. This is also consistent with the large literature on within-firm transfer of knowledge (e.g., [Markusen, 1984, Antràs, 2003, Keller and Yeaple, 2013, Atalay et al., 2014, Antràs and Yeaple, 2014, Ding et al., 2022, Boehm et al., 2022, Ding, 2023, Cai and Xiang, 2025]). For consistency with our covering of the entire manufacturing sector, we stay agnostic on the concrete nature of the energy-biased technology, such that it can be thought of as any firm-wide considerations that reduces the firm's energy intensity. Examples include carbon capture, utilization and storage technologies, scrap recycling, which type of energy source (fossil or renewable) the firm is using in its different plants, re-using of excess process fuels generated at upstream stages, or even the firm manager's environmental awareness, or R&D teams shared expertise.

and $q_{ij}x_{ij}(\hat{\omega})$ represents the quality-adjusted volume of input $\hat{\omega}$ that is used by firm j. In particular, q_{ij} depends on which firm produces input $\hat{\omega}$. We normalize $q_{jj} = 1$ if the input is produced in-house by firm j. This means that if $q_{ij} > 1$, outsourcing benefits, such as specialization gains, dominate and outsourcing increases the quality of the input while if $q_{ij} < 1$, outsourcing costs, such as the loss of control, dominate, and outsourcing decreases the quality of the input. Since q_{ij} reflects all the factors outside of the energy-biased technology that make in-house production and outsourcing differ, we refer to it as outsourcing capabilities.[40] γ represents the degree of substitutability in how the qualitative characteristics of the inputs frame the final quantity of the product.

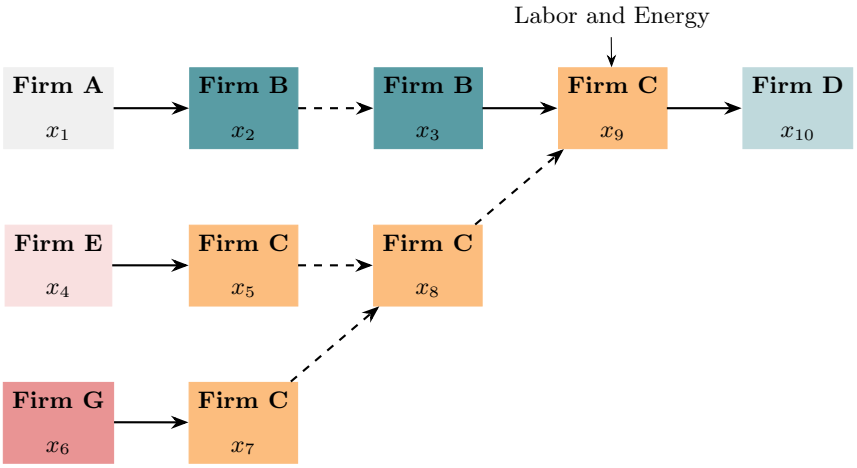


Figure 6: Illustration of the production network

Firm production. Firms produce modules according to the production function (3). Firms sell only one product, which we call their *core* product. The recipe of this core product determines which modules the firm requires. All firms have to perform the last step of their production in-house. In other words, a firm selling an input ω cannot outsource the production module ω. This last step is characterized by a firm-specific TFP ψ_j.[41] Additionally, they can buy each

[40] By staying agnostic on the exact composition of the outsourcing capabilities, and especially on their firm-level correlation with the energy-biased technology, our flexible model is equally capable to generate most of the empirical predictions of more complex models. This lack of empirical predictions would present a quantitative challenge in most settings. We are able to keep this generality because we calibrate each of our model firms based on a firm in our data, meaning that we don't have to match a distribution using aggregate parameters. The limitation of this approach is, of course, that it is not helpful to explain the microfoundations of firms' make-or-buy decisions and the potential trade-off between their outsourcing capabilities and their energy-biased technology. We explore this question in a companion paper [Marullaz, 2025] through the lens of energy-biased technology adoption, and nonetheless briefly discuss some potential avenues in Appendix **??**. It also means that our quantification is not immune to the Lucas' critique: if the energy-biased technology and the outsourcing capapbilities are not exogenous primitives of the firms, as we postulate, but respond to energy price changes, we will be missing this channel. We see our measure of misallocation as a lower bound, because it only captures the potential gains from reallocation at the current technology levels.

[41] Adding firm heterogeneity through a firm-specific TFP implies that the outsourcing capabilities q_{ij} only summarize outsourcing-specific factors but don't include TFP consideration, as the latter will be absorbed in the costs. In other words, it isolates the supplier productivity factor from the other outsourcing factors. This allows to reduce the importance of the calibrated q_{ij} to build the counterfactual, thus mitigating the impact of our identifying assumptions.

input $\hat{\omega} \in \Omega_\omega$ from a supplier,[42] or produce it in-house according to the same production function. If they produce $\hat{\omega} \in \Omega_\omega$ in-house, they can also produce the required inputs $\tilde{\omega} \in \Omega_{\hat{\omega}}$ in-house, and so on and so forth. In other words, the sequentiality of the production process implies that they can only produce consecutive stages in-house.[43] We define the set $\mathcal{I}_j$ as all inputs produced in-house by firm j, and the set $\mathcal{O}_j$ as all inputs that are directly purchased by the firm. Therefore, the production function of a firm selling module ω can be rewritten as

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

Where $x_j^{\mathcal{O}}(\omega) = \int_{\hat{\omega} \in \mathcal{O}_j} (q_{ij} x_{ij}(\hat{\omega}))^\gamma d\hat{\omega}$ is the quality-adjusted aggregate of all intermediate inputs that are directly outsourced by firm j ($\mathcal{O}_j$), while

$$x_j^{\mathcal{I}}(\omega) = \left(\int_{\hat{\omega} \in \mathcal{I}_j} \underbrace{(\beta_{\hat{\omega}}^{1-\rho(\hat{\omega})} (\nu_j e_j(\hat{\omega}))^{\rho(\hat{\omega})} + (1-\beta_{\hat{\omega}})^{1-\rho(\hat{\omega})} l_j(\hat{\omega})^{\rho(\hat{\omega})})^{\frac{\gamma}{\rho(\hat{\omega})}}}_{\equiv x_j^{\mathcal{I}}(\hat{\omega})^\rho} d\hat{\omega} \right)^{\frac{1}{\gamma}}$$

Importantly, because the energy-labor nest differs across input, we model the cost minimization problem for each input separately. In appendix A.2.1, we show that 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.

Timing. As it is standard in the misallocation literature, the model is static, and consists of two stages: a make-or-buy decision stage and a production stage. Firms observe their energy-biased technology 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 $i \in S_{j\hat{\omega}}$ of the module $\hat{\omega}$.[44] With this information, they decide whether to produce each module in-house or to outsource it.[45] Conditional on outsourcing it, they choose their preferred supplier among the set of potential suppliers. Once those make-or-buy decisions have been made, all firms produce and the household consumes.

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

$$c_j^{\mathcal{I}}(\hat{\omega}) = \left(c_j^{el}(\hat{\omega})^{\frac{-\gamma}{1-\gamma}} + \int_{\tilde{\omega} \in \Omega_{\tilde{\omega}}} c_j(\tilde{\omega})^{\frac{-\gamma}{1-\gamma}} d\tilde{\omega} \right)^{\frac{-(1-\gamma)}{\gamma}}$$

[42] The literature has consistently shown that firms are largely undiversified, that is, they have very few suppliers - usually only one - for each of their input (e.g., [Antras et al., 2017, Alfaro et al., 2025]). Diversification has even declined over 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 to a tariff shock.

[43] [Alfaro et al., 2019] show that firms integrate stages that are bunched together.

[44] Since the model is static, firms draw their outsourcing capabilities q_{ij} only once, and this set is fixed across different scenarios.

[45] Due to the sequentiality of production, firm j cannot produce module $\hat{\omega}$ in house if it has decided to outsource module ω located downstream to $\hat{\omega}$ (for instance, if $\hat{\omega} \in \Omega_\omega$). However, its decision is still characterized for each module, since this decision determines the potential in-house cost of the downstream module.

Where $c_j^{el}(\hat{\omega}) = \left(\beta_{\hat{\omega}}\left(\frac{p_e}{\nu_j}\right)^{\frac{-\rho(\hat{\omega})}{1-\rho(\hat{\omega})}} + (1-\beta_{\hat{\omega}})w^{\frac{-\rho(\hat{\omega})}{1-\rho(\hat{\omega})}}\right)^{\frac{-(1-\rho(\hat{\omega}))}{\rho(\hat{\omega})}}$ and $c_j(\tilde{\omega})$ is the effective cost of module $\tilde{\omega}$, that we now characterize. For each module that can be outsourced, the firm chooses separately[46] whether to produce it in-house or to outsource it to minimize its effective unit cost. Therefore,

$$\begin{aligned} c_j(\hat{\omega}) &= \min\left\{c_j^{\mathcal{I}}(\hat{\omega}), c_j^{\mathcal{O}}(\hat{\omega})\right\} \\ &= \min\left\{\left(\left(\beta_{\hat{\omega}}\left(\frac{p_e}{\nu_j}\right)^{\frac{-\rho(\hat{\omega})}{1-\rho(\hat{\omega})}} + (1-\beta_{\hat{\omega}})w^{\frac{-\rho(\hat{\omega})}{1-\rho(\hat{\omega})}}\right)^{\frac{(1-\rho(\hat{\omega}))\gamma}{(1-\gamma)\rho(\hat{\omega})}} + \int_{\tilde{\omega}\in\Omega_{\hat{\omega}}} c_j(\tilde{\omega})^{\frac{-\gamma}{1-\gamma}} d\tilde{\omega}\right)^{\frac{-(1-\gamma)}{\gamma}}, \right. \\ &\qquad \left. \min_{i\in\mathcal{S}_{\hat{\omega}}}\left(\frac{c_i(\hat{\omega})}{q_{ij}}\right)\right\} \end{aligned}$$

Since ν_j decreases $c_j^{\mathcal{I}}(\hat{\omega})$, firms with higher energy-efficient technology are less likely to outsource. In contrast, firms with high outsourcing capability draws will tend to outsource more.

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

$$c_j(\omega) = \frac{1}{\psi_j}\left(\left(\beta_\omega\left(\frac{p_e}{\nu_j}\right)^{\frac{-\rho(\omega)}{1-\rho(\omega)}} + (1-\beta_\omega)w^{\frac{-\rho}{1-\rho(\omega)}}\right)^{\frac{(1-\rho(\omega))\gamma}{(1-\gamma)\rho(\omega)}} + \int_{\hat{\omega}\in\Omega_\omega} c_j(\hat{\omega})^{\frac{-\gamma}{1-\gamma}} d\hat{\omega}\right)^{\frac{-(1-\gamma)}{\gamma}}$$

Energy intensity. We now summarize the key equilibrium conditions related to energy at the firm and at the aggregate level. Firm j's demand for energy and labor to produce any input ω is:

$$\begin{aligned} \frac{e_j(\omega)}{x_j(\omega_j)} &= \psi_j^{\frac{\gamma}{1-\gamma}}\beta_\omega\left(\frac{c_j(\omega_j)}{c_j^{el}(\omega)}\right)^{\frac{1}{1-\gamma}}\left(\frac{c_j^{el}(\omega)}{p_e/\nu_j^{\rho(\omega)}}\right)^{\frac{1}{1-\rho(\omega)}} \\ \frac{l_j(\omega)}{} &= \psi_j^{\frac{\gamma}{1-\gamma}}(1-\beta_\omega)\left(\frac{c_j(\omega_j)}{c_j^{el}(\omega)}\right)^{\frac{1}{1-\gamma}}\left(\frac{c_j^{el}(\omega)}{w}\right)^{\frac{1}{1-\rho(\omega)}} \end{aligned}$$

Where ω_j is the core product sold by firm j. Therefore, firm-level demands for energy and labor are:

$$\begin{aligned} \frac{e_j}{x_j(\omega_j)} &= \psi_j^{\frac{\gamma}{1-\gamma}}\int_{\omega\in\mathcal{I}_j}\beta_\omega\left(\frac{c_j(\omega_j)}{c_j^{el}(\omega)}\right)^{\frac{1}{1-\gamma}} c_j^{el}(\omega)^{\frac{1}{1-\rho(\omega)}}\left(\frac{p_e}{\nu_j^{\rho(\omega)}}\right)^{\frac{-1}{1-\rho(\omega)}} d\omega \\ \frac{l_j}{x_j(\omega_j)} &= \psi_j^{\frac{\gamma}{1-\gamma}}\int_{\omega\in\mathcal{I}_j}(1-\beta_\omega)\left(\frac{c_j(\omega_j)}{c_j^{el}(\omega)}\right)^{\frac{1}{1-\gamma}} c_j^{el}(\omega)^{\frac{1}{1-\rho(\omega)}} w^{\frac{-1}{1-\rho(\omega)}} d\omega \end{aligned}$$

[46] Since there is no fixed cost to outsource or to produce in-house, the decision for each module can be taken independently. This is somewhat equivalent to a world in which production is purely sequential in a snake form, and each establishment (stage) of a firm makes its own input decision independently, taking the upstream decisions as given. However, the decision is governed by the common firm-level energy-biased technology and the firm-level component in the outsourcing capabilities, representing technology diffusion within firm.

And in expenditure share:

$$\frac{p_e e_j}{c_j(\omega_j)x_j(\omega_j)} = \psi_j^{\frac{\gamma}{1-\gamma}} \int_{\omega\in\mathcal{I}_j} \beta_\omega \left(\frac{c_j(\omega_j)}{c_j^{el}(\omega)}\right)^{\frac{\gamma}{1-\gamma}} c_j^{el}(\omega)^{\frac{\rho}{1-\rho(\omega)}} \left(\frac{p_e}{\nu_j}\right)^{\frac{-\rho(\omega)}{1-\rho(\omega)}} d\omega$$

$$\frac{wl_j}{c_j(\omega_j)x_j(\omega_j)} = \psi_j^{\frac{\gamma}{1-\gamma}} \int_{\omega\in\mathcal{I}_j} (1-\beta_\omega) \left(\frac{c_j(\omega_j)}{c_j^{el}(\omega)}\right)^{\frac{\gamma}{1-\gamma}} c_j^{el}(\omega)^{\frac{\rho(\omega)}{1-\rho(\omega)}} w^{\frac{-\rho(\omega)}{1-\rho(\omega)}} d\omega$$

While for outsourced intermediate inputs:

$$\frac{c_j(\hat{\omega})x_{ij}(\hat{\omega})}{c_j(\omega_j)x_j(\omega_j)} = \psi_j^{\frac{\gamma}{1-\gamma}} \left(\frac{c_j(\omega_j)}{c_j(\hat{\omega})}\right)^{\frac{\gamma}{1-\gamma}}$$

Those expenditures correspond to the input expenditures shares that are observed in the data.

The total expenditure shares in energy and intermediate inputs show that their relationship can be driven by two forces: first, they might be positively correlated through scale, as both depend on $c_j(\omega)$. Secondly, they can be negatively correlated through substitution: ceteris paribus, $\frac{p_e e_j}{c_j(\omega_j)x_j(\omega_j)}$ increases with the set of inputs produced in-house, while total expenditure shares on all intermediate inputs will logically decrease with the number of inputs produced in-house. Which force dominates is ultimately an empirical question.

The aggregate energy required to produce one unit of GDP (energy intensity) amount to:

$$\begin{aligned}
\frac{E}{PY} &= \frac{1}{PY}\int_0^N \psi_j^{\frac{\gamma}{1-\gamma}} p_j(\omega_j)x_j(\omega_j) \int_{\omega\in\mathcal{I}_j} \left(\frac{p_e}{\nu_j^{\rho(\omega)}}\right)^{\frac{-1}{1-\rho(\omega)}} \beta_\omega \left(\frac{c_j(\omega_j)}{c_j^{el}(\omega)}\right)^{\frac{\gamma}{1-\gamma}} c_j^{el}(\omega)^{\frac{\rho(\omega)}{1-\rho(\omega)}} d\omega dj \\
&= \int_0^N \psi_j^{\frac{\gamma}{1-\gamma}} \left(\delta_{\omega_j}\left(\frac{P}{P_\omega}\right)^{\eta-1}\left(\frac{P_\omega}{c_j(\omega_j)}\right)^{\varepsilon-1} + \int_{i:j\in\mathcal{S}_i} \frac{c_i(\hat{\omega}_i)x_i(\hat{\omega}_i)}{PY}\left(\frac{c_i(\hat{\omega}_i)}{c_j(\omega_j)/q_{ij}}\right)^{\frac{\gamma}{1-\gamma}} di\right) \int_{\omega\in\mathcal{I}_j} \beta_\omega \left(\frac{c_j(\omega_j)}{c_j^{el}(\omega)}\right)^{\frac{\gamma}{1-\gamma}} c_j^{el}(\omega)^{\frac{\rho(\omega)}{1-\rho(\omega)}} \left(\frac{p}{\nu_j^{\rho}}\right. \\
&= \underbrace{\int_0^N \psi_j^{\frac{\gamma}{1-\gamma}} \delta_{\omega_j}\left(\frac{P}{P_\omega}\right)^{\eta-1}\left(\frac{P_\omega}{c_j(\omega_j)}\right)^{\varepsilon-1} \int_{\omega\in\mathcal{I}_j} \beta_\omega \left(\frac{c_j(\omega_j)}{c_j^{el}(\omega)}\right)^{\frac{\gamma}{1-\gamma}} c_j^{el}(\omega)^{\frac{\rho(\omega)}{1-\rho(\omega)}} \left(\frac{p_e}{\nu_j^{\rho(\omega)}}\right)^{\frac{-1}{1-\rho(\omega)}} d\omega dj}_{\text{Energy used to directly produce for the household}} \\
&+ \underbrace{\int_0^N \psi_j^{\frac{\gamma}{1-\gamma}} \int_{i:j\in\mathcal{S}_i} \frac{c_i(\hat{\omega}_i)x_i(\hat{\omega}_i)}{PY}\left(\frac{c_i(\hat{\omega}_i)}{c_j(\omega_j)/q_{ij}}\right)^{\frac{\gamma}{1-\gamma}} di \int_{\omega\in\mathcal{I}_j} \beta_\omega \left(\frac{c_j(\omega_j)}{c_j^{el}(\omega)}\right)^{\frac{\gamma}{1-\gamma}} c_j^{el}(\omega)^{\frac{\rho(\omega)}{1-\rho(\omega)}} \left(\frac{p_e}{\nu_j^{\rho(\omega)}}\right)^{\frac{-1}{1-\rho(\omega)}} d\omega dj}_{\text{Energy used for intermediates}}
\end{aligned} \tag{4}$$

Where $\mathcal{S}_i$ is the set of suppliers of firm i for the inputs it outsources (i.e., $\omega_j \in \mathcal{O}_i \; \forall j \in \mathcal{S}_i$).

Aggregate energy intensity ($\frac{E}{GDP}$) increases if inefficient make-or-buy decisions allocates production to low $\nu-$firms. At the firm level, both the firm boundaries ($\mathcal{I}_j$) and the suppliers selection distort the final cost $c_j(\omega_j)$ 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 $\int_{\omega\in\mathcal{I}_j} \beta_\omega c_j^{el}(\omega)^{\frac{\rho(\omega)-\gamma}{(1-\rho(\omega))(1-\gamma)}} \left(\frac{p_e}{\nu_j^{\rho(\omega)}}\right)^{\frac{-1}{1-\rho(\omega)}} d\omega$. It is therefore crucial to account for firm boundaries to compare the effective energy intensity of the firms. These results show that if the equilibrium set of inputs has very different energy-labor cost relative to the optimal set of inputs, misallocation is amplified. Even if all inputs are identically

produced ($\beta_\omega = \beta, \rho(\omega) = \rho \; \forall \omega$), the make-or-buy decision still amplifies misallocation if the *number* of inputs produced in-house vary.

Importantly, a counterfactually higher energy price affects energy intensity through the usual intensive margin effects by reallocating resources across factors, within firm, and across firms, as some firms are relatively more negatively affected. However, it also has an extensive margin effect that additionally reallocates production across firms, as firm boundaries extend or contract.[47] Therefore, we cannot use the popular exact hat algebra method used in international trade [Dekle et al., 2008] to quantify the change in energy intensity or welfare.

Finally, we close the model with labor market clearing, which implies that $X = L - L^m$, where $L^m = \int_0^N l_j dj$.

Multi-dimensional heterogeneity. The model incorporates multiple dimensions of heterogeneity. This is required by the absence of systematic empirical patterns in the data that would allow us to characterize firms' choices along one dimension of heterogeneity. We discuss each dimension case-by-case:

1. The energy-biased technology ν_j is firm-specific. This reflects the large within-product dispersion in energy intensity, that we associate to both a large dispersion in energy-biased technology, and different in-house inputs set $\mathcal{I}_j$.

2. The outsourcing capabilities q_{ij} are relationship-specific. In particular, we express them as a combination of buyer (z_j) and supplier (z_i) fixed effects, and some buyer-supplier controls (X_{ij}) such as the distance between the two firms: $q_{ij} = f(z_i, z_j, X_{ij})$

   (a) The buyer fixed effect z_j captures large within-product dispersion in in-house inputs set. Moreover, it allows to match the joint distribution between firms' effective energy efficiency (ν_j) and their make-or-buy practices (characterized by the set $\mathcal{I}_j$). To the best of our knowledge, we are the first to identify this empirical joint distribution.

   (b) The supplier fixed effect z_i matches the observation that buyers choose different suppliers to source the same input.[48]

3. The input energy share β_ω or the energy to labor elasticity of substitution $\rho(\omega)$ accommodate for different production process across inputs. This mitigates the implications of assuming that, within a firm, all inputs are produced with the same technology. One can thus think about ν_j as the overall effective energy efficiency of a firm that is applied to all inputs (e.g., fuel mix used, manager's preferences for sustainability, etc.), but its action is limited by the input-specific energy requirements driven by β_ω or $\rho(\omega)$.[49]

[47] In Appendix A.2.3, we provide an analytical decomposition of the different channels through which energy prices affect energy intensity.

[48] Should we observe in the data that there is a systematic empirical correlation between the suppliers size and their number of buyers, conditional on their costs, we will replace this fixed effect by a supplier-specific fixed cost to create links.

[49] In the current version of the calibration, we can only identify one product-specific parameter, so we set $\beta_\omega = \beta \; \forall \omega$ and capture product heterogeneity through $\rho(\omega)$. However, we are investigating different strategies to calibrate a varying β_ω across inputs.

4. The firm-specific TFP ψ_j is isolated from the outsourcing capabilities q_{ij} simply to clarify the intuition behind each parameter: the model is isomorphic to absorbing the firm-specific TFP within the outsourcing capabilities, but the separation allows to interpret the latter as outsourcing-specific factors only.

5 Calibration

The calibration is still work-in-progress, therefore, this draft does not contain any result. However, we describe our strategy below. Appendix B.2 provides more details on how we prepare the data.

We want to characterize the network of the Italian economy in terms of its energy intensity, which is driven by the allocation of energy consumption across firms that differ in their energy-biased technology ν_j and outsourcing capabilities q_{ij}. Assuming that those parameters would remain the same in the counterfactual scenario, we can then use them to characterize the efficient allocation and quantify the energy and welfare loss that arise due to the misallocation in the outsourcing decision.[50] Both parameters require identifying the set $\mathcal{I}_j$ of inputs produced in-house by each firm, but this set is not directly observed. Thus, in section 5.1, we first develop a novel strategy to infer it from our transaction data. As opposed to a measure of vertical span, a standard proxy for in-house production, our strategy recovers exactly *which* input is produced in-house. We leverage this feature to include input heterogeneity in our model. In section 5.2, we describe the identification of this parameter. Characterizing in-house production is key to accurately identify the effective energy efficiency of a firm. Armed with this measure, we can thus turn to calibrating the energy-biased technology ν_j in section 5.3. In section 5.4, we describe the strategy that allows us to retrieve the outsourcing capabilities for both the observed and unobserved (potential) pairs of buyers and suppliers. Finally, we calibrate the hicks-neutral productivity ψ_j in section 5.5, the demand shifter in section 5.6 and discuss how we build the counterfactual allocation in section 5.7.

As discussed above, we make standard assumptions on the network that facilitate the numerical resolution of the model.

1. The network is acyclic, which guarantees the uniqueness of the equilibrium [Dhyne et al., 2023].[51]
2. The number of production module is finite.
3. All firms only sell *one* core product (but may produce several of their inputs in-house), and this core product is always produced in-house.

[50] In our framework, the factors that drive the make-or-buy decision, namely the effective unit cost of outsourcing $\{\frac{c_i(\omega)}{q_{ij}}\}_{i\in\mathcal{S}_j}$ and the energy-biased technology ν_j are not perfectly correlated. Indeed, one of our contribution is to empirically retrieve the joint distribution between outsourcing capabilities and energy efficiency, which, to our knowledge, is unknown. This means that we cannot characterize the make-or-buy decision based on a continuous object (for instance, the share of in-house inputs out of all required inputs). In turn, this prevents us to use standard welfare evaluation tools such as the exact hat algebra [Dekle et al., 2008]. Instead, we leverage our rich firm-level data to exactly calibrate the firm-level parameters, taking each empirical observation as one model unit.

[51] [Dhyne et al., 2023] show that the Belgium production network is consistent with an acyclic network structure for more than 80% of the observed transactions, and that dropping inconsistent transactions has little impact on sectoral linkages and aggregate measures.

4. Firms only have one supplier for each input.[52]

We start by exogenously calibrating $\eta-1=\varepsilon-1=1.5$ as in [Boehm and Oberfield, 2023]. Following [Baqaee and Farhi, 2020], we set $\gamma=-4$ based on both [Boehm et al., 2014] and [Atalay, 2017]'s estimations.[53] We set $\alpha=0.75$ to match the price elasticity of demand in manufacturing [Kreinin, 1967]. We normalize the wage to 1 in all scenarios. Finally, we set the equilibrium energy price price to the average energy price faced by manufacturing firms.[54]

5.1 Measuring in-house production

The set $\mathcal{I}_j$ summarizes the set of inputs that is produced in-house by the firm, among all inputs that are required to produce the firm's output. Intuitively, it is inversely related to the total cost share of intermediate input (net of energy) expenditures, which corresponds to the inputs that are directly outsourced by the firms. However, the intermediate input expenditure share of total costs is a naive statistics for the degree of outsourcing, because it abstracts from firm-level heterogeneity. Indeed, recall that it can be expressed as $\frac{p_j^{\mathcal{O}} x_j^{\mathcal{O}}}{c_j x_j} = \left(\frac{c_j}{p_j^{\mathcal{O}}}\right)^{\frac{\gamma}{1-\gamma}}$. One needs to make the restrictive assumptions that $\nu_j=\nu\ \forall j$, $q_{ij}=q\ \forall ij$, and shut down input heterogeneity ($\beta_\omega=\beta, \rho(\omega)=\rho\ \forall\omega$), which result in identical prices for all firms, for the total cost share of intermediate input expenditures to correspond to the number of outsourced inputs ($\frac{p_j^{\mathcal{O}} m_j^{\mathcal{O}}}{c_j x_j} = |\mathcal{O}_j|$). However, as our stylized facts in section 2 demonstrate, firms are extremely heterogeneous in terms of both energy efficiency and firm boundaries.

