

Cleaning Up or Passing the Buck? Environmental Spillovers in Production Networks*

Ezequiel Garcia-Lembergman Leticia Juarez
Ignacio Marra de Artiñano Christian Volpe Martincus

September 23, 2026

Abstract

Most environmental regulations target individual firms, even though production spans networks of upstream suppliers. Do clean firms induce cleaner production upstream, or shift pollution onto their suppliers? We develop a model of endogenous production networks and environmental upgrading in which suppliers tailor the energy intensity of intermediates to buyer demand. Two opposing forces shape the response of suppliers: outsourcing of dirty activities and transmission of clean technology. Their relative strength is governed by a single parameter, which we can estimate directly. Using novel data linking the universe of firm-to-firm transactions in Uruguay to official energy records, we find that suppliers increase their energy intensity by about 5 percent after connecting to an ISO-certified clean buyer, implying that outsourcing dominates. This pattern holds when buyers are defined as clean by ISO certification, by low energy intensity, or by multinational status. Calibrating the model to Uruguay, we find a supply-chain leakage rate of 44 percent. Extending the scope of environmental regulation to suppliers of regulated firms doubles the resulting welfare gains.

JEL Codes: F14, F18, F23, Q56, L14

Keywords: Environmental standards, supply chains, energy intensity, ISO certification, multinational enterprises, production networks

*García-Lembergman, Pontificia Universidad Católica de Chile (eglembergman@berkeley.edu); Juarez, IADB (leticiaaj@umich.edu); Marra de Artinano, ECARES (ignacio.marra.de.artinano@ulb.be); Volpe-Martincus, IADB and CESifo (christianv@iadb.org). The views and interpretations expressed in this paper are strictly those of the authors and should not be attributed to the organizations with which the authors are or have been affiliated, their Executive Boards, management, or their member countries. Other usual disclaimers also apply. We thank Francisco Temossi for excellent research assistance.

1 Introduction

Environmental sustainability has become central to industrial and trade policy. Governments have introduced carbon taxes and other measures that raise the price of energy, and firms have increasingly adopted voluntary environmental standards: ISO certifications alone now cover nearly three million firms in over 170 countries.¹ Multinational enterprises have gone further, adopting corporate sustainability programs and public environmental commitments. Yet most of these policies apply only to the regulated firm, not to its suppliers.² Since a large share of a firm’s environmental footprint originates upstream, the effectiveness of these policies depends critically on how their effects travel through supply chains.

Do the environmental gains of regulated firms extend beyond their boundaries? When a clean buyer enters a supplier’s network, does the supplier become cleaner, or does it absorb the dirty activities the buyer no longer performs? Two competing channels are at work. On the one hand, cleaner buyers may transmit clean technology, knowledge, or contractual requirements to their suppliers, inducing environmental upgrading along the supply chain. On the other hand, the very process by which a firm reduces its own emissions may involve outsourcing pollution-intensive production stages upstream. Despite the importance of this question for understanding the aggregate effectiveness of environmental policies, two obstacles have prevented researchers from addressing it: the absence of a quantitative model that connects firm-level regulation to supply-chain emissions through buyer–supplier linkages, and the lack of data combining firm-to-firm networks with energy variables.

Using a unique database of firms in Uruguay, we study the effects of connecting to a clean buyer on the environmental performance of suppliers. We develop a quantitative model of firm-to-firm networks in which regulated buyers optimally outsource production to suppliers who can tailor the energy intensity of their intermediates to each buyer’s demand. The model nests two opposing channels—cost-driven outsourcing of dirty tasks and green supply-chain management—through a single parameter whose sign is disciplined by the data. Using an event-study design on the universe of firm-to-firm transactions matched with energy consumption records, we find that suppliers increase their energy intensity by approximately 5 percent after connecting to a clean buyer, indicating that the outsourcing channel dominates. The effect is general for different definitions of cleaner buyers, appearing for ISO-certified firms and multinational buyers alike. Calibrating the model to Uruguay data, we estimate a supply-chain leakage rate of 44%, implying that firm-level environmental policies are partially offset by the reallocation of dirty activities upstream.

Our analysis proceeds in five steps. First, we develop a quantitative model of pro-

¹ International Organization for Standardization, *The ISO Survey of Management System Standard Certifications*, 2024.

² A recent exception is the European Union’s Carbon Border Adjustment Mechanism (CBAM), which begins to account for emissions embodied in supply chains.

duction networks and the environment. Heterogeneous buyers and suppliers form links endogenously, with each buyer producing using labor, in-house energy, and outsourced intermediates sourced from multiple suppliers. A key feature of the model is that suppliers tailor their production process for its product to each buyer’s specification: the energy intensity of the intermediate that the seller produces depends on what the buyer demands. If the buyer outsources its most energy-intensive activities orders a dirtier intermediate, while a buyer that transmits clean production practices orders a cleaner one. The supplier adjusts its input mix accordingly, tilting toward energy or labor depending on the specification.

The model nests both channels in a single structural parameter, the *demand-composition elasticity* ν . When a buyer faces higher in-house energy costs—from certification, regulation, or corporate sustainability commitments— ν determines whether its suppliers become dirtier ($\nu > 0$) or cleaner ($\nu < 0$). We derive a structural equation that maps ν directly to the reduced-form event-study coefficient, allowing the data to resolve the theoretical ambiguity.

Answering whether cleaner buyers green their suppliers or dirty them and estimating the structural parameter ν requires overcoming a central challenge: combining detailed firm-to-firm transaction data with reliable information on energy use, data typically absent from standard administrative sources such as VAT records. In the second step, We build a unique database for Uruguay that merges several administrative datasets covering the 2007–2018 period: the universe of firm-to-firm transactions (2013–2016), a firm-level panel with employment, sales, wages, and imports, official energy-use registries from the national energy directorate, and records of ISO 14001 and ISO 50001 certifications at the firm level. This combined dataset allows us to observe who trades with whom, how much energy each firm uses, and which buyers are environmentally regulated.

In the third step, we use the data to document a series of novel stylized facts about environmental performance along the production network that are in line with the model predictions. We find evidence of *negative assortative matching* in energy intensity: buyers with relatively low energy intensity are disproportionately linked to suppliers with relatively high energy intensity. In the light of the model, this is a first glance that clean buyers might outsource dirtier inputs to their suppliers.

In the fourth step, we implement our main event-study design to estimate the effects of starting to supply a clean buyer on suppliers energy intensity, which also helps to recover the key parameter ν of the model. Our main empirical strategy is an event-study design that tracks the evolution of supplier energy intensity before and after the first time a supplier starts selling to a clean buyer. This design exploits the staggered timing of new supply-chain relationships: different suppliers form their first connection with a clean buyer in different years, allowing us to compare treated suppliers to not-yet-treated and never-treated firms, while controlling for supplier and year fixed effects. In our preferred specification, we identify clean buyers as firms holding ISO 14001 or ISO 50001 certification. The event-study estimates reveal a sharp and persistent increase in supplier energy intensity of about 5 percent points following the connection. In terms of the model, the positive sign of the event-study coefficient implies $\hat{\nu} > 0$: the cost-driven

outsourcing of dirty tasks dominates whatever green supply-chain management certified buyers practice.

Although ISO certification provides a natural way to identify clean firms, a key question is whether the upstream pollution shift generalizes to other type of clean buyers. We explore this through a quantitatively important case: multinational firms. MNEs are substantially cleaner than domestic firms, yet no causal evidence exists on how they affect the emissions of their domestic suppliers—a significant gap, since supply-chain emissions account for roughly half of multinationals’ total environmental footprint (Steenbergen and Saurav, 2023). We find that the pattern extends to MNEs: suppliers that form a first connection with a multinational become dirtier. Importantly, the effect is not driven by foreign ownership itself: when we match MNE buyers to equally clean domestic firms, the differential effect vanishes, indicating that the mechanism is the buyer’s cleanliness, not its multinational status.

Furthermore, in a series of falsification tests we show that connecting to non-clean buyers have no effect on suppliers energy intensity. We show that connecting to a random buyer, a buyer at the median of the energy-intensity distribution, or a buyer with Quality ISO certifications produces no significant effect on suppliers energy intensity. This set of results is exactly what the model predicts: only cleaner buyers with high effective energy costs can generate the demand-composition shift that raises upstream energy intensity.

In the final step, we combine the reduced-form estimates with the structural model to recover $\hat{\nu}$ and quantify the aggregate consequences of environmental regulation along the supply chain. We compare the estimated economy with a counterfactual economy in which certification is unavailable. This exercise allows us to decompose the environmental effect of certification into certified buyers’ own energy reductions, the response of non-certified buyers, and the change in energy use by suppliers. We also use the model to evaluate the welfare consequences of certification and to identify which suppliers absorb the upstream increase in energy use.

We use the model to run two counterfactual exercises. The first quantifies the general equilibrium effects of ISO certification as currently deployed. Certification’s supply-chain leakage rate is 44%: for every unit of energy that certified buyers abate, 0.44 units reappear at their suppliers. Two channels generate this leakage. The first is an energy price channel: as certified buyers cut their energy demand, the energy price falls and unregulated firms expand. More importantly, the second is the task reallocation channel that originates due to the positive ν : suppliers additionally absorb the dirty tasks that certified buyers outsource. Both channels operate on suppliers, while non-certified firms outside the value chain are exposed only to the first, which is why the increase in suppliers’ energy use is twice as large.

The second asks what would be gained from a policy designed to also reach the suppliers of regulated firms. We evaluate a policy design that extends the standard to the supply chain, in which suppliers of certified buyers also face a compliance wedge on energy. The extension closes the margin through which abatement leaks. Aggregate abatement doubles and upstream leakage is close to zero, as suppliers clean up alongside

their buyers rather than inheriting their dirty tasks. Welfare gains nearly doubles. The effectiveness of firm-level environmental policy thus depends not only on how much regulated firms reduce their own emissions, but also on whether the policy reaches their supply chain.

Related Literature

First, we contribute to the literature on firm-to-firm production networks (Bernard et al., 2019; Eaton et al., 2022; Oberfield, 2018; Chaney, 2014; Arkolakis et al., 2025; Huneeus et al., 2021; Demir et al., 2024). This literature studied how productivity, geography, and wages shape supply-chain relationships, and how shocks propagate through firm-to-firm linkages. Yet the environmental dimension in networks remains unexplored: we know little about how energy costs or environmental regulation shape buyer-supplier linkages, or how pollution transmits along the supply chain. We fill this gap both theoretically and empirically. We propose a quantitative model of production networks in which environmental regulation can affect both which firms trade with each other and how much energy their suppliers use. We then provide the first evidence that forming a link with a clean buyer raises the energy intensity of the supplier.

Second, we contribute to a growing literature on trade and the environment. Existing work studies how trade openness, comparative advantage, and trade policy interact with environmental outcomes, operating primarily at the sector or across countries (Antweiler et al., 2001; Copeland and Taylor, 2004; Shapiro and Walker, 2018; Shapiro, 2021; Levinson and Taylor, 2008; Cherniwchan et al., 2017; Garcia-Lembergman et al., 2026b). A central theme in this literature is carbon leakage across jurisdictions: the concern that emissions migrate from regulated to unregulated countries (Fowlie and Reguant, 2018). Recent work traces how carbon pricing leaks across borders: Colmer et al. (2025) and Coster et al. (2024) show that firms subject to EU ETS regulation adjust their import sourcing toward inputs from unregulated countries. We complement this work by documenting leakage *within* a country, from clean buyers to their unregulated suppliers. We trace how environmental regulation propagates through individual buyer-supplier relationships and quantify the aggregate effects through a model of production networks with endogenous energy decisions.

Third, we also contribute to the literature on multinational firms and the environment. Existing work finds that MNE affiliates are cleaner than domestic plants (Eskeland and Harrison, 2003; Brucal et al., 2019; Garcia-Lembergman et al., 2026a) and examines whether environmental regulation affects MNE location decisions (Keller and Levinson, 2002; Hanna, 2010; Ben-David et al., 2021).³ However, we know little about how MNEs affect the energy intensity of their local suppliers. This is an important gap, since supply-chain emissions account for roughly half of multinationals' total environmental footprint (Steenbergen and Saurav, 2023), and the direction of the spillover is theoretically ambiguous. We provide the first estimates of how MNE buyers affect the energy intensity of their domestic suppliers.

Our MNE results also speak to a recent literature on responsible sourcing. Alfaro-

³ See Cole et al. (2017) and Garetto et al. (2025) for comprehensive reviews.

Ureña et al. (2022) study the labor-market effects of MNE responsible-sourcing codes in Costa Rica. We study a complementary margin—the environmental dimension—and examine the general population of MNE buyers, not only those with formal codes. For this broader population of MNE, the outsourcing channel dominates, suggesting that without binding environmental supply-chain requirements, clean multinationals shift dirty production upstream.

Fourth, we contribute to the literature on voluntary environmental standards (Potoski and Prakash, 2005; Arimura et al., 2008; King and Lenox, 2005). This literature mostly evaluates whether ISO 14001 certification reduces the certified firm’s own emissions. Recent empirical evidence confirms this direct effect; for instance, Wang and Xu (2025) estimate that ISO certification reduces a manufacturing firm’s own energy intensity by approximately 11%. We shift the focus from the certified firm to its supply chain: even if certification reduces the buyer’s own energy intensity, the relevant question for aggregate outcomes is how much of the reduction is amplified or attenuated by responses of firms upstream. Building on this perspective, our event-study design provides a direct answer to how emissions and production are reallocated across the network, and the structural model allows us to quantify the aggregate impact of ISO certifications.

2 Model

We develop a model of production networks in which downstream buyers produce final goods using intermediates sourced from upstream suppliers. Buyers differ in their environmental status: some face a higher effective cost of in-house energy—whether through voluntary certification, corporate mandates, or government regulation—which we call environmental upgrading. The model features a nested production technology in which upgraded buyers substitute away from in-house energy toward outsourced intermediates, and suppliers can tailor the energy intensity of their production to each buyer’s specifications. Two opposing forces shape how a buyer’s environmental commitments propagate upstream: a cost-driven outsourcing channel, through which upgraded buyers shift their dirtiest tasks onto suppliers, and a green supply-chain channel, through which they impose cleaner production standards. A single structural parameter, the demand-composition elasticity ν , governs the net direction. The model endogenizes the formation of buyer–supplier links and the upgrading decision itself, and delivers a structural equation that maps ν directly onto an estimable event-study coefficient.

2.1 Environment

Preferences and Social Welfare. A representative consumer has CES preferences over differentiated final goods:

$$U = \left(\int_{\omega \in \Omega} q(\omega)^{\frac{\sigma-1}{\sigma}} d\omega \right)^{\frac{\sigma}{\sigma-1}}, \quad (1)$$

where $\sigma > 1$ is the elasticity of substitution. Demand for variety ω is

$$q(\omega) = \frac{Y}{P} \left(\frac{p(\omega)}{\delta_\omega P} \right)^{-\sigma}, \quad P = \left[\int_{\omega \in \Omega} \left(\frac{p(\omega)}{\delta_\omega} \right)^{1-\sigma} d\omega \right]^{\frac{1}{1-\sigma}}, \quad (2)$$

where Y is aggregate income, P is the ideal price index, and $\delta_\omega \geq 1$ is a green demand premium for environmentally upgraded varieties (defined below), equal to 1 for varieties that do not upgrade. At a given posted price, an upgraded variety captures a larger share of demand.

Social welfare includes the environmental externality and is given by:

$$W = U - SCC \cdot E^{\text{total}},$$

where $SCC > 0$ is the social cost of carbon and E^{total} is aggregate energy consumption across buyers and suppliers. The last term is the environmental externality.

Firms. There is a fixed mass N_B of downstream final good producers, which we call *buyers*. There is a fixed mass N_S of upstream intermediate input producers, which we call *suppliers*. Buyer and seller productivities are drawn from Pareto distributions. We assume that firms do not internalize the environmental externality from energy use so the competitive equilibrium features excessive energy use.

Factors. There are two factors: labor (L) and energy (E). Labor is the numeraire ($w = 1$). Energy is supplied competitively with an upward-sloping supply curve:

$$E^S(p_E) = \bar{E} p_E^\epsilon, \quad (3)$$

where p_E is the energy price and $\epsilon > 0$ is the supply elasticity. Price of energy is an equilibrium result that depends on total supply and demand of energy.

Environmental upgrading. Each buyer makes a binary environmental-upgrading decision $\chi_b \in \{0, 1\}$. Upgrading ($\chi_b = 1$) raises the effective cost of the firm's own energy use to $\tau_b p_E$ —through compliance investments, monitoring, and process redesign—and earns its good the green demand premium δ_b in (2), where

$$\tau_b = \tau^{\chi_b}, \quad \delta_b = \delta^{\chi_b}, \quad (4)$$

so that $(\tau_b, \delta_b) = (\tau, \delta)$ for upgraded buyers and $(1, 1)$ otherwise, with $\tau > 1$ and $\delta > 1$ fixed parameters of the upgrading technology. Upgrading also carries a fixed cost f_b^{up} , in units of labor, heterogeneous across firms. The decision χ_b is endogenous; τ , δ , and the cost distribution are primitives, and the upgrading rate $\mu = \frac{1}{N_B} \sum_b \chi_b$ is an equilibrium outcome. We the optimal environmental upgrading decision in Section 2.4, once revenue and the supplier network are defined.

Environmental upgrading takes different forms across settings, all sharing the same

structure—a higher effective cost of direct energy use in exchange for a market benefit: (i) voluntary environmental standards such as ISO 14001/50001 certification, which require firms to monitor, audit, and reduce their energy use; (ii) multinational headquarters policies, where corporate-wide emissions targets, environmental reporting, and sustainability commitments raise the effective energy cost at affiliates; and (iii) government regulation, including carbon taxes, emissions trading, and energy efficiency mandates that raise the cost of direct energy use for covered firms.

The first is a firm choice; the latter two impose χ_b exogenously. Our formulation nests both: the predictions apply to any buyer that faces a higher effective energy cost than its suppliers, whether environmental upgrading is chosen or imposed. In our empirical analysis, we use ISO certification as the primary marker for environmental upgrading and extends the results to multinational firms, but the model's conclusions apply equally to policies such as carbon taxes.

2.2 Production Technologies

We now describe the production technologies of buyers and suppliers.

2.2.1 Buyer Technology

Each buyer b produces a differentiated final good using a nested production technology with three levels.

Upper nest: clean vs. dirty production (elasticity σ_b). At the top level, the buyer combines a clean input C_b that uses only labor and a dirtier composite D_b :

$$y_b = z_b \left[\gamma^{\frac{1}{\sigma_b}} C_b^{\frac{\sigma_b-1}{\sigma_b}} + (1-\gamma)^{\frac{1}{\sigma_b}} D_b^{\frac{\sigma_b-1}{\sigma_b}} \right]^{\frac{\sigma_b}{\sigma_b-1}}, \quad (5)$$

where z_b is buyer productivity, $\gamma \in (0, 1)$ is the weight on clean production, and $\sigma_b > 0$ is the elasticity of substitution between clean and dirty production. Clean production uses only labor: $C_b = \ell_b^C$, with unit cost w , which we set as the numeraire.

Middle nest: in-house energy vs. outsourced composite (elasticity η). Dirty production combines two inputs: in-house energy use e_b and an outsourced composite M_b assembled from intermediates purchased from multiple suppliers:

$$D_b = \left[\alpha^{\frac{1}{\eta}} e_b^{\frac{\eta-1}{\eta}} + (1-\alpha)^{\frac{1}{\eta}} M_b^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}, \quad (6)$$

where $\alpha \in (0, 1)$ is the weight on in-house energy and $\eta > 0$ is the elasticity of substitution between in-house energy use and the outsourced composite. When the buyer upgrades ($\tau_b > 1$), its effective in-house energy cost rises to $\tau_b p_E > p_E$. If $\eta > 1$, this raises the

relative cost of in-house energy and induces the buyer to substitute toward the outsourced composite—the core outsourcing channel.

Lower nest: outsourced composite (elasticity ζ). The outsourced composite M_b aggregates intermediates from all active suppliers:

$$M_b = \left[\sum_{s \in \mathcal{S}_b} m_{bs}^{\frac{\zeta-1}{\zeta}} \right]^{\frac{\zeta}{\zeta-1}}, \quad (7)$$

where $\mathcal{S}_b \subseteq \mathcal{M}_b$ is the endogenous set of active suppliers (determined in Section 2.3), m_{bs} is the quantity purchased from supplier s , and $\zeta > 1$ is the elasticity of substitution across suppliers.

Summary of nesting structure. Table 1 summarizes the three nests.

Table 1: Nesting Structure of the Buyer's Production Technology

Nest	Aggregates	Elasticity	Key parameter
Upper	Clean (C_b) vs. dirty (D_b)	σ_b	γ
Middle	In-house energy (e_b) vs. outsourced (M_b)	η	α
Lower	Across suppliers (m_{bs})	ζ	—

2.2.2 Supplier Technology

Each supplier s simultaneously produces intermediates for all buyers in its customer set $\mathcal{B}_s$. The intermediates are tailored to each buyer. For example, a supplier serving an upgraded buyer and a non-upgraded buyer might produce two distinct variants of its product, allocating different input mixes to each order. A contract manufacturer, for instance, might run an energy-intensive deep-forging line for one buyer and a lighter stamping process for another, using the same factory floor and workforce. The supplier's productivity $\tilde{a}_s$ is common across all orders; what varies is the specification demanded by each buyer, which can affect the energy mix.

Production function. When producing the intermediate for buyer b , supplier s uses labor ℓ_{sb} and energy ε_{sb} with a CES technology:

$$m_{sb} = \tilde{a}_s \left[\ell_{sb}^{\frac{\rho-1}{\rho}} + \psi_b^{\frac{1}{\rho}} \varepsilon_{sb}^{\frac{\rho-1}{\rho}} \right]^{\frac{\rho}{\rho-1}}, \quad (8)$$

where $\rho > 0$ is the elasticity of substitution between labor and energy, $\tilde{a}_s$ is the supplier's productivity.

We define ψ_b as the energy tailoring parameter. ψ_b is a buyer-specific parameter that governs the importance of energy to produce the product for the specific buyer. Given a

set of prices and wages, a higher ψ_b tilts the supplier's optimal input mix toward energy, raising the energy intensity of the intermediate.

The parameter ψ_b is determined by what the buyer demands. The environmental wedge τ_b can affect ψ_b through two opposing channels. On the one hand, an upgraded buyer faces higher in-house energy costs, creating an incentive to outsource its most energy-intensive activities. This pushes ψ_b up. On the other hand, the same buyer might also evaluate suppliers' environmental performance, require responsible sourcing policies, and communicate environmental requirements to their suppliers. This pushes ψ_b down.

To fix ideas, consider a food processing company that produces ready-to-eat meals using two activities: cooking (energy-intensive, requiring industrial ovens burning natural gas) and packaging (labor-intensive). After facing a higher energy cost—whether from obtaining ISO certification, complying with a corporate sustainability mandate from headquarters, or being subject to a carbon tax—its in-house energy cost rises to τp_E . Cooking, which accounts for most of the energy bill, becomes disproportionately expensive. The company outsources the cooking to a contract manufacturer operating at market energy prices, specifying high-temperature, fossil-fuel intensive processes (ψ_b high). The food company's measured energy intensity drops—it now mostly packages. The contract manufacturer's energy intensity rises—it inherited the dirtiest stage. Alternatively, the food company could require the manufacturer to install energy-efficient ovens (ψ_b low), greening its supply chain but at higher cost.

We parameterize the net effect of these opposing forces as:

$$\psi_b = \psi_0 \cdot \tau_b^\nu, \quad (9)$$

where $\psi_0 > 0$ is the baseline energy weight for non-upgraded buyers and ν is the *demand-composition elasticity*, which is the key structural parameter of the model. At $\tau_b = 1$ (not upgraded, $\chi_b = 0$), $\psi_b = \psi_0$ regardless of ν .

The sign of ν determines the direction of the composition shift: when $\nu > 0$, the outsourcing channel dominates—upgraded buyers demand dirtier intermediates and suppliers become more energy-intensive; when $\nu < 0$, the green supply-chain channel dominates and upgraded buyers impose cleaner production; and when $\nu = 0$, the two channels exactly offset, so upgrading changes the volume of outsourcing but not the composition of the outsourced bundle. Overall, the net effect of connecting to an upgraded buyer is ambiguous. Our event-study estimates which scenario prevails.

Supplier's unit cost. The unit cost of producing the intermediate with specification ψ_b is:

$$r_s(\psi_b) = \frac{1}{\tilde{a}_s} \left[w^{1-\rho} + \psi_b p_E^{1-\rho} \right]^{\frac{1}{1-\rho}}. \quad (10)$$

2.3 Network Formation

We model the formation of buyer–supplier links as the outcome of a random meeting process combined with an endogenous activation decision: a buyer trades with a supplier only if the cost savings from adding that supplier to its outsourced composite exceed the fixed cost of establishing the relationship.

2.3.1 Meeting Process

Each buyer b encounters potential suppliers through a Poisson process with arrival rate $\lambda > 0$, reflecting the intensity of the buyer’s search activity. Conditional on buyer b meeting supplier s , the pair draws a match-specific bilateral friction $\xi_{bs} \geq 1$ from a distribution H on $[1, \bar{\xi}]$, independently across pairs. The friction captures geographic distance, product incompatibility, contracting costs, and all other pair-specific obstacles to trade. The meeting set $\mathcal{M}_b$ is the set of suppliers that buyer b has encountered.

Delivered cost. The delivered cost of supplier s ’s intermediate to buyer b is:

$$\tilde{r}_{bs} = r_s(\psi_b) \cdot \xi_{bs}, \quad (11)$$

where $r_s(\psi_b)$ is the supplier’s unit cost defined in equation (10). Frictions are of the iceberg form: delivering one unit to the buyer requires the supplier to produce ξ_{bs} units, so suppliers earn zero rents and the friction is a real resource cost.

2.3.2 Link Formation

Each buyer–supplier link requires paying a fixed cost $f > 0$ paid in units of labor, representing contracting, coordination, and monitoring costs. Given its meeting set $\mathcal{M}_b$, buyer b activates the subset of met suppliers that maximizes profit net of these link costs.

Rank the met suppliers by delivered cost, $\tilde{r}_{b(1)} \leq \tilde{r}_{b(2)} \leq \dots \leq \tilde{r}_{b(n_b)}$, where $n_b = |\mathcal{M}_b|$. Because the outsourced composite is CES with $\zeta > 1$, a cheaper supplier always commands a larger expenditure share—and hence delivers a larger profit gain—than a more expensive one. The buyer would therefore never activate a supplier while leaving a cheaper met supplier idle, which gives the following characterization.

Proposition 1. *The optimal active supplier set consists of the K cheapest met suppliers for some $K \in \{0, 1, \dots, n_b\}$. The buyer’s problem reduces to*

$$K_b^* = \arg \max_{K \in \{0, 1, \dots, n_b\}} \left[\Pi_b(K) - Kf \right], \quad (12)$$

where $\Pi_b(K) = R_b(K)/\sigma$ is operating profit when the K cheapest suppliers are active and $R_b(K) \propto c_b(K)^{1-\sigma}$.