To account for firm-level heterogeneity, we instead need to understand *which* input is outsourced *by which* firm and *to which* firm.[55] To achieve this, we reconstruct the complete set of inputs — both direct and indirect - that enter the production of each product (see figure 7 for an illustration). This set, which we refer to as a product's *recipe* (Ω_ω), is inferred from our firm-to-firm transaction data structured as a directed acyclic graph (DAG), where edges represent the use of an input in the production of another good. Each node corresponds to a product, and a directed edge from product $\hat{\omega}$ to product ω indicates that $\hat{\omega}$ is used in the production of ω.

The key underlying assumption behind this strategy is that all firms producing product ω use the same recipe (i.e., need the same inputs) (as in e.g., [Alfaro et al., 2019, Boehm and Oberfield, 2020, Rachapalli, 2021, Alfaro et al., 2024]).[56] This allows us to infer all firms' input requirements, and therefore to deduce, from their observed intermediate inputs

[52] The literature has consistently shown that firms are largely undiversified, that is, they have very few suppliers - usually only one - for each of their input (e.g., [Antras et al., 2017, Alfaro et al., 2025]). Diversification has even declined over 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 to a tariff shock.

[53] This low elasticity of substitution is a conservative measure, because, as shown by [Baqaee and Farhi, 2020], the effect of reallocation is amplified by the elasticity of substitution between firms.

[54] In a robustness exercise, we compare our results to allowing for firm-specific wage, measured as wage bill divided by total labor, and to building a firm-specific energy price based on a weighted average of national fuel prices, weighted by the fuel share in the total energy expenditures of the firms.

[55] The literature (e.g., [Alfaro et al., 2019, Boehm and Oberfield, 2020]) has often proxied for in-house production by measuring the vertical span of the firm. However, it is likely that inputs are not identically produced (i.e., they have different β_ω or $\rho(\omega)$). In that case, the number of stages that are produced in-house (the vertical span) is not enough: we need to identify exactly which input is produced in-house to understand how it affects the firms' energy demand.

[56] In other words, the underlying technology to produce a product is exogenous and identical across firms, but the decision on how to produce, that is, which inputs to produce in-house, is endogenous, thereof determining the firms' production function. Note that our CES framework

expenditures which input is produced in-house and which one is outsourced, out of all inputs entering the recipe.[57]

Recipes. To characterize the recipes, we use our firm-to-firm transaction data and construct a DAG that links all products to their direct and indirect inputs, avoiding cycles. Starting from the observed direct input relationships, we construct a *forward graph* mapping each product to its immediate inputs. We then recursively traverse this graph to recover the entire transitive closure of inputs for each product. Importantly, we ensure acyclicity both ex ante and ex post: first, by filtering transaction-level data to remove feedback loops (through eliminating weakly connected components, as measured by average transaction shares, that create cycles, and defining a set of terminal nodes); and second, 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.

Standard DAG understands 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 might buy input $\hat{\omega}$ to produce $\tilde{\omega}$ in-house, and use $\tilde{\omega}$ to produce ω_j. 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's implied by a longer path.

Using this in-house-consistent DAG, we can recover two crucial sets of inputs for each product j. Defining $tr_{ij} = 1$ if input i is connected to product j in the DAG and l_{ij} as the number of links between input i and product j, we first build the recipe $\Omega_j = \{i : tr_{ij} = 1\}$ as the set of inputs that are required, directly or indirectly, to produce input j (upstream set). Then, we build $\mathcal{D}_j = \{i : tr_{ji} = 1\}$ as the set of products that can be produced, directly or indirectly, using input j (downstream set). Finally, we build $\mathcal{U}_j^{direct} = \{i : tr_{ij} = 1 \ \& \ l_{ij} = 1\} \subseteq \Omega_j$ as the set of inputs that only *directly* enter the production of input j. With these different elements, combined with the set $\mathcal{O}_j$ of inputs purchased by the firms, we can identify the set of inputs that are produced in-house by our firms.

In-house product set. The set of inputs produced in-house by a firm producing product ω is the intersection between the set of inputs that are upstream to product ω (i its recipe Ω_ω) and the set of products that are downstream the products that are purchased by the firm ($\cup_{\hat{\omega} \in \mathcal{O}_j} \mathcal{D}_{\hat{\omega}}$), or in a branch that is entirely produced in-house. To ensure that we are capturing all inputs produced in-house, we move input-by-input.[58] In effect, we proceed as follow. We build an algorithm that moves upstream the supply chains starting from the firms' product:

1. For each firm j, we start from product ω_j and consider all inputs $\hat{\omega}$ that are directly upstream of product ω_j

implies that the expenditure shares of each input varies across firms. Therefore, we only impose that the *set* of inputs in a recipe is the same for all firms producing the same product. This also means that we build recipes as a set of inputs, and not as a vector of expenditure shares.

[57] Data limitations force us to characterize a product by its ATECO six-digit code. This empirical limitation is unavoidable and is finer than the standard four-digit level used in the literature (e.g., [Alfaro et al., 2019, Alfaro et al., 2024, Huneeus et al.,]). Our framework relies on the crucial assumption that the same product produced by different firms always require the same direct and indirect inputs (recipes). Six-digit codes might contain more granular products that require different recipes. Nonetheless, if those products, within a six-digit classification, are sufficiently similar, their recipes will overlap sufficiently to limit the measurement errors. We provide empirical evidence supporting the large overlap within more granular categories in Appendix B.3.1.

[58] In appendix B.3, we illustrate the potential limitations that may arise if we don't move recursively.

($\hat{\omega} \in \mathcal{U}_{\omega_j}^{direct}$. These inputs are produced in-house (i.e., added to the set $\mathcal{I}_j$) if they are not purchased by the firm ($\hat{\omega} \notin \mathcal{O}_j$).

2. After having identified the in-house inputs $\hat{\omega} \in \mathcal{I}_j$, we move upstream. We 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 $\tilde{\omega} \in \mathcal{U}_{\hat{\omega}}^{direct}$). Again, these inputs are produced in-house (i.e., added to the set $\mathcal{I}_j$) if they are not purchased by the firm ($\tilde{\omega} \notin \mathcal{O}_j$). In effect, we stop our algorithm every time we reach a purchased input or a terminal node.

3. We continue to move upstream until we are not adding any new input to the set of in-house inputs.

This recursive algorithm guarantees that we cover all nodes of the supply chain and that we don't miss out any input.

5.2 Product-specific parameter

In the current version of the calibration, we can only identify one product-specific parameter, so we set $\beta_\omega = \beta \; \forall \omega$ and capture product heterogeneity through $\rho(\omega)$. However, we are investigating different strategies to calibrate a varying β_ω across inputs. $\rho(\omega)$ corresponds to the energy to labor elasticity of substitution of a given product ω. [Bretschger and Jo, 2021a] 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 strategies to identify the elasticity parameter $\rho(\omega)$.

Strategy I: Using [Bretschger and Jo, 2021b] estimates. [Bretschger and Jo, 2021b] report the elasticity of substitution for 19 NACE 2−digit manufacturing industries. We use their estimates as a first calibration of our $\rho(\omega)$, 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, of course, highly restricts product heterogeneity. Moreover, [Bretschger and Jo, 2021b] ignore in-house production. Formally, they estimate, for each industry separately

$$log\frac{e_{jt}}{l_{jt}} = \delta_0 + \underbrace{\frac{1}{1-\rho(\omega)}}_{\equiv \delta_1} log\Big(\frac{w_{jt}}{p_{e_{jt}}}\Big) + \delta_j + \delta_t + \xi_{jt}$$

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

$$log\frac{e_j}{l_j} = log\frac{\int_{\omega\in\mathcal{I}_j}(c_j^{el}(\omega))^{\frac{1}{1-\gamma}}\left(\frac{c_j^{el}}{p_e}\right)^{\frac{1}{1-\rho(\omega)}}\nu_j^{\frac{\rho(\omega)}{1-\rho(\omega)}}d\omega}{\int_{\omega\in\mathcal{I}_j}(c_j^{el}(\omega))^{\frac{1}{1-\gamma}}\left(\frac{c_j^{el}}{w}\right)^{\frac{1}{1-\rho(\omega)}}d\omega}$$

Which is non-linear in $\frac{w}{p_e}$. Moreover, 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 only identifies the true elasticity of substitution if all inputs ω produced in-house by their firms share a common elasticity of substitution. Due to the large variations in the types of inputs that are produced in-house even by similar firms, this assumption does not seem very credible. Therefore, we use insights from our model to develop a strategy that estimates model-consistent elasticity of substitution $\rho(\omega)$ for each of our products, accounting

for in-house production.

Strategy II: Estimating $\rho(\omega)$ on outsourcing firms. Recall that in our model, the energy to labor expenditure ratio takes the form:

$$\frac{p_e e_j}{w l_j} = \frac{\int_{\omega \in \mathcal{I}_j} \left(\frac{c_j(\omega_j)}{c_j^{el}(\omega)}\right)^{\frac{\gamma}{1-\gamma}} \left(\frac{c_j^{el}}{p_e/\nu_j}\right)^{\frac{\rho(\omega)}{1-\rho(\omega)}} d\omega}{\int_{\omega \in \mathcal{I}_j} \left(\frac{c_j(\omega_j)}{c_j^{el}(\omega)}\right)^{\frac{\gamma}{1-\gamma}} \left(\frac{c_j^{el}}{w}\right)^{\frac{\rho(\omega)}{1-\rho(\omega)}} d\omega}$$

This is the moment we will use to calibrate ν_j (see section 5.3), but this requires calibrating the elasticity $\rho(\omega)$ first, and ν_j and $\rho(\omega)$ interact in a highly non-linear way. However, for the subsample of single-product firms that only produce their core product ω_j in-house ($\mathcal{I}_j = \{\omega_j\}$), this simplifies to

$$log\frac{p_e e_j}{w l_j} = \frac{\rho(\omega)}{1-\rho(\omega)} log\left(\frac{w}{p_e}\right) + \frac{\rho(\omega)}{1-\rho(\omega)} log\nu_j$$

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 2021 to avoid the COVID-pandemic shock), so we 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 $\mathbf{1}\{\mathcal{I}_j = \{\omega\}\}$, 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 $\mathbf{1}\{\mathcal{I}_j = \{\omega\}\}$ over the two periods, under the (strong) assumption that the energy-biased technology remain fixed over time. We therefore additionally control for important fuel mix changes in the hope of partially capturing any energy-biased technology changes. In effect, we estimate

$$\Delta log\frac{p_{e_{jt}} e_{jt}}{w_{jt} l_{jt}} = \frac{\rho(\omega)}{1-\rho(\omega)} \Delta log\left(\frac{w_{jt}}{p_{e_{jt}}}\right) + \delta' \mathbf{1}\{\text{fuel mix change}\} + \eta_{jt}$$

Where η_{jt} is an iid error term which is not correlated with any component driving the make-or-buy decision.

Strategy III: Robustness using French data. We use the Annual production survey (EAP) over the years 2010 to 2020 (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 $\mathcal{I}_j = \{\omega_j\}$ 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 5.1. With longer time periods, we can include firm and time fixed effects. To run our regression, we complement this dataset with three 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, Fontaine and Marullaz, 2024]. We also use the balanced-sheet data (FARE-FICUS), from which we retrieve wage and labor expenditures. Additionally, we can use the Environmental protection studies and investments (ANTIPOL), that provides information on energy efficiency investment for a representative sample of firms. We use the latter as an additional control in our regression, as it proxies for energy-biased technology investment. However, because the EAP and ANTIPOL are both non-exhaustive, representative samples, with a medium overlap, the inclusion of this control greatly reduces our sample.

5.3 Energy-biased technology

We can then estimate the energy-biased technology ν_j for all firms using the energy-to-labor expenditure share,[59]. In particular, we calibrate each ν_j to match:

$$\frac{p_e e_j}{w l_j} = \frac{\int_{\omega \in \mathcal{I}_j} \left(\frac{c_j(\omega_j)}{c_j^{el}(\omega)}\right)^{\frac{\gamma}{1-\gamma}} \left(\frac{c_j^{el}}{p_e/\nu_j}\right)^{\frac{\rho(\omega)}{1-\rho(\omega)}} d\omega}{\int_{\omega \in \mathcal{I}_j} \left(\frac{c_j(\omega_j)}{c_j^{el}(\omega)}\right)^{\frac{\gamma}{1-\gamma}} \left(\frac{c_j^{el}}{w}\right)^{\frac{\rho(\omega)}{1-\rho(\omega)}} d\omega} \tag{5}$$

Equation (5) highlights the importance of taking in-house production into account.[60] Intuitively, we are able to identify ν_j as the energy-biased technology only if we control for exactly all the products that are *produced* by the firm, and not only the final product that is used to classify the firm into its sector.[61]

5.4 Outsourcing capabilities

After having recovered ν_j and $\mathcal{I}_j$ for every firm, we can characterize the distribution of the outsourcing capabilities q_{ij}. Recall that we observe, for each input that is outsourced or produced in-house (i.e., located downstream of the outsourced inputs), the chosen supplier among the set of all potential suppliers, which includes the buyer herself through in-house production. This is essentially an issue of selection, reminiscent to the labor literature [Heckman, 1990, Heckman and Honore, 1990] and an important focus of network econometrics [Mayer and Puller, 2008, Fafchamps and Gubert, 2007] However, three main challenges arise in our setting.[63] First, to build the counterfactual efficient network, we need to recover the outsourcing capabilities not only for each observed pairs ij that trade, but also all potential $ij-$pairs that might trade under another allocation, since counterfactual allocations might connect different customer-supplier

[59] In practice, we use manufacturing labor and exclude management labor to be consistent with our model.

[60] Our model highlights that estimates that don't account for in-house production (e.g., [Carraro and De Cian, 2013, Hawkins-Pierot and Wagner, 2023]) confound the energy-biased technology with the degree of outsourcing.

[61] This is still limited by the fact that we only observe product code at the 6−digit level, which might hide quality or input requirement heterogeneity. Unfortunately, a finer level of product classification doesn't exist in firm-to-firm transaction data, and we follow the literature by assuming no within- 6−digit product heterogeneity.

[62] Dyadic models of network formation conceptualize network formation as the result of an optimal choices of individuals who maximize their objective function. The statistical analysis of network formation then relies on a dyadic non-linear regression, where the decision to establish a link between two agents is modeled as a binary choice driven exclusively by their pair-specific attributes. [Graham, 2015]

[63] We also discuss finite sample considerations in Appendix B.5.

pairs relative to the observed allocation.[64] Inputs for which we observe the outsourcing decision (i.e., inputs that are outsourced or downstream to those) are well-characterized by dyadic models with agent heterogeneity and homophily [Graham, 2017, Dzemski, 2019, Yan et al., 2019]. However, the sequentiality of our production function implies that, for inputs located upstream of the outsourced inputs, we don't observe the outsourcing decision of the buyer firm, preventing us to use these pairs to identify any parameter. Secondly, as opposed to standard dyadic models, each buyer is limited to only one supplier (including herself for in-house production) per input. This implies that the outside option that determines the payoff of forming a link is the next cheapest potential supplier, which is unknown for the supplier selected in the data. Those two challenges prevent us to use the usual tools and econometric results developed in network econometrics. The last challenge is that there are two levels of selection: the supplier selection, conditional on outsourcing, and the decision to outsource or produce in-house. These distinct decisions break the symmetry between the different options, which renders the use of the usual tools of discrete-choice models [McFadden, 1974, Berry, 1994, Eaton and Kortum, 2002] inapplicable (see Appendix **??** for a more thorough discussion). Moreover, because the outside option is relationship-specific, the selection term is also relationship-specific, meaning that it would be absorbed by any supplier and buyer fixed effects. This means that we cannot consistently build the outsourcing capabilities without accounting for the selection issue.

We experiment with two methods. As a benchmark, we assume that selection is purely based on observables. Alternatively, we allow for selection on unobservables and follow a revealed preference argument to extract some information on the unobserved outsourcing capabilities. Albeit it is more complex, we start by describing the latter strategy (section 5.4.1) as some key concepts are used again in the former (section 5.4.2).

5.4.1 Selection on unobservables

Our limited set of observables can cast doubts on the validity of our identifying assumption that selection is based purely on observables. Thus, we implement a strategy that allows for selection on unobservables.

Before describing our strategy to deal with selection, it is useful to summarize the different types of relationships that we try to characterize:

1. **Observed transactions:** Firm j buys input ω from firm i and the transaction is directly observed in the data. We will calibrate the q_{ij} characterizing those relationships directly using expenditure shares.[65]

2. **Observed decisions:** Firm s produces 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. We will calibrate the *distribution* of the q_{sj} characterizing those relationships by leveraging the properties of the Pareto distribution, and draw a q_{sj} from these distributions.

[64] Since the extend of misallocation induced by the outsourcing decision depends on the energy-biased technology of the buyer relative to its supplier's, it is necessary to identify the exact supplier in each allocation.

[65] With a hicks-neutral productivity ψ_j, we cannot disentangle q_{ij} from ψ_j using expenditure shares. However, we can still build the suppliers' costs for each buyer firm, as q_{ij} and ψ_j enter in the same way. This allows us to build the cost of the core product of each firm, which we use to identify ψ_j using household sales. Once we've identified the ψ_j, we can retrieve q_{ij} net of the ψ_j.

3. **Unobserved decisions:** Firm s produces input $\hat{\omega} \in \Omega_\omega$, but firm j outsources input ω. Therefore, firm j's internal make-or-buy decision on input $\hat{\omega}$ is not observed. We will calibrate the q_{sj} characterizing those relationships after having estimated the buyer and supplier fixed effects using the q_{ij} and q_{sj} recovered in case the previous steps.

Moreover, we need a sufficiently flexible characterization of the q_{ij} such that it can match its empirical joint distribution with the buyer and supplier's energy-biased technology ν_j, ν_i.

We now explain our strategy in words, before laying out the formulae we use. A more detailed description of our algorithm can be found in appendix B.4. Our strategy is inspired by the consumer concept of revealed preference, applied to firms' choice. The observed trade flows and internal production choices are informative on the distribution of the outsourcing capabilities that characterize the counterfactual transactions that we don't observe. We utilize this information to build a consistent representative sample of outsourcing capabilities, that we use to recover the q_{ij} for each potential pair of buyer and supplier.[66] Formally, we leverage the properties of truncated Pareto distributions to recover the firm-input specific distribution of outsourcing capabilities q_{ij} that match the observed fragmentation decisions of the firms. This functional form assumption allows us to recover the outsourcing capabilities for all potential buyer-supplier links of each input for which we *observe* the make-or-buy decision. In turn, this protects us from the sparse network issue [Graham, 2024] and from sample selection, meaning that we can identify the outsourcing capabilities for all the buyer-supplier pairs of each input for which we *don't observe* the make-or-buy decision, without relying on a large set of observables. In particular, we assume that $\frac{1}{q}$ is Pareto distributed with a supplier-buyer specific scale parameter and a common shape parameter θ.[67]

In particular, we model $\frac{1}{q_{ij}} = z_{ij} = e^{z_i^s + z_j^b + \mathbf{X}'_{ij}\delta} e^{\epsilon_{ij}}$ for $i \neq j$,[68] where $\mathbf{X}_{ij}$ controls for observables (e.g., distance), $\epsilon_{ij} \sim Exp(\theta)$ is an iid random variable uncorrelated across firms, and z_i^s, z_j^b are respectively supplier and buyer fixed effects.[69] This specification implies that $\frac{1}{q} \sim Pareto(e^{z_i^s + z_j^b + \mathbf{X}'_{ij}\delta}, \theta)$[70] and relies on the assumption that the controls, supplier and buyer fixed effects capture all the interdependence across input outsourcing capabilities. We start by exter-

[66] A key feature of our approach is that we calibrate our parameters to exactly match all model decisions for which we observe a corresponding decision in the data. For those buyer-supplier relationships, a model unit corresponds to an actual empirical relationships. In other words, our data *is* the model, and not one realization of our model. This allows us to fully match the empirical joint dispersion of Italian firms in energy use and firm boundary decisions, with limited parameter restrictions on the underlying distributions of their primitives, while at the same time discussing predictions for specific individual Italian firms.

[67] The Pareto distribution assumption might generate chosen suppliers with much larger q_{ij} than the second best choice. This has counteracting implications on our aggregate measure of misallocation. On the one hand, it reduces the differences in decisions across the equilibrium and the counterfactual scenario, since very large energy-biased technology differences would be required to compensate the large differences in relative outsourcing capabilities. This downplays the measured misallocation. On the other hand, if the equilibrium and the counterfactual decisions do differ, the loss in q_{ij} will be large, negatively affecting total production, which amplifies the overall differences in aggregate energy intensity across the two scenarios. Overall, the direction of the bias is unclear. We mitigate it by taking multiple draws and keeping the average q_{ij}.

[68] We impose that $q_{jj} = 1 \ \forall j$, meaning that it cannot come from the same distribution. This asymmetry, also present in [Boehm and Oberfield, 2023], highlights that in-house production can differ in many ways from outsourcing. Empirically, $q_{jj} = 1 \ \forall j$ is required since we don't observe any input-specific transactions for in-house production.

[69] This specification allows for $z_j^b \neq z_j^s$, that is the fixed effect of j as a supplier, in transaction of its product with its buyers, can differ from its fixed effect as a buyer, in the transactions with its suppliers.

[70] In other words, this specification allows for the distribution of outsourcing capabilities to be buyer-supplier specific, while still being random

nally calibrating $\theta = 2$ following the literature.

If we estimate the fixed effects using only the observed transactions q_{ij} (case 1), selection would arise. Thus, we need to first identify the q_{ij} for all the potential suppliers of an input, that is, the q_{sj} for the observed decisions (case 2). To do so, we calibrate the *distribution* of those outsourcing capabilities, from which we can subsequently draw a value for each q_{sj}. The intuition is as follows. If input ω is produced in-house by firm j, we know that $\frac{1}{q_{ij}} > \frac{c_j^{\mathcal{I}}(\omega)}{c_i(\omega)}\ \forall i \in \mathcal{S}_j(\omega)$. We also know that, if $\frac{1}{q} \sim Pareto(\bar{z}, \theta)$, then $\frac{1}{q}|\frac{1}{q} > \frac{1}{q_m}$, where $\frac{1}{q_m} > \bar{z}$, is Pareto distributed with scale parameter q_m and the same shape parameter θ. In other words, the *conditional* distribution of $\frac{1}{q}$ is $Pareto(\frac{1}{q_m}, \theta)$. Therefore, conditional on observing that the firm is producing input ω in-house, we know that the actual distribution of outsourcing capabilities $\frac{1}{q}$ of that input ω is $Pareto\left(\frac{c_j^{\mathcal{I}}(\omega)}{c_i(\omega)}, \theta\right)$. We can then draw the inverse outsourcing capability $\frac{1}{q_{ij}}$ from that distribution to characterize the potential relationship between supplier i and buyer j. Since the q_{ij} and the distributions are fixed across scenarios, we can use those drawn q_{ij} to assess whether firm j's outsourcing decision changes under a different energy price, which changes the relative cost ratio. A similar reasoning applies to an input that is outsourced to firm i. It implies that $\frac{1}{q_{ij}} > \frac{c_i(\omega)/q_{ij}}{c_s(\omega)} \forall i \in \mathcal{S}_j(\omega)$. Therefore, we know that all the distribution of outsourcing capabilities of that input is Pareto with shape parameter $\frac{c_i(\omega)/q_{ij}}{c_s(\omega)}$. By characterizing the distribution using the input producer's cost, we fully capture the correlation between outsourcing capabilities and energy-biased technology, and ensure that our simulations are always consistent with the data.[71]

As a final note, our strategy relies on measuring $c_j^{\mathcal{I}}(\omega)$ and $c_i^{\mathcal{O}}(\omega)$ for all inputs, 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.

5.4.2 Selection on observables

As a benchmark comparison, we also estimate the fixed effects assuming that selection is purely based on observables. In particular, firms observe their suppliers' costs and their outsourcing capability q_{ij}, which we express as

$$q_{ij} = exp(z_i^s + z_j^b + \mathbf{X}_{ij}\delta)$$

Where z_i^s, z_j^b are supplier and buyer fixed effects, X_{ij} controls for distance and degree of input specialization. Thus, we can estimate $\{\{z_i^s\}_i, \{z_j^b\}_j, \delta\}$ using a simple OLS regression of the

$$log(q_{ij}) = z_i^s + z_j^b + \mathbf{X}_{ij}\delta + \varepsilon_{ij}$$

thanks to the ϵ_{ij} component. q_{ij} will be higher if z_i^s or z_j^b are small, that is for firms that tend to have an advantage in outsourcing. This implies that we allow for some correlation between the firm-specific energy-biased technology ν_j and their relative advantage in outsourcing as a buyer or as a supplier.

[71] The truncated Pareto distribution guarantees that our drawn outsourcing capabilities always match the observed decisions, except for the decision to outsource an input relative to producing it in-house, since we don't calibrate a $q_{jj}(\omega)$ for these decisions. However, the decision relies on the potential in-house unit cost $c_j^{\mathcal{I}}(\omega)$ of that input, which is unknown, since it depends on unobserved make-or-buy decision for upstream inputs. In that sense, our simulations cannot be empirically invalid for that decision either.