Proposition 1 collapses the problem from 2^{n_b} subsets to $n_b + 1$ nested candidates, with the prefix price indices $p_b^M(K)$, $p_b^D(K)$, $c_b(K)$ given by the closed forms in Section 2.5 (the prefix property and algorithm are in Appendix C.2).

It is important that the buyer cannot simply add suppliers one at a time and stop at the first unprofitable one. The marginal gain from activating supplier s is

$$\Delta\Pi_b(s \mid \mathcal{S}_b) = \frac{1}{\zeta - 1} \omega_{bs} X_b^{D,O}, \quad (13)$$

where ω_{bs} is s 's expenditure share (18) and $X_b^{D,O}$ is total outsourced dirty spending (23). But $X_b^{D,O}$ itself rises as more suppliers are activated: each additional supplier lowers the outsourced price index p_b^M , which raises outsourced spending and thus the marginal value of *every* other supplier. Link profitability is therefore supermodular in the active set—a block of suppliers can be jointly worthwhile even when none clears the threshold $\omega_{bs} X_b^{D,O} > (\zeta - 1)f$ in isolation—so the per-supplier threshold describes the marginal supplier at a given set but is not a valid stopping rule. The global maximization in (12) is required.

Endogenous network and environmental upgrading. The activation problem (12) makes the production network endogenous. When a buyer upgrades ($\chi_b = 1$, so $\tau_b = \tau$), total outsourced dirty spending $X_b^{D,O}$ rises, raising the marginal value of every potential supplier and hence the optimal K_b^* . Upgraded buyers therefore activate more suppliers than non-upgraded buyers.

2.4 The Upgrading Decision

Alongside its choice of suppliers, each buyer chooses whether to upgrade. Let $\Pi_b^{\text{op}}(\chi_b)$ be buyer b 's operating profit under upgrading status χ_b , evaluated at its optimal supplier set $\mathcal{S}_b(\chi_b)$ from Section 2.3 and net of link costs:

$$\Pi_b^{\text{op}}(\chi_b) = \frac{1}{\sigma} \delta_b^{\sigma-1} \left(\frac{p_b}{P} \right)^{1-\sigma} Y - |\mathcal{S}_b(\chi_b)| f, \quad p_b = \frac{\sigma}{\sigma-1} c_b(\chi_b), \quad (14)$$

where the unit cost $c_b(\chi_b)$ and revenue follow the closed forms derived in Section 2.5, and $(\tau_b, \delta_b) = (\tau, \delta)$ if $\chi_b = 1$, and $(1, 1)$ otherwise. The supplier set is re-optimized under each status: an upgraded buyer faces a higher in-house energy cost, outsources more, and generally activates a different set of suppliers. The gain from upgrading is the difference in operating profit across the two statuses,

$$\Delta\pi_b \equiv \Pi_b^{\text{op}}(1) - \Pi_b^{\text{op}}(0), \quad (15)$$

which trades off the green premium δ against the higher in-house energy cost and the additional link costs. Buyer b upgrades if and only if this gain exceeds its idiosyncratic

fixed cost:

$$\chi_b = \mathbf{1}\{\Delta\pi_b > f_b^{\text{up}}\}. \quad (16)$$

Because $\Delta\pi_b$ increases in revenue, larger and more productive buyers (high z_b) are the most likely to upgrade, consistent with the size gap between certified and matched uncertified firms in Table 14.

Equilibrium upgrading. The gain $\Delta\pi_b$ depends on the aggregates (P, Y, p_E) , which in turn depend on *which* buyers upgrade. Equilibrium upgrading is therefore a fixed point: an upgrading set $\mathcal{C}^* = \{b : \chi_b = 1\}$ and aggregates (P^*, Y^*, p_E^*) such that every buyer's decision satisfies (16) given the aggregates, and the aggregates are consistent with $\mathcal{C}^*$ and clear the energy market. We solve for this fixed point jointly with p_E in Section 2.7.

2.5 Equilibrium Prices and Shares

Given the active supplier set $\mathcal{S}_b$, the CES structure delivers closed-form expressions for all prices and expenditure shares.

Outsourced composite price index. The price index of the outsourced composite is:

$$p_b^M = \left[\sum_{s \in \mathcal{S}_b} \tilde{r}_{bs}^{1-\zeta} \right]^{\frac{1}{1-\zeta}}. \quad (17)$$

Expenditure share on supplier s . The share of outsourced spending allocated to supplier s is:

$$\omega_{bs} = \frac{\tilde{r}_{bs}^{1-\zeta}}{\sum_{s' \in \mathcal{S}_b} \tilde{r}_{bs'}^{1-\zeta}} = \left(\frac{\tilde{r}_{bs}}{p_b^M} \right)^{1-\zeta}. \quad (18)$$

Dirty composite price index. The price index of the dirty composite combines the in-house energy price and the outsourced composite price:

$$p_b^D = \left[\alpha (\tau_b p_E)^{1-\eta} + (1-\alpha) (p_b^M)^{1-\eta} \right]^{\frac{1}{1-\eta}}. \quad (19)$$

In-house share of dirty spending. The share of dirty spending allocated to in-house energy is:

$$s_b^{D,H} = \alpha \left(\frac{\tau_b p_E}{p_b^D} \right)^{1-\eta}, \quad s_b^{D,O} = 1 - s_b^{D,H}. \quad (20)$$

When τ_b rises, the effective price of in-house energy increases relative to the outsourced composite, and the in-house share falls: the buyer substitutes toward outsourcing.

Buyer's total unit cost. The buyer's total unit cost from the upper nest is:

$$c_b = \frac{1}{z_b} \left[\gamma w^{1-\sigma_b} + (1-\gamma) (p_b^D)^{1-\sigma_b} \right]^{\frac{1}{1-\sigma_b}}, \quad (21)$$

with the dirty share in total costs:

$$s_b^D = (1-\gamma) \left(\frac{p_b^D}{c_b \cdot z_b} \right)^{1-\sigma_b}. \quad (22)$$

Total outsourced dirty spending. Combining the upper and middle nests, the buyer's total outsourced dirty spending—which enters the marginal gain (13) and the upgrading decision (15)—is:

$$X_b^{D,O} = s_b^{D,O} \cdot s_b^D \cdot C_b, \quad C_b = \frac{\sigma-1}{\sigma} R_b, \quad (23)$$

where $C_b = c_b y_b$ is the buyer's total variable cost. Outsourced spending is a share of cost, not revenue: the markup $R_b - C_b = R_b/\sigma$ is profit and is not paid out to suppliers. This is the cost base used in the profit-gain derivation in Appendix C.1.

2.6 Energy Intensity

We now derive the energy intensity of buyers and suppliers. These expressions map directly to the reduced-form estimates in the empirical section.

2.6.1 Buyer's Energy Intensity

The buyer's energy expenditure (valued at market price p_E) relative to revenue is:

$$EI_b = \frac{s_b^D \cdot s_b^{D,H}}{\tau_b}. \quad (24)$$

Upgraded buyers have lower measured EI for two reinforcing reasons: (i) $s_b^{D,H}$ is lower because the wedge $\tau_b > 1$ raises the relative price of in-house energy, inducing substitution toward outsourcing, and (ii) the $1/\tau_b$ factor further compresses their measured energy-to-revenue ratio.

2.6.2 Supplier's Energy Intensity

The supplier's energy expenditure share when serving buyer b is:

$$\text{EI}_s(b) = \frac{\psi_b p_E^{1-\rho}}{w^{1-\rho} + \psi_b p_E^{1-\rho}}. \quad (25)$$

This is increasing in ψ_b : a dirtier specification raises the supplier's measured energy intensity.

Since supplier s serves multiple buyers with different specifications, its aggregate energy intensity is the ratio of total energy expenditure to total revenue:

$$\text{EI}_s = \sum_{b \in \mathcal{B}_s} \omega_{sb} \text{EI}_s(b), \quad (26)$$

where ω_{sb} is the share of supplier s 's total revenue coming from buyer b :

$$\omega_{sb} = \frac{X_{sb}}{\sum_{b' \in \mathcal{B}_s} X_{sb'}}, \quad X_{sb} = \omega_{bs} \cdot X_b^{D,O}. \quad (27)$$

Here X_{sb} is the total purchases of buyer b from supplier s and $X_b^{D,O}$ is buyer b 's total outsourced dirty spending defined in equation (23). Both ω_{sb} and X_{sb} are directly observable in the firm-to-firm transaction data.

Note that a supplier whose buyer portfolio shifts toward upgraded buyers will see its aggregate energy intensity change through two channels: an intensive margin, through which each unit produced for an upgraded buyer has a higher energy share $\text{EI}_s(b)$ when $\nu > 0$, because the buyer demands a dirtier specification (ψ_b); and an extensive margin, through which upgraded buyers outsource more (higher $X_b^{D,O}$) and may activate more suppliers, shifting the revenue weights ω_{sb} toward upgraded buyers.

When $\nu > 0$, both margins reinforce each other: the supplier produces dirtier intermediates *and* earns a larger share of its revenue from upgraded buyers. When $\nu < 0$, the two margins work in opposite directions: each unit is cleaner, but the supplier earns more from upgraded buyers, making the net effect on aggregate EI_s ambiguous.

2.7 Equilibrium

Definition 1. An equilibrium is a set of prices $\{p_E, \{p_b\}_{b=1}^{N_B}, \{r_s(\psi_b)\}_{s,b}\}$, allocations $\{\ell_b^C, e_b, \{m_{bs}, \ell_{sb}, \varepsilon_{sb}\}_{s \in \mathcal{S}_b}\}$ for each buyer b and supplier s , an upgrading profile $\{\chi_b\}_{b=1}^{N_B}$, and a network $\mathcal{S}^* \subseteq \{1, \dots, N_B\} \times \{1, \dots, N_S\}$, such that:

- (i) Consumers maximize utility (1) given prices $\{p_b\}$ and income $Y = wL + p_E E^S + \sum_b \pi_b$, where energy is a household-owned factor paying rent $p_E E^S$, and $\pi_b = R_b - c_b y_b - |\mathcal{S}_b|f -$

$\chi_b f_b^{up}$ is buyer b 's profit net of link and upgrading costs.

- (ii) Each buyer b chooses $\{\ell_b^C, e_b, \{m_{bs}\}_{s \in \mathcal{S}_b}\}$ to minimize cost given τ_b, p_E , and $\{\tilde{r}_{bs}\}_{s \in \mathcal{S}_b}$, and sets $p_b = \frac{\sigma}{\sigma-1} c_b$.
- (iii) Each supplier s chooses $\{\ell_{sb}, \varepsilon_{sb}\}$ to minimize the cost of producing m_{sb} given p_E, w , and ψ_b , for each $b \in \mathcal{B}_s$.
- (iv) $(b, s) \in \mathcal{S}^*$ if and only if s is among the K_b^* cheapest suppliers met by b , where K_b^* solves the activation problem (12).
- (v) Each buyer's upgrading decision is optimal: $\chi_b = \mathbf{1}\{\Delta\pi_b > f_b^{up}\}$, evaluated at the equilibrium aggregates (16).
- (vi) Markets clear:

$$\sum_{b=1}^{N_B} e_b + \sum_{s=1}^{N_S} \sum_{b \in \mathcal{B}_s} \varepsilon_{sb} = \bar{E} p_E^\epsilon \quad (\text{energy})$$

$$\sum_{b=1}^{N_B} \ell_b^C + \sum_{s=1}^{N_S} \sum_{b \in \mathcal{B}_s} \ell_{sb} + \sum_{b=1}^{N_B} (\tau_b - 1) p_E e_b + \sum_{b=1}^{N_B} |\mathcal{S}_b| f + \sum_{b=1}^{N_B} \chi_b f_b^{up} = L \quad (\text{labor})$$

Given the CES structure, the equilibrium is characterized by a single price. Buyer and supplier optimization yield cost functions that depend on p_E and w through the nesting equations (17)–(21). The network $\mathcal{S}^*$ is determined by the link formation condition (??) given p_E . With $w = 1$ as numeraire, the labor market clears by Walras' law, and the equilibrium reduces to finding the energy price p_E^* that satisfies the energy market clearing condition, jointly with the upgrading fixed point of Section 2.4.

2.8 Identifying the Demand-Composition Elasticity ν

The central object of interest is the demand-composition elasticity ν , which determines whether upgraded buyers make their suppliers dirtier ($\nu > 0$) or cleaner ($\nu < 0$). The model delivers a structural equation that maps ν directly to an estimable regression coefficient.

From the supplier's energy intensity equation (25) and the demand composition $\psi_b = \psi_0 \tau_b^\nu$, the change in log energy intensity experienced by supplier s upon connecting to an upgraded buyer is:

$$\Delta \ln \text{EI}_s \approx \omega_{sb} \cdot (1 - \text{EI}_s) \cdot \nu \ln \tau, \quad (28)$$

where ω_{sb} is the revenue share of the upgraded buyer in the supplier's total sales, $(1 - \text{EI}_s)$ is the elasticity of supplier energy intensity with respect to ψ_b , and τ is the environmental wedge.

This implies the following regression equation:

$$\ln \text{EI}_{st} = \underbrace{\ln \text{EI}_s(\psi_0, p_E, w)}_{\alpha_s} + \underbrace{f(p_{Et}, w_t)}_{\lambda_t} + \underbrace{(1 - \text{EI}_s) \cdot \nu \ln \tau \cdot \omega_{st}^{\text{clean}}}_{\beta_1} + \varepsilon_{st}, \quad (29)$$

where α_s is absorbed by a supplier fixed effect, λ_t by sector $\times$ year fixed effects, and $\omega_{st}^{\text{clean}}$ is the share of the supplier's revenue from upgraded buyers—which jumps from zero to a positive value when the link forms.

The event-study design in the next section estimates a binary version of this equation, yielding:

$$\hat{\beta}_1 = \bar{\omega}^{\text{clean}} \cdot (1 - \overline{\text{EI}}_s) \cdot \nu \ln \tau. \quad (30)$$

To recover ν from $\hat{\beta}_1$, we need one additional object: the environmental wedge τ . We identify τ from the observed gap in energy shares between propensity-score matched upgraded and non-upgraded buyers. The model implies $s_b^{D,H} = \alpha(\tau_b p_E / p_b^D)^{1-\eta}$, so:

$$\tau \approx \left(\frac{\bar{s}_{\text{clean}}^{D,H}}{\bar{s}_{\text{non}}^{D,H}} \right)^{\frac{1}{1-\eta}}. \quad (31)$$

With $\hat{\beta}_1$ and $\hat{\tau}$ in hand, we recover ν :

$$\hat{\nu} = \frac{\hat{\beta}_1}{\bar{\omega}^{\text{clean}} \cdot (1 - \overline{\text{EI}}_s) \cdot \ln \hat{\tau}}. \quad (32)$$

Since $\bar{\omega}^{\text{clean}} > 0$, $(1 - \overline{\text{EI}}_s) > 0$, and $\ln \tau > 0$, the sign of $\hat{\beta}_1$ directly reveals the sign of ν . The event study therefore provides a sharp test: $\hat{\beta}_1 > 0$ implies the outsourcing channel dominates; $\hat{\beta}_1 < 0$ implies the green supply-chain channel dominates. Combined with the structural equation, it also allows us to measure ν and conduct counterfactual analysis.

3 Empirical Analysis: Spillovers of Connecting to a Clean Buyer

The model delivers a sharp testable prediction: the sign of the demand-composition elasticity ν determines whether connecting to a clean buyer raises or lowers the energy intensity of the supplier, and Section 2.8 shows that ν maps directly onto the coefficient of an event-study regression. We now take this prediction to the data. We assemble a novel dataset for Uruguay that links the universe of firm-to-firm transactions with official energy-use registries and records of ISO 14001/50001 environmental certifications. Using

the staggered timing of new connections between suppliers and ISO-certified buyers, we estimate equation (30) in an event-study design that compares the energy intensity of newly connected suppliers against a matched control group of suppliers that never form such a link. Section 3.1 describes the data sources and the construction of the key variables; Section 3.2 discusses the identification strategy based on ISO certifications; Section 3.3 lays out the event-study specification; and Section 3.4 presents the main results. We then probe robustness through placebo tests and alternative definitions of clean buyers—including firms in the lowest within-sector decile of energy intensity and subsidiaries of multinational enterprises—in Sections 3.5 and 3.6.

3.1 Data

Our analysis combines four datasets. We start from the universe of firm-to-firm transactions in Uruguay, which we use to build the domestic production network, and link it to a firm-level panel with annual information on sales, employment, wages, and trade. We then merge official energy-use registries from Uruguay’s national energy directorate and, finally, records of ISO 14001 and ISO 50001 environmental certifications from the Instituto Uruguayo de Normas Técnicas (UNIT) and the International Accreditation Forum. We describe each source in turn before defining the analysis sample.

Firm-to-Firm Transaction Data. Our primary data source is the universe of domestic business-to-business (B2B) transactions recorded in Uruguay’s electronic invoicing system over 2013–2016. Each observation identifies a buyer–supplier pair at the annual level, along with the total value of transactions between them. The system captures all formal firms, providing a comprehensive view of the domestic production network. On average, the data cover 272,232 firms and 3,246,824 bilateral links per year. For each firm and year, we observe the full set of active buyers and suppliers and the value transacted with each, which lets us track the year-by-year evolution of every relationship and pin down the year in which it is formed or dissolved.

Firm-Level Economic Data. We complement the network with a firm-level panel spanning 2007–2018, built from corporate tax returns filed with the Dirección General Impositiva (DGI). The panel provides annual information on total sales, employment, and wages for the universe of formal firms, and we link it to transaction-level customs records that add each firm’s annual imports and exports. The wide time coverage is central to our design: the panel begins six years before the first year of the network data (2013), so we can track supplier outcomes well in advance of any new buyer–supplier link—essential for assessing parallel trends in our event study—and it extends two years past the last network year (2016), allowing us to capture medium-run post-treatment dynamics.

Energy Use. Uruguayan firms report their energy consumption to the national energy directorate (*Dirección Nacional de Energía*), and we use these records to construct our main outcome, *Energy Intensity* (EI). A naive measure—total energy expenditure over sales—would capture direct (scope 1) emissions but miss the indirect (scope 2) emissions

embedded in purchased electricity. To address this, we exploit the VAT data to weight utility expenditures by the carbon intensity of each firm’s electricity supplier. Specifically, for firm i in year t :

$$EI_{it} = \frac{\text{Fuel}_{it} + \text{Utilities}_{it} \times \lambda_{it}}{\text{Sales}_{it}}, \quad (33)$$

where Fuel_{it} is firm i ’s direct fossil fuel expenditure, Utilities_{it} is its electricity purchases, and λ_{it} is a firm-specific *fuel content factor*, defined as the ratio of fuel expenditure to total energy costs at firm i ’s electricity supplier in year t . This factor proxies for the carbon intensity of the generation mix the firm draws from, so firms sourcing from cleaner generators receive a correspondingly lower effective energy intensity than those supplied by fossil-fuel-intensive producers.

We observe energy intensity for 758,934 firm-year observations. Although this is a subset of the universe of formal firms, it captures the economically relevant core of the economy: our sample accounts for 100% of total sales, 99% of total wages, and 97% of total imports (Table 11). Firms with missing energy data are substantially smaller, with few employees and negligible sales; Appendix A.0.1 provides a detailed comparison.

ISO Certification Records. ISO 14001 (environmental management) and ISO 50001 (energy management) both require organizations to implement a formal management system that identifies environmental or energy impacts, sets quantitative objectives for improvement, monitors performance on an ongoing basis, and is subject to periodic management review. Both standards involve third-party auditing and periodic recertification, providing a transparent and externally verified signal of a firm’s environmental commitments. We obtain certification records from two complementary sources: the public directory of the Instituto Uruguayo de Normas Técnicas (UNIT), Uruguay’s national standardization body, cross-referenced with the IAF CertSearch global database, which aggregates verified data from accredited certification bodies worldwide. Combining the two lets us accurately date each firm’s initial certification and track its active status over time.

Institutional Setting of ISO Certification ISO 14001 and ISO 50001 are voluntary international standards for environmental and energy management. Rather than imposing static emission caps, these standards require organizations to implement a dynamic Environmental Management System (EMS) or Energy Management System (EnMS) based on the Plan-Do-Check-Act cycle. Under this framework, firms must formally identify their “Significant Environmental Aspects” (SEAs) or “Significant Energy Uses” (SEUs)—such as in-house electricity consumption or fossil fuel combustion—and establish measurable reduction targets.

To maintain certification, firms undergo a continuous auditing cycle conducted by accredited third-party bodies. This cycle includes a rigorous initial certification audit followed by mandatory annual on-site surveillance audits to verify continuous improvement. Auditors rigorously review documentation, utility bills, and emissions inventories against the firm’s stated objectives. Failure to adequately monitor significant aspects or achieve continuous improvement targets results in formal “Non-Conformities,” which, if left unresolved, can lead to the immediate suspension or withdrawal of the certificate.

Table 2: Summary Statistics: Firm Characteristics, Trade, and Energy Profile

Variable	Firms	Mean	Std. Dev.	p25	Median	p75
<i>Firm characteristics</i>						
Total sales (USD)	42,924	2,799,416	39,200,000	147,368	366,940	1,048,369
Employment (N)	42,924	19.97	143.01	1.00	5.00	14.00
<i>Energy profile</i>						
Energy intensity	42,924	0.0857	2.4838	0.0096	0.0217	0.0501
Fuel intensity	42,924	0.0425	0.8456	0.0000	0.0048	0.0200
<i>International trade</i>						
Total exports (USD)	3,178	4,431,290	20,300,000	11,378	77,974	775,150
Total imports (USD)	13,806	998,405	5,010,903	23,946	106,471	422,339
<i>Domestic network</i>						
Domestic buyers (N)	30,107	186.67	1,249.14	5.00	17.00	86.00
Domestic suppliers (N)	33,262	112.29	182.01	24.00	50.00	128.00

Notes: Sample of formal firms in Uruguay used in the empirical analysis. The *Firms* column reports the number of unique firms with a positive value of the corresponding variable in at least one observed year. Distributional statistics for firm characteristics and energy variables are computed across all firm-year observations for the 42,924 firms in the baseline sample. For trade and network variables, statistics are conditional on positive values in a given year. Monetary variables are expressed in USD. *Energy intensity* is defined as the ratio of total energy expenditures to total sales, adjusted for the carbon content of the firm's electricity supplier as described in Section 3.1. The sample also includes 655 firms identified as subsidiaries of multinational enterprises.

The sample comprises 42,924 unique formal firms observed over 2007–2018, spanning the full size and sectoral distribution of the Uruguayan economy. Three features of the data are particularly relevant for the empirical analysis that follows. First, energy intensity is substantially dispersed across firms: the median value is 2.2% of sales, the mean is 8.6%, and a long right tail of energy-intensive firms accounts for much of the cross-sectional spread. This dispersion provides the cross-firm variation in environmental performance that we use to define clean and dirty buyers and to trace the transmission of energy intensity along the supply chain. Second, firms are densely connected in the production network: the median firm reports 17 buyers and 50 suppliers, and the firm-to-firm panel records the formation and dissolution of these relationships at annual frequency, providing the variation underlying our event-study design. Third, the sample contains 655 subsidiaries of multinational enterprises, which we use as an alternative definition of clean buyers in Section 3.6. As is standard in firm-level data, the size distribution is highly skewed—the median firm has total sales of USD 367K and 5 employees, against means of USD 2.8M and 20 employees, respectively—and only a small share of firms is active in international trade (13,806 firms record positive imports and 3,178 record positive exports, with median annual values of USD 106K and USD 78K).

3.2 Empirical Strategy: ISO Certifications

Having derived the model’s structural equation and its mapping to the demand-composition elasticity ν , we now turn to the data. Our empirical strategy has two goals: first, to test and resolve the key theoretical tension: whether connecting to a clean buyer raises or lowers supplier energy intensity. Second, to estimate $\hat{\beta}_1$ to recover the demand composition parameter $\hat{\nu}$ to be able to conduct counterfactual analysis.

We use ISO 14001 and ISO 50001 certifications as a marker of buyer cleanliness and exploit the staggered timing of new buyer–supplier relationships in an event-study design. The first subsection introduces the certification data and characterizes the treated sample of suppliers; the second formalizes the event-study specification. We then present robustness analysis and show that conclusions of the paper go beyond ISO certifications and expand to connecting with other type of clean firms like multinationals and exporters.

We classify firms as “clean” based on their ISO 14001 (environmental management) and ISO 50001 (energy management) certifications. As discussed in Section 3.1, these standards require firms to implement and audit formal management systems for monitoring and reducing environmental and energy impacts; together, they provide a transparent, third-party-verified signal of a firm’s environmental commitments.

We compile a list of 255 unique certified firms by aggregating records from three sources: the Uruguayan Institute of Technical Standards (UNIT), the International Accreditation Forum (IAF), and the certification body LSQA. The distribution is heavily concentrated in ISO 14001 (228 firms), while 27 firms hold ISO 50001, with an overlap of 4 firms holding both. Matching these records against our firm-level data over 2007–2018, we identify 181 ISO-certified buyers in our sample, which corresponds to roughly 89% of the 203 active certifications officially reported by the IAF for Uruguay as of 2018. Of these, 159 satisfy the data requirements of our baseline sample (positive sales, employment, and energy intensity).

ISO-certified firms are concentrated in construction, manufacturing, and wholesale and retail—sectors with the highest environmental footprint and the largest scope for management-system gains. Table 3 reports descriptive statistics for ISO-certified buyers, non-certified buyers, and a matched control group of non-certified firms. In the unadjusted sample, ISO-certified buyers have a mean energy intensity of 0.027, roughly half the 0.051 average among non-certified buyers, and Figure 4(b) shows that their entire distribution of energy intensity lies to the left of the non-certified one. ISO-certified buyers are also substantially larger in sales and employment and maintain larger supplier networks. To account for these compositional differences, we construct a matched control group of non-certified firms using one-nearest-neighbor propensity-score matching on firm size and network degree. Even after matching, certified buyers remain measurably cleaner (0.027) than their non-certified matched counterparts (0.033), confirming that the ISO premium on energy intensity is not entirely driven by size or network structure. Appendix A.1 reports the energy-intensity distribution on the matched sample.