Where our outcome variable q_{ij} has been identified using the expenditure share of the observed transactions, and ε_{ij} is an iid measurement error term *unobserved* by the firm during selection process. Formally, ε_{ij} is uncorrelated with the selection decision and is mean-zero, conditional on observables: $E(\varepsilon_{ij}|i,j,X_{ij},i \text{ selected by } j) = E(\varepsilon_{ij}|i,j,X_{ij}) = 0$. This means that

$$\begin{aligned} E(log(q_{ij})|i,j,X_{ij},i \text{ selected by } j) &= \hat{z}_i^s + \hat{z}_j^b + X_{ij}\hat{\delta} + E(\varepsilon_{ij}|i,j,X_{ij},i \text{ selected by } j) \\ &= \hat{z}_i^s + \hat{z}_j^b + X_{ij}\hat{\delta} \end{aligned}$$

We then build the remaining unobserved q_{sj} as

$$q_{sj} = exp(\hat{z}_s^s + \hat{z}_j^b + \mathbf{X}_{sj}\hat{\delta})$$

This method is valid if the regression delivers a high R^2, as it means that selection can be largely explained by our observables.

5.5 Hicks-neutral productivity

Equipped with a firm's energy-biased technology, the outsourcing capabilities, the product-specific parameters and the set of input produced in-house, we can recursively build the unit cost for its core product, as described in appendix B.4.. This correspond to the price it charges.

Once we've build the costs of all firms, we identify the hicks-neutral productivity ψ_j[72] using sales share to household. In particular, the household sales ratio of two firms producing the same product and selling it to the household is

$$\frac{p_j(\omega)x_j(\omega)^{hh}}{p_i(\omega)x_i(\omega)^{hh}} = \left(\frac{p_j(\omega)}{p_i(\omega)}\right)^{1-\varepsilon} = \left(\frac{c_j(\omega)}{c_i(\omega)}\right)^{1-\varepsilon}$$

Therefore, we calibrate $\frac{\psi_i}{\psi_j}$ using

$$\frac{\psi_i}{\psi_j} = \left(\frac{p_j(\omega)x_j(\omega)^{hh}}{p_i(\omega)x_i(\omega)^{hh}}\right)^{\frac{-1}{(1-\varepsilon)}} \frac{\left(\int_{\hat{\omega}\in\mathcal{I}_j} c_j^{el}(\hat{\omega})^{\frac{-\gamma}{1-\gamma}} d\hat{\omega} + \int_{\hat{\omega}\in\mathcal{O}_j} \left(\frac{c_i(\hat{\omega})}{q_{ij}}\right)^{\frac{-\gamma}{1-\gamma}} d\hat{\omega}\right)^{\frac{-(1-\gamma)}{\gamma}}}{\left(\int_{\hat{\omega}\in\mathcal{I}_j} c_j^{el}(\hat{\omega})^{\frac{-\gamma}{1-\gamma}} d\hat{\omega} + \int_{\hat{\omega}\in\mathcal{O}_j} \left(\frac{c_i(\hat{\omega})}{q_{ij}}\right)^{\frac{-\gamma}{1-\gamma}} d\hat{\omega}\right)^{\frac{-(1-\gamma)}{\gamma}}}$$

As discussed above, this means that we actually identify $\tilde{q}_{ij} = q_{ij}\psi_j$ for the observed transactions. We can then use our measure of ψ_j to filter it out from $\tilde{q}_{ij}$ and retrieve the true q_{ij}. 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 (e.g., transport costs, specialization gains, loss of controls over the supply chains, relationship building, etc.). Therefore, after having identified the observed q_{ij} but before running the PPML regression, we identify the ψ_j using our

[72] Additionally to the outsourcing capabilities q_{ij} and the energy-biased technology ν_j, we can either identify the firm-specific hicks-neutral technology (TFP), or the firm-specific in-house-inputs-biased technology q_{jj}. This is because in-house input transactions are never observed, thus they must be characterized by the same technology. We choose to identify TFP and to normalize q_{jj} to 1, as its makes the interpretation of the q_{ij} as gains or losses from outsourcing more intuitive.

measure of costs, and clear the observed (case 1) and drawn q_{ij} (case 2) from their ψ_j. This reduces the importance of the value of the supplier fixed effect in the q_{ij} for aggregate object, hence mitigates the impact of our identifying assumption on the q_{ij} calibration.

One could interpret ψ_j as the q_{hj}-equivalent parameter that drives the relationship between a firm and the household. Using ψ_j as another outsourcing capability q_{hj} in the PPML regression would increase variations, but also impose that the supplier fixed effect z_j^s is the same in ψ_j and the other outsourcing capability. To give us more flexibility, we prefer to split the two parameters and interpret ψ_j as the general hicks-neutral productivity, while the q_{ij} represent factors specific to outsourcing. Again, this means that the firm fixed effects are clear of the hicks-neutral productivity, and summarize only outsourcing-specific factors. Moreover, isolating the firm-specific hicks-neutral productivity also allows us to specifically analyze how it correlates with our other firm characteristics, namely the energy-biased technology and the outsourcing capabilities.

5.6 Demand shifter

To calibrate the product-specific demand shifter $\delta(\omega)$, we use the aggregate sales of product ω to the representative "rest-of-the-economy" firm.[73] In particular, firm j's sales $sales_j^R$ are equal to total sales minus sales to other manufacturing firms, which are then aggregated at the product level across firms selling ω. In the model, the latter correspond to:

$$\frac{sales^R(\omega)}{PY} = \delta(\omega)\left(\frac{P}{P_\omega}\right)^{\eta-1}$$

Where PY aggregate $sales^R(\omega)$ across all manufacturing products and $P_\omega = \left(\int_{j\in N_\omega}\left(\frac{\varepsilon}{\varepsilon-1}\right)^{1-\varepsilon} c_j(\omega)^{1-\varepsilon} dj\right)^{\frac{1}{1-\varepsilon}}$.

$$\frac{sales^R(\omega')}{sales^R(\omega)} = \frac{\delta(\omega')}{\delta(\omega)}\left(\frac{P_\omega}{P_{\omega'}}\right)^{\eta-1} \quad \forall \omega \neq \omega'$$

In equilibrium, the domestic "rest-of-the-economy" demand implies that $P = (PY - EXP)^{\frac{\alpha-1}{\alpha}}$. Therefore, we retrieve $\delta(\omega)$ from

$$\frac{P^{1-\eta}}{\delta(\omega)} = \sum_{\omega'}\frac{\delta_{\omega'}}{\delta_\omega}P_{\omega'}^{1-\eta} \Rightarrow \delta(\omega) = (PY - EXP)^{\frac{(\alpha-1)(1-\eta)}{\alpha}}\left[\sum_{\omega'\neq\omega}\frac{\delta_{\omega'}}{\delta_\omega}P_{\omega'}^{1-\eta} + P_\omega^{1-\eta}\right]^{-1}$$

5.7 Optimal allocation

Once we have calibrated our parameters for all firms, we can simulate the efficient allocation under the efficient *feasible* price $p_e = p_e^{SCC}$. The key assumption is that the parameters $\{\nu_j, \{q_{ij}\}_{i\in\mathcal{S}_\omega,\omega\in\Omega_{\omega_j}}\}_{j\in N}$ are fundamentals of the domestic firms, that is, they don't change with the energy price. We consider a counterfactual energy price that applies only to Italy, and keep the foreign firm prices constant. This will allow us to relate our results to the effect of a domestic carbon tax, and Italy is a small open economy in terms of energy.

[73] This residual firm includes exports.

Unit costs and make-or-buy decision. We compute the unit cost of outsourcing and of in-house production for each firm and each input in its recipe, under the new optimal energy price. The unit cost of outsourcing is the minimum *effective* price across all potential suppliers. As before, we proceed recursively moving downstream, starting from inputs that only use leaf-inputs. For each input, we then compare the outsourcing cost to the in-house cost to infer the outsourcing decision of the firm for each input. Note that we characterize this decision for all inputs but due to the sequentiality of production, only the most downstream outsourced input would effectively be purchased by the firm under the SCC scenario.

Sales. We also need each firm's sales shares (to GDP) in the optimal allocation to compute the optimal aggregate energy intensity.[74] To compute the sales to GDP ratio under the efficient allocation, we proceed recursively moving upstream.

1. With our costs, we compute the product price indices $P_\omega = \left(\int_{j\in N_\omega} p_j(\omega)^{1-\varepsilon} dj\right)^{\frac{1}{1-\varepsilon}}$ and the aggregate price index $P = \left(\sum_\omega \delta_\omega P_\omega^{1-\eta}\right)^{\frac{1}{1-\eta}}$

2. We start with the firms who only sell to the "rest-of-the-economy" sector, $j \in \mathcal{H}$. Those firms' sales (net of markup) are:
$$\frac{p_j(\omega_j)x_j(\omega_j)}{PY} = \delta_{\omega_j}\left(\frac{\varepsilon-1}{\varepsilon}\right)^{\varepsilon-1} c_j(\omega_j)^{1-\varepsilon} P^{\varepsilon-1}\left(\frac{P}{P_\omega}\right)^{\eta-\varepsilon}$$

3. We move upstream to firms that sell only to firms for which we have computed the sales ($j \in \mathcal{H}$). Those firms' sales are
$$\frac{p_j(\omega_j)x_j(\omega_j)}{PY} = \delta_{\omega_j}\left(\frac{\varepsilon-1}{\varepsilon}\right)^{\varepsilon-1} c_j(\omega_j)^{1-\varepsilon} P^{\varepsilon-1}\left(\frac{P}{P_\omega}\right)^{\frac{\eta(\eta-\varepsilon)}{\varepsilon}} + \int_{i\in\mathcal{H}:j\in\mathcal{O}_i} \frac{p_i(\hat{\omega}_i)x_i(\hat{\omega}_i)}{PY}\left(\frac{c_i(\hat{\omega}_i)}{c_j(\omega_j)/q_{ji}}\right)^{\frac{\gamma}{1-\gamma}} di$$

4. We add those firms to the set $\mathcal{H}$ of firms for which we have computed the sales.

5. We start again at step 3 and repeat until $\mathcal{H}$ contains all our firms.

Counterfactual analysis. We are finally equipped to quantify the extend of misallocation induced by distorted energy prices in an economy in which firms are heterogeneous in their energy-biased technology and outsourcing capabilities, which result in heterogeneous organization structures.

For our counterfactual analysis, we use the sample of firms present in the balanced-sheet data, since we need information on labor to calibrate ν_j. However, we account for the transactions with all domestic firms in the B2B to compute firms' counterfactual sales. We evaluate $\frac{E}{PY}$ at different energy prices,[75] such that we can decompose the different effects:

- **Factor substitution** For a given network allocation, firms reallocate resources between energy and labor for each of the inputs they produce in-house. They also reallocate resources across intermediate inputs. We measure this effect by keeping the equilibrium allocation fixed.

[74] In the equilibrium allocation, we observe total sales in the data. We match them using the demand shifter.

[75] Because welfare is dependent on the assumptions on the form of the utility function, as well as on the use of tax revenue, our preferred evaluation is based on aggregate energy intensity.

- **Reallocative efficiency:** This directly captures the effect of misallocation. We measure this effect as the residual effect net out of the factor substitution effect.

As a final exercise, we also re-calibrate our model and build the feasibly efficient allocation when $\eta = \varepsilon = \rho = \gamma \to -\infty$ (Leontieff). As shown by [Baqaee and Farhi, 2020], a change in factor prices should have no reallocation effect in this economy. However, the make-or-buy decision of the firms effectively introduces an additional margin of substitution: potential supplier, including the firm itself, are always perfectly substitutable, regardless of the other elasticity of substitution. Therefore, this setting isolates the contribution due to reallocation that is solely arising from make-or-buy decisions.

6 Results

7 Conclusion

How does the make-or-buy decisions of firms affect the aggregate energy demand of production? The large heterogeneity in energy efficiency across similar firms that we document implies that the micro-allocation of energy demand across the supply chain is highly relevant for aggregate variables. We show that when energy is under-priced, the make-or-buy decisions of firms under-weight the importance of energy efficiency, leading to misallocation of production across firms. In turn, this misallocation increases aggregate energy intensity and decreases welfare. To quantify the importance of this misallocation, we develop and calibrate a quantitative model of production organized in supply chains, in which firms differ in their energy-biased technology and the extent to which they gain from outsourcing (outsourcing capabilities). We use the universe of Italian firm-to-firm transactions to build the recipes behind the production of each product, namely the set of all inputs that enter directly or indirectly the production of a product. We use these recipes and the transaction data to infer the set of inputs produced in-house by each firm, which allows us to identify the energy-biased technology as residuals. We then recover the outsourcing capabilities for all possible customer-supplier pairs, using the ranking of potential options revealed by the observed fragmentation choice of the firms. Equipped with those firm-level parameters, we can build the counterfactual efficient allocation and measure the welfare loss arising from the energy inefficient make-or-buy decisions. This exercise is still work-in-progress.

References

[Acemoglu and Azar, 2020] Acemoglu, D. and Azar, P. D. (2020). Endogenous production networks. *Econometrica*, 88(1):33–82.

[Acemoglu and Tahbaz-Salehi, 2025] Acemoglu, D. and Tahbaz-Salehi, A. (2025). The macroeconomics of supply chain disruptions. *Review of Economic Studies*, 92(2):656–695.

[Aghion et al., 2024] Aghion, P., Barrage, L., Hémous, D., and Liu, E. (2024). Transition to green technology along the supply chain.

[Agostino et al., 2016] 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 et al., 2023] Akcigit, U., Baslandze, S., and Lotti, F. (2023). Connecting to power: political connections, innovation, and firm dynamics. *Econometrica*, 91(2):529–564.

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

[Alfaro et al., 2025] Alfaro, L., Brussevich, M., Minoiu, C., and Presbitero, A. F. (2025). Bank financing of global supply chains. Technical report, National Bureau of Economic Research.

[Alfaro and Charlton, 2009] Alfaro, L. and Charlton, A. (2009). Intra-industry foreign direct investment. *American Economic Review*, 99(5):2096–2119.

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

[Allen and Arkolakis, 2014] Allen, T. and Arkolakis, C. (2014). Trade and the topography of the spatial economy. *The Quarterly Journal of Economics*, 129(3):1085–1140.

[Alpino et al., 2024] Alpino, M., Citino, L., and Frigo, A. (2024). Firms in the energy crisis: evidence from 2021-22. *Available at SSRN 5251730*.

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

[Andersson, 2019] Andersson, J. J. (2019). Carbon Taxes and CO2 Emissions: Sweden as a Case Study. *American Economic Journal: Economic Policy*, 11(4):1–30.

[Antonietti et al., 2017] Antonietti, R., De Marchi, V., and Di Maria, E. (2017). Governing offshoring in a stringent environmental policy setting: Evidence from italian manufacturing firms. *Journal of cleaner production*, 155:103–113.

[Antràs, 2003] Antràs, P. (2003). Firms, contracts, and trade structure. *The Quarterly Journal of Economics*, 118(4):1375–1418.

[Antràs, 2005] Antràs, P. (2005). Incomplete contracts and the product cycle. *American economic review*, 95(4):1054–1073.

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

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

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

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

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

[Antràs and Yeaple, 2014] Antràs, P. and Yeaple, S. R. (2014). Multinational firms and the structure of international trade. *Handbook of international economics*, 4:55–130.

[Arkolakis et al., 2023] Arkolakis, C., Eckert, F., and Shi, R. (2023). Combinatorial discrete choice: A quantitative model of multinational location decisions. Technical report, National Bureau of Economic Research.

[Arkolakis et al., 2025] Arkolakis, C., Huneeus, F., and Miyauchi, Y. (2025). Production network formation, trade, and welfare. Technical report, CESifo Working Paper.

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

[Atalay et al., 2019] Atalay, E., Hortaçsu, A., Li, M. J., and Syverson, C. (2019). How wide is the firm border? *The Quarterly Journal of Economics*, 134(4):1845–1882.

[Atalay et al., 2014] Atalay, E., Hortaçsu, A., and Syverson, C. (2014). Vertical integration and input flows. *American Economic Review*, 104(4):1120–1148.

[Baily, 1978] Baily, M. N. (1978). Some aspects of optimal unemployment insurance. *Journal of public Economics*, 10(3):379–402.

[Baqaee and Farhi, 2019] Baqaee, D. R. and Farhi, E. (2019). The macroeconomic impact of microeconomic shocks: Beyond hulten's theorem. *Econometrica*, 87(4):1155–1203.

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

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

[Bernard and Kichian, 2021] Bernard, J.-T. and 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.

[Berry, 1994] Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, pages 242–262.

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

[Boehm et al., 2014] Boehm, C., Flaaen, A., and Pandalai-Nayar, N. (2014). Complementarities in multinational production and business cycle dynamics. Technical report, Technical report, Working paper, University of Michigan.

[Boehm, 2022] 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 et al., 2022] Boehm, J., Dhingra, S., and Morrow, J. (2022). The comparative advantage of firms. *Journal of Political Economy*, 130(12):3025–3100.

[Boehm and Oberfield, 2020] Boehm, J. and 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 and Oberfield, 2023] Boehm, J. and Oberfield, E. (2023). Growth and the fragmentation of production. Technical report, Tech. rep., mimeo. 7.

[Bolz et al., 2024] Bolz, S. J., Naumann, F., and Richter, P. M. (2024). Unilateral environmental policy and offshoring.

[Boudreau et al., 2023] Boudreau, L., Cajal-Grossi, J., and 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 and Jo, 2021a] Bretschger, L. and Jo, A. (2021a). Complementarity Between Labor and Energy: A Firm-Level Analysis. *SSRN Electronic Journal*.

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

[Cai and Xiang, 2025] Cai, S. and Xiang, W. (2025). Multinational production, technology diffusion, and economic growth. *Journal of International Economics*, page 104137.

[Capelle et al., 2023] Capelle, D., Kirti, D., Pierri, N., and Villegas Bauer, G. (2023). Mitigating climate change at the firm level: Mind the laggards.

[Cappariello and Felettigh, 2015] Cappariello, R. and 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.

[Caron, 2022] Caron, J. (2022). Empirical evidence and projections of carbon leakage: some, but not too much, probably. In *Handbook on trade policy and climate change*, pages 58–74. Edward Elgar Publishing.

[Carradori et al., 2023] Carradori, O., Giuzio, M., Kapadia, S., Salakhova, D., and Vozian, K. (2023). Financing the low-carbon transition in europe. *Available at SSRN 4510461*.

[Carraro and De Cian, 2013] Carraro, C. and De Cian, E. (2013). Factor-augmenting technical change: an empirical assessment. *Environmental Modeling & Assessment*, 18:13–26.

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

[Carvalho et al., 2021] Carvalho, V. M., Nirei, M., Saito, Y. U., and Tahbaz-Salehi, A. (2021). Supply chain disruptions: Evidence from the great east japan earthquake. *The Quarterly Journal of Economics*, 136(2):1255–1321.

[Cherniwchan et al., 2017] Cherniwchan, J., Copeland, B. R., and Taylor, M. S. (2017). Trade and the environment: New methods, measurements, and results. *Annual Review of Economics*, 9(1):59–85.

[Cherniwchan and Taylor, 2022] Cherniwchan, J. M. and Taylor, M. S. (2022). International trade and the environment: Three remaining empirical challenges. Technical report, National Bureau of Economic Research.

[Chetty, 2006] Chetty, R. (2006). A general formula for the optimal level of social insurance. *Journal of Public Economics*, 90(10-11):1879–1901.

[Choi et al., 2023] Choi, J., Hyun, J., Kim, G., and Park, Z. (2023). The cleanup of us manufacturing through pollution offshoring. *Available at SSRN 4345442*.

[Chor and Ma,] Chor, D. and Ma, L. Contracting frictions in global sourcing: Implications for welfare.

[Chor et al., 2021] Chor, D., Manova, K., and Yu, Z. (2021). Growing like china: Firm performance and global production line position. *Journal of International Economics*, 130:103445.

[Cole et al., 2014] Cole, M. A., Elliott, R. J., and Okubo, T. (2014). International environmental outsourcing. *Review of World Economics*, 150:639–664.

[Cole et al., 2021] Cole, M. A., Elliott, R. J., Okubo, T., and Zhang, L. (2021). Importing, outsourcing and pollution offshoring. *Energy Economics*, 103:105562.

[Copeland and Taylor, 2004] Copeland, B. R. and Taylor, M. S. (2004). Trade, growth, and the environment. *Journal of Economic literature*, 42(1):7–71.

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

[Cui et al., 2021] Cui, J., Wang, C., Zhang, J., and Zheng, Y. (2021). The effectiveness of china's regional carbon market pilots in reducing firm emissions. *Proceedings of the National Academy of Sciences*, 118(52):e2109912118.

[Dai et al., 2021] Dai, R., Duan, R., Liang, H., and Ng, L. (2021). Outsourcing climate change. *European Corporate Governance Institute–Finance Working Paper*, (723).

[De Lyon and Dechezleprêtre, 2025] De Lyon, J. and Dechezleprêtre, A. (2025). The great dispersion in energy productivity between firms. Technical report, OECD.

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

[Dechezleprêtre et al., 2019] Dechezleprêtre, A., Koźluk, T., Kruse, T., Nachtigall, D., and 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.

[Dekle et al., 2008] Dekle, R., Eaton, J., and Kortum, S. (2008). Global rebalancing with gravity: Measuring the burden of adjustment. Technical report, National Bureau of Economic Research.

[Devulder and Lisack, 2020] Devulder, A. and Lisack, N. (2020). Carbon tax in a production network: Propagation and sectoral incidence.

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

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

[Di Giovanni et al., 2024] Di Giovanni, J., Levchenko, A. A., and Mejean, I. (2024). Foreign shocks as granular fluctuations. *Journal of Political Economy*, 132(2):000–000.

[Ding, 2023] Ding, X. (2023). *Industry linkages from joint production*. US Census Bureau, Center for Economic Studies.

[Ding et al., 2022] Ding, X., Fort, T. C., Redding, S. J., and Schott, P. K. (2022). Structural change within versus across firms: Evidence from the united states. Technical report, National Bureau of Economic Research.

[Dingel and Tintelnot, 2020] Dingel, J. I. and Tintelnot, F. (2020). Spatial economics for granular settings. Technical report, National Bureau of Economic Research.

[Donald et al., 2025] Donald, E., Fukui, M., and Miyauchi, Y. (2025). Unpacking aggregate welfare in a spatial economy. Technical report, National Bureau of Economic Research.

[Dugoua and Dumas, 2021] Dugoua, E. and Dumas, M. (2021). Green product innovation in industrial networks: A theoretical model. *Journal of Environmental Economics and Management*, 107:102420.

[Dussaux et al., 2023] Dussaux, D., Vona, F., and Dechezleprêtre, A. (2023). Imported carbon emissions: Evidence from french manufacturing companies. *Canadian Journal of Economics/Revue canadienne d'économique*, 56(2):593–621.

[Dzemski, 2019] Dzemski, A. (2019). An empirical model of dyadic link formation in a network with unobserved heterogeneity. *Review of Economics and Statistics*, 101(5):763–776.

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

[Eaton et al., 2022] Eaton, J., Kortum, S. S., and Kramarz, F. (2022). Firm-to-firm trade: Imports, exports, and the labor market. Technical report, National Bureau of Economic Research.

[EPA, 2022] EPA, E. P. A. (2022). Environmental protection agency: Report on the social cost of greenhouse gases: Estimates incorporating recent scientific advances.

[Fafchamps and Gubert, 2007] Fafchamps, M. and Gubert, F. (2007). The formation of risk sharing networks. *Journal of development Economics*, 83(2):326–350.

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

[Fetzer et al., 2024] Fetzer, T., Lambert, P. J., Feld, B., and Garg, P. (2024). Ai-generated production networks: Measurement and applications to global trade. Technical report.

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

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

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

[Fort, 2023] 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.

[Gabaix, 2011] Gabaix, X. (2011). The granular origins of aggregate fluctuations. *Econometrica*, 79(3):733–772.

[Garg et al., 2019] Garg, S., Ghosh, P., and Tan, B. (2019). Vertical integration and supply chains.

[Gourieroux et al., 1984] Gourieroux, C., Monfort, A., and Trognon, A. (1984). Pseudo maximum likelihood methods: Applications to poisson models. *Econometrica: Journal of the Econometric Society*, pages 701–720.

[Graham, 2015] Graham, B. S. (2015). Methods of identification in social networks. *Annu. Rev. Econ.*, 7(1):465–485.

[Graham, 2017] Graham, B. S. (2017). An econometric model of network formation with degree heterogeneity. *Econometrica*, 85(4):1033–1063.

[Graham, 2024] Graham, B. S. (2024). Sparse network asymptotics for logistic regression under possible misspecification. *Econometrica*, 92(6):1837–1868.

[Grossi et al., 2023] Grossi, J. C., Del Prete, D., and Macchiavello, R. (2023). Supply chain disruptions and sourcing strategies.

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

[Hawkins-Pierot and Wagner, 2023] Hawkins-Pierot, J. T. and Wagner, K. R. (2023). *Technology lock-in and costs of delayed climate policy*. US Census Bureau, Center for Economic Studies.

[Head and Mayer, 2014] Head, K. and Mayer, T. (2014). Gravity equations: Workhorse, toolkit, and cookbook. In *Handbook of international economics*, volume 4, pages 131–195. Elsevier.

[Heckman, 1990] Heckman, J. (1990). Varieties of selection bias. *The American Economic Review*, 80(2):313–318.

[Heckman and Honore, 1990] Heckman, J. J. and Honore, B. E. (1990). The empirical content of the roy model. *Econometrica: Journal of the Econometric Society*, pages 1121–1149.

[Helper and Munasib, 2022] Helper, S. and Munasib, A. (2022). *Economies of scope and relational contracts: exploring global value chains in the automotive industry*. US Department of Commerce, Bureau of Economic Analysis.

[Holland and Leinhardt, 1981] Holland, P. W. and Leinhardt, S. (1981). An exponential family of probability distributions for directed graphs. *Journal of the american Statistical association*, 76(373):33–50.

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

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

[Huneeus et al.,] Huneeus, F., Koike-Mori, Y., and Martner, A. Growth accounting in open economies with distortions.

[Jochmans, 2018] Jochmans, K. (2018). Semiparametric analysis of network formation. *Journal of Business & Economic Statistics*, 36(4):705–713.

[Känzig, 2023] Känzig, D. R. (2023). The unequal economic consequences of carbon pricing. Technical report, National Bureau of Economic Research.

[Keller and Yeaple, 2013] Keller, W. and Yeaple, S. R. (2013). The gravity of knowledge. *American Economic Review*, 103(4):1414–1444.

[Kim et al., 2022] Kim, J., Sovacool, B. K., Bazilian, M., Griffiths, S., Lee, J., Yang, M., and 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.