Table 3: Descriptive Statistics: ISO, Non-ISO, and Matched Buyers

Variable	Non-Certified (Full sample)			ISO-Certified (Treated group)			Matched Non-Certified (Control group)		
	Mean	Median	SD	Mean	Median	SD	Mean	Median	SD
Energy intensity	0.051	0.021	0.094	0.027	0.006	0.065	0.033	0.015	0.063
Total sales (USD)	2.47M	0.38M	29.9M	66.9M	11.4M	342.0M	39.4M	6.97M	126.0M
Employment	16.9	5.0	91.8	261.7	85.5	676.6	257.5	77.0	571.5
N of suppliers	44.5	29.0	56.1	230.5	176.0	207.9	217.1	155.0	211.6
Observations	29,714			159			151		

Notes: The table compares unadjusted non-certified buyers, ISO-certified buyers, and a matched control group of non-certified buyers. The matched control group is constructed using one-nearest-neighbor propensity-score matching (caliper 0.05) based on the log number of supplier links, log employment, and total sales. Energy intensity is the share of fuel and utility expenditures over total sales. Total sales are expressed in millions of USD.

Connecting to an ISO-certified buyer is an economically meaningful event for suppliers. During the network sample period 2013–2016, 13,446 domestic suppliers were linked with an ISO-certified buyer. Table 4 reports the sectoral distribution of both the certified buyers and the suppliers they connect to. Certified buyers are concentrated in construction (22.0%), manufacturing (20.8%), wholesale and retail (20.1%), and transport (17.6%), while the suppliers they source from span the full economy—the highest shares are in wholesale and retail (50.3%), manufacturing (16.7%), and transport (9.2%). The sectoral distribution of treated suppliers closely mirrors that of the broader economy, suggesting that the treatment we exploit is not confined to a particular industry.

Table 4: Sectoral Distribution of ISO-Certified Buyers and their Connected Suppliers

Sector	ISO buyers		Connected suppliers	
	Firms	(%)	Firms	(%)
Construction	35	22.01	643	4.8
Manufacturing	33	20.75	2,246	16.7
Wholesale and retail	32	20.13	6,767	50.3
Transport	28	17.61	1,237	9.2
Professional, scientific and admin. activities	9	5.66	1,084	8.1
Real estate	6	3.77	292	2.2
Primary activities	6	3.77	228	1.7
Energy and water supply	4	2.52	60	0.4
Information and communication	2	1.26	362	2.7
Public administration and other activities	2	1.26	387	2.9
Financial and insurance activities	2	1.26	140	1.0
Total	159	100.00	13,446	100.0

Notes: The table reports the cross-sectional sectoral distribution of the 159 ISO-certified buyers and the 13,446 suppliers that form at least one observed B2B linkage with an ISO-certified buyer during the sample period. Sectors are defined according to 2-digit industry codes: primary activities (1–9); manufacturing (10–34); energy and water supply (35–40); construction (41–43); wholesale and retail (45–47); transport (48–52); information and communication (58–63); financial and insurance activities (64–66); real estate (68); professional, scientific and administrative activities (69–82); and public administration and other activities (84–99).

These bilateral relationships are also economically meaningful for the suppliers. The ISO-certified buyer accounts on average for 16.0% of a supplier’s total B2B sales in the year of the first connection (SD 0.281) and remains a substantial revenue source thereafter, with the share stabilizing at an average of 14.9% across all connected years (SD 0.239). The magnitude of this revenue dependence suggests that suppliers face strong financial incentives to meet the operational and environmental requirements imposed by their ISO-certified clients, motivating our empirical focus on these specific network linkages.

3.3 Event-Study Design

We use the staggered timing of new buyer–supplier connections in an event-study design to identify how supplying an ISO-certified buyer affects a supplier’s energy intensity. The *treated* group consists of domestic suppliers that form their first link with an ISO-certified buyer between 2014 and 2016; the *control* group consists of suppliers that never form such a link. Because the network data begin in 2013, we drop connections formed that year, which we cannot identify as new. We merge this sample with the firm-level panel (2007–2018) and the energy-use registry, restricting attention to firms with strictly positive, non-missing values for total sales, employment, and energy intensity. In our preferred specification, a clean buyer is one that holds an ISO 14001 or ISO 50001 certification at any point during the sample period. In section 3.6 we consider alternative measures of

cleanliness.

Formally, we estimate the following event-study:

$$\ln(\text{EI}_{st}) = \alpha_s + \lambda_{mjt} + \sum_{k=C_-}^{C_+} \theta_k D_{st}^k + X'_{st}\beta + \varepsilon_{st}, \quad (34)$$

where $\ln(\text{EI}_{st})$ is log energy intensity for supplier s in year t , defined as the ratio of total energy expenditure to total revenue; α_s are supplier fixed effects, which absorb all time-invariant differences across firms (including sector, location, and baseline productivity); and λ_{mjt} are municipality $\times$ sector $\times$ year fixed effects, which control for time-varying sector- and location-specific shocks such as changes in demand conditions or industry-wide technological adoption. The vector X_{st} includes time-varying controls—lagged employment and the change in the supplier’s number of active connections—to account for the possibility that changes in firm scale or network size mechanically affect energy intensity.

Let t_s^0 denote the first year in which supplier s forms a link with a clean buyer. We define event-time indicators as

$$D_{st}^k \equiv \mathbb{1}\{t = t_s^0 + k\} \quad \text{for } k \in \{C_- + 1, \dots, -2, 0, 1, \dots, C_+ - 1\},$$

with endpoint bins that pool distant periods:

$$D_{st}^{C_-} \equiv \mathbb{1}\{t \leq t_s^0 + C_-\}, \quad D_{st}^{C_+} \equiv \mathbb{1}\{t \geq t_s^0 + C_+\}.$$

We normalize $\theta_{-1} = 0$, taking the year immediately prior to the connection as the reference period, and set $C_- = -6$ and $C_+ = 4$. The coefficients $\{\theta_k\}$ trace out the dynamic effect of connecting to a clean buyer. The pre-treatment coefficients $\theta_{-6}, \dots, \theta_{-2}$ provide a direct test of the parallel-trends assumption: if treated and control suppliers were on similar energy-intensity trajectories before the connection, these coefficients should be close to zero. The post-treatment coefficients $\theta_0, \theta_1, \dots, \theta_4$ measure the causal effect of the connection at each horizon and identify our parameter of interest, ν .

The panel covers 2007–2018, providing up to six years of pre-treatment data and up to four years of post-treatment data for each treated supplier. The identifying assumption is that, absent the connection to a clean buyer, treated suppliers would have followed the same energy-intensity trajectory as control suppliers, conditional on the fixed effects and controls. Two features of the design support this assumption. First, the timing of the first connection is driven primarily by the buyer’s side—when the buyer decides to source from a new supplier—rather than by changes in the supplier’s energy intensity. Second, the pre-treatment coefficients $\theta_{-6}, \dots, \theta_{-2}$ allow us to directly inspect whether treated and control suppliers were trending differently before the event.

3.4 Results

This section presents the main empirical findings. We use ISO 14001 (environmental management) and ISO 50001 (energy management) certifications to identify clean buyers and report event-study estimates of equation (34) for the sample of domestic suppliers described in Section 3.2.

Figure ??, Panel A, presents the event-study estimates from equation (3), where treatment is defined as a supplier’s first connection with an ISO-certified buyer. The dependent variable is the supplier’s log energy intensity. We additionally interact year fixed effects with the supplier’s pre-treatment energy intensity and log sales to account for differential trends based on initial environmental performance and firm size. This provides further support for the pollution-shifting mechanism while ruling out alternative explanations driven by suppliers’ baseline characteristics. Throughout, we report three estimators that are robust to staggered treatment timing: two-way fixed effects (TWFE), Callaway and Sant’Anna (2021) (CS), and Sun and Abraham (2021) (SA).

The figure highlights three main findings. First, the pre-treatment coefficients are economically small and statistically indistinguishable from zero, providing support for the parallel trends assumption. Second, energy intensity increases sharply upon forming a relationship with an ISO-certified buyer ($k = 0$) and continues to rise over subsequent years, reaching approximately 5–8 log points above the pre-treatment level by $k = 3$ to $k = 4$. This gradual increase is consistent with suppliers progressively adapting their production processes to meet the demand of ISO-certified buyers. Third, the three estimators yield nearly identical point estimates and confidence intervals, indicating that the results are not driven by heterogeneous treatment effects, negative weighting, or contamination from staggered treatment adoption.

In the context of the model, the positive post-treatment effect rejects the null hypothesis that $\nu \leq 0$ and suggests that ISO-certified buyers induce suppliers to produce more pollution-intensive intermediate inputs, increasing their energy intensity rather than improving supply-chain environmental performance through the diffusion of cleaner technologies or more responsible sourcing practices.

To ensure that our estimates isolate the causal effect of the connection and are not driven by selection on observables, we propose another specification that combines the event-study design with propensity-score matching. The potential concern for identification is that ISO-certified buyers may select suppliers based on observable characteristics that predict the future evolution of energy intensity. While the event-study design absorbs time-invariant differences across suppliers through fixed effects, certified buyers might preferentially link to suppliers whose energy intensity is expected to be on a particular trajectory—for example, suppliers specialized in production processes they anticipate becoming more (or less) energy-intensive over time.


We address this concern by matching each treated supplier to its nearest neighbor within the same 2-digit sector, based on pre-treatment log energy intensity, sales, log employment, and log wages, all measured in the year prior to the connection. The premise is

that, conditional on this strict set of observables, certified buyers are unable to predict future changes in the suppliers' trajectories. Appendix B.1 reports the balance diagnostics, confirming that treated and matched control suppliers are statistically indistinguishable on every pre-treatment characteristic.

Figure 1 Panel B reports the matched-sample event-study estimates of equation (34), where treatment is the formation of a first link with a buyer holding an ISO 14001 or ISO 50001 certification. The dependent variable is log energy intensity at the supplier level, and the omitted period is the year immediately before the connection ($k = -1$).

Three patterns stand out. First, the pre-treatment coefficients are statistically indistinguishable from zero and economically small, supporting the parallel-trends assumption underlying the event-study design. Second, energy intensity rises sharply upon connection ($k = 0$) and continues to increase over the following years, reaching roughly 6–10 log points above its pre-treatment level by $k = 3$ –4. The gradual buildup—rather than a one-time jump at the moment of the link—is consistent with suppliers progressively adapting their production processes to meet the requirements of certified clients, whether through scaled-up volumes, tighter delivery schedules, or product specifications that are inherently more energy-intensive to produce. Third, the three estimators yield nearly identical point estimates and confidence intervals, indicating that the result is robust to heterogeneous treatment effects, negative weighting, and contamination from staggered adoption.

Figure 1: Effect of Connecting to an ISO-Certified Buyer on Supplier Energy Intensity



Notes: Event-study coefficients $\hat{\theta}_k$ from equation (34). The dependent variable is $\ln(EI_{st})$, the log ratio of energy expenditure to total revenue, with $k = -1$ as the omitted baseline period. Panel (a) augments baseline specification by including interactions between year fixed effects and the supplier's pre-treatment energy intensity and log sales, isolating the treatment effect from natural efficiency trends and scale-driven dynamics. Panel (b) reports estimates on the propensity-score matched sample (nearest neighbor within the same 2-digit sector based on pre-treatment energy intensity, sales, log employment, and log wages). Three estimators are shown: two-way fixed effects (TWFE, red), Callaway and Sant'Anna (2021) (CS, green), and Sun and Abraham (2021) (SA, purple). Vertical bars indicate 95% confidence intervals based on standard errors clustered at the supplier level. Balance diagnostics for the matched sample are reported in Table 15.

In terms of magnitude, an increase of 6–10 log points implies that suppliers connected to an ISO-certified buyer become roughly 6–10% more energy-intensive than otherwise comparable, observationally equivalent suppliers within four years of the connection. The sign of the effect is informative on its own: connecting to a certified buyer pushes suppliers toward dirtier production rather than catalyzing the diffusion of cleaner technologies or responsible-sourcing practices. This pattern is consistent with the upstream pollution-shifting mechanism, in which certified buyers outsource the more energy-intensive stages of their production to non-certified suppliers. In light of the model presented in Section 2.8, this is evidence consistent with $\nu > 0$.

3.5 Beyond Certification: Multinationals and Upstream Pollution

The ISO results raise a natural question: is the upstream pollution shift specific to environmental certification, or does it reflect a more general pattern in which any clean buyer transmits higher energy intensity to its suppliers? If the mechanism is cost-driven outsourcing of dirty tasks—as the model predicts whenever $\nu > 0$ —then the effect should not depend on the particular source of the buyer’s cleanliness. Any firm facing a higher effective cost of in-house energy, whether through voluntary certification, corporate sustainability mandates, or simply operating with a cleaner technology, should generate the same upstream response. We explore this by extending the analysis to two economically distinct types of clean buyers.

The first extension uses a purely data-driven definition: we classify a buyer as clean if it falls in the lowest decile of energy intensity within its 2-digit sector, measured in the year prior to the connection. This captures firms that are objectively energy-efficient regardless of whether they hold any certification, and allows us to test whether the outsourcing channel operates for any low-EI buyer or is instead tied to something specific about the ISO audit process.


The second extension turns to multinational enterprises, connecting our findings to a long-standing debate in the trade-and-environment literature. MNEs are well documented to be cleaner than comparable domestic firms in developing countries (Eskeland and Harrison, 2003; Brucal et al., 2019; Garcia-Lembergman et al., 2026a)—in Uruguay they are on average 86% cleaner in terms of energy intensity—and supply-chain emissions account for roughly half of multinationals’ total environmental footprint (Steenbergen and Saurav, 2023). Yet whether this firm-level advantage propagates upstream remains an open question. On the one hand, MNEs may transmit environmental best practices to their suppliers through technology transfer, training, and contractual requirements—the green supply-chain channel. On the other hand, the same cost-driven logic that applies to ISO-certified firms applies here: an MNE whose headquarters imposes emissions targets on its affiliates faces an effective energy cost $\tau p_E > p_E$, creating incentives to outsource the most energy-intensive stages to domestic suppliers operating at market prices. Our framework provides a direct test. We identify MNE buyers as firms with foreign

ownership exceeding 10% of equity; the 2013–2016 network includes 1,206 such buyers.⁴


Figure 2 reports the event-study estimates for both definitions side by side. The pattern is qualitatively identical to the ISO baseline: flat pre-trends followed by a statistically significant and growing increase in supplier energy intensity after the connection. The outsourcing channel is not specific to certification—it operates for any sufficiently clean buyer, including multinationals.