[King et al., 2019] King, M., Tarbush, B., and Teytelboym, A. (2019). Targeted carbon tax reforms. *European Economic Review*, 119:526–547.

[Klenow et al., 2024] Klenow, P. J., Pastén, E., and Ruane, C. (2024). Carbon taxes and misallocation in chile. Technical report, Mimeo, Stanford.

[Köveker et al., 2025] Köveker, T., Richter, P. M., Schiersch, A., and Sogalla, R. (2025). Clean production, dirty sourcing: How embodied emissions alter the environmental footprint of exporters. Technical report, DIW Berlin, German Institute for Economic Research.

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

[Lafontaine and Slade, 2013] Lafontaine, F. and Slade, M. (2013). Inter-firm contracts. *The handbook of organizational economics*, 3.

[Lafrogne-Joussier et al., 2023] Lafrogne-Joussier, R., Martin, J., and Mejean, I. (2023). Supply shocks in supply chains: Evidence from the early lockdown in china. *IMF economic review*, 71(1):170.

[Lanteri and Rampini, 2023] Lanteri, A. and Rampini, A. (2023). Financing the adoption of clean technology.

[La'O and Tahbaz-Salehi, 2022] La'O, J. and Tahbaz-Salehi, A. (2022). Optimal monetary policy in production networks. *Econometrica*, 90(3):1295–1336.

[Le Moigne et al., 2024] Le Moigne, M., Lepot, S., Ossa, R., Ritel, M., and Simon, D. Z. (2024). Greening ricardo: Environmental comparative advantage and the environmental gains from trade. Technical report, WTO Staff Working Paper.

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

[Li and Zhou, 2017] Li, X. and Zhou, Y. M. (2017). Offshoring pollution while offshoring production? *Strategic Management Journal*, 38(11):2310–2329.

[Lim, 2018] Lim, K. (2018). Endogenous production networks and the business cycle. *Work. Pap.*

[Liu, 2019] Liu, E. (2019). Industrial policies in production networks. *The Quarterly Journal of Economics*, 134(4):1883–1948.

[Lyubich et al., 2018] Lyubich, E., Shapiro, J. S., and Walker, R. (2018). Regulating mismeasured pollution: Implications of firm heterogeneity for environmental policy. In *AEA Papers and Proceedings*, volume 108, pages 136–142. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

[Marin and Vona, 2021] Marin, G. and 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.

[Markusen, 1984] Markusen, J. R. (1984). Multinationals, multi-plant economies, and the gains from trade. *Journal of international economics*, 16(3-4):205–226.

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

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

[Martin et al., 2024] Martin, R., Muûls, M., and Stoerk, T. (2024). The effects of carbon pricing along the production network. Technical report.

[Marullaz, 2025] Marullaz, C. (2025). Firm boundaries and green innovation. *Work In Progress*.

[Mayer and Puller, 2008] Mayer, A. and Puller, S. L. (2008). The old boy (and girl) network: Social network formation on university campuses. *Journal of public economics*, 92(1-2):329–347.

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

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

[Metcalf, 2019] Metcalf, G. E. (2019). On the economics of a carbon tax for the united states. *Brookings Papers on Economic Activity*, 2019(1):405–484.

[Metcalf and Stock, 2020] Metcalf, G. E. and Stock, J. H. (2020). Measuring the macroeconomic impact of carbon taxes. In *AEA papers and Proceedings*, volume 110, pages 101–106. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

[Michel, 2013] Michel, B. (2013). *Is Offshoring Driven by Air Emissions?: Testing the Pollution Haven Effect for Imports of Intermediates*. Federal Planning Bureau.

[Monarch, 2014] 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 and Teece, 1982] Monteverde, K. and Teece, D. J. (1982). Supplier switching costs and vertical integration in the automobile industry. *The Bell Journal of Economics*, pages 206–213.

[Nordhaus, 2017] Nordhaus, W. D. (2017). Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences*, 114(7):1518–1523.

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

[Panzar and Willig, 1981] Panzar, J. C. and Willig, R. D. (1981). Economies of scope. *The American economic review*, 71(2):268–272.

[Petrick et al., 2011] Petrick, S., Rehdanz, K., and 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, 2019] Pindyck, R. S. (2019). The social cost of carbon revisited. *Journal of Environmental Economics and Management*, 94:140–160.

[Rachapalli, 2021] Rachapalli, S. (2021). Learning between buyers and sellers along the global value chain. Technical report, Working paper.

[Ramírez-Santos et al., 2018] Ramírez-Santos, Á. A., Castel, C., and 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, 1999] Rauch, J. E. (1999). Networks versus markets in international trade. *Journal of international Economics*, 48(1):7–35.

[Sahay, 2003] Sahay, B. S. (2003). Understanding trust in supply chain relationships. *Industrial Management & Data Systems*, 103(8):553–563.

[Schenker et al., 2018] Schenker, O., Koesler, S., and 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, 2021] Shapiro, J. S. (2021). The environmental bias of trade policy. *The Quarterly Journal of Economics*, 136(2):831–886.

[Silva and Tenreyro, 2006] Silva, J. S. and Tenreyro, S. (2006). The log of gravity. *The Review of Economics and statistics*, pages 641–658.

[Stadler et al., 2018] Stadler, K., Wood, R., Bulavskaya, T., Södersten, C.-J., Simas, M., Schmidt, S., Usubiaga, A., Acosta-Fernández, J., Kuenen, J., Bruckner, M., et al. (2018). Exiobase 3: Developing a time series of detailed environmentally extended multi-regional input-output tables. *Journal of Industrial Ecology*, 22(3):502–515.

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

[Wagner et al., 2020] 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 et al., 2008] Wang, C., Larsson, M., Ryman, C., Grip, C.-E., Wikström, J.-O., Johnsson, A., and Engdahl, J. (2008). A model on co2 emission reduction in integrated steelmaking by optimization methods. *International journal of energy research*, 32(12):1092–1106.

[Yan et al., 2019] Yan, T., Jiang, B., Fienberg, S. E., and Leng, C. (2019). Statistical inference in a directed network model with covariates. *Journal of the American Statistical Association*, 114(526):857–868.

A Theoretical Appendix

A.1 Two-stage Model

This section derives the proof of the two-stage model results presented in the main text. It then derives additional results.

Proof of Proposition 1 If $p_e = p_e^{SCC}$, there is no distortion in the economy.

1. **The equilibrium allocation is efficient**, given the allocation across firms (first Welfare Theorem).

   Proof. The social planner maximizes an objective function $W(y, e)$ where $\frac{\partial W}{\partial e} < 0$. subject to $y_j = q_{kj}\left(\beta^{1-\rho}(\nu_k e_{ij})^{\rho} + (1-\beta)^{1}\right.$ $\{i, j\}$ and $l_j = \bar{L}$. For the decentralized equilibrium to match the social planner allocation, $p_e = \frac{\partial W/\partial e}{\partial W/\partial y} = SCC$ and $w = \frac{\mu_l}{\partial W/\partial y}$ where μ_l is the Lagrange multiplier associated to the labor endowment constraint.

2. **The outsourcing decision is efficient.**

   Proof. Since energy decreases welfare, the social planner wants to minimize energy used to produce one unit of output when choosing which firm should produce. The efficient allocation of energy and labor yields the following resources needed to produce one unit of output:

   $$1 \leq q_{kj}^{-1}\left(\beta\Big(\frac{\partial W/\partial e}{\partial W/\partial y}\Big)^{\frac{-\rho}{1-\rho}} + (1-\beta)\Big(\frac{\mu_l}{\partial W/\partial y}\Big)^{\frac{-\rho}{1-\rho}}\right)^{\frac{-(1-\rho)}{\rho}}$$

   Which is minimized by choosing the firm $k = \{i, j\}$ with the lowest unit cost of production c_{ij}, that is to solve for $\min\left\{c_{ij}^{\mathcal{I}}, \frac{c_{ij}^{\mathcal{O}}}{q_{ij}}\right\}$.

If $p_e < p_e^{SCC}$:

1. **Insufficient weight is given to ν in the outsourcing decision when $\rho < 0$.**

   Proof. The make-or-buy decision can be rewritten as

   $$\left(\left(\frac{q_{ij}}{q_{jj}}\right)^{\frac{-\rho}{1-\rho}}\nu_j^{\frac{\rho}{1-\rho}} - \nu_i^{\frac{\rho}{1-\rho}}\right)\beta p_e^{\frac{-\rho}{1-\rho}} + \left(\left(\frac{q_{ij}}{q_{jj}}\right)^{\frac{-\rho}{1-\rho}} - 1\right)(1-\beta)w^{\frac{-\rho}{1-\rho}} < (>)\ 0 \text{ if } \rho < (>)0$$

   To simplify exposure, we normalize $\nu_j = 1$ and let q_{jj} vary in this proof. However, the results hold with any choice of normalization. We can derive $\frac{dLHS}{dlog\nu_i/\nu_j} = -\frac{\rho}{1-\rho}\beta_\omega\left(\frac{p_e}{\nu_i}\right)^{\frac{-\rho}{1-\rho}}$ and $\frac{dLHS}{dlogq_{ij}/q_{jj}} = \frac{-\rho}{1-\rho}\left(\frac{q_{ij}}{q_{jj}}\right)^{\frac{-\rho}{1-\rho}}\left(\beta p_e^{\frac{-\rho}{1-\rho}} + (1-\beta)w^{\frac{-\rho}{1-\rho}}\right)$. Thus,

   $$\frac{\frac{dLHS}{dlogq_{ij}/q_{jj}}}{\frac{dLHS}{dlog\nu_i/\nu_j}} = \left(\frac{q_{ij}}{q_{jj}\nu_i}\right)^{\frac{-\rho}{1-\rho}}\left(1 + \frac{1-\beta}{\beta}\Big(\frac{w}{p_e}\Big)^{\frac{-\rho}{1-\rho}}\right)$$

   Which increases with p_e if $\rho < 0$.[76] The decision is more sensitive to the capabilities ratio $\frac{q_{ij}}{q_{jj}}$ (relative to its sensitivity to the energy-biased technology ratio $\frac{\nu_j}{\nu_i}$) the lower p_e is. If there is a trade-off between $\frac{q_{ij}}{q_{jj}}$ and $\frac{\nu_j}{\nu_i}$, the consideration will be tilted toward the outsourcing capabilities at lower p_e. Formally, $|\frac{\partial^2 LHS}{\partial\nu_i/\nu_j\partial p_e}| \geq |\frac{\partial^2 LHS}{\partial q_{ij}/q_{jj})\partial p_e}|$.

[76] The opposite holds if $\rho > 0$. However, this still increases energy intensity, as ν_j becomes an energy-augmenting technology which amplifies energy consumption. Therefore, *too much* weight is given to ν_j in that case. Overall, it remains that too little consideration is given to the energy component, and its externality.

Effect of carbon tax when the allocation is fixed. Suppose that firm j initially decided to outsource input ω and this allocation fixed. We consider the effect of an increase in energy price on the costs of the firm.

To express the marginal effect, we consider the (potentially fictitious) inputs for which the firm was initially indifferent ($\frac{c_j^{\mathcal{I}}(\omega)}{q_{jj}} = \frac{c_i^{\mathcal{O}}}{q_{ij}}$). They are interesting, because the make-or-buy decision for these inputs might switch even for a marginal change in energy price. The firm is outsourcing input ω, thus its effective cost changes by

$$\frac{dln\frac{c_i^{\mathcal{O}}(\omega)}{q_{ij}}}{dlnp_e} = \beta\Big(\frac{c_i^{\mathcal{O}}(\omega)}{p_e}\Big)^{\frac{\rho}{1-\rho}}\nu_i^{\frac{\rho}{1-\rho}}$$

If instead, it was producing ω in-house, its effective unit cost would change by

$$\frac{dln\frac{c_j^{\mathcal{I}}(\omega)}{q_{jj}}}{dlnp_e} = \beta\Big(\frac{c_j^{\mathcal{I}}(\omega)}{p_e}\Big)^{\frac{\rho}{1-\rho}}\nu_j^{\frac{\rho}{1-\rho}}$$

Initial indifference implies that

$$\Delta = \frac{dln\frac{c_i^{\mathcal{O}}(\omega)}{q_{ij}}}{dlnp_e} - \frac{dln\frac{c_j^{\mathcal{I}}(\omega)}{q_{jj}}}{dlnp_e} = \beta\Big(\frac{c_j^{\mathcal{I}}(\omega)}{p_e}\Big)^{\frac{\rho}{1-\rho}}\left(\left(\frac{q_{ij}\nu_i}{q_{jj}\nu_j}\right)^{\frac{\rho}{1-\rho}} - 1\right)\nu_j^{\frac{\rho}{1-\rho}}$$

This is necessarily positive if the firm would have like to change its make-or-buy decision under the higher energy price. In other words, the firm loses more (i.e., its unit cost increases more) when outsourcing than it would have if it was previously producing in-house.

Extensive and intensive margin decomposition. As we have done at the firm level, we can decompose the change in aggregate energy intensity into extensive and intensive margin adjustments if we consider a continuum N of downstream firms, and augment again our simple model with the multiple-input specification. We express optimal value with a SCC index. First, we derive a simple expression for aggregate energy intensity, which takes into account the energy used both directly (scope 1 and 2) and indirectly (scope 3) by the downstream firms. Recall that:

$$\text{Scope 1 and 2: } e_j^{\mathcal{I}}(\omega) = \beta\Big(\frac{c_j^{\mathcal{I}}}{p_e/\nu_j^{\rho}}\Big)^{\frac{1}{1-\rho}}\Big(\frac{c_j}{c_j^{\mathcal{I}}(\omega)/q_{jj}^{\gamma}}\Big)^{\frac{1}{1-\gamma}}y_j \quad \text{if } \omega \text{ in-house}$$

$$\text{Scope 3: } e_j^{\mathcal{O}}(\omega) = \beta\Big(\frac{c_i^{\mathcal{O}}(\omega)}{p_e/\nu_i^{\rho}}\Big)^{\frac{1}{1-\rho}}\Big(\frac{c_j}{c_i^{\mathcal{O}}(\omega)/q_{ij}^{\gamma}}\Big)^{\frac{1}{1-\gamma}}y_j \quad \text{if } \omega \text{ outsourced}$$

Thus,

$$\begin{aligned}E &= \int_{j\in N}\left(\int_{\omega\in\mathcal{I}_j}\beta\Big(\frac{c_j^{\mathcal{I}}}{p_e/\nu_j^{\rho}}\Big)^{\frac{1}{1-\rho}}\Big(\frac{c_j}{c_j^{\mathcal{I}}/q_{jj}^{\gamma}}\Big)^{\frac{1}{1-\gamma}}y_j d\omega + \int_{\omega\in\mathcal{O}_j}\beta\Big(\frac{c_i^{\mathcal{O}}(\omega)}{p_e/\nu_i^{\rho}}\Big)^{\frac{1}{1-\rho}}\Big(\frac{c_j}{c_i^{\mathcal{O}}(\omega)/q_{ij}^{\gamma}}\Big)^{\frac{1}{1-\gamma}}y_j d\omega\right)dj \\ &= \int_{j\in N}\frac{c_j y_j}{p_e}\left(\int_{\omega\in\mathcal{I}_j}\beta\Big(\frac{c_j^{\mathcal{I}}}{p_e/\nu_j}\Big)^{\frac{\rho}{1-\rho}}\Big(\frac{c_j}{c_j^{\mathcal{I}}/q_{jj}}\Big)^{\frac{\gamma}{1-\gamma}}d\omega + \int_{\omega\in\mathcal{O}_j}\beta\Big(\frac{c_i^{\mathcal{O}}(\omega)}{p_e/\nu_i}\Big)^{\frac{\rho}{1-\rho}}\Big(\frac{c_j}{c_i^{\mathcal{O}}(\omega)/q_{ij}}\Big)^{\frac{\gamma}{1-\gamma}}d\omega\right)dj\end{aligned}$$

To simplify exposure, suppose that the final demand is characterized by $\frac{y_j}{Y} = D\Big(\frac{c_j}{P}\Big)$, and that $\gamma = \rho$. Then,

$$\frac{E}{Y} = \int_{j\in N}\beta D\Big(\frac{c_j}{P}\Big)\Big(\frac{c_j}{p_e}\Big)^{\frac{1}{1-\rho}}\left(|\mathcal{I}_j|(\nu_j q_{jj}))^{\frac{\rho}{1-\rho}} + \int_{i:\omega_i\in\mathcal{O}_j}(\nu_i q_{ij})^{\frac{\rho}{1-\rho}}di\right)dj$$

We can decompose the changes in aggregate energy intensity into extensive and intensive margin adjustments:

$$\begin{aligned}
\frac{E}{Y} - \Big(\frac{E}{Y}\Big)^{SCC} =& \int_{j\in N} \beta D\Big(\frac{c_j}{P}\Big)\Big(\frac{c_j}{p_e}\Big)^{\frac{1}{1-\rho}} \left(|\mathcal{I}_j|(\nu_j q_{jj}))^{\frac{\rho}{1-\rho}} + \int_{i:\omega_i\in\mathcal{O}_j} (\nu_i q_{ij})^{\frac{\rho}{1-\rho}} di\right) \\
& - \beta D\Big(\frac{c_j^{SCC}}{P^{SCC}}\Big)\Big(\frac{c_j^{SCC}}{p_e^{SCC}}\Big)^{\frac{\rho}{1-\rho}} \left(|\mathcal{I}_j^{SCC}|(\nu_j q_{jj}))^{\frac{\rho}{1-\rho}} + \int_{i:\omega_i\in\mathcal{O}_j^{SCC}} (\nu_i q_{ij})^{\frac{\rho}{1-\rho}} di\right) dj \\
=& \beta \int_{j\in N} D\Big(\frac{c_j}{P}\Big)\left[\Big(\frac{c_j}{p_e}\Big)^{\frac{1}{1-\rho}} - \Big(\frac{c_j^{SCC}}{p_e^{SCC}}\Big)^{\frac{1}{1-\rho}}\right]\left(|\mathcal{I}_j|(\nu_j q_{jj}))^{\frac{\rho}{1-\rho}} + \int_{i:\omega_i\in\mathcal{O}_j} (\nu_i q_{ij})^{\frac{\rho}{1-\rho}} di\right) \\
& + \left(D\Big(\frac{c_j}{P}\Big) - D\Big(\frac{c_j^{SCC}}{P^{SCC}}\Big)\right)\Big(\frac{c_j^{SCC}}{p_e^{SCC}}\Big)^{\frac{1}{1-\rho}} \left(|\mathcal{I}_j^{SCC}|(\nu_j q_{jj}))^{\frac{\rho}{1-\rho}} + \int_{i:\omega_i\in\mathcal{O}_j^{SCC}} (\nu_i q_{ij})^{\frac{\rho}{1-\rho}} di\right) \\
& + D\Big(\frac{c_j^{SCC}}{P^{SCC}}\Big)\Big(\frac{c_j^{SCC}}{p_e^{SCC}}\Big)^{\frac{1}{1-\rho}} \Big((|\mathcal{I}_j| - |\mathcal{I}_j^{SCC}|)(\nu_j q_{jj}))^{\frac{\rho}{1-\rho}} \\
& + \int_{i:\omega_i\in\mathcal{O}_j \& \omega_i\notin\mathcal{O}_j^{SCC}} (\nu_i q_{ij})^{\frac{\rho}{1-\rho}} di - \int_{i:\omega_i\notin\mathcal{O}_j \& \omega_i\in\mathcal{O}_j^{SCC}} (\nu_i q_{ij})^{\frac{\rho}{1-\rho}} di\Big) dj
\end{aligned}$$

The first line corresponds to the standard, within-firm, intensive margin adjustment. It is positive for all firms, because $c_j\Big(\frac{c_j}{p_e}\Big)^{\frac{1}{1-\rho}} - c_j^{SCC}\Big(\frac{c_j^{SCC}}{p_e^{SCC}}\Big)^{\frac{1}{1-\rho}} > 0$. Indeed, the relative price of energy is higher under the social cost of carbon, which reallocates resources away from energy and toward labor within the firm, decreasing energy intensity.

The second line corresponds to the intensive margin, between-firm reallocation of sales shares. $\left(D\Big(\frac{c_j}{P}\Big) - D\Big(\frac{c_j^{SCC}}{P^{SCC}}\Big)\right)$ is positive for firms that suffer the most from a higher energy price, that is for firms with a lower energy-biased technology ν_j. Those firms see their relative costs increase, and lose market share. Those firms also produce with a higher aggregate energy intensity ($\Big(|\mathcal{I}_j^{SCC}|(\nu_j q_{jj}))^{\frac{\rho}{1-\rho}} + \int_{i:\omega_i\in\mathcal{O}_j^{SCC}} (\nu_i q_{ij})^{\frac{\rho}{1-\rho}} di\Big)$ is larger). Therefore, the overall effects of reallocation through market share is also positive, i.e., it decreases energy intensity under the SCC.

Finally, the last two lines correspond to the extensive margin adjustments that arise from different make-or-buy decisions. Equilibrium and optimal make-or-buy decisions imply that $\Big(\frac{q_{ij}\nu_i}{q_{jj}\nu_j}\Big)^{\frac{-\rho}{1-\rho}} > 1$ for the inputs that are produced in-house under p_e^{SCC} but outsourced under p_e ($\omega \in \mathcal{I}_j^{SCC}$ & $\omega_i \in \mathcal{O}_j$), and that $\Big(\frac{q_{ij}\nu_i}{q_{jj}\nu_j}\Big)^{\frac{-\rho}{1-\rho}} < 1$ for the inputs that are produced in-house under p_e but outsourced under p_e^{SCC} ($\omega \in \mathcal{I}_j$ & $\omega_i \in \mathcal{O}_j^{SCC}$). Thus, the extensive margin is also positive and reinforce the intensive margin effects. Overall $\frac{E}{Y} - \Big(\frac{E}{Y}\Big)^{SCC} > 0$.

A.1.1 Misallocation as a wedge.

In this section, we use the toy model to characterize our concept of *energy* misallocation in terms of wedges, as it is standard in the misallocation literature. Recall that a firm's energy intensity is

$$\frac{e_j}{c_j y_j} = \frac{1}{p_e}\left(\int_{\omega\in\mathcal{I}_j} \beta\Big(\frac{c_j^{\mathcal{I}}(\omega)}{p_e/\nu_j}\Big)^{\frac{\rho}{1-\rho}}\Big(\frac{c_j}{c_j^{\mathcal{I}}(\omega)/q_{jj}}\Big)^{\frac{\gamma}{1-\gamma}} d\omega + \int_{\omega\in\mathcal{O}_j} \beta\Big(\frac{c_i^{\mathcal{O}}(\omega)}{p_e/\nu_i}\Big)^{\frac{\rho}{1-\rho}}\Big(\frac{c_j}{c_i^{\mathcal{O}}(\omega)/q_{ij}}\Big)^{\frac{\gamma}{1-\gamma}} d\omega\right)$$

To simplify exposure, suppose that $\gamma = \rho$. We also normalize $\nu_i = 1$ and $q_{jj} = 1$, such that we can interpret q_{ij} as relative outsourcing capabilities, and ν_j as relative energy-biased technology. With only one parameter varying across inputs, we can rank all the inputs $\omega \in \Omega$ in increasing order of the corresponding q_{ij} associated to their suppliers, for each

downstream firm j. Therefore, we can characterize the cutoff input $\tilde{q}_j$, such that inputs associated to a $q_{ij} < (>)\ \tilde{q}_j$ are produced in-house (outsourced) by firm j. In particular, the private equilibrium cutoff is

$$\tilde{q}_j^{eq} \equiv \left(\frac{\nu_j^{\frac{\rho}{1-\rho}} + \frac{1-\beta}{\beta}\left(\frac{w}{p_e}\right)^{\frac{-\rho}{1-\rho}}}{1 + \frac{1-\beta}{\beta}\left(\frac{w}{p_e}\right)^{\frac{-\rho}{1-\rho}}} \right)^{\frac{(1-\rho)}{\rho}} \tag{6}$$

If $\nu_j > (<)\ 1$, $1 < (>)\ \tilde{q}_j^{eq} < (>)\ \nu_j$.

Suppose also that q_{ij} are uniformly distributed between 0 and $q^m > 1$. Then, with a continuum of inputs normalized to 1, $|\mathcal{I}|_j = \frac{\tilde{q}_j}{q^m}$ Therefore,

$$\frac{e_j}{c_j y_j} = \beta \frac{1}{p_e}\left(\frac{c_j}{p_e}\right)^{\frac{\rho}{1-\rho}} \left(\int_{\omega \in \mathcal{I}_j} \nu_j^{\frac{\rho}{1-\rho}} d\omega + \int_{\omega \in \mathcal{O}_j} q_{ij}^{\frac{\rho}{1-\rho}} d\omega \right) = \beta \frac{1}{p_e}\left(\frac{c_j}{p_e}\right)^{\frac{\rho}{1-\rho}} \underbrace{\left(\frac{\tilde{q}_j}{q^m}\nu_j^{\frac{\rho}{1-\rho}} + q_m^{\frac{\rho}{1-\rho}}(1-\rho)(1-\left(\frac{\tilde{q}_j}{q^m}\right)^{\frac{1}{1-\rho}}) \right)}_{\equiv Q_j^e}$$

Where Q_j^e represents the supply chain-level of technology. For a given cost, $\frac{e_j}{c_j y_j}$ is minimized at $\tilde{q}_j = \nu_j$ (for $\rho < 0$).[77] In words, for a given cost, energy intensity is minimized by giving the same weights to each type of technology, because both play the same role for energy intensity. Thus, firms should only outsource the inputs associated to a higher q_{ij} than their relative energy efficiency. However, there is also a scale effect through costs. We know that c_j is minimized at $\tilde{q}_j$, since firms minimize private costs. With $\rho < (>)\ 0$, energy intensity is minimized at large (small) costs due to this scale effect. Using the definition of $\tilde{q}_j$ in equation (6), it is easy to show that $1 < (>)\ \tilde{q}_j < (>)\ \nu_j$ if $\nu_j > (<)\ 1$, because the cutoff that minimizes the costs puts less emphasis on the energy-biased technology, since it balances reducing both the energy and the labor cost. However, this result also tells us that some inputs with relatively high ($q_{ij} > 1$) outsourcing capability will still be produced in-house, because there is still some weight given to the relative energy-biased technology. In other words, the private firm considers both the relative outsourcing capability and the relative energy-biased technology in its decision, with a weight on each element determined by the relative price of energy. Taking the technology and the scale effect into account, the cutoff $\tilde{q}_j^E$ that minimizes energy intensity satisfies if $\nu_j > 1$:[78]

$$\left(\left(\frac{\nu_j}{\tilde{q}_j^{eq}}\right)^{\frac{\rho}{1-\rho}} - \left(\frac{\tilde{q}_j^E}{\tilde{q}_j^{eq}}\right)^{\frac{\rho}{1-\rho}} \right) = \frac{\frac{\tilde{q}_j^E}{q^m}\nu_j^{\frac{\rho}{1-\rho}} + (q^m)^{\frac{\rho}{1-\rho}}(1-\rho)(1-\left(\frac{\tilde{q}_j^E}{q^m}\right)^{\frac{1}{1-\rho}})}{\frac{\tilde{q}_j^E}{q^m}(\tilde{q}_j^{eq})^{\frac{\rho}{1-\rho}} + (q^m)^{\frac{\rho}{1-\rho}}(1-\rho)(1-\left(\frac{\tilde{q}_j^E}{q^m}\right)^{\frac{1}{1-\rho}})} \left(1 - \left(\frac{\tilde{q}_j^E}{\tilde{q}_j^{eq}}\right)^{\frac{\rho}{1-\rho}}\right) \tag{7}$$

The first line represents the direct technology effect (pushing $\tilde{q}_j^E$ toward ν_j), while the second line represents the effect on costs (pushing $\tilde{q}_j^E$ toward $\tilde{q}_j$). From private cost minimization, the second line is equal to zero under $\tilde{q}_j$. Thus, $\tilde{q}_j$ is only

[77] If $\rho > 0$, the direct effect of technology is to increase energy intensity. In that case, the minimizing cutoff is 0 if $\nu_j > \tilde{q}_j > 1$ (which implies that $\nu_j q_j > \nu_i q_i\ \forall i < |\mathcal{I}_j|$, so the downstream firm is overall more productive than its potential suppliers, hence uses more energy through the direct technology effect), and q^m if $\nu_j < \tilde{q}_j < 1$. This is because the minimizing cutoff simply allocates production to the least productive firm.

[78] This only holds if $\frac{e_j}{c_j y_j}$ is convex in $\tilde{q}_j$. It is almost always satisfied if $\nu_j > 1$, but not necessarily with $\nu_j < 1$. In that case, the optimal cutoff is at a corner solution.

optimal if $\tilde{q}_j = \nu_j$, which requires no across-firm heterogeneity ($\nu_j = 1$). Therefore, if the objective of the social planner is to minimize firm-level energy intensity[79], the firm-level measure that summarizes misallocation is $\frac{\tilde{q}_j^E}{\tilde{q}_j}$. Equation (7) implicitly determines misallocation $\frac{\tilde{q}_j^E}{\tilde{q}_j^{eq}}$ as a decreasing function of ν_j. Misallocation also decreases in p_e. In other words, the gap between $\tilde{q}_j^E$ and $\tilde{q}_j$ contains all the inputs that could be produced using less energy. Those inputs would be produced at a higher costs, due to their lower associated outsourcing capabilities. However, the energy externality that they would spare more than compensate that higher cost.

Thus, a social planner wants to reduce misallocation by enforcing $\tilde{q}_j^{eq} = \tilde{q}_j^E$. This requires that $\left(\frac{\nu_j}{\tilde{q}_j^{eq}}\right)^{\frac{\rho}{1-\rho}} = 1$. Suppose that she can use a downstream-firm-specific tax τ_j on the price of energy to enforce $\tilde{q}_j^{eq} = \tilde{q}_j^E$. This effectively transforms the requirement to $\left(\frac{\nu_j/(1+\tau_j)}{\tilde{q}_j^{eq}}\right)^{\frac{\rho}{1-\rho}} = 1$, Thus, the optimal tax should exactly be set to compensate the firm-specific energy-biased technology ν_j. In standard misallocation jargon, ν_j reflects the firm-specific wedge.[80]

Thus, one can think of $1 + \tau_j$ as a missing tax, or equivalently of ν_j as the wedge, that induces misallocation, which is measured by $\tilde{q}_j^{eq}$. Indeed, if $\left(\frac{\nu_j/(1+\tau_j)}{\tilde{q}_j^{eq}}\right)^{\frac{\rho}{1-\rho}} = 1$, $\tilde{q}_j^E = 1$. Therefore, $\tilde{q}_j^{eq} \neq 1$ measures misallocation. When $1 + \tau_j = \nu_j$, the optimal tax neutralizes firm-level heterogeneity in energy-biased technology. As a result, $\tilde{q}_j^{eq} = 1$, which means that the allocation of production is solely based on q.

We can relate firm-level energy intensity to the wedge. First, notice that at $\tilde{q}_j^{eq} = \frac{\nu_j}{1+\tau_j}$, we get

$$\left(\frac{p_e e_j}{c_j y_j}\right)^{\underline{E}} = \left(1 + \frac{1-\beta}{\beta}\left(\frac{w}{p_e}\right)^{\frac{-\rho}{1-\rho}}\right)^{-1} = \left(1 + \frac{w l_j}{p_e e_j}\nu_j^{\frac{-\rho}{1-\rho}}\right)^{-1}$$

Where $p_e e_{int}^{\underline{E}} \equiv \left(\frac{p_e e_j}{c_j y_j}\right)^{\underline{E}}$ is the expenditure share of firm j on energy under the optimal allocation. It is equalized across firms. This means that we can rewrite energy intensity as

$$\frac{e_j}{c_j y_j} = e_{int}^{\underline{E}} \frac{\left(\frac{\nu_j}{\tilde{q}_j^{eq}}\right)^{\frac{\rho}{1-\rho}} + (1-\rho)\left(\left(\frac{\tilde{q}_j^{eq}}{q^m}\right)^{\frac{-1}{1-\rho}} - 1\right)}{\left(1 + (1-\rho)\left(\left(\frac{\tilde{q}_j^{eq}}{q^m}\right)^{\frac{-1}{1-\rho}} - 1\right)\right)}$$

First, if $\tau_j = 0$, $\frac{e_j}{c_j y_j} = e_{int}^{\underline{E}}$ because this firm does not misallocate its production. Secondly, if $\tilde{q}_j^{eq} > 1$, $\frac{\partial log \frac{e_j/c_j y_j}{(e_j/c_j y_j)^*}}{\partial log \tilde{q}_j^{eq}} > 0$. Increasing misallocation increases energy intensity. If $\tilde{q}_j^{eq} < 1$, $\frac{\partial log \frac{e_j/c_j y_j}{(e_j/c_j y_j)^*}}{\partial log \tilde{q}_j^{eq}} < 0$. Decreasing misallocation (by increasing

[79] This implicitly assumes that energy and output enters the welfare function proportionally. This simple objective also abstracts from aggregation across firms, which creates an additional margin since the allocation of sales share across firms determine total output, meaning that it should also be considered in the choice of $\tilde{q}_j^E$.

[80] Without any fiscal externalities, the optimal tax simply closes the gap in the marginal product of energy. If a second-best world in which the planner is constrained, the tax will trade off the reduction in dispersion against generating fiscal externalities that discourage production, reminiscent the optimal unemployment insurance literature and, in particular, the Baily-Chetty formula [Baily, 1978, Chetty, 2006].

$\tilde{q}_j^{eq}$ toward 1) decreases energy intensity.

If we aggregate across firms, assuming that the final demand is such that firms have constant sales share $c_j y_j = PY$, we obtain

$$\frac{E/PY}{(E/PY)^{\underline{E}}} = \frac{1}{N}\int_j \frac{\left(\frac{\nu_j}{\tilde{q}_j^{eq}}\right)^{\frac{\rho}{1-\rho}} + (1-\rho)\left(\left(\frac{\tilde{q}_j^{eq}}{q^m}\right)^{\frac{-1}{1-\rho}} - 1\right)}{\left(1 + (1-\rho)\left(\left(\frac{\tilde{q}_j^{eq}}{q^m}\right)^{\frac{-1}{1-\rho}} - 1\right)\right)} dj$$

$\nu_j \neq \nu_i = 1$ induces misallocation, increasing aggregate energy intensity. In other words, energy intensity in equilibrium, relative to the optimal allocation, is an average of the contribution of each firm to energy intensity, which can be expressed as a function of the firm-level wedges (as misallocation measures are also a function of the wedge $\tilde{q}_j^{eq}(\nu_j)$). This function balances the cost effect and the technology effect and, as before, always increases (decreases) with $\tau > (<) 1$. The term $\left(\frac{\nu_j}{\tilde{q}_j^{eq}}\right)^{\frac{\rho}{1-\rho}} \neq 1$ induces misallocation: if $\left(\frac{\nu_j}{\tilde{q}_j^{eq}}\right)^{\frac{\rho}{1-\rho}} = 1$, $\frac{E}{PY} = \left(\frac{E}{PY}\right)^*$. Indeed, misallocation arises because we are considering $\frac{E}{PY}$ (wile output is maximized). More (less) dispersion in ν_j increases aggregate energy intensity for $\rho > (<) 0$. Moreover, a higher (lower) correlation between ν_j and q_j^m increases aggregate energy intensity for $\rho > (<) 0$, because it implies a larger trade-off: firms that are effectively more energy efficient (that is, that use less energy) also have higher outsourcing capabilities.

As a final note, as long as firm-level sales $c_j y_j$ are maximized when the firm-level unit cost is minimized, aggregate GDP PY is maximized in equilibrium: $PY = \int_j c_j y_j dj = \int_j f(c_j) dj$. Since the allocation minimizes c_j for all firms, PY is maximized.[81]

A.1.2 Why does the make-or-buy decision of the firms matter for aggregate outcomes?

In this section, we use the toy model to illustrate why the make-or-buy decision matters. If firms are heterogeneous in their energy-biased technology, which firm is using energy matters. This is determined by the allocation of production across firms. Fundamentally, this allocation is driven by the selection of suppliers by each firm. However, the selection of suppliers include the firm itself, as many firms make their own intermediate inputs. Therefore, the make-or-buy decision, defined as both whether *and* whom to outsource, determines the equilibrium allocation of energy use across heterogeneous firms. Beyond supplier selection, the make-or-buy decision itself is crucial to properly identify the energy-biased technology of each firm, which is key when talking about allocation. In the following, we show how the aggregate response of energy intensity to an energy price shock depends on the make-or-buy decision of each firm, due to firm heterogeneity. The reallocation of energy use across firms that make-or-buy adjustments induces occurs through the extensive margin, which can generate significant, first-order effect on aggregate outcomes.

[81] The rationale is mirroring spatial equilibrium models in which agents make their location choices based on utility considerations [Donald et al., 2025]. Therefore, the allocation does not necessarily maximize aggregate real GDP. In our case, the allocation maximizes output but not welfare, because firms *don't* consider the energy externality in their make-or-buy decision.

To illustrate that the make-or-buy decision plays a key role as soon as firms differ along at least one dimension, we strip down the toy model to its simplest form and only compare the model with and without heterogeneity in energy-biased technology. In other words, we assume that $\gamma = \rho$ and that $q_{ij} = q \; \forall i, j$.

Homogeneous firms. First, suppose that there is no heterogeneity across firms: $\nu_j = \nu_i = \nu \; \forall i, j$ and $q_{ij} = q \; \forall i, j$. This directly implies that all firms are indifferent between making or buying their inputs. [82] Moreover, since all firms are identical, $cy = \frac{PY}{N}$.

The aggregate energy intensity response of to a common shock on energy price p_e is

$$\frac{\partial log E/PY}{\partial log p_e} = \underbrace{\frac{-1}{1-\rho}}_{\text{Substitution}} + \underbrace{\frac{\rho}{1-\rho}\left(\frac{p_e}{\nu q}\right)^{\frac{-\rho}{1-\rho}} \beta (c^{el})^{\frac{\rho}{1-\rho}}}_{\text{Scale}} = \frac{-1}{1-\rho} + \frac{\rho}{1-\rho}\frac{p_e E}{PY}$$

Heterogeneous Firms - No make-or-buy. We next introduce heterogeneity across firms, that is, $\nu_j \neq \nu_i \; \forall i \neq j$. However, we still shut down the make-or-buy decision: $\mathcal{I}_j$ is fixed and exogenous. The aggregate energy intensity response of to a common shock on energy price p_e is

$$\frac{\partial log E/PY}{\partial log p_e} = \underbrace{\frac{-1}{1-\rho}}_{\text{Substitution effect}} + \int_j \frac{e_j}{E} \left[\underbrace{\frac{\partial log c_j y_j / PY}{\partial log p_e}}_{\text{Reallocation: sales share}} + \underbrace{\frac{-\rho}{(1-\rho)} \frac{\frac{1}{|\mathcal{I}_j|}\int_{\omega_i \in \mathcal{O}_j} \frac{p_e e_i}{c_i x_i} - \frac{p_e e_j}{c_j y_j} d\omega_i}{1 + \frac{1}{|\mathcal{I}_j|}\int_{\omega_i \in \mathcal{O}_j} \frac{c_i x_i}{c_j^{el} x_j^{el}} d\omega_i}}_{\text{Substitution (intermediates) + scale effects}} \right] dj$$

We retrieve the same substitution effect between labor and energy (within-firm substitution) as in the homogeneous model. Additionally, the initial scale effect is also still present. It is entangled with the substitution between intermediate inputs, since they are priced differently. Moreover, the effect is no weighted by the energy share of each firm, that is by $\frac{e_j}{E}$, since firms are heterogeneous. One dimension of heterogeneity is enough to generate this result: it holds if $\mathcal{I}_j = \mathcal{I} \; \forall j$.

Finally, since firms are heterogeneous, their cost shares also respond to the shock: there is reallocation between firms at the intensive margin. We can shut down this margin by supposing that firms face a log-linear demand that imposes $c_j y_j = \frac{PY}{N}$. As we will show below, there will still be a reallocative effect induced by the make-or-buy decision.

Before introducing the make-or-buy decision, we consider three extreme cases.

Only in-house production: Suppose that firms only produce in-house, that is $|\mathcal{I}_j| = 1 \; \forall j$. Then,

[82] Indeed, the make-or-buy decision of the firms is

$$q^{-1}\left(\beta\left(\frac{p_e}{\nu}\right)^{\frac{-\rho}{1-\rho}} + (1-\beta) w^{\frac{-\rho}{1-\rho}}\right)^{\frac{-(1-\rho)}{\rho}} = q^{-1}\left(\beta\left(\frac{p_e}{\nu}\right)^{\frac{-\rho}{1-\rho}} + (1-\beta) w^{\frac{-\rho}{1-\rho}}\right)^{\frac{-(1-\rho)}{\rho}}$$

$$\frac{\partial logE/PY}{\partial logp_e} = \frac{-1}{1-\rho} + \frac{\rho}{(1-\rho)} \int_j \frac{e_j}{E} \frac{p_e e_j}{c_j y_j} dj$$

The response of aggregate energy intensity is a weighted average of the firm-level response to the shock that we characterized in the homogeneous case. The last line shows us that the difference with the homogeneous case boils down to $\int_j \frac{e_j}{E} \frac{p_e e_j}{c_j y_j} dj$, which is a weighted average of firm-level energy expenditure share, weighted by energy share. Similar to the Hulten theorem, the aggregate response simply depends on the firm (energy) shares. In the homogeneous case, we retrieve $\int_j \frac{e}{E} \frac{p_e e}{cy} dj = \frac{Ne}{E} \frac{p_e e}{PY/N} = \frac{p_e E}{PY}$. Importantly, $\frac{p_e e_j}{c_j y_j} = \left(\frac{p_e}{\nu_j q}\right)^{\frac{-\rho}{1-\rho}} \beta (c_j^{el})^{\frac{\rho}{1-\rho}}$. Thus, it is crucial to accurately estimate ν_j to measure the aggregate response. As we've discussed repeatedly, this requires controlling for the set of inputs produced in-house.

Only outsourcing: Suppose that all downstream firms outsource all their inputs, that is $\mathcal{I}_j = \emptyset \ \forall j$. In that case, the downstream firms heterogeneity is irrelevant

$$\frac{\partial logE/PY}{\partial logp_e} = N \left[\underbrace{\frac{-1}{1-\rho}}_{\text{Substitution effect}} + \underbrace{\frac{-\rho}{(1-\rho)} \frac{\int_{\omega_i} \frac{p_e e_i}{c_i x_i} d\omega_i}{(c/q)^{\frac{-\rho}{1-\rho}}}}_{\text{Substitution (intermediates) + scale effects}} \right]$$

The firm-level response is identical for all downstream firms, since they outsource to the same suppliers. Since they don't produce in-house, their own energy-biased technology ν_j is irrelevant. Therefore, they face the same prices and have the same marginal cost, and all respond in the same way.

Finally, consider the case where firms use only one input. We define the set of firms $\mathcal{F}_I$ as the firms that produce that inputs in-house, while the set $\mathcal{F}_O$ includes the firms that outsource it. Then,

$$\frac{\partial logE/PY}{\partial logp_e} = \underbrace{\frac{-1}{1-\rho}}_{\text{Substitution effect}} + \underbrace{\frac{\rho}{(1-\rho)} \left[\int_{j \in \mathcal{F}_I} \frac{e_j}{E} \frac{p_e e_j}{q^{\frac{\rho}{1-\rho}} c_j y_j} dj + \int_{j \in \mathcal{F}_O} \frac{e_j}{E} \frac{p_e e_i}{c_i x_i} dj \right]}_{\text{Scale effect}}$$

The scale effect depends on which firm along the supply chain is producing the input.

Heterogeneous Firms with make-or-buy. We now endogenize the make-or-buy decision. In particular, we let the sets $\mathcal{I}_j$ and $\mathcal{O}_j$ respond to the shock. As we will show, this adds extensive margin responses. Therefore, we can no longer consider $\partial logE/PY$, but we have to compute $\Delta logE/PY$. To simplify exposure, we again consider the one-input case. The aggregate energy intensity change is

$$\frac{\Delta E/PY}{E/PY} = \frac{-1}{1-\rho}\left(\frac{(E/PY)'}{E/PY}\right)\partial log p_e + \frac{-\rho}{1-\rho}\left(\int_{j\in\mathcal{F}_{OO}}\left(\frac{e_j}{E}\right)\frac{p_e e_i}{c_i x_i}dj + \int_{j\in\mathcal{F}_{II}}\left(\frac{e_j}{E}\right)\frac{p_e e_j}{q^{\frac{\rho}{1-\rho}}c_j x_j}dj\right)\partial log p_e$$
$$\underbrace{+\int_{j\in\mathcal{F}_{IO}}\left(\frac{e_j}{E}\right)\frac{(c_i^{el})^{\frac{-\rho}{1-\rho}} - (c_j^{el})^{\frac{-\rho}{1-\rho}}}{(c_j^{el})^{\frac{-\rho}{1-\rho}}}dj + \int_{j\in\mathcal{F}_{OI}}\left(\frac{e_j}{E}\right)\frac{(c_j^{el})^{\frac{-\rho}{1-\rho}} - (c_i^{el})^{\frac{-\rho}{1-\rho}}}{(c_i^{el})^{\frac{-\rho}{1-\rho}}}dj}_{\text{Extensive margin: make-or-buy adjustments}}$$

The second line captures the effect on aggregate energy intensity when firms adjust their make-or-buy decision as a response to the energy price shock. The extensive margin responses it induces have first-order reallocative effects.

In sum, it is crucial to account for the make-or-buy decision of the firms, because it induces extensive margin responses - through reallocation of energy use across firms - that have first-order effects in the aggregate response. However, the make-or-buy responses are very difficult to measure empirically. First, in-house production is not directly observed. Secondly, the make-or-buy decisions are persistent and likely joint with other decisions. Therefore, this paper does not intend to capture the true aggregate response of energy intensity to an energy price change. Instead, it aims at measuring the changes that would come from reallocation through the make-or-buy decisions, keeping all other decisions constant, with the intention to measure the importance of this channel. We thus use a structural model to identify the factors behind these decisions in the most flexible way, which allows us to isolate the changes in allocation induced by the make-or-buy decision under a counterfactual higher energy price.

A.2 Multiple-stage Model

This section provides more details on the multiple-stage model.

A.2.1 Independent decision for each input

We model the cost minimization problem for the energy and labor separately for each input within a firm. In other words, the firms choose each $e_j(\omega)$ and $l_j(\omega)$ separately. In this section, we show that this is equivalent to having firms choosing their total demand for energy and labor, and then choose how to allocate it across their different inputs. Indeed, we can write this second decision as firms choosing each $e_j(\omega)$ and $l_j(\omega)$ separately, but under the constraints that $e_j = \int_{\mathcal{I}_j} e_j(\omega)d\omega$ and $l_j = \int_{\mathcal{I}_j} l_j(\omega)d\omega$. Denote the respective Lagrange multipliers as λ_e and λ_l. Then, the cost minimization problem for any $e_j(\omega)$ yields $e_j(\omega) = x^{el}(\hat{\omega})^{\frac{\gamma-\rho}{1-\rho}}\beta_\omega\left(\frac{p_e+\lambda_e}{\lambda_c}\right)^{\frac{-1}{1-\rho}} y_j \nu_j^{\frac{\rho}{1-\rho}}$, where λ_c is the Lagrange multiplier on costs. Plugging it into $e_j = \int_{\mathcal{I}_j} e_j(\omega)d\omega$, we obtain $e_j(\omega) = e_j\frac{x^{el}(\hat{\omega})^{\frac{\gamma-\rho}{1-\rho}}\beta_\omega}{\int_{\mathcal{I}} x^{el}(\hat{\omega})^{\frac{\gamma-\rho}{1-\rho}}\beta_\omega d\omega}$. Therefore, we can actually rewrite the production function as a function of e_j and l_j:

$$y_j = \Big(\int_{\hat{\omega}\in\mathcal{I}_j}(\beta_{\hat{\omega}}^{1-\rho}(\nu_j e_j\frac{x^{el}(\hat{\omega})^{\frac{\gamma-\rho}{1-\rho}}\beta_{\hat{\omega}}}{\int_{\mathcal{I}} x^{el}(\hat{\omega})^{\frac{\gamma-\rho}{1-\rho}}\beta_\omega d\omega})^\rho + (1-\beta_{\hat{\omega}})^{1-\rho}(l_j\frac{x^{el}(\hat{\omega})^{\frac{\gamma-\rho}{1-\rho}}(1-\beta_{\hat{\omega}})}{\int_{\mathcal{I}} x^{el}(\hat{\omega})^{\frac{\gamma-\rho}{1-\rho}}(1-\beta_\omega)d\omega})^\rho)^{\frac{\gamma}{\rho}}d\hat{\omega} + x_j^{\mathcal{O}}(\omega)^\gamma\Big)^{\frac{1}{\gamma}}$$

The cost minimization problem with respect to the aggregate e_j and l_j yields $e_j = \left(\frac{c_j}{p_e}\right)^{\frac{1}{1-\rho}} y_j^{\frac{1-\gamma}{1-\rho}}\nu_j^{\frac{\rho}{1-\rho}}(\int_{\mathcal{I}} x^{el}(\hat{\omega})^{\frac{\gamma-\rho}{1-\rho}}\beta_\omega d\omega)$

(and similarly for l_j), where $x^{el}(\omega) = c_j^{\frac{1}{(1-\gamma)}} y_j \left(\beta_\omega \left(\frac{p_e}{\nu_j}\right)^{\frac{-\rho}{1-\rho}} + (1-\beta_\omega) w^{\frac{-\rho}{1-\rho}}\right)^{\frac{1-\rho}{\rho(1-\gamma)}}$. Thus, we obtain the usual unit cost:

$$c_j = \left(\int_{\mathcal{I}_j} \left(\beta_\omega \left(\frac{p_e}{\nu_j}\right)^{\frac{-\rho}{1-\rho}} + (1-\beta_\omega) w^{\frac{-\rho}{1-\rho}}\right)^{\frac{\gamma(1-\rho)\gamma}{\rho(1-\gamma)}} d\omega + (c_j^{\mathcal{O}})^{\frac{-\gamma}{1-\gamma}}\right)^{\frac{-(1-\gamma)}{\gamma}}$$

Which means that each make-or-buy decision can be taken independently as it boils down to comparing the unit costs $(c_j^{\mathcal{I}})^{\frac{-\gamma}{1-\gamma}}$ and $\left(\frac{c_i}{q_{ij}}\right)^{\frac{-\gamma}{1-\gamma}}$. Indeed, if $\hat{\omega}$ is outsourced, the total unit cost of the firm is

$$c_j = \left(\int_{\mathcal{I}_j^{other}} \left(\beta_\omega \left(\frac{p_e}{\nu_j}\right)^{\frac{-\rho}{1-\rho}} + (1-\beta_\omega) w^{\frac{-\rho}{1-\rho}}\right)^{\frac{\gamma(1-\rho)\gamma}{\rho(1-\gamma)}} d\omega + \left(\frac{c_i(\hat{\omega})}{q_{ij}}\right)^{\frac{-\gamma}{1-\gamma}} + \int_{\mathcal{O}_j^{\text{other}}} \left(\frac{c_i}{q_{ij}}\right)^{\frac{-\gamma}{1-\gamma}} d\omega\right)^{\frac{-(1-\gamma)}{\gamma}}$$

While if $\hat{\omega}$ is produced in-house,

$$c_j = \left(\int_{\mathcal{I}_j^{other}} \left(\beta_\omega \left(\frac{p_e}{\nu_j}\right)^{\frac{-\rho}{1-\rho}} + (1-\beta_\omega) w^{\frac{-\rho}{1-\rho}}\right)^{\frac{\gamma(1-\rho)\gamma}{\rho(1-\gamma)}} d\omega + \int_{\mathcal{I}_j^{\text{for } \hat{\omega}}} \left(\beta_\omega \left(\frac{p_e}{\nu_j}\right)^{\frac{-\rho}{1-\rho}} + (1-\beta_\omega) w^{\frac{-\rho}{1-\rho}}\right)^{\frac{\gamma(1-\rho)\gamma}{\rho(1-\gamma)}} d\omega \right.$$
$$\left. + \int_{\mathcal{O}_j^{\text{for } \hat{\omega}}} \left(\frac{c_i}{q_{ij}}\right)^{\frac{-\gamma}{1-\gamma}} d\omega + \int_{\mathcal{O}_j^{\text{other}}} \left(\frac{c_i}{q_{ij}}\right)^{\frac{-\gamma}{1-\gamma}} d\omega\right)^{\frac{-(1-\gamma)}{\gamma}}$$

Where $(c_j^{\mathcal{I}})^{\frac{-\gamma}{1-\gamma}} = \int_{\mathcal{I}_j^{\text{for } \hat{\omega}}} \left(\beta_\omega \left(\frac{p_e}{\nu_j}\right)^{\frac{-\rho}{1-\rho}} + (1-\beta_\omega) w^{\frac{-\rho}{1-\rho}}\right)^{\frac{\gamma(1-\rho)\gamma}{\rho(1-\gamma)}} d\omega + \int_{\mathcal{O}_j^{\text{for } \hat{\omega}}} \left(\frac{c_i}{q_{ij}}\right)^{\frac{-\gamma}{1-\gamma}} d\omega$ and the *for* $\hat{\omega}$ index represents the inputs in the recipe of $\hat{\omega}$ that would be produced in-house if $\hat{\omega}$ is produced in-house.