Figure 2: Effect of Connecting to a Clean or Multinational Buyer on Supplier Energy Intensity

(a) Clean buyer (cleanest 10% within sector)



(b) Multinational buyer



Notes: Event-study coefficients $\hat{\theta}_k$ from equation (34). Panel (a) defines a clean buyer as a firm in the lowest decile of energy intensity within its 2-digit sector, measured in the year prior to the connection. Panel (b) defines a multinational buyer as a firm with foreign ownership exceeding 10% of equity. The dependent variable, omitted period, and confidence intervals are defined as in Figure 1.

These findings carry two broader implications. First, the upstream pollution shift is not an artifact of the ISO label or its audit process. It is consistent with any policy or norm that raises costs at the firm boundary but not upstream—carbon taxes, the EU emissions trading scheme, and corporate net-zero pledges all share the same asymmetric structure ($\tau > 1$ for the regulated firm, $\tau = 1$ for its suppliers). Some share of the intended abatement leaks upstream through the production network. Second, on the multinationals debate, our results complement [Alfaro-Ureña et al. \(2022\)](#), who document labor-market spillovers of responsible-sourcing codes in Costa Rica, by showing that on the environmental margin—where monitoring is harder and standards less codified—MNEs behave like other clean buyers, outsourcing rather than transmitting their environmental advantage. To disentangle the role of foreign ownership from cleanliness per se, in Appendix B.4 we match MNE buyers to equally clean domestic buyers and show that the differential MNE effect vanishes: the upstream shift is driven by the buyer's energy profile, not its multinational status. Appendix B.2 reports analogous results for the cleanest 5% and 15% within-sector thresholds.

⁴ No causal evidence exists on how MNEs affect their suppliers' energy intensity—in part because data combining firm-to-firm transaction records with energy-use registries have until recently been unavailable.

3.6 Robustness: Falsification Tests

We assess the robustness of our main finding—that connecting to a clean buyer raises supplier energy intensity—along two dimensions. We run two falsification tests that re-estimate the event study with placebo connections: to a generic new buyer and to a buyer certified only for quality management (ISO 9001) (Section 3.6).

A natural concern is that any new buyer connection could mechanically raise supplier energy intensity—if, for example, forming a new relationship requires scaling up production with rising returns on dirty inputs. Although all specifications already control for changes in the supplier’s number of active connections, we conduct two falsification tests to verify that the increase in energy intensity is specific to connections with clean buyers rather than a generic consequence of forming new business relationships.

Falsification Test: Random Buyers. To ensure that the observed upstream pollution shifting is driven specifically by the environmental stringency of the buyer, we conduct a falsification test. We restrict the pool of potential placebo buyers to domestic firms that are not in the lowest decile of sectoral energy intensity and neither ISO-certified, nor multinational corporations.


Aside from the treatment definition, the empirical design for this placebo—including the dependent variable, omitted period, controls, and fixed effects—is constructed identically to our baseline event-study model. If the increase in upstream energy intensity were merely a byproduct of scaling up production to satisfy a new, average client, we should observe similar post-treatment dynamics in this alternative group. Conversely, Figure 3(a) shows that the post-treatment coefficients for this placebo remain statistically indistinguishable from zero. This null effect contrasts sharply with the baseline ISO-certified trajectory, confirming that pollution shifting is a unique feature of connecting to environmentally stringent buyers.

ISO 9001 as a placebo certification. We test whether the result is driven by the certification process itself rather than its environmental content. ISO 9001 (quality management) shares the rigorous third-party auditing, documentation, and operational discipline of ISO 14001 and 50001 but, crucially, lacks any specific mandate to control or monitor environmental impacts within the firm or along the supply chain. If our main result were driven by a generic “certification effect”—larger, better-managed buyers imposing more demanding requirements on their suppliers—we should observe a similar upstream response when suppliers connect to ISO 9001-only firms.

We identify 403 ISO 9001 certificates in our data, of which 295 firms hold ISO 9001 exclusively (the strict placebo group used in the event study) and 108 also hold ISO 14001 or ISO 50001 and are excluded. For reference, the IAF reports 1,026 active ISO 9001 certifications in Uruguay for 2018, so our compiled sample covers roughly 40% of the universe. Panel (b) of Figure 3 reports the event-study estimates with the placebo group as treated buyers. The estimated post-treatment effect is essentially zero across all three estimators, with only a small and largely insignificant rise in the immediate year following the con-

nection. Given that ISO 9001 imposes comparable auditing and documentation burdens to ISO 14001/50001, this contrast confirms that what drives the upstream pollution shift is the environmental content of the certification, not the certification process or selection on firm size or managerial quality.

Figure 3: Falsification: Connecting to a Random Buyer or to an ISO 9001–Only Buyer



Notes: Event-study coefficients $\hat{\theta}_k$ from equation (34). Panel (a) defines treatment as forming a first connection with a buyer that is not in the lowest decile of sectoral energy intensity, is not ISO-certified, and is not a multinational (mean baseline energy intensity: 0.041; median: 0.025). Panel (b) defines treatment as forming a first connection with a buyer exclusively certified for ISO 9001, meaning not also certified for ISO 14001 or ISO 50001 (mean energy intensity: 0.032; median: 0.012); the treated group consists of 295 placebo buyers. For baseline comparison, the primary environmentally certified buyers (ISO 14001/50001) evaluated in Figure 1 exhibit a mean energy intensity of 0.027 (median: 0.007). The dependent variable, omitted period, and confidence intervals are defined as in Figure 1.

4 Quantitative analysis

This section takes the model of Section 2.8 to the Uruguayan data. We first describe the quantitative implementation of the model (Section 4.1); we then present the algorithm used to solve for equilibrium (Section 4.2); and we finally set out the parameterization and estimation of the structural parameters (Section 4.3). We use the estimated model to quantify the aggregate and distributional consequences of environmental upgrading in production networks.

4.1 Model Set-Up

To bring the model to the data, we specify the distributional assumptions, fix the calibration sample and the units in which the model is expressed, and assign firms to the two sides of the bipartite network.

Distributional assumptions. Buyer and supplier productivities are drawn from bounded Pareto distributions. Buyer productivity z_b has tail (shape) parameter θ_z and supplier productivity $\tilde{a}_s$ has tail parameter θ_a , each on a bounded support so that all revenue and cost moments are finite. We impose $\theta_z > \sigma - 1$, so that buyer revenue has a finite mean.

Each buyer meets potential suppliers through the Poisson process of Section 2.3: a buyer–supplier pair meets with probability λ . Conditional on a meeting, the pair draws a bilateral friction ξ_{bs} from a Pareto distribution with shape parameter θ_ξ , which enters the delivered cost $\tilde{r}_{bs} = r_s(\psi_b) \xi_{bs}$ in equation (11). Each active link costs f_b units of labor; we allow this cost to vary across buyers, with location f_{link} and dispersion $f_{\text{link},\sigma}$, nesting the homogeneous-cost case of Section 2.3.

Finally, the fixed cost of environmental upgrading f_b^{up} is drawn from a log-normal distribution with location f_μ and dispersion f_σ , so that buyers upgrade according to the threshold rule (16). Energy is supplied competitively along the upward-sloping schedule $E^S(p_E) = \bar{E} p_E^\epsilon$ of equation (3), with supply elasticity ϵ and scale $\bar{E}$.

Calibration sample, period, and units. We parameterize the model on the universe of energy-reporting active firms with annual revenue above \$10,000.⁵ Firm counts, the firm-size distribution, and certification status are measured in the 2016 cross-section. All financial variables are expressed in US dollars.

We normalize the wage to one ($w = 1$), so that labor is the numeraire, and set the labor endowment L equal to total employment in Uruguay in 2016 in the calibration sample, expressed in thousands of workers. Because observed energy expenditure is denominated in currency while model expenditure is in numeraire units, we calibrate the energy-supply scale $\bar{E}$ so that, in the baseline equilibrium, energy spending relative to the wage bill equals its data counterpart; thus preserving the real ratio of energy expenditure to labor compensation.

Assignment of firms to buyers and suppliers. The model is bipartite: buyers are downstream final-good producers and suppliers are upstream input producers. We assign firms to the two sides by their two-digit sector of activity. We define as buyers sectors those with the highest share of environmental upgrading as measured by the relative share of ISO 14001 and 90001 certificates. Buyers thus comprise manufacturing, construction, and transport; suppliers comprise all remaining sectors.⁶ This yields $N_B = 5,760$ buyers and $N_S = 16,812$ suppliers. The assignment fixes the masses N_B and N_S and hence the meeting rate λ consistent with the observed network density.

Parameter partition. The model’s parameters fall into three groups, which organize the estimation in Section 4.3. A first group is set *externally*, using estimates from the

⁵ The threshold removes inactive and shell entities that report negligible activity; the results are not sensitive to the specific revenue threshold

⁶ More precisely buyer sectors are thus manufacturing (CIIU 10–34), construction (41–43), and transport (48–52); supplier sectors are trade, energy and water, primary activities, professional services, finance, information and communication, real estate, and other sectors.

literature. A second group is recovered *directly from the data* in closed form, outside the estimation, from cost-share identities and from the reduced-form estimates of Section 3.4. The remaining parameter, which govern the certification-cost distribution, the network, and the productivity tails, are estimated jointly by *simulated method of moments*. The equilibrium objects that enter these moments are computed using the algorithm we describe next.

4.2 Solution Algorithm

Given the CES nested structure, the equilibrium of Section 2.7 reduces to a nested fixed point. By Walras' law, once the energy market clears the labor market clears as well, so the problem collapses to finding the energy price p_E and the upgrading profile $\{\chi_b\}$ consistent with optimal firm behavior. We solve it with the three nested layers summarized in Algorithm 1: an outer loop on energy prices p_E , a middle loop on the environmental upgrading fixed point, and an inner per-buyer supplier-set choice. The full solution algorithm, including the supplier-set subroutine and the convergence criteria, is explained in more detail in Appendix D.1.

Algorithm 1 Solving for equilibrium

- 1: Guess an energy price p_E
 - 2: **repeat** (*outer: clear the energy market*)
 - 3: **repeat** (*middle: find the upgrading equilibrium*)
 - 4: Taking (p_E, P, Y) as given, buyers choose supplier set and whether to upgrade
 - 5: Recompute the price index P and aggregate income Y
 - 6: **until** P , Y , and the set of upgrading buyers converge
 - 7: Compare aggregate energy demand with energy supply $\bar{E} p_E^\epsilon$ at p_E
 - 8: **until** the energy market clears.
 - 9: Report equilibrium prices, supplier network, upgrading profile, and welfare.
-

Innermost layer: supplier-set activation. Given the energy price p_E , the aggregates (P, Y) , and a buyer's upgrading status χ_b , each buyer chooses its active suppliers. By Proposition 1, the optimal set consists of the K_b cheapest met suppliers, so the buyer ranks its meeting set by delivered cost $\tilde{r}_{bs}$ and evaluates the $|\mathcal{M}_b| + 1$ nested prefixes, selecting the K_b^* that maximizes profit net of link costs in (12). Each candidate uses the closed-form prefix price indices and expenditure shares of Section 2.5, so the inner problem requires no further iteration. Because link profitability is supermodular in the active set, the buyer solves this global problem rather than adding suppliers until the first one is unprofitable.

Middle layer: the upgrading fixed point. At a given energy price, the upgrading profile is itself a fixed point. Holding the aggregates (P, Y) fixed, each buyer computes its optimal supplier set under both statuses, evaluates the upgrading gain $\Delta\pi_b$ in (15), and upgrades if and only if the gains are higher than the fixed costs from upgrading

$\Delta\pi_b > f_b^{\text{up}}$. The resulting upgrading set determines a new price index P and a new level of aggregate income Y ; we iterate on $(\{\chi_b\}, P, Y)$ until the upgrading profile and the aggregates are mutually consistent.

Outer layer: energy-market clearing. The outer loop searches for the energy price that clears the energy market, equating aggregate energy demand (buyers' in-house use e_b plus suppliers' use ε_{sb}) to supply $\bar{E} p_E^\epsilon$. We solve the excess-demand equation with a root-finder, re-solving the new upgrading fixed point at each trial price.

4.3 Parameterization

We assign the model's parameters in three ways. A first group is set externally, from estimates in the literature. A second group is recovered directly from Uruguayan data in closed form, outside the estimation, using cost-share identities and the reduced-form estimates of Section 3.4. The remaining parameters are estimated jointly by simulated method of moments.

The set of parameters. The model's parameters are the elasticities of substitution $\{\sigma, \sigma_b, \eta, \rho, \zeta\}$; the cost-share weights of the buyer nests $\{\gamma, \alpha\}$; the environmental-upgrading block $\{\tau, \nu, \delta, \psi_0, f_\mu, f_\sigma\}$; the network block $\{\lambda, f_{\text{link}}, f_{\text{link}, \sigma}, \theta_\xi\}$; the productivity tails $\{\theta_z, \theta_a\}$; the scale and energy parameters $\{L, \bar{E}, \epsilon\}$; the social cost of carbon SCC; and the wage numeraire w . The elasticities are disciplined externally, the cost-share and scale objects are recovered from accounting identities and reduced-form estimates, and the remaining objects are estimated by simulated method of moments.