This results arises because more in-house inputs correspondingly implies less outsourced inputs. Thus, there is no love-for-variety in the number of inputs produced in-house. To give intuitions, we can shut down input heterogeneity by assuming that $\beta_\omega = \beta \ \forall \omega$, and that $\frac{c_i}{q_{ij}} = \frac{c}{q} \ \forall i$. In that case, $e_j(\omega) = \frac{e_j}{|\mathcal{I}_j|} \ \forall \omega$, $l_j(\omega) = \frac{l_j}{|\mathcal{I}_j|} \ \forall \omega$, and $x_j^{\mathcal{O}} = \int_{\omega \notin \mathcal{I}_j} (q\tilde{x}_j^{\mathcal{O}}(\omega))^\gamma d\omega = (1-|\mathcal{I}_j|)\left(q\frac{x_j^{\mathcal{O}}}{1-|\mathcal{I}_j|}\right)^\gamma$ where $\tilde{x}_j^{\mathcal{O}} = \frac{x_j^{\mathcal{O}}}{q}$. We also normalize the number of inputs in a recipe to 1.

$$\begin{aligned} y_j &= \left(|\mathcal{I}_j|(\beta^{1-\rho}\left(\nu_j \frac{e_j}{|\mathcal{I}_j|}\right)^\rho + (1-\beta)^{1-\rho}\left(\frac{l_j}{|\mathcal{I}_j|}\right)^\rho)^{\frac{\gamma}{\rho}} + (x_j^{\mathcal{O}})^\gamma\right)^{\frac{1}{\gamma}} \\ &= \left(|\mathcal{I}_j|^{1-\gamma}\left(\beta^{1-\rho}\left(\nu_j e_j\right)^\rho + (1-\beta)^{1-\rho} l_j^\rho\right)^{\frac{\gamma}{\rho}} + (1-|\mathcal{I}_j|)^{1-\gamma}(q\tilde{x}_j^{\mathcal{O}})\right)^{\frac{1}{\gamma}} \\ \Rightarrow c_j &= \left(|\mathcal{I}_j|\left(\beta\left(\frac{p_e}{\nu_j}\right)^{\frac{-\rho}{1-\rho}} + (1-\beta) w^{\frac{-\rho}{1-\rho}}\right)^{\frac{(1-\rho)\gamma}{\rho(1-\gamma)}} + (1-|\mathcal{I}_j|)\left(\frac{c}{q}\right)^{\frac{-\gamma}{1-\gamma}}\right)^{\frac{1-\gamma}{-\gamma}} \end{aligned}$$

Thus, there is no love-for-variety in the number of inputs produced in-house $|\mathcal{I}_j|$, despite diminishing returns in each individual inputs. This is because $|\mathcal{I}_j|$ affects all inputs, which jointly exhibit CRS, since more in-house inputs correspondingly implies less outsourced inputs.

A.2.2 Value-added and outsourcing

The labor expenditure share highlights that the *a priori* intuitive use of the value-added to sale ratio to characterize the extend to which the firm outsources might be misleading. Indeed, the VA to sales ratio is a sufficient statistics for

outsourcing, only once we condition on intermediate input effective costs and energy-biased technology that is, on the heterogeneous parameters of the firms.[83] To simplify exposure, suppose that $\rho(\omega) = \gamma \ \forall \omega$ and $\beta_\omega = \beta \ \forall \omega$. We have:

$$\begin{aligned}\frac{VA}{sales} &= \frac{wl_j}{c_j(\omega_j)x_j(\omega_j)} = \left(\frac{w}{c_j(\omega_j)}\right)^{\frac{-\rho}{1-\rho}} |\mathcal{I}_j|(1-\beta) \\ &= \frac{(1-\beta)w^{\frac{-\rho}{1-\rho}}}{\left(\beta\left(\frac{p_e}{\nu_j}\right)^{\frac{-\rho}{1-\rho}} + (1-\beta)w^{\frac{-\rho}{1-\rho}}\right) + \frac{1}{|\mathcal{I}_j|}\int_{\hat{\omega}\in\mathcal{O}_j}\left(\frac{c_i(\hat{\omega})}{q_{ij}}\right)^{\frac{-\rho}{1-\rho}} d\hat{\omega}}\end{aligned}$$

A firm can have a large $\frac{VA}{sales}$ because it produces a large set of in-house (large $|\mathcal{I}_j|$), but also because it faces low effective intermediate input costs, that is low $\frac{c_i}{q_{ij}}$ (for $\rho < 0$). However, higher q_{ij} induces less in-house production (small $|\mathcal{I}_j|$), since the q_{ij} are correlated across i, within a firm j. This means that the raw correlation between $\frac{VA}{sales}$ and $\mathcal{I}_j$ is ambiguous.

A.2.3 Decomposition

To shed light on the mechanisms through which energy prices influence energy intensity, we introduce a decomposition framework analogous to the canonical scale-composition-technique decomposition in environmental economics [Copeland and Taylor, 2004], or to the pure technology effect and the allocative efficiency effect in [Baqaee and Farhi, 2020]. We first briefly discuss the different effects at play at the firm-level (equation (8)), then study how they unfold at the aggregate level (equation (9)). To simplify exposure, we impose that $\rho(\omega) = \rho = \gamma$ and that $\eta = \varepsilon$, and define $\mathcal{I}_j^*$ and $\mathcal{I}_j^{SCC}$ as the sets of in-house inputs of firm j under the initial and final energy price, respectively, and similarly for $\mathcal{O}_j$. We obtain:

$$\begin{aligned}\frac{\partial ln\left(\frac{e_j}{c_j(\omega_j)x_j(\omega_j)}\right)}{\partial ln(p_e)} = &\underbrace{-1}_{\text{direct price effect}} + \underbrace{\frac{\rho}{1-\rho}\frac{p_e e_j}{c_j(\omega_j)x_j(\omega_j)}\left(1 - \frac{p_e e_j}{c_j(\omega_j)x_j(\omega_j)}\right)}_{\text{cost effect}} \\ &+ \underbrace{\int_{\hat{\omega}\notin\mathcal{I}_j^* \,\&\, \hat{\omega}\in\mathcal{I}_j^{SCC}} \frac{p_e e_j(\hat{\omega})}{c_j(\omega_j)x_j(\omega_j)} d\hat{\omega} - \int_{\hat{\omega}\in\mathcal{I}_j^* \,\&\, \hat{\omega}\notin\mathcal{I}_j^{SCC}} \frac{p_e e_j(\hat{\omega})}{c_j(\omega_j)x_j(\omega_j)} d\hat{\omega}}_{\text{in-house effect}}\end{aligned} \tag{8}$$

The first two effects are the most standard economic mechanisms. First, higher energy prices depress firms' energy demand, thereby lowering their energy intensity. Second, higher energy prices increase the cost of all firms, which mitigates or amplifies the direct price effect, depending on ρ. Third, higher energy prices incentivize firms to restructure their firm boundaries. They substitute away from producing in-house some inputs for which they can find suppliers that are relatively more energy efficient and can therefore offer a lower effective cost, and integrate the production of some inputs for which they are relatively more energy efficient than their potential suppliers. Related mechanisms arise at the aggregate level, but reallocation *across* firms now play a key role. Recall that aggregate energy intensity is expressed as

[83] The same intuition applies to energy and labor costs ("in-house" costs): $\frac{p_e e_j + wl_j}{c_j(\omega_j)x_j(\omega_j)} = \frac{\beta\left(\frac{p_e}{\nu_j}\right)^{\frac{-\rho}{1-\rho}} + (1-\beta)w^{\frac{-\rho}{1-\rho}}}{\beta\left(\frac{p_e}{\nu_j}\right)^{\frac{-\rho}{1-\rho}} + (1-\beta)w^{\frac{-\rho}{1-\rho}} + \frac{1}{|\mathcal{I}_j|}\int_{\hat{\omega}\in\mathcal{O}_j}\left(\frac{c_i(\hat{\omega})}{q_{ij}}\right)^{\frac{-\rho}{1-\rho}} d\hat{\omega}}$.

$$\frac{E}{PY} = \frac{1}{p_e PY}\int_0^N \left(\left(\frac{p_e/\nu_j}{c_j(\omega_j)}\right)^{\frac{-\rho}{1-\rho}} \int_{\omega\in\mathcal{I}_j} \beta_\omega d\omega\right) c_j(\omega_j)x_j(\omega_j)dj$$
$$= \frac{1}{p_e}\int_0^N \left(\left(\frac{p_e/\nu_j}{c_j(\omega_j)}\right)^{\frac{-\rho}{1-\rho}} \int_{\omega\in\mathcal{I}_j} \beta_\omega d\omega\right)\left(\delta_{\omega_j}\left(\frac{\varepsilon-1}{\varepsilon}\right)^{\varepsilon-1}\left(\frac{P}{c_j(\omega_j)}\right)^{\eta-1} + \int_{i:j\in\mathcal{S}_i} \frac{c_i(\hat{\omega}_i)x_i(\hat{\omega}_i)}{PY}\left(\frac{c_i(\hat{\omega}_i)}{c_i(\omega_j)}\right)^{\frac{\gamma}{1-\gamma}} di\right) dj$$

Therefore,

$$\begin{aligned}
\frac{dln\left(\frac{E}{PY}\right)}{dln(p_e)} &= \int_0^N \frac{e_j}{E}\left(\frac{\partial ln\left(\frac{e_j}{c_j(\omega_j)x_j(\omega_j)}\right)}{\partial ln(p_e)} + \frac{\partial ln\left(\frac{c_j(\omega_j)x_j(\omega_j)}{PY}\right)}{\partial ln(p_e)}\right) dj \\
&= \underbrace{-1}_{\text{direct price effect}} + \frac{p_e}{E/PY}\int_0^N \Bigg(\underbrace{\frac{\rho}{1-\rho}\frac{p_e e_j}{c_j(\omega_j)x_j(\omega_j)}\left(1-\frac{p_e e_j}{c_j(\omega_j)x_j(\omega_j)}\right)}_{\text{cost effect}} \\
&+ \underbrace{\int_{\hat{\omega}\notin\mathcal{I}_j^*\,\&\,\hat{\omega}\in\mathcal{I}_j^{SCC}} \frac{p_e e_j(\hat{\omega})}{c_j(\omega_j)x_j(\omega_j)}d\hat{\omega} - \int_{\hat{\omega}\in\mathcal{I}_j^*\,\&\,\hat{\omega}\notin\mathcal{I}_j^{SCC}} \frac{p_e e_j(\hat{\omega})}{c_j(\omega_j)x_j(\omega_j)}d\hat{\omega}}_{\text{in-house effect}} \\
&+ \underbrace{\frac{sales_j^{hh}(\omega_j)}{PY}(\eta-1)\left(\frac{dlnP}{dlnp_e} - \frac{dlnc_j(\omega_j)}{dlnp_e}\right)}_{\text{household demand effect}} + \underbrace{\int_{i:j\in\mathcal{S}_i^*\cap\mathcal{S}_i^{SCC}} \frac{sales_j^i(\omega_j)}{PY}\left(\frac{\partial ln\left(\frac{c_i(\omega_i)x_i(\omega_i)}{PY}\right)}{\partial ln(p_e)} + \frac{\partial ln\left(\frac{sales_j^i(\omega_j)}{c_i(\omega_i)x_i(\omega_i)}\right)}{\partial ln(p_e)}\right) di}_{\text{intensive firm demand effect}} \\
&+ \underbrace{\int_{i:j\notin\mathcal{S}_i^*\,\&\,j\in\mathcal{S}_i^{SCC}} \frac{sales_j^i(\omega_j)}{PY}di - \int_{i:j\in\mathcal{S}_i^*\,\&\,j\notin\mathcal{S}_i^{SCC}} \frac{sales_j^i(\omega_j)}{PY}di}_{\text{outsourcing effect}}\Bigg) dj
\end{aligned} \tag{9}$$

Equation (9) formalize how higher energy prices reduces aggregate energy intensity through distinct channels. The effects identified at the firm level are also present at the aggregate level, since they directly determine the firms' demand for energy (second and third line). But this demand is weighted by the firms' sales, which are also impacted. Firms' sales are affected because the firms' relative cost vary as energy prices increase, depending on their relative energy efficiency. Sales change at the intensive margin (fourth line): both types of customers (household and other firms) respond to the relative cost change by adjusting their demand for the firm's product. The customer firms' and household own revenue are affected by the energy price change directly, through changes in their own costs and labor income, respectively, and indirectly, through changes in their suppliers' costs, which changes the relative cost of the firm. In other words, energy prices alters the composition of demand by shifting expenditure away from energy-inefficient firms and towards energy-efficient ones. Third, and uniquely to a setting that considers firm boundaries, higher energy prices incentivize firms to substitute away from energy-inefficient suppliers (themselves included) in favor of more energy-efficient ones (fifth line). We term this the outsourcing effect. At the firm-level, it concerns the reallocation of production at the *extensive* margin: firms decide to shift the production of their input from one supplier to another. While the first two effects are standard in the environmental literature, the third is inherently tied to the cross-firms allocation of intermediate input production, and highlights the environmental benefits or losses of firm boundaries. This decomposition also highlights that, abstracting

from firm boundaries, one might misinterpret heterogeneous responses in firms' energy demand per unit of output as resulting only from differences in prices or energy-biased technology, while they effectively also arise from differences in firm boundaries.

Two final comments are worth making. First, this decomposition abstract from a technique effect because we keep the energy-biased technology (ν_j) and outsourcing capabilities (q_{ij}) fixed, in accordance with our measure of static misallocation. In a companion working paper [Marullaz, 2025], the interaction between firm boundaries and innovation in energy-biased technology is explored.

Secondly, the results on allocative efficiency on GDP in [Baqaee and Farhi, 2020] extend to our setting with a distorted energy price, since the effective cost $\frac{p_e}{\nu_j}$ is firm-specific.[84] Our allocative mechanism specifically introduces an extensive margin channel through the choice of firm boundaries and supplier selection for each input[85] that is worth pointing out. In the make-or-buy decision, potential suppliers (including the firm itself) are perfectly substitutable, regardless of the elasticity of substitution between inputs (γ). This means, for instance, that, when there is no gains or losses from reallocation following a change in factor prices when $\gamma \to -\infty$ (Leontief) in [Baqaee and Farhi, 2020]'s setting, our extensive margin is separate and still active, allowing for important reallocation gains or losses.[86]

A.2.4 Discussion of assumptions

Perfect competition. We assume perfect competition across firms to abstract from any distortion other than the environmental externality, which is the focus of our paper. However, our results would be qualitatively similar in a monopolistically competitive setting, as it will deliver constant markups. Heterogeneity in markups would ne absorbed by the general outsourcing capability q_{ij}, which is assumed to be a primitive of the firms. Thus, it wouldn't affect our results as long as markups don't respond to aggregate energy price changes.

Acyclicality of production. We assume that production is acyclic and exclude round-about production. As shown by [Dhyne et al., 2023], this assumption guarantees a unique equilibrium. Moreover, the impact of firm-level make-or-buy decisions on aggregate energy intensity depends on the energy-biased technology of the firm *relative* to the one of its (potential) supplier. Therefore, we need to explicitly characterize all potential suppliers in terms of their energy-biased technology, and we need to solve for the optimal supplier selection in the counterfactual decision. This step requires knowing the price charged by all these potential suppliers before solving for the fragmentation decision of a given firms, which involves solving the model recursively moving downstream. Along with our calibration results, we will provide several statistics to evaluate the cyclicality of our network.

[84] Their framework can be generalized to include heterogeneous factor-biased technology by modeling each factor with such technology as produced by a distinct (fictitious) producer. This means that the distortion on p_e would only apply to those "fictitious" energy producers. Therefore, changes in p_e impact firms differently and can lead to reallocation.

[85] This is different from an extensive margin that would operate through entry of new products, as the latter affects macroeconomic outcome through love-for-variety.

[86] The authors show that misallocation increases with the elasticity of substitution between producers.

Single-product firms. Our firms produce only one single product, characterized by their ATECO code. This limitation is imposed by the data, but might bias our recipes and the set of in-house inputs if the firm's purchases actually blend together inputs for multiple products. Fortunately, the Italian manufacturing sector is characterized by its large number of small firms [Agostino et al., 2016], which are less likely to be able to diversify into making multiple products. Moreover, we identify multi-product firms using export data[87] and exclude them to build the recipes (see section B.2.1). We then build their firm-level in-house input set based on all the inputs they require to produce all their products. Finally, we consider them as potential suppliers of all their products, and refine their observed outflows to correspond to an input that enters the buyer's recipe, among the set of products sold by the multi-product firms.

One supplier per input. To be able to retrieve the set of inputs produced in-house, we are forced to assume that each firm has only one supplier (potentially itself) per input. Absent this restriction, an input could be produced in-house and outsourced at the same time. By construction, we cannot evaluate whether the latter holds true in the data. However, we do observe the number of suppliers per input for each firm and, reassuringly, the number of buyer-input pair with multiple suppliers is very limited.

No interdependencies. As in [Eaton et al., 2022], we assume that there is no fixed cost of outsourcing.[88] This exclude any interdependencies in the fragmentation decision of each input.[89] We chose this specification for two reasons. First, it allows to relate our model to generic frameworks of distorted production network (e.g., [Baqaee and Farhi, 2020, Liu, 2019]) by considering the production of each input as undertaken by a fictitious producer. Secondly, inderdependencies would give rise to a complex combinatorial discrete choice problems. However, the tools recently developed to solve those problems [Antras et al., 2017, Arkolakis et al., 2023] cannot be applied to our setting. Indeed, the mapping between the energy-biased technology, which characterizes our firms, and the make-or-buy decision, is not necessarily monotonic if fixed costs to outsourcing are introduced.[90] Large ν_j-firms benefit more from producing in-house, but might also be larger depending on factor prices and the level of ν_j itself[91] hence would tend to outsource more. Furthermore, fragmentation decisions are far from being solely determined by energy-related considerations, and are known to be highly idiosyncratic [Alfaro and Charlton, 2009, Lafontaine and Slade, 2013, Alfaro et al., 2019, Atalay et al., 2019, Boudreau et al., 2023, Grossi et al., 2023]. Therefore, our firms are also characterized by their broadly-defined outsourc-

[87] This implicitly assumes that all firms that only sell domestically are single-product firms. Standard trade models 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.

[88] Relatedly, we also abstract from any economy of scope in in-house production that would arise if e.g., firms could improve their energy-biased technology. It would similarly result in firms maximizing total profits, not minimizing costs in any single input decision separately [Panzar and Willig, 1981, Boehm and Oberfield, 2020].

[89] This is somewhat equivalent to a world in which production is purely sequential in a snake form, and each establishment (stage) of a firm makes its own input decision independently, taking the upstream decisions as given. However, the decision is governed by the common firm-level energy-biased technology and the firm-level component in the outsourcing capabilities, representing technology diffusion within firm.

[90] Formally, this means that the condition to solve for the policy function, SCD-T, in [Arkolakis et al., 2023] is not satisfied.

[91] As discussed below, we identify the energy-biased technology *relative* to the labor-biased technology $\frac{\nu_j}{A_j}$. A change in the *relative* energy-biased technology affects the firm's profits positively if $\beta_\omega \left(\frac{p_e}{\nu_j/A_j}\right)^{\frac{-\rho}{1-\rho}} > (1-\beta_\omega)(w)^{\frac{-\rho}{1-\rho}}$, where A_j is a labor-biased technology.

ing capabilities q_{ij}, which vary across suppliers to match the empirical observations that firms outsource the same input to different suppliers. A monotonic ranking of $(\nu_j, \{q_{ij}\}_{\omega \in \Omega_{\omega_j}, i \in \mathcal{S}_\omega})-$firms would require to assume a perfect correlation between those characteristics, or to consider only a few types in one of the component (e.g., reducing the support of ν_j to a binary categorization that splits firms into using a clean and a polluting technology), such that the ranking can be done relative to the other component within each type. However, a key contribution of our paper is to specifically identify the joint distribution of the energy-biased technology and these outsourcing capabilities, and the discrete type alternative cannot match the large heterogeneity that we observe in both characteristics. A third alternative would be to model those two components as affecting different types of costs. For instance, ν_j enters the marginal costs, but the outsourcing capabilities are all summarized into a fixed cost.[92] However, not only does it prevent to discuss specialization gains unrelated to energy, but the non-monotonicity issue implied by ν_j remains. Overall, because those interdependencies are not the focus of our analysis, we deem reasonable to abstract from them in light of the complexities they would generate. In a companion working paper [Marullaz, 2025], the interaction between firm boundaries and innovation in energy-biased technology is explored, which directly creates interdependencies.

Identical recipes. When building the recipes (see section 5), we assume that products require the same inputs regardless of which firm is producing it. In other words, this assume away any within-product specificities or quality differentiations. For instance, a car needs tires, which need rubber, regardless of whether it is a Lamborghini, or an Alfa Romeo. Our recipes are built as sets, thus don't impose any restriction on the quantities needed. Quantity and quality differences, or customized features, are therefore captured by the relationship-specific capabilities q_{ij}. One might fear that this might bias our results. Consider a very specialized supplier, selling to only a handful of buyers. It will be attributed low q_{ij} for all its other potential links in our calibration strategy. However, high specialization usually implies higher prices, meaning that this supplier's observed q_{ij} will accordingly be high. Overall, the estimated fixed effect of this supplier will capture these two counteracting forces, landing somewhere in the middle. In our counterfactual allocation, some buyers might be induced to select this supplier, even if effectively, the specialized input is not suited for their product (i.e., this supplier does not actually belong to their set of potential supplier). On the other hand, its mid-value of q_{ij} might incentivize buyers not to choose it, when they would effectively benefit a lot from its specialized input. These bias will average out if the "wrongly-lost buyers" and the "wrongly-gained buyers" are not systematically different. Since firms requiring specialized inputs could be more productive, or more specialized, we will carry a robustness exercise to mitigate this concern. In particular, we keep the make-or-buy decision of specialized inputs - based on [Rauch, 1999]'s classification - fixed across all scenarios, with the rationale that those inputs are too specialized to be produced by any other supplier regardless of the economic conditions.

Identical factor prices. We assume identical factor prices across firms, to focus on the heterogeneity in energy-biased technology. The assumption of identical labor prices is not as restrictive as it seems: any comparative advantage in wages that enters the make-or-buy decision will be embedded in the outsourcing capability term in our identification. However,

[92] [Antras et al., 2017] show how to solve for a model with heterogeneity along variable and fixed costs of sourcing, with the fixed cost generating interdependencies.

it won't affect the actual decision on relative quantity or energy to labor in the production phase. Since our focus in on the make-or-buy decision rather than the production decision, we deem reasonable to abstract from this heterogeneity. Similarly, identical energy prices is innocuous if we think of ν_j not exclusively as a technology in the mechanical sense, but as capturing any residual firm-level heterogeneity related to energy (relative to labor). Energy prices[93] and energy-biased technology enter all relevant firm-level decisions in the same way, meaning that variations in one component is isomorphic to variations in the other. The fixed factor prices assumption, however, does play a role when we aggregate firm-level energy use to measure aggregate energy intensity. Indeed, this measure is in quantity, and not in value terms, since this is the relevant metrics to think about emissions.[94] In equation (4), energy prices p_e and energy-biased technology ν_j enter with different exponents, so heterogeneity directly modeled in p_e would have slightly different impact relative to the dispersions absorbed by ν_j. This difference increases as ρ moves away from 1. We will allow for some energy price heterogeneity in two robustness exercises. In the first exercise, we build total energy prices using the different fuel mix of the firms, and in the second exercise, we restrict our sample to firms for which we observe energy contracts, using the Indagine sulle imprese industriali e dei servizi (Inquiry into investments of industrial and services firms) survey conducted by the Bank of Italy on a representative sample of firms with 20 employees or more, and in which an ad-hoc section collecting firm-level energy information was added in 2021 [Alpino et al., 2024].

A.3 Comparison with other models

In this section, we discuss different approaches taken by the recent literature, and explain how our framework compares to them, or why we had to diverge from them.

A.3.1 Theoretical Framework

Firm boundary literature. The recent firm boundary literature has significantly advanced our understanding of the drivers and consequences of the organizational structure of the firms, by conceptualizing those decisions through a firm-level relational framework that characterizes supply chain linkages by contractual frictions, relationship-specific investments, and customized inputs [Grossman and Helpman, 2002, Sahay, 2003, Alfaro et al., 2019, Antràs, 2020, Boudreau et al., 2023]. These highly idiosyncratic, relationship-specific drivers guide our modelization of the outsourcing capabilities. However, as the focus of these analysis was to understand the determinants of firm boundary decisions, they have abstracted from the production network. They usually study the equilibrium decisions at only one stage of the supply chain, or bilateral links formation. However, when firms are heterogeneous, supplier selection is crucial to quantify the extent of misallocation that is induced by inefficient boundary decisions. Indeed, the extent of misallocation is not only driven by *whether* the firm is producing in-house or outsourcing, but also *to which firm* it would outsource. As our identification strategy highlights, an accurate measure of the energy-biased technology can only be retrieved if the boundaries of the firms are properly taken into account. Therefore, our analysis calls for a different framework that embeds the fragmentation decision into a full production network in which firms at each stage choose their boundaries.

[93] Since we normalize the wage to 1 in all scenarios, variation in labor prices are absorbed by energy prices.

[94] Unlike changes in productivity, when a factor revenues are rebated to the household, a change in its price only affects *physical* quantities by changing the input allocation.

Production network. The production network literature has developed a large set of generic tools to study inefficient networks (e.g., [Baqaee and Farhi, 2019, Liu, 2019, La'O and Tahbaz-Salehi, 2022]). Those model restrict to firms producing only one input or stage in-house. However, some intuitions and results still apply to our setting, because the latter is a special case of these generic model, thanks to our assumption that there is no fixed cost to outsource or produce in-house. Indeed, because the make-or-buy decisions of the firms are taken independently for each input, one can decompose firms into as many fictitious producers as there are inputs in the firms' recipes. This effectively creates a network of single-input producers. The remaining difference between our setting and those general frameworks is that energy is a primary factor whose supply is not fixed. However, this can be resolved by modeling energy as another intermediate input, that also requires labor (and/or other intermediate inputs) in its production. In this world, our setting resembles [Liu, 2019]'s model, because our distortion ($p_e < p_e^{SCC}$) is mathematically equivalent to his "quasi-rent" wedges: their "revenue" are not rebated to the household, and they can be seen as a deadweight-loss in the sense that they leave the economy in terms of the consumption good.[95] With labor as the unique factor of production, we could use [Liu, 2019]'s non-parametric results to think about the effect of different carbon tax schedules (applied differently to e.g., different sectors) on total consumption, without making any functional form assumption.[96] However, the insights for our analysis are limited: indeed, they simply tell us that energy-using firms should be taxed, independently of their underlying characteristics. This is because the distortion centrality is built in terms of costs, meaning that the energy-biased *technology* ν_j is irrelevant, and all firms using energy face the same wedge. However, we are interested in a physical quantity term, that is, aggregate energy E, for which the distribution of energy use by energy-biased technology matters.[97] To fully understand the effect of different carbon schedules on aggregate energy (intensity), we thus need to go beyond this framework and impose some structure that allows us to identify the production technologies.[98] This echoes our discussion about using [Baqaee and Farhi, 2020]'s generic framework. Because their results are ex-post, they require observing the counterfactual allocations. This again confirms that we need to impose some structure. As discussed in section 5, we nonetheless leverage our rich firm-level dataset to avoid as much as possible to make functional form assumptions on the underlying distributions of the firms' primitives in our model ($\nu_j, \{q_{ij}\}$), on which little is known.

[95] To see the equivalence, suppose that emissions affect the utility indirectly by damaging production: $PC = (1 - D(E))PY$. In that case, the welfare object of interest is PC, and not PY, and $D(E)PY$ corresponds to the DWL, as it is a resource that simply leaves the economy. The same applies to the case where energy enters the household's disutility directly, as again, we are interested in PC, which is equal to total income through the budget constraint. Total income is based on p_e, not on p_e^{SCC}, therefore the distortions leave the economy without benefiting any group.

[96] Note that those results solely apply to an economy with no initial carbon tax.

[97] As [Liu, 2019] emphasizes, one cannot generally infer the allocation response to policy solely from empirical objects observed in the pre-intervention economy. Allocation effects are determined by the structural properties of production technologies, which are not directly observable. By contrast, price responses can always be characterized by reduced-form objects identified from the pre-policy equilibrium. The reason is that producers optimize on costs and, by the envelope theorem, policy-induced changes in production elasticities affect prices only at the second order — even in environments with distortions.

[98] In particular, with the numeraire $w = 1$, $d\frac{PY/E}{dp_e} = 1 + \frac{dL^m/E}{dp_e}$, where $L^m = \int_j^N l_j dj = \int_j^N \frac{dlny_j}{dlnl_j} p_j y_j dj$ and $E = \frac{1}{p_e} \int_j^N \frac{dlny_j}{dlne_j} p_j y_j dj$, which emphasizes that we need to know how the physical allocation of resources across firms change in the counterfactual economy and that for this, we need to know the structural characteristics of the production technologies of each firm. Note that without any heterogeneity in energy-biased technology across firms ($\nu_j = \nu$), the firm boundaries would not matter, as we could aggregate the network into a representative firm, and the substitution between aggregate labor and aggregate energy would be a sufficient statistics.

Relatedly, a recent literature has endogenized the production network by modeling firms choosing their number of inputs through a love-for-variety and matching costs trade-off [Huneeus, 2018, Lim, 2018, Acemoglu and Azar, 2020, Bernard et al., 2022, Dhyne et al., 2023, Acemoglu and Tahbaz-Salehi, 2025]. In that context, one could draw a parallel with firm boundary decisions by understanding in-house production as the outside option to changing the number of inputs, such that the overall recipes remain fixed. However, applying these frameworks to our question would mechanically result in larger firms (i.e., more energy efficient in our context) outsourcing more, a correlation that is not necessarily supported in the data. We avoid this restriction by staying agnostic on the relative values of the different drivers behind the outsourcing decision $(\nu_j, \{\{\nu_i, q_{ij}\}_{i\in\mathcal{S}_\omega}\}_{\omega\in\Omega_{\omega_j}})$. Some of these models abstract from relationship specificities, or model a continuum of suppliers with identical prices. This is not an appropriate framework to quantify misallocation, which depends on the alternative options that is, the energy use of the potential suppliers. Most of these models also characterize the equilibrium in terms of parametric distribution, while our rich data allows us to build the price that the buyer faces, and to recover the effective cost and the outsourcing capabilities for each firm, hence their effective make-or-buy decision in the counterfactual scenario. Therefore, we can go beyond parametric distributions and calibrate each of our model unit using a specific firm in our data.

[Boehm and Oberfield, 2023]'s model. [Boehm and Oberfield, 2023] provide the closest framework to ours, as they also embed a model of make-or-buy decisions within a multiple-stage production network.[99] They characterize their equilibrium by making novel functional form assumptions that deliver a closed-form solution for the probability to outsource at each node. However, their framework doesn't allow for firm-specific components driving the outsourcing capability distribution. This is problematic for our setting, in which we are particularly interested in the joint distribution between the firms' energy-biased technology and their outsourcing capabilities, a moment for which we have little empirical evidence and that drives the extent of energy misallocation induced by inefficient make-or-buy decisions. Moreover, our data is richer in the sense that we directly observe the identity of the buyer and the supplier for each pair. As previously discussed, this allows us to go beyond distributions to calibrate our model.

A.3.2 Empirical Strategy

Network Econonometrics. A natural approach to our calibration would be to follow the network econometrics literature, which sees network formation as the result of an optimal choices of individuals who maximize their objective

[99] [Eaton et al., 2022] also include in-house production within a network of buyers and sellers. However, the in-house and outsourcing production process differ beyond the firm-specific technologies. Their production network exhibits a spider-like structure: each input (labeled task) enters the firm's production horizontally. This means that, within each firm, an input is produced in-house using a primary factor only. In contrast, if this input is outsourced, it corresponds to the supplier's final output, which is produced by combining all inputs. This specification implies that downstream in-house input costs are not a function of upstream input costs, meaning that the distribution of the former can be exogenously determined by the underlying factor-biased technology distribution. This is a primitive of the model conveniently set to match the distribution of the outsourcing costs, such that the options can easily be compared and the outcome of the decision can be expressed in terms of probability. However, this framework is not appropriate to compare the allocation of energy across buyer and supplier through the make-or-buy decision, as their production processes differ beyond their factor-biased technology.

function, exactly like ours. The statistical analysis of network formation then relies on a dyadic non-linear regression, where the decision to establish a link between agents i and j is modeled as a binary choice driven exclusively by their pair-specific attributes. This specification is straightforward to implement and has been widely adopted in empirical work (e.g., [Mayer and Puller, 2008, Fafchamps and Gubert, 2007]). Although such models abstract from strategic interdependencies in network formation, they nonetheless succeed in replicating key stylized facts of observed networks [Jochmans, 2018]. To account for unobserved agent-level heterogeneity,[100] researchers often introduce high-dimensional fixed effects [Holland and Leinhardt, 1981],[101].

Our setting differs from the standard dyadic model, in which each link formation decision consists of only two mutually exclusive alternatives, because we consider *multiple*, mutually exclusive alternatives. This implies that the outside option that determines the payoff of forming a link is the cheapest potential supplier. Indeed, only one supplier per input can be chosen by each buyer.[102] This implies that the link formation rule for the effectively chosen relationship (in-house production included) differs from the rule that applies to the other potential relationships, since the outside options differ. The outside option of the former is the second cheapest supplier, who is unknown, while the outside option of all the other links is the chosen relationship. However, the chosen relationship must also be included in the estimation, to create variations in the outcome variable. Therefore, this approach is not applicable to our setting. We thus turn to discuss Discrete Choice Models, which are closer to our setting.

Discrete Choice Models. Our setting is reminiscent of a discrete choice model (DCM) [McFadden, 1974], widely used in the IO (e.g. [Berry, 1994]) and Trade (e.g., [Eaton and Kortum, 2002, Head and Mayer, 2014]) literature. As in network econometrics, DCM are based on the idea that agents choose among a discrete number of mutually exclusive alternative, selecting the one that provides them with the highest perceived benefit. This means that the econometrician only observes the actually chosen alternative. The benefits are modeled as depending on different attributes, including an idiosyncratic component that follows a Type I extreme value distribution. The estimation strategy leverage the properties of those distributions to control for selection, for which a closed-form expression can be derived. In our context, this would mean expressing the outsourcing capabilities of all potential alternative (in-house included) as a function of buyer and supplier fixed effects, some controls, and a Gumbel iid noise.

We cannot follow this strategy due to the presence of the make-or-buy decision that renders the costs stochastic. Indeed, this strategy requires all alternatives to be characterized by an idiosyncratic Gumbel term, including the in-house option (q_{jj}) which is also an alternative. This means that it relies on the expectation:

$$E\left(\max_{i\in\mathcal{S}_\omega}\frac{q_{ij}}{c_i}\right)=E\left(\frac{q_{ij}}{c_i}|\frac{q_{ij}}{c_i}>\frac{q_{sj}}{c_s},\ldots,\frac{q_{ij}}{c_i}>\frac{q_{jj}}{c_j^{\mathcal{I}}}\right)$$

This random variable is logit distributed only if all the unit costs c_i are deterministic. Only in that that case, we can use:

[100] The likely presence of individual effects bias the identification of the homophily parameters if not accounted for, because gregarious agents will create more links indiscriminately of their similarities on particular attributes with other agents [Graham, 2017].

[101] In directed network such as ours, agents may differ both in their propensity to send ties and in their attractiveness as receivers. Accordingly, two separate fixed-effects must be introduced to capture the heterogeneity in the buyer's tendency to initiate outgoing links (that is, outsourcing) and in the supplier's propensity to receive incoming links.

[102] If we allow for multiple suppliers per input, we can no longer infer the in-house set from the set of inputs in a recipe that are not outsourced.

$$E\left(\frac{q_{ij}}{c_i}|\frac{q_{ij}}{q_{sj}} > \frac{c_i}{c_s}, \ldots, \frac{q_{ij}}{q_{jj}} > \frac{c_i}{c_j^{\mathcal{I}}}\right)$$

To illustrate the issue, suppose that $q_{kj} = exp(z_k + \varepsilon_{ij})$ and that $q_{jj} = exp(z_j + \varepsilon_{ij})$. Notice that there is no buyer fixed effect in q_{ij}, which is already a limitation relative to our baseline strategy.[103] Suppose that we run our fixed effects regression product by product, meaning that we can build each q_{ij} sequentially. We can then build c_k after having identified q_{jj} from the sales to the household (instead of recovering ψ_j). Then, consider input ω that is outsourced to supplier i. Exploiting our directed acyclical graph assumption, all (potential) suppliers' unit costs $c_s(\omega)$ are known (hence deterministic) before we consider estimating the fixed effects behind the q_{sj} at this stage. Similarly, all the fixed effects z_k^s of the (potential) suppliers of even more upstream inputs $\hat{\omega} \in \Omega_\omega$, which enter the outside option $c_j^{\mathcal{I}}(\omega)$ are also known. However, $c_j^{\mathcal{I}}(\omega)$ remain stochastic, because it contains q_{jj} that must be distributed Gumbel.

$$\frac{c_j^{\mathcal{I}}(\omega)}{q_{jj}} = \left(\left(c_j^{el}(\omega)^{\frac{-\gamma}{1-\gamma}} + \int_{\hat{\omega}\in\mathcal{I}_j\cap\Omega_\omega} c_j^{el}(\hat{\omega})^{\frac{-\gamma}{1-\gamma}} d\hat{\omega}\right) + q_{jj}^{\frac{\gamma}{1-\gamma}} \int_{\hat{\omega}\in\mathcal{O}_j\cap\Omega_\omega} \left(\frac{c_i^{\mathcal{O}}(\hat{\omega})}{q_{ij}}\right)^{\frac{-\gamma}{1-\gamma}} d\hat{\omega}\right)^{\frac{1-\gamma}{-\gamma}}$$

With $q_{ij} = exp(z + \varepsilon_{ij})$ multiplying only the "outsourcing" part of the cost, $\frac{c_j^{\mathcal{I}}(\omega)}{q_{jj}}$ is not Gumbel distributed, thus $\frac{q_{ij}c_j^{\mathcal{I}}}{q_{jj}} > c_i$ is not logit-distributed, and the usual Type I extreme value property cannot be applied. Moreover, this strategy completely ignores the information coming from the inputs that are produced in-house, as these decisions could only be used once all q_{ij} have been estimated (and made deterministic), jointly at the firm-level. This would allow to identify the buyer fixed effect z_j, but this is useless since q_{jj} has already been characterized and z_j doesn't enter anywhere else. If it would enter q_{ij}, then the upstream suppliers' effective costs could not be computed.[104] The difficulty arises from the combination of the production network and the unobserved within-firm transactions on in-house inputs. Intuitively, it is because i) q_{jj} only applies to in-house inputs, but to all of them and ii) because in-house production is an option, q_{jj} must be stochastic.[105]

More generally, our model specification does not allow to characterize the shares of firms outsourcing a certain input. Indeed, even if the outsourcing capability distribution was e.g., sector-specific, instead of relationship-specific (through the buyer and supplier fixed effects), the cutoffs remain they depend on the in-house cost $c_j^{\mathcal{I}}$ and the supplier cost $c_i^{\mathcal{O}}$. As opposed to [Boehm and Oberfield, 2023], we are *computing* these costs, not characterizing them through a distribution (which would require a Cobb-Douglas production function and their distributional assumptions), so it is unique to each buyer and each supplier. Thus, even if the underlying distribution was the same for all firms within a sector, we would still

[103] A buyer fixed effect would make the upstream q_{kj} still stochastic even after estimating z_k.

[104] As a final comparison to our chosen strategy, recall that one of its distinctive feature is that it is directly calibrated using our firm-level data, that is a model unit corresponds exactly to an Italian firm. We can therefore envisage the data as corresponding exactly to our model, and not as *one* realization of our model. Accordingly, our preferred strategy always guarantees that the calibration is consistent with the data, that is, a potential supplier that is not chosen in equilibrium will always obtain an outsourcing capability draw that match this empirical observation. The DCM approach would deliver draws that might not be consistent with this observation, since the distribution of the idiosyncratic error term is not characterized to fit the firm-level decisions.

[105] Otherwise, we could identify it using household sales, following the same method that we apply to calibrate ψ_j.

not be able to express the shares of firms outsourcing a certain input ω in closed-form, because this closed-form depends on the cutoff, and that cutoff is firm-specific.

Symmetric outsourcing capabilities. Outsourcing capabilities could be simplified to be symmetric for each pair (i.e., $q_{ij} = q_{ji}$). In effect, this assumes that the buyer and the supplier fixed effects are equal for each firm ($z_j^s = z_j^b$). In that case, one can recover any unobserved outsourcing capability between two firms by using the observed q_{ij} that link each intermediary step. Formally,

$$\begin{aligned} q_{aj} &= e^{z+z_a+z_j+\varepsilon_{aj}} \\ \hat{q}_{aj} &= \frac{q_{ab}}{q_{bc}}\frac{q_{cd}}{q_{de}}\cdots\frac{q_{gh}}{q_{hi}}q_{ij} \\ &= e^{z_a-z_c+\varepsilon_{ab}-\varepsilon_{bc}}e^{z_c-z_e+\varepsilon_{cd}-\varepsilon_{de}}\cdots e^{z_g-z_i+\varepsilon_{gh}-\varepsilon_{hi}}e^{z+z_i+z_j+\varepsilon_{ij}} \\ &= e^{z+z_a+z_j}e^{\varepsilon_{ab}-\varepsilon_{bc}+\varepsilon_{cd}-\varepsilon_{de}\cdots+\varepsilon_{gh}-\varepsilon_{hi}+\varepsilon_{ij}} \end{aligned}$$

However, this identification strategy only holds if all the error terms ε, conditional on the choices made by the firms at each intermediary step, sum to zero in expectation (i.e., $E(\varepsilon_{ab} - \varepsilon_{bc} + \varepsilon_{cd} - \varepsilon_{de}\cdots + \varepsilon_{gh} - \varepsilon_{hi} + \varepsilon_{ij}|\mathcal{I}_s, \mathcal{O}_s \forall s \in \{a \to j\}) = E(\varepsilon_{aj}|\mathcal{I}_s, \mathcal{O}_s \forall s \in \{a \to j\}) = 0$, where $\mathcal{I}_s, \mathcal{O}_s \forall s \in \{a \to j\}$ represents the in-house and outsourcing choices of all the firms for which we observe the transaction along the supply chain that goes from a to j). This is unlikely considering that the observed transactions are a selected sample of relationships. In other words, the selection issue arises in that context as well, and one cannot apply the Type I Extreme Value toolbox to deal with it. Indeed, it would again require to first compute the unit costs, which depends on the suppliers' q_s and the buyer q_j (for the in-house cost), which are specifically the coefficients that are being identified. Moreover, this strategy again delivers a simulated model that might not be fully consistent with the data.

B Empirical Appendix

B.1 Energy expenditures

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 energy. The corresponding ATECO and TARIC codes are listed in tables 1 and 2 respectively. For our manufacturing firms, which exclude firms that process and resale energy, these purchases correspond to the actual fuel used to power their production.

Table 1: List of ATECO codes of energy-generating activities

ATECO Code	**Activity**
05.20.00	Mining of lignite
06.10.00	Extraction of crude petroleum
06.20.00	Extraction of natural gas
09.10.00	Support activities for petroleum and natural gas extraction
16.26.00	Manufacture of solid fuels from vegetable biomass
19.10.00	Manufacture of 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
20.11.00	Manufacture of industrial gases
20.51.00	Manufacture of liquid biofuels
24.46.00	Processing of nuclear fuel
28.29.20	Manufacture of distilling or rectifying plant for petroleum refineries and chemical industries
35.11.00	Production of electricity from non-renewable sources
35.12.00	Production of electricity from renewable sources
35.13.00	Transmission of electricity
35.14.00	Distribution of electricity
35.15.00	Trade of electricity
35.16.00	Storage 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
35.24.00	Storage of gas as part of network supply services
35.30.00	Steam and air conditioning supply
38.22.00	Energy recovery (from waste disposal)
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
46.81.00	Wholesale of solid, liquid and gaseous fuels and related products
49.50.10	Transport of gases via pipeline

Table 2: List of TARIC codes of energy-generating activities

TARIC Code	Product
Coal	Briquettes, ovoids and similar solid fuels manufactured from coal
2702	Lignite, whether or not agglomerated, excluding jet
2704	Coke and semi-coke of coal, of lignite or of peat, whether or not agglomerated; retort carbon
2705	Coal gas, water gas, producer gas and similar gases, other than petroleum gases and other gaseous hydrocarbons
2706	Tar distilled from coal, from lignite or from peat, and other mineral tars, whether or not dehydrated or partially distilled, including reconstituted tars
2707	Oils and other products of the distillation of high temperature coal tar; similar products with more aromatic than non-aromatic constituents
2708	Pitch and pitch coke, obtained from coal tar or from other mineral tars
2709	Petroleum oils and oils obtained from bituminous minerals, crude
2710	Petroleum oils and oils obtained from bituminous minerals, other than crude; certain preparations and waste oils
2711	Petroleum gases and other gaseous hydrocarbons
27131	Petroleum coke, petroleum bitumen and other residues of petroleum oils or of oils obtained from bituminous minerals
27132	Petroleum coke, petroleum bitumen and other residues of petroleum oils or of oils obtained from bituminous minerals
2716	Electrical energy
2803	Carbon
280410	Hydrogen
290512	Propan
290513	Butan
3826	Biodiesel and mixtures thereof with a basis of vegetable oils, animal fats or other products
44011	Fuel wood, in logs, in billets, in twigs, in faggots or in similar forms
4402	Wood charcoal
3824999266	Other chemical products for use as fuel

This ignore in-house energy generation. However, we evaluate its prevalence in our sample by looking at transactions of equipment directly used to generate energy (see tables 3 and 4 for ATECO and TARIC codes, respectively), and find that they seem to be insignificant.

Table 3: List of ATECO codes of energy-generating equipment

Code	Sector
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
38.22.00	Energy recovery (from waste disposal)
42.22.00	Construction of utility projects for electricity and telecommunications
43.21.01	Installation of lighting and photovoltaic systems on buildings
43.22.01	Installation of geothermal systems
43.22.06	Installation of gas distribution systems

Table 4: List of TARIC codes of energy-generating equipment

Code	Sector
8511	Electrical ignition or starting equipment of a kind used for spark-ignition or compression-ignition internal combustion engines (for example, ignition magnetos, magneto-dynamos, ignition coils, sparking plugs and glow plugs, starter motors); generators (for example, dynamos, alternators) and cut-outs of a kind used in conjunction with such engines
8402	Steam or other vapor generating boilers (for example, central heating boilers, super-heated water boilers)
85017	Photovoltaic DC generators
85018	Photovoltaic AC generators
7311	Containers for compressed or liquefied gas, or iron or steel

B.2 Mapping the data to the model

We take the model to the manufacturing data by defining each product ω as a six-digit code.[106] We assume that each firm's activity code corresponds to the *single* product it sells (see Appendix B.2.1 for a discussion on how we deal with multi-product firms).[107] For a few firms, our data only reports a more aggregated code.[108] 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 6−digit code within their category, based on the transaction shares of those 6−digit codes within these broader categories.

To construct the production tree, we proceed in the following steps. First, 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. This set includes primary material, but we exclude all energy inputs. Indeed, energy expenditures are already accounted for separately when we

[106] As emphasized by [Fort, 2023], manufacturing production is constituted of three distinct stages: 1) product design and innovation, 2) a sequence of physical transformation steps and 3) sales, marketing and distribution. We solely focus on the second stage, as it is the most relevant in terms of energy use.

[107] Firms are characterized by their six-digit activity code in the Italian ATECO classification. The ATECO classification is based on the NACE classification, and similarly organizes economic activity into 615 detailed four-digit industries, further subdivided into 1241 six-digit subcategories of product. This level of disaggregation is finer than other firm-to-firm transaction datasets that are usually at the 4−digit level [Huneeus et al.,]. 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 20.59.11 ("Manufacture of matches"), 20.59.12 ("Manufacture of explosives"), 20.59.20 ("Manufacture of glues"), 20.59.30 ("Manufacture of essential oils"), 20.59.91 ("Manufacture of e-liquids used in vaping devices") and finally 20.59.99 ("Manufacture of all other various chemical products n.e.c.").

[108] This seems to be due to misreporting, and appears to be random.

build firms' cost. Following [Boehm and Oberfield, 2023], we also select products classified as "residual" ("not elsewhere classified", "others" etc), or whose definition is very broad.[109] Firms producing these products likely have different input mixes because they produce different goods, rather than because they adopted different organizational forms.
Next, we don't have a firm identifier for foreign firms.[110] However, they remain potential suppliers. We aggregate foreign firms into one country Seller per input×country. Since we observe prices in the custom data, the unit cost of this national seller[111] is computed as the average price across the buyers of that input, weighted by their importance in the sales of the Seller. We similarly build their energy intensity using Exiobase.[112]

The model assumes that the recipe (the input-output linkages) for each product is a directed acyclical graph, i.e. a product is not used as an input in its own recipe. We therefore manually clean the data to eliminate cycles:

1. 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.
2. Following [Carvalho, 2014] and [Rachapalli, 2021], we only consider transactions that compose more than 1% of the total input expenditures of a firm, in order to avoid picking up unrelated products due to misreporting.

We then construct the recipes as a directed acyclical graph (DAG) of input-output linkages between products.

In the production network, we also model firms that are not present in the balanced-sheet data (i.e., for which we don't observe labor expenditures and total sales) as a terminal node.

Finally, the model assumes that each firm has only one supplier per input.[113]If we observe multiple suppliers per input for one firm, we keep the supplier with the highest transaction value.

[109]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 additive manufacturing machinery n.e.c., Manufacture of all other various special-purpose machinery n.e.c., Manufacture of other parts and accessories for motor vehicles, Manufacture of other transport equipment n.e.c., and Manufacture of other various articles n.e.c.

[110]Non-EU foreign suppliers don't have a firm identifier in our data. Therefore, we cannot retrieve any information on their production. EU foreign suppliers have a unique firm identifier. In the future, we will combine data from Orbis IP and EUTL to characterize their energy efficiency.

[111]In our model, the heterogeneity in effective prices that buyers face only comes from q_{ij}.

[112]Exiobase 3 is a globally consistent, environmentally extended multi-regional input–output (MRIO) database covering 44 individual countries and a residual "rest of the world" aggregate. The construction of the database is documented in [Stadler et al., 2018], and has been widely applied in the environmental literature (e.g., [Shapiro, 2021]). Exiobase reports input–output tables for 200 product groups (163 sectors) annually from 1995 to 2019, and augments these with detailed environmental extensions, including water and energy use, pollutant release, and carbon dioxide emissions per unit of output. Our analysis exploits energy use to measure the energy intensity of our World Seller. To match Exiobase with customs data, we employ the official crosswalk to CPA 2002 product codes, which enables us to map each customs observation to its corresponding aggregate category in Exiobase.

[113]The literature has consistently shown that firms are largely undiversified, that is, they have very few suppliers - usually only one - for each of their input (e.g., [Antras et al., 2017, Alfaro et al., 2025]). Diversification has even declined over years. In particular, [Alfaro et al., 2025]

B.2.1 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. We then combine these sets to create the firms' total in-house set which allows us to accurately identify their energy-biased technology ν_j. Indeed, to be consistent with our model, we assume that the energy-biased technology 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, we pick the input that has the highest within-supplier export share.[114] We also carry a robustness exercise in which we assign first the 6-digit code of the supplier, if the latter enters the buyer's produce recipe. Finally, in the counterfactual we define multi-product firms as potential suppliers of all their products (for products with a within-firm export share larger than 10%), but have to assume that all products are sold at the same price, which corresponds to the average unit cost of the firm.