Parameters set from the literature. Table 5 lists the parameters set outside the model, with their sources. We first set the elasticity of substitution across final-good varieties to $\sigma = 6.03$, the median structural-gravity trade elasticity in [Head and Mayer \(2014\)](#). The implied markup $\sigma/(\sigma - 1) \approx 1.20$ is within standard estimates. The elasticity of substitution across suppliers in the outsourced composite, ζ , is a CES variety nest of the same kind, and we set it to the same value, $\zeta = 6.03$. The outsourcing elasticity between in-house energy and the outsourced composite is set to $\eta = 2.79$, using the benchmark sourcing parameter of [Antràs et al. \(2017\)](#). The elasticity between clean and dirty production in the buyer's upper nest is set to $\sigma_b = 0.88$, following the value-added-versus-intermediate-input elasticity estimated by [Atalay \(2017\)](#). The elasticity of substitution between labor and energy in supplier production is set to $\rho = 0.64$, the firm-level estimate of [Bretschger and Jo \(2024\)](#). The energy-supply elasticity is set to $\epsilon = 0.20$, in the inelastic range of the short-run energy-supply literature ([Bornstein et al., 2023](#)) and consistent with Uruguay's hydroelectric-heavy, short-run-inelastic grid. Finally, the social cost of carbon, which enters only the welfare evaluation, is set to $\text{SCC} = \$185$ per ton of CO_2 ([Rennert et al., 2022](#)).

Table 5: Parameters Set Externally (Literature)

Parameter	Description	Source	Value
σ	Elast. of subst. across final goods	Head and Mayer (2014)	6.03
ζ	Elast. of subst. across suppliers	Head and Mayer (2014)	6.03
η	Elast., in-house energy vs. outsourcing	Antràs et al. (2017)	1.79
σ_b	Elast., clean vs. dirty (upper nest)	Atalay (2017)	0.88
ρ	Elast., labor vs. energy (suppliers)	Bretschger and Jo (2024)	0.64
ϵ	Energy-supply elasticity	Bornstein et al. (2023)	0.20
SCC	Social cost of carbon (\$/tCO ₂)	Rennert et al. (2022)	185

The substitution elasticities are set externally because the cross-sectional moments available in our data do not identify CES curvature: absent credible relative-price or policy variation, a given energy-share gap between upgraded and non-upgraded buyers can be rationalized by a high wedge with low substitutability or a lower wedge with high substitutability. We therefore fix the elasticities from external estimates and report sensitivity to them. The resulting hierarchy $\sigma \approx \zeta > \eta > \sigma_b > \rho$ is empirically grounded: substitution is highest across varieties of similar goods and suppliers, intermediate for outsourcing, and lowest across broad production factors and input bundles.

Parameters calibrated directly from data. Table 6 lists the parameters recovered directly from the data, each matched to a single observable in closed form, outside the estimation. The wage is the numeraire, $w = 1$. The clean-labor weight in the upper nest, γ , equals the buyer labor cost share (the wage bill divided by total variable cost). Similarly, the in-house-energy weight in the middle nest, α , equals the in-house-energy share of dirty spending, that is, energy expenditure over the sum of energy expenditure and outsourced intermediate purchases. The baseline supplier energy weight ψ_0 is recovered by inverting the supplier energy-intensity equation (25) at the baseline supplier energy intensity on links serving non-upgraded buyers. The construction of these cost shares is detailed in Appendix D.2.

The environmental wedge τ is recovered in closed form from the gap in in-house-energy shares between propensity-score-matched upgraded and non-upgraded buyers, using equation (31) given η . The demand-composition elasticity ν , which is the central structural parameter, is recovered from the event-study coefficient through the structural equation (32), combining the estimated supplier energy-intensity response to connecting with an upgraded buyer, the average exposure of suppliers to upgraded buyers, the baseline supplier energy intensity, and τ . Because ν is read off the empirical coefficient, the sign of the estimated coefficient determines the sign of ν ; our estimate $\nu \approx 1.10 > 0$ indicates that the outsourcing channel dominates. Finally, L is total employment in the calibration sample and $\bar{E}$ is set so that baseline energy expenditure relative to the wage bill matches its data counterpart.

Table 6: Parameters Calibrated Directly from Data

Parameter	Description	Data target / identity	Value
w	Wage (numeraire)	Normalization	1
γ	Clean-labor weight (upper nest)	Buyer clean-labor cost share	0.200
α	In-house-energy weight (mid. nest)	In-house share of dirty spending	0.026
ψ_0	Baseline supplier energy weight ^a	Baseline supplier EI (non-upgraded links)	0.037
τ	Environmental wedge	Matched in-house-energy-share gap	1.227
ν	Demand-composition elasticity	Event-study coeff. via eq. (32)	1.099
L	Labor endowment (000s)	Employment in calibration sample	498.0
$\bar{E}$	Energy-supply scale	Baseline energy expenditure / wage bill	46.83

^a Provisional: the supplier energy-intensity moment is being rebuilt on a value-added denominator (energy over labor plus energy), which raises ψ_0 relative to the sales-based value reported here. See Appendix D.2.

Identification. Each parameter in this group maps to a single observable through an accounting identity or a closed-form model inversion, and is imposed exactly at the baseline equilibrium rather than chosen to minimize a criterion. The cost-share weights γ and α are pinned by the corresponding expenditure shares; ψ_0 , τ , and ν by the supplier energy intensity, the matched energy-share gap, and the event-study coefficient, respectively; and L and $\bar{E}$ by the sample’s labor and energy accounts.

Parameters estimated via simulated method of moments. The remaining parameters, which govern the certification-cost distribution, the network, and the productivity tails are estimated jointly by simulated method of moments. Table 7 reports each parameter with its primary identifying moment.

In the certification block, the green demand premium δ is disciplined by the cross-sectional ISO sales premium. Since demand scales with $\delta^{\sigma-1}$, the premium is the target moment to identify δ . The location and dispersion of the certification-cost distribution, f_μ and f_σ , are disciplined by the share of buyers that are certified and by the standard deviation of log sales among certified buyers, the latter being informative about how widely certification spreads across the size distribution: when certification costs are concentrated, upgrading is a near-deterministic function of size and certified firms occupy a narrow size band, whereas dispersed costs admit certified firms across a broad range of sizes.

In the network block, the meeting rate λ is set by network density; the link-cost location f_{link} by the slope of buyer degree with respect to buyer size; the link-cost dispersion $f_{\text{link},\sigma}$ by the spread of degree among buyers of the same size; and the friction dispersion θ_ξ by within-buyer spending concentration at a fixed supplier count. The separation between λ and f_{link} is clean because meetings are size-independent, so the entire degree–size slope is governed by activation of linkages.

Finally, the buyer and supplier productivity tails θ_z and θ_a are disciplined by the top-decile shares of buyer and supplier sales, respectively.

Table 7: Parameters Estimated via Simulated Method of Moments

Parameter	Description	Key identifying moment	Estimate
δ	Green demand premium	ISO sales premium (cross-sectional)	1.171 (0.018)
f_μ	Certification-cost location	Share of buyers certified	0.209 (0.080)
f_σ	Certification-cost dispersion	S.d. of log sales among certified buyers	1.446 (0.243)
λ	Meeting rate	Network density	0.051 (0.003)
f_{link}	Link-cost location	Degree-size slope	1.3×10^{-6} (1.1×10^{-7})
$f_{\text{link},\sigma}$	Link-cost dispersion	Residual degree spread (given size)	0.882 (0.025)
θ_ξ	Bilateral-friction dispersion	Within-buyer concentration at fixed K_b	3.292 (0.662)
θ_z	Buyer productivity tail	Buyer top-10% sales share	5.569 (0.127)
θ_a	Supplier productivity tail	Supplier top-10% sales share	5.114 (0.058)

Standard errors in parentheses, from the exactly identified SMM formula $(1 + 1/S) G^{-1} \widehat{\Omega} (G^{-1})'$, with $\widehat{\Omega}$ estimated by a nonparametric firm bootstrap of the data moments ($B = 500$) and G the numerical Jacobian of the simulated moments at the estimate; the location parameter is reported as e^{f_μ} with a delta-method standard error. See Appendix D.3.

Estimation. The parameters are estimated by minimising a weighted distance between model and data moments, computed on the in-sample universe.⁷ We give equal weight to all moments. While, as discussed above, each moment is the primary source of identification for its parameter; the moments are not perfectly orthogonal. For instance, mean degree also falls as f_{link} rises and concentration depends on ζ . The parameters are thus identified from a joint estimation. The construction of every moment is detailed in Appendix D.2. Appendix D.3 documents identification—the criterion around the estimate, one parameter at a time—and the construction of the standard errors.

Optimization. To minimize the criterion we use a derivative-free, multistart procedure.⁸ More specifically we use a local trust-region optimiser at multiple independent start points (BOBYQA, which builds a local quadratic approximation of the criterion from function values alone) and keep the best point found. We then re-centre the starting values on that point, shrink the perturbations, and repeat, in the spirit of the TikTok multistart algorithm of [Arnoud et al. \(2019\)](#).

⁷ We use common random numbers across trial parameter vectors, so the same primitive draws underlie every evaluation and estimate the parameters jointly.

⁸ The optimization problem is not globally smooth due to the buyers' discrete choices of supplier set and of whether to upgrade create kinks and jumps in the simulated moments. Gradient-based methods are therefore unreliable, and we instead minimise the criterion with a derivative-free method.

4.4 Model Fit

Table 8 compares the nine targeted moments in the data and at the estimated parameters. The model reproduces the certification block — the share of buyers certified, the ISO sales premium, and the dispersion of sales within the certified set — together with the network moments, including the within-buyer spending concentration: eight of the nine moments are matched to within roughly two percent. The one larger gap is the supplier-sales concentration, which the model under-predicts by about five percent: it generates a somewhat more even distribution of sales across suppliers than the data. The sensitivity of each estimate to each moment is reported in Appendix D.3.

Table 8: Model Fit: Targeted Moments

Moment	Data	Model
ISO sales premium	0.988	0.997
Share of buyers certified	0.0148	0.0148
S.d. of log sales, certified buyers	1.602	1.620
Network density	0.0176	0.0176
Degree–size slope	0.425	0.426
Residual degree spread	0.789	0.789
Within-buyer concentration ($K_b = 2$)	0.772	0.756
Buyer top-10% sales share	0.790	0.793
Supplier top-10% sales share	0.876	0.836

4.5 Counterfactual Exercises

Table 9 reports the aggregate effects of environmental upgrading. We compare two policy designs against the same benchmark economy in which certification is unavailable ($\chi_b = 0$ for every buyer): ISO certification as deployed (column 1), and certification extended to the supply chain (column 2), under which certified buyers also impose the compliance wedge τ on the energy their suppliers use to serve them, as under emerging Scope 3 requirements (Appendix C.2.1); certification remains endogenous in both. The first three rows report each group’s own percent change in energy use; the aggregate row weights these changes by the groups’ baseline energy shares (8.2, 51.8, and 40.0 percent for certified buyers, non-certified buyers, and suppliers), and, since emissions are proportional to energy use, it is also the change in emissions.

Certification as deployed both cleans up and passes the buck. Certified buyers cut their energy use by 31 percent—2.5 percentage points of aggregate energy—but 44 cents of every unit they abate reappear at their suppliers, and non-certified buyers expand as the energy price falls. The two offsets differ in nature. Non-certified buyers respond only to cheaper energy (+1.4 percent); suppliers face the same price incentive and in addition absorb the dirty tasks certified buyers outsource—the demand-composition channel at the core of the model—so their increase is twice as large (+2.8 percent). Netting out both

margins, only a quarter of certified buyers' abatement survives, and aggregate emissions fall by 0.67 percent.

Extending the standard to the supply chain closes the margin to which abatement leaks. Aggregate abatement doubles, to 1.35 percent, and upstream leakage turns slightly negative: suppliers clean up alongside their buyers instead of inheriting their dirty tasks. Certified buyers abate somewhat less—outsourcing is no longer an escape valve—but the value chain as a whole abates far more.⁹ Welfare rises under both designs, by 0.05 and 0.14 percent of benchmark welfare: because the extension roughly doubles abatement at a modest incremental resource cost, it delivers nearly three times the welfare gain. Leakage is thus a consequence of policy design rather than a fate of firm-level environmental policy: covering the supply chain roughly doubles the environmental effectiveness of certification and raises its welfare gain, which is the sense in which our results speak to supply-chain-wide instruments such as Scope 3 disclosure and border-adjustment mechanisms.

Table 9: Counterfactuals: ISO Certification as Deployed and Extended to Suppliers

	(1) ISO certification	(2) ISO extended to suppliers
Certified buyers' energy (%)	−30.6	−27.4
Non-certified buyers' energy (%)	+1.4	+2.1
Suppliers' energy (%)	+2.8	−0.5
Aggregate energy = emissions (%)	−0.67	−1.35
Upstream leakage ($\Delta E_{\text{sup}}/ \Delta E_{\text{cert}} $)	0.44	−0.09
Δ welfare (%)	+0.05	+0.14

Notes: Each column compares an economy under the indicated policy to the same benchmark economy in which certification is unavailable ($\chi_b = 0$ for every buyer), at the estimated baseline parameters, holding the energy-supply scale $\bar{E}$ fixed and clearing the energy market in both economies. Groups are defined by baseline certification status (certified share 1.5% of buyers); the same firms are tracked in each counterfactual. Column (1) is ISO certification as deployed. In column (2), certified buyers additionally impose the compliance wedge τ on the energy their suppliers use on certified links (Appendix C.2.1); certification remains endogenous. The first three rows report percent changes in each group's own energy use; they do not sum to the aggregate row, which is their baseline-energy-share-weighted sum (weights 8.2, 51.8, and 40.0 percent), computed before rounding. Emissions are proportional to aggregate energy use. Upstream leakage is the change in supplier energy per unit of energy abated by certified buyers; a negative value means suppliers reduce energy alongside their buyers. Δ welfare is the change in real consumption net of climate damages, $W = U - \text{SCC} \cdot E$, in percent of its benchmark level, with consumption valued at quality-neutral prices and abatement valued at the social cost of carbon. Entries are averages across 16 simulation replications of the full network ($N_B = 5,760$, $N_S = 16,812$).