Since the firm-to-firm transaction data assigns only one product per firm, we take several measures to identify multi-product firms. First, we use export data to evaluate the prevalence of multi-product firms. We define multi-product firms as firms with product export shares (out of their total exports) of several distinct 6−digit products above 10%. This implicitly assumes that all firms that only sell domestically are single-product firms.[115]

We mitigate this assumption with a robustness exercise that flag firms based on their *likelihood* to produce multiple products. In particular, we infer the likelihood that the a firm produces multiple products the recipes built with [Fetzer et al., 2024]'s product×input dataset.[116] In practice, we evaluate the share of distinct inputs (weighted by their transaction shares)[117] purchased by a firm that enters 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 is, we exclude them when building

find that firms drop their supplier before switching to a new one, instead of diversifying, when hit by to a tariff shock.

[114] The implications 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 will biased the estimation of the energy-biased technology of the buyer, since we control for the in-house input set. However, this control is built by weighting each in-house input by its energy share β_ω, which we don't observe at a very granular level due to data limitations. Therefore, the bias, in our setting, will only arise if A and B strongly differ in their energy share, which is unlikely in our data.

[115] 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.

[116] Before building the recipe, we convert HS 6−digit code to ATECO 6−digit code using a crosswalk provided by the Bank of Italy.

[117] The weights give greater importance to economically significant input flows, ensuring that large, relevant purchases dominate the metric rather than small, incidental ones.

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 $\mathcal{R}_\omega$ denote the technological recipe of product ω, defined as the set of inputs required to produce ω either directly or indirectly. Intuitively, a firm's input bundle $\mathcal{O}_f$ should closely overlap with the recipe $\mathcal{R}_\omega$ 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 $\mathcal{R}_\omega$. We compute a weighted Jaccard index that accounts for the intensity of each input in the firm's purchases:

$$s^w_{f\omega} = \frac{\sum_{\hat{\omega} \in \mathcal{O}_f \cup \mathcal{R}_\omega} \min(v_{if}, v_{\hat{\omega}p})}{\sum_{\hat{\omega} \in \mathcal{O}_f \cup \mathcal{R}_\omega} \max(v_{if}, v_{ip})},$$

where $v_{\hat{\omega}f}$ denotes the share of total transaction value that firm f allocates to input $\hat{\omega}$, normalized so that $\sum_{i \in \mathcal{O}_f} v_{\hat{\omega}f} = 1$, and $v_{\hat{\omega}p}$ represents the relative importance of input $\hat{\omega}$ in the recipe of product p (set to 1 if unknown). This measures the proportion of common inputs relative to the total number of distinct inputs observed in either set, 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:

$$P(\omega \mid f) = \frac{s_{fp}}{\sum_{\omega'} s_{f\omega'}}$$

For each firm, additional to its reported product, we retain the top 3 (with robustness using up to 5) products with the highest posterior probabilities $P(\omega \mid f)$, 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.3 Measuring in-house production

The set of inputs produced in-house by a firm producing product ω is the intersection between the set of inputs that are upstream to product ω (i its recipe Ω_ω) and the set of products that are downstream the products that are purchased by the firm ($\cup_{\hat{\omega} \in \mathcal{O}_j} \mathcal{D}_{\hat{\omega}}$), or in a branch that is entirely produced in-house. Formally, this would be $\mathcal{I}_j = \Omega_\omega \cup \{\cup_{\hat{\omega} \in \mathcal{O}_j} \mathcal{D}_{\hat{\omega}}\}$. Alternatively, we could also use $\mathcal{I}_j = \Omega_\omega \backslash \mathcal{O}^e_j$, where $\mathcal{O}^e_j = \mathcal{O}_j \cup \left(\cup_{\hat{\omega} \in \mathcal{O}_j} \Omega_{\hat{\omega}}\right)$ is the set of all inputs outsourced by the firm either directly or indirectly. However, both strategies fail to fully capture the true $\mathcal{I}_j$ due to inputs being used in multiple supply chains. Figure 7 illustrates the issues. It provides an illustrative production process for some product ω. Stages indexed by h corresponds to inputs that are produced in-house, stages indexed by o corresponds to inputs that are directly outsourced that is, *observed in the data*, and stages indexed by oe are inputs that are indirectly outsourced as they enter the production of an outsourced product. $\mathcal{I}_j = \Omega_\omega \backslash \mathcal{O}^e_j$ will wrongly exclude $c(h)$ from $\mathcal{I}_j$ due to $c(oe)$, since input c will enter $\mathcal{O}^e_j$. $\mathcal{I}_j = \Omega_\omega \cup \{\cup_{\hat{\omega} \in \mathcal{O}_j} \mathcal{D}_{\hat{\omega}}\}$ will wrongly include $k(o)$ and $m(oe)$ in $I(j)$ due to $b(o)$. Therefore, we need to move stage by stage, recursively, to accurately identify all inputs produced in-house.

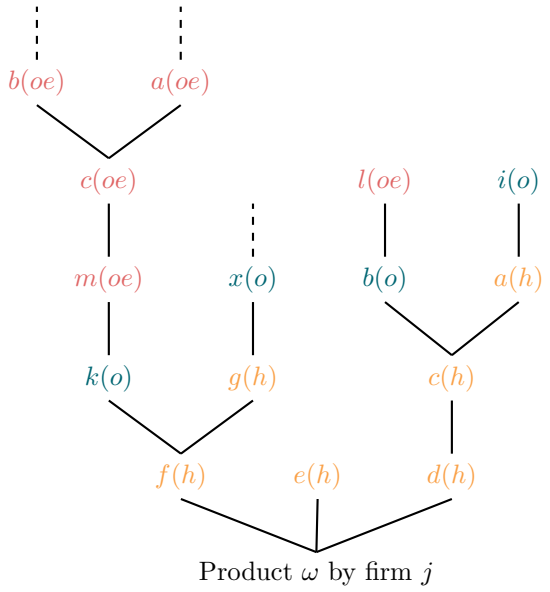


Figure 7: Illustration of the production process of product ω by firm j

B.3.1 Overlap of 6-digit Product Recipes within 4-digit codes

One key assumption behind our framework is that products share the same recipes. To evaluate the importance of the granularity level to describe the recipes, we assess whether the recipes of even more granular products within the same four-digit classification tend to overlap.[118] In particular, we use [Fetzer et al., 2024]'s production network to build the recipes for 5439 products classified at the HS 6−digit level. Using generative AI technology, the authors have created a network that connects 5439 unique products to their inputs in the production process. The network does not report transaction values, and simply report whether an input enters the production of a product. Therefore, our recipes algorithm breaks cycles based on the highest number of in-degree connections. Table 5 reports summary statistics on the size of the recipes across all 5439 products.

We then group the products at the HS 4−digit level[119] and evaluate how much the recipes overlap within each group, using the Jaccard similarity index between each product.[120] Figure 8 reports the distribution of the average Jaccard similarity index within the 4−digit group. Most 4−digit group have a very high similarity index, which confirms that aggregating products at the 4−digit level does not hide much variation in recipes.

[118] We have to focus on overlaps within four-digit codes, instead of within six-digit codes, because the most disaggregated HS product codes in [Fetzer et al., 2024]'s dataset are also at the six-digit level.

[119] The granularity level of the HS and the NACE 4−digit categories generally coincide. For instance, the 4−digit HS code 1006 corresponds to rice, while the 4−digit NACE code 0112 classifies "Growing of rice".

[120] The Jaccard similarity index between two recipes is the share of inputs that appear in both recipes, out of all the inputs contained in both recipes: $Jaccard^s = \frac{|A \cap B|}{|A \cup B|}$

Table 5: Summary Statistics on the HS 6-digit Product Recipes

Percentiles	**Value**	**Summary**	**Value**
1%	0	Observations	5439
5%	84		
10%	310	Min	0
25%	1058	Max	4118
50%	1893	Mean	1955.28
75%	3024	Standard deviation	1128.26
95%	3746	Skewness	−0.1341
99%	3746	Excess Kurtosis	−1.238

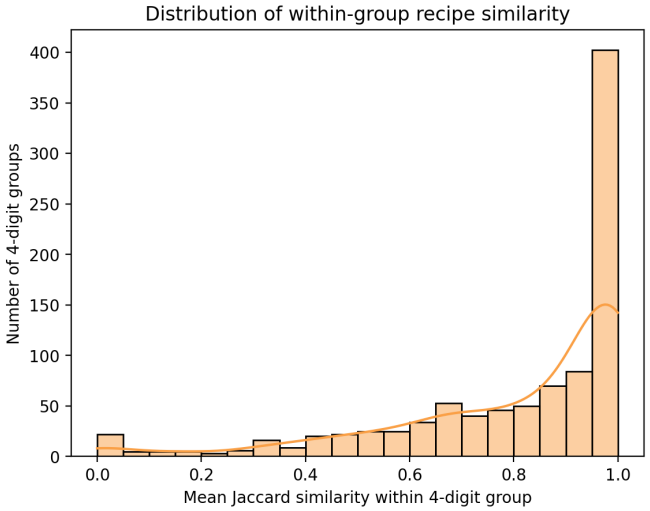


Figure 8: Distribution of within-4-digit codes recipe similarity

Notes: The graph plots the distribution of the within HS 4-digit cpde average Jaccard similarity index between the recipes of the HS 6-digit products. Recipes are built using the production network of [Fetzer et al., 2024].

B.4 Algorithm to calibrate the outsourcing capabilities

Our strategy relies on measuring $c_j^{\mathcal{I}}(\omega)$ and $c_i^{\mathcal{O}}(\omega)$ for all inputs, 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.

1. For each input ω and domestic firm j, we compute $c_j^{el}(\omega)$.
2. For each input ω and foreign seller f,[121] we use the (weighted) average import price across all buyers from that seller at the seller's unit cost $c_f(\omega)$, since the product price is reported in the custom data.

[121] Recall that our foreign seller are at the country×input level.

3. For each leaf-input ω^{leaf}, we use $c_i(\omega^{leaf}) = c_i^{el}(\omega)$ if supplier i is a domestic firm (and the unit cost $c_f(\omega^{leaf})$ from foreign prices if the supplier is a foreign firm).

4. We define the set $\mathcal{R}$ as the set of inputs for which we have identified all suppliers' c_i. We then compute the unit costs of each firm recursively. The last module is always produced in-house. Therefore, $c_j(\omega_j) = c_j^{\mathcal{I}}(\omega_j)$.

   (a) Using all the transactions of input $\omega \in \mathcal{R}$, we identify q_{ij} for all suppliers i of $\omega \in \mathcal{R}$ and buyer j of $\omega \in \mathcal{R}$ as

   $$log(\tilde{q}_{ij}) = \frac{1-\gamma}{-\gamma}\left(log\frac{p_e e_j + wl_j}{p_{ij}(\omega)x_{ij}(\omega)} + \frac{\gamma}{1-\gamma}log\Big(\frac{c_j^{el,tot}}{c_i(\omega)}\Big)\right) \tag{10}$$

   Where $c_j^{el,tot} = \int_{\omega\in\mathcal{I}_j} c_j^{el}(\omega)^{\frac{-\gamma}{1-\gamma}}d\omega$ and $\tilde{q}_{ij} = q_{ij}\psi_j$.[122] Note that for transactions with foreign firms, using the aggregate foreign seller unit cost (c_j^f) and a relationship-specific q_{ij} in the expression above is equivalent to directly using the firm-level foreign price (c_{ij}^f) that we also observe in the data, if we assume that $q_{ij} = \frac{c_j^f}{c_{ij}^f}$.

   (b) We move downstream to the firms that only purchase inputs $\omega \in \mathcal{R}$. The unit cost of producing ω in-house is:

   $$c_j(\omega) = c_j^{\mathcal{I}}(\omega) = \left(c_j^{el}(\omega)^{\frac{-\gamma}{(1-\gamma)}} + \int_{\hat{\omega}\in\Omega_\omega \& \hat{\omega}\in\mathcal{I}_j} c_j^{el}(\hat{\omega})^{\frac{-\gamma}{(1-\gamma)}}d\hat{\omega} + \int_{\hat{\omega}\in\Omega_\omega\subseteq\mathcal{R}\&\hat{\omega}\in\mathcal{O}_j}\left(\frac{c_{\hat{\omega}}}{q_{ij}}\right)^{\frac{-\gamma}{1-\gamma}}d\hat{\omega}\right)^{\frac{-(1-\gamma)}{\gamma}} \tag{11}$$

   Where i is the observed supplier of $\hat{\omega}$ to buyer j. We do this for all firms that only purchase inputs $\omega \in \mathcal{R}$. If this includes all firms producing input ω. We can then add ω to $\mathcal{R}$

   (c) We repeat steps 4a and 4b until we have identified all unit costs for all suppliers.

5. For the outsourcing capabilities of the transactions that include inputs that are produced in-house, we also require $c_j^{\mathcal{I}}(\omega)$ for all in-house inputs, and not only for the final product of firm. Thus, for all $\omega \in \mathcal{I}_j$, for all firms, we compute

   $$c_j^{\mathcal{I}}(\omega) = \left(c_j^{el}(\omega)^{\frac{-\gamma}{(1-\gamma)}} + \int_{\hat{\omega}\in\Omega_\omega \& \hat{\omega}\in\mathcal{I}_j} c_j^{el}(\hat{\omega})^{\frac{-\gamma}{(1-\gamma)}}d\hat{\omega} + \int_{\hat{\omega}\in\Omega_\omega\&\hat{\omega}\in\mathcal{O}_j}\left(\frac{c_{\hat{\omega}}}{q_{ij}}\right)^{\frac{-\gamma}{1-\gamma}}d\hat{\omega}\right)^{\frac{-(1-\gamma)}{\gamma}} \tag{12}$$

6. 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.

   (a) **Outsourced inputs:** For the potential supplier s of an input that is *outsourced* by firm j to supplier i,[123] the distribution of outsourcing capability is $\frac{1}{q_{sj}} \sim Pareto\left(\frac{c_i(\omega)/q_{ij}}{c_s(\omega)}, \theta\right)$.

   We characterize in this way the distributions of all buyer-supplier outsourcing capabilities for the suppliers of

[122] With a hicks-neutral productivity ψ_j, we cannot disentangle q_{ij} from ψ_j using expenditure shares. However, we can still build the suppliers' costs for each buyer firm, as q_{ij} and ψ_j enter in the same way. This allows us to build the cost of the core product of each firm, which we use to identify ψ_j using household sales. Once we've identified the ψ_j, we can retrieve q_{ij} net of the ψ_j. We do so before using them for our PPML regression, to reduce the importance of the supplier fixed effect.

[123] Notice that in that case, the counterfactual in-house cost used in the fragmentation decision is

$$c_j^{\mathcal{I}}(\omega)^{\frac{-\gamma}{1-\gamma}} = c_j^{el}(\omega)^{\frac{-\gamma}{1-\gamma}} + \int_{\hat{\omega}\in\Omega_\omega} \min\left\{c_j^{\mathcal{I}}(\omega), \min_{i\in\mathcal{S}_j(\hat{\omega})}\left(\frac{c_i(\omega)}{q_{ij}}\right)\right\}^{\frac{-\gamma}{1-\gamma}}$$

Based on the decision for all upstream inputs. We recover it in step 6c.

inputs that are outsourced by firm j, and draw their respective q_{sj} from these distributions.[124] We do the same for all buyers.

(b) **In-house inputs:** For the potential supplier s of an input $\hat{\omega}$ that is produced *in-house* by firm j, the distribution of outsourcing capability is $\frac{1}{q_{ij}} \sim Pareto\left(\frac{c_j^{\mathcal{I}}(\hat{\omega})}{c_i(\hat{\omega})}, \theta\right)$.
We characterize in this way the distributions of all buyer-supplier outsourcing capabilities for the suppliers of inputs that are produced in-house by firm j, and draw their respective q_{ij} from these distributions. We do the same for all buyers.

(c) **Unobserved fragmentation decision:** Finally, we need the q_{ij} for products produced further upstream the supply chain, for which we *don't observe* the fragmentation decision of firm j. Indeed, if firm j outsources input ω under p_e but produces it in-house under p_e^{SCC}, we need to know its fragmentation decision for $\hat{\omega} \in \Omega_\omega$ under p_e^{SCC}, hence need the q_{ij} for those pairs as well. As stated above, $\frac{1}{q_{ij}} = z_{ij} = e^{z_i^s + z_j^b + \mathbf{X}'_{ij}\delta} e^{\epsilon_{ij}}$ where $\mathbf{X}_{ij}$ controls for observables (e.g., distance), $\epsilon_{ij} \sim Exp(\theta)$ is iid across firms and inputs, and where z_i^s and z_j^b are respectively supplier and buyer fixed components. We then use the observed and drawn q_{ij}[125] to recover z_i^s, z_j^b as the buyer and supplier fixed effects in a Poisson pseudo Maximum Likelihood (PPML) estimation of $z_{ij} = exp(z_i^s + z_j^b + \epsilon_{ij})$.[126] Since we have drawn the q_{ij} for all potential suppliers, they constitute a representative sample of observations for each input, so we no longer have to account for selection. Formally, we estimate

$$\ell(\mathbf{z}^b, \mathbf{z}^s, \mathbf{X}_{ij}) = \sum_{i=1}^{N_b} \sum_{j=1}^{N_s} \Big\{ z_{ij} \times (z + z_i^s + z_j^b + \mathbf{X}'_{ij}\delta) - \exp(z + z_i^s + z_j^b + \mathbf{X}'_{ij}\delta) \Big\}$$

Where $z = log\frac{\theta}{\theta - 1}$ when we normalize the fixed effects to be centered (instead of using reference-category coding), that is imposing that $\sum_i z_i^s = 0$ and $\sum_j z_j^b = 0$.[127]
Since the PPML regression does not fully account for the potential heteroskedasticity of the error terms - likely to occur under our exponential distribution assumption - inference is based on an Eicker–Huber–White robust covariance matrix estimator as recommended by [Silva and Tenreyro, 2006], and cluster at the buyer-level. Once we have estimated all the z_i^s, z_j^b, we build the unobserved q_{ij} using them and drawing ϵ_{ij} randomly from the exponential distribution, for all inputs for which we don't observe the fragmentation decision, for all firms. We will use the standard error of the fixed effects to bound our measure of misallocation.

[124] The entire support of q_{sj} in the identified distribution is consistent with the observed allocation, so we cannot pin it down uniquely in our cross-sectional data. In effect, this means that some drawn q_{sj} will be higher than the true one, while some will be lower. However, these differences are random, hence independent from the characteristics of the firms. This means that they will even out for sufficiently many potential suppliers for each buyer, and that the resulting fragmentation decisions won't be systematically biased.

[125] In practice, we draw the q_{sj} of potential suppliers many times and run our regression for each draw, using the mean fixed effects identified in those different simulated sample as our final estimates.

[126] As it is well-known, PPML estimators are consistent as long as the conditional mean is correctly specified. The data do not have to actually be Poisson distributed, and it even allows for our outcome variable z_{ij} not to be an integer [Gourieroux et al., 1984]. The fixed effects estimated with our PPML estimator won't directly depend on our distributional assumption for the outsourcing capabilities q_{ij}, which only becomes relevant to build the q_{ij} themselves. PPML specifications are also recommended to deal consistently with heteroskedastic errors - likely to occur under our exponential distribution assumption - in log-linearized equations [Silva and Tenreyro, 2006].

[127] Note that this strategy still requires to externally calibrate θ beforehand, since we need to draw the q_{ij}.

B.5 Estimation in granular setting

The predictions of our quantitative model are sensitive to the baseline outsourcing capapbilities q_{ij} that we calibrate. [Dingel and Tintelnot, 2020] (hereafter DT) highlight a fundamental limitation in calibrating high-dimensional models with observations reflecting a discrete, idiosyncratic choice: because each agent reveals only a single choice among a vast set of alternatives, many of these alternatives end up with very few effective observations. In particular, DT show that calibration procedures that recover a bilateral idiosyncratic parameter for each alternative to exactly match observed outcomes perform very poorly in these high-dimensional settings. This reflects an overfitting problem,[128] which becomes particularly severe when the number of observations is small relative to the number of estimated parameters, leading to large measurement errors. A covariate-based approach that assumes away any unobserved idiosyncratic component and expresses the parameter as a function of pair-specific covariates yields better counterfactual predictions in this high-dimensional setting, except in scenarios in which the dispersion in the unobserved idiosyncratic component is particularly important. Ultimately, the respective performance of each approach depends on the size of the noises coming from the granular setting, relative to the importance of the unobserved components in driving the decisions.

Our hybrid method does not directly correspond to either of these approaches, but the setting is close enough for these finite sample considerations to be addressed. In step 4a of our calibration strategy, we recover the outsourcing capabilities q_{ij} for each observed transactions. This resembles DT's first procedure, in the sense that it matches each observed outcome by fitting an initially unobserved parameter q_{ij}. However, at this stage, we are not inferring any other outsourcing capabilities, meaning that we exactly use T observed transactions to retrieve T parameters. DT's procedure uses I observations, from I individual's choices, to fit N^2 parameters bilaterally connecting N cities. In granular settings, $I \ll N^2$, which leads to measurement errors. Steps 6a and 6b diverge from the usual applications considered by DT. On the one hand, by leveraging the properties of the Pareto distribution, we exactly match the *binary* observed decision on whether to outsource, and to which firm. On the other hand, we do not exactly match the effective unit cost, which we don't observe. Doing so would require building a cost distribution based on functional form assumption (e.g., using a DCM and transaction shares, or [Boehm and Oberfield, 2023]'s specifications), but as discussed in Appendix A.3, those tools do not suit our framework. This means that our strategy does not identify a unique set of parameters - although all possible calibrations will be data-consistent - but it also escapes the risk of high measurement errors induced by our finite number of observed decisions. In step 6c, we infer the outsourcing capabilities for all pairs for which we don't observe any decision. To do so, we retrieve a fixed effect for each N firm, that is $2N$ fixed effects.[129] We use both the calibrated observed-transactions q_{ij} from step 4a, and the drawn observed-decision q_{ij} from steps 6a and 6b. Since we have at least one q_{ij} per firm, this sample is larger than N. However, there might be some firms for which this sample contains only very few q_{ij}, resulting in very noisy fixed effects. This will be the case if some firms outsource their most direct inputs (or inputs very close to their products), have only few of those direct inputs, and only a handful of potential suppliers make those inputs. These cases are, fortunately, quite limited in our sample. Overall, we see that despite the resemblance with the methods discussed by DT, the concerns don't exactly apply to our setting, since we are not fitting each idiosyncratic

[128] A more flexible parameterization enhances the in-sample fit but deteriorates the model's predictive accuracy out of sample.

[129] We actually estimate a slightly lower number, since we cannot estimate a buyer (supplier) fixed effect for the firms that never buy (sell).

component q_{ij} exactly in steps 6a and 6b. Moreover, our last step (6c) differs from the usual applications, since no decision is observed for these pairs. Our novel strategy still allows to both consistently calibrate the q_{ij} to match *all* the observed binary decisions, while recovering the fixed effects in a way that avoids any selection and overfitting issue.

A different but related concern arising in our granular setting is the importance of idiosyncracies for aggregate outcome [Gabaix, 2011]. Our Pareto distribution assumption can generate very large differences in outsourcing capabilities across suppliers. As discussed in section 5, this feature has counteracting implications on our aggregate measure of misallocation. On the one hand, it reduces the differences in decisions across the equilibrium and the counterfactual scenario, since very large energy-biased technology differences would be required to compensate the large differences in relative outsourcing capabilities. This downplays the measured misallocation. On the other hand, if the equilibrium and the counterfactual decisions do differ, the loss in q_{ij} will be large, negatively affecting total production, which amplifies the overall differences in aggregate energy intensity across the two scenarios. Overall, the direction of the bias is unclear. To mitigate the impact of this specification, we estimate the fixed effects several times with different draws of q_{ij} and take the average across the resulting estimates. We do the same for our resulting aggregate measures. It is worth noting that DT show that the two calibration procedures they consider yield similar *aggregate* results in a model with a finite number of individuals, despite delivering very different allocations across alternatives. In other words, aggregate outcomes are not too sensitive to individual idiosyncracies.

In addition to parameter uncertainty, a finite population gives rise to a distribution of equilibrium outcomes that reflects the distribution of individual idiosyncracies. In their application, DT find that the variation around the counterfactual *individual-level* choices arising from idiosyncratic heterogeneity in a finite population exceeds the variation in counterfactual predictions attributable to parameter uncertainty in the continuum model. Again, their counterfactual aggregate measures are not sensitive to idiosyncracies, which is reassuring as our analysis is mainly interested in quantifying macroeconomic outcomes. Nonetheless, we do caution that our firm-level quantitative results should be understood as applying specifically to the context of the Italian manufacturing sector, for the given distribution of firms - and their associated idiosyncratic characteristics - active in 2019.

B.6 Calibrating the elasticity of substitution between intermediate inputs

We will confirm the robustness of our externally calibrated value for the elasticity of substitution between intermediate inputs $\frac{1}{1-\gamma}$, by estimating it using our detailed custom data. Custom data contains both prices and quantities of inputs purchased at the monthly level. Therefore, assuming that the firm boundaries don't change at such frequency, we have that

$$\Delta log \frac{c_j(\omega)x_{ij}(\omega)}{c_j(\hat{\omega})x_{sj}(\hat{\omega})} = \frac{-\gamma}{1-\gamma}\Delta log \frac{c_j(\omega)}{c_j(\hat{\omega})} + \frac{\gamma}{1-\gamma}\underbrace{\Delta log \frac{q_{ij}}{q_{sj}}}_{=0}$$

Where $\frac{c_j(\omega)}{c_j(\hat{\omega})}$ is the relative price imported input ω and $\hat{\omega}$, and $\frac{q_{ij}}{q_{sj}}$ is the ratio of buyer's j 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 monthly variation in markups (that our model assumes away) and in relative outsourcing capabilities

$\frac{q_{ij}}{q_{sj}}$. We run this regression on all intermediate input imported, excluding energy imports, over the 12 months available in 2019. Assuming that all intermediate inputs face the same elasticity of substitution, we can compare the estimated γ to our externally calibrated value of $\gamma = -4$.