⁹ Delivering the same two abatement levels with a uniform energy tax on all firms would require rates of roughly 4.5 and 9.0 percent, respectively.

5 Conclusion

This paper studies whether firm-level environmental policies clean up production networks or merely relocate pollution upstream. We develop a model of production networks and the environment in which two opposing forces—cost-driven outsourcing of dirty tasks and green supply-chain management—shape how a buyer’s environmental commitments propagate to its suppliers. A single structural parameter, the demand-composition elasticity ν , determines which channel dominates.

To answer this question empirically, we assemble a novel dataset for Uruguay that merges the universe of firm-to-firm transactions with official energy-use registries and ISO certification records. Our event-study estimates reveal that suppliers increase their energy intensity by approximately 5 percent after connecting with a clean buyer. The result is robust across multiple definitions of buyer cleanliness—ISO-certified firms, firms in the lowest within-sector decile of energy intensity, and multinational enterprises—and survives placebo and matching exercises. The outsourcing channel dominates: clean buyers shift dirty activities onto their suppliers rather than greening them.

These findings carry direct implications for the design of environmental policy. Regulations that target individual firms without accounting for supply-chain reallocation may overstate their true abatement, as part of the reduction in regulated firms’ energy use is offset by increases upstream. Our structural estimates of the leakage rate provide a first quantification of this margin.

Several questions remain open. First, our analysis focuses on a single country; whether the outsourcing channel dominates in economies with different industrial structures, stricter environmental enforcement, or deeper integration into global value chains is an important question for future work. Second, while we document the upstream shift in energy intensity, we do not observe whether suppliers adopt dirtier technologies or simply expand output of already energy-intensive products—disentangling these margins would sharpen the welfare implications. Third, our model abstracts from supplier bargaining power and long-term relational contracting, which may moderate or amplify the outsourcing channel over longer horizons. Finally, whether supply-chain-wide regulations—such as the EU’s Carbon Border Adjustment Mechanism or Scope 3 disclosure mandates—can effectively close the leakage we document remains a natural next step for both theoretical and empirical investigation.

References

- Alfaro-Ureña, A., B. Faber, C. Gaubert, I. Manelici, and J. P. Vasquez (2022). Responsible sourcing? Theory and evidence from Costa Rica. Working Paper 30683, National Bureau of Economic Research.
- Andrews, I., M. Gentzkow, and J. M. Shapiro (2017). Measuring the sensitivity of parameter estimates to estimation moments. *Quarterly Journal of Economics* 132(4), 1553–1592.
- Antràs, P., T. C. Fort, and F. Tintelnot (2017). The margins of global sourcing: Theory and evidence from US firms. *American Economic Review* 107(9), 2514–2564.
- Antweiler, W., B. R. Copeland, and M. S. Taylor (2001). Is free trade good for the environment? *American Economic Review* 91(4), 877–908.
- Arimura, T. H., A. Hibiki, and H. Katayama (2008). Is a voluntary approach an effective environmental policy instrument? A case for environmental management systems. *Journal of Environmental Economics and Management* 55(3), 281–295.
- Arkolakis, C., F. Huneeus, and Y. Miyauchi (2025). Production network formation, trade, and welfare. Working Paper 11833, CESifo.
- Arnoud, A., F. Guvenen, and T. Kleineberg (2019). Benchmarking global optimizers. Working Paper 26340, National Bureau of Economic Research.
- Atalay, E. (2017). How important are sectoral shocks? *American Economic Journal: Macroeconomics* 9(4), 254–280.
- Ben-David, I., Y. Jang, S. Kleimeier, and M. Viehs (2021). Exporting pollution: Where do multinational firms emit CO₂? *Economic Policy* 36(107), 377–437.
- Bernard, A. B., A. Moxnes, and Y. U. Saito (2019). Two-sided heterogeneity and trade. *Review of Economics and Statistics* 101(3), 425–439.
- Bornstein, G., P. Krusell, and S. Rebelo (2023). A world equilibrium model of the oil market. *The Review of Economic Studies* 90(1), 132–164.
- Bretschger, L. and A. Jo (2024). Complementarity between labor and energy: A firm-level analysis. *Journal of Environmental Economics and Management* 124, 102934.
- Brucal, A., B. S. Javorcik, and I. Love (2019). Good for the environment, good for business: Foreign acquisitions and energy intensity. *Journal of International Economics* 121, 103247.
- Callaway, B. and P. H. C. Sant’Anna (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics* 225(2), 200–230.
- Chaney, T. (2014). The network structure of international trade. *American Economic Review* 104(11), 3600–3634.

- Cherniwchan, J., B. R. Copeland, and M. S. Taylor (2017). Trade, the environment, and development: A survey of recent advances and future directions. *Annual Review of Economics* 9, 537–565.
- Cocci, M. D. and M. Plagborg-Møller (2025). Standard errors for calibrated parameters. *Review of Economic Studies*. doi:10.1093/restud/rdae099.
- Cole, M. A., R. J. Elliott, and L. Zhang (2017). The impact of FDI on the environment. *Routledge Handbook of Trade and the Environment*, 193–207.
- Colmer, J., R. Martin, M. Muûls, and U. J. Wagner (2025). Does pricing carbon mitigate climate change? Firm-Level evidence from the European Union Emissions Trading System. *Review of Economic Studies* 92(3), 1625–1660.
- Copeland, B. R. and M. S. Taylor (2004). Trade, growth, and the environment. *Journal of Economic Literature* 42(1), 7–71.
- Coster, P., J. di Giovanni, and I. Mejean (2024). Firms’ supply chain adaptation to carbon taxes. Discussion Paper 19644, CEPR.
- Demir, B., A. C. Fieger, D. Y. Xu, and K. K. Yang (2024). O-ring production networks. *Journal of Political Economy* 132(1), 200–247.
- Eaton, J., S. S. Kortum, and F. Kramarz (2022). Firm-to-firm trade: Imports, exports, and the labor market. Working Paper 29685, National Bureau of Economic Research.
- Eskeland, G. S. and A. E. Harrison (2003). Moving to greener pastures? Multinationals and the pollution haven hypothesis. *Journal of Development Economics* 70(1), 1–23.
- Fowlie, M. and M. Reguant (2018). Challenges in the measurement of leakage risk. *AEA Papers and Proceedings* 108, 124–129.
- Garcia-Lembergman, E., N. Ramondo, A. Rodríguez-Clare, and J. S. Shapiro (2026a). The carbon footprint of multinational production. Technical report. Working paper.
- Garcia-Lembergman, E., N. Ramondo, A. Rodríguez-Clare, and J. S. Shapiro (2026b). Global value chains and the design of carbon taxes. Technical report. Working paper.
- Garetto, S., N. Pavcnik, N. Ramondo, V. Alviarez, J. Fan, N. Pandalai-Nayar, N. Limodio, I. Manelici, N. Morales, E. Dardati, E. Garcia-Lembergman, G. Gu, G. Hale, D. Hemous, R. Martin, F. Farrokhi, H. Pellegrina, P. L. Vezina, L. Boudreau, and J. Vasquez (2025). Foreign direct investment and development. *VoxDevLit* 13(1).
- Hanna, R. (2010). Us environmental regulation and FDI: Evidence from a panel of US-based multinational firms. *American Economic Journal: Applied Economics* 2(3), 158–189.
- Head, K. and T. Mayer (2014). Gravity equations: Workhorse, toolkit, and cookbook. In G. Gopinath, E. Helpman, and K. Rogoff (Eds.), *Handbook of International Economics*, Volume 4, pp. 131–195. Elsevier. Meta-analysis of 435 trade-elasticity estimates from 32 papers; median trade elasticity 5.03.

- Huneus, F., K. Kroft, and K. Lim (2021). Earnings inequality in production networks. Working Paper 28424, National Bureau of Economic Research.
- Keller, W. and A. Levinson (2002). Pollution abatement costs and foreign direct investment inflows to US states. *Review of Economics and Statistics* 84(4), 691–703.
- King, A. A. and M. J. Lenox (2005). Does it really pay to be green? An empirical study of firm environmental and financial performance. *Journal of Industrial Ecology* 5(1), 105–116.
- Levinson, A. and M. S. Taylor (2008). Unmasking the pollution haven effect. *International Economic Review* 49(1), 223–254.
- Milewska, B. and D. Milewski (2023). The impact of energy consumption costs on the profitability of production companies in poland in the context of the energy crisis. *Energies* 16(18), 6519.
- Oberfield, E. (2018). A theory of input–output architecture. *Econometrica* 86(2), 559–589.
- Potoski, M. and A. Prakash (2005). Do voluntary programs reduce pollution? Examining ISO 14001’s effectiveness across countries. *Journal of Policy Analysis and Management* 24(4), 745–769.
- Rennert, K., F. Errickson, B. C. Prest, L. Rennels, R. G. Newell, W. Pizer, C. Kingdon, J. Wingenroth, R. Cooke, B. Parthum, et al. (2022). Comprehensive evidence implies a higher social cost of CO₂. *Nature* 610(7933), 687–692.
- Shapiro, J. S. (2021). The environmental bias of trade policy. *Quarterly Journal of Economics* 136(2), 831–886.
- Shapiro, J. S. and R. Walker (2018). Why is pollution from US manufacturing declining? The roles of environmental regulation, productivity, and trade. *American Economic Review* 108(12), 3814–3854.
- Steenbergen, V. and A. Saurav (2023). The environmental footprint of multinational enterprises. Technical report, World Bank Group.
- Sun, L. and S. Abraham (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics* 225(2), 175–199.
- Wang, E.-Z. and J. Xu (2025). Something for nothing? the effect of voluntary environmental program on energy intensity. *Resource and Energy Economics* 83, 101511.

Appendix

A Appendix: Data & Facts

A.0.1 Sample Coverage and Missing Energy Data

Table 10 compares firms with and without observed energy intensity. The 1,199,575 firm-year observations with missing energy data correspond overwhelmingly to micro and informal firms: they report zero sales, employ on average fewer than 2 workers, and account for negligible shares of aggregate wages and imports. By contrast, firms with observed energy intensity average 10 employees and over USD 1.4 million in annual sales.

Table 11 quantifies the economic coverage of our analytical sample. Restricting attention to firms with non-missing energy intensity retains 100% of total sales, 99% of total wages, and 97% of total imports in the economy. Even after further excluding firms with zero energy intensity, the sample still captures 68% of sales and 82% of wages. Our analysis therefore covers the economically meaningful core of the Uruguayan formal sector, with sample attrition concentrated entirely among the smallest firms.

Table 10: Summary Statistics by Energy Intensity Availability

Variable	Energy Intensity Observed			Energy Intensity Missing		
	Mean	SD	Obs.	Mean	SD	Obs.
Total sales (USD)	1,461,387	31,007,030	758,934	0	0	1,199,575
Employment	10	105	758,934	2	119	1,199,575
Wages (USD)	136,260	3,105,207	758,934	908	132,724	1,199,575
Total imports (USD)	140,168	7,190,234	758,934	2,429	197,284	1,199,575

Notes: This table compares summary statistics for firm-year observations with observed versus missing energy intensity. Monetary values are expressed in USD.

Table 11: Share of Aggregate Economic Activity Retained by Energy Intensity Filter

Variable	Excluding Missing Values (%)	Excluding Missing and Zero Values (%)
Total Sales	100.00	67.58
Total Wages	98.96	81.73
Employment	72.67	50.46
Total Imports	96.87	61.80

Notes: Each cell reports the percentage of the economy-wide aggregate that is retained in the analytical sample after applying the indicated filter. For example, firms with non-missing energy intensity account for 100% of total sales in Uruguay.

A.0.2 Data Consistency and the Energy Cost Share

A critical variable in our analysis is the energy intensity of manufacturing firms. In our main descriptive statistics, we report an unweighted average energy intensity of approximately 4.3% (0.043) across the sample of manufacturing firms. To ensure the reliability of our data and the robustness of our measure, it is essential to benchmark this figure against international evidence and macro-level aggregates.

The unweighted mean of 4.3% is fully consistent with international micro-level evidence when considering the distribution of firm sizes and sector composition. For instance, recent studies analyzing the cost structure of European manufacturing firms (Milewska and Milewski, 2023) indicate that the average share of energy costs to total revenues is approximately 4.1%. Similarly, reports from the Italian industrial confederation (?) state that the historical pre-crisis incidence of energy costs over total operating costs in Italian manufacturing was around 3.9%, rising to 4.9% during the recent energy shocks. Our unweighted mean of 4.3% falls squarely within this documented empirical range for typical industrial structures.

Micro vs. Macro Aggregation and the Role of Firm Size

In highly skewed distributions—characteristic of the manufacturing sector—a small number of massive corporations account for a disproportionate share of total national sales but exhibit very low energy intensities. Conversely, the vast majority of firms are small and medium enterprises (SMEs) or operate in traditional, energy-intensive sectors (e.g., agro-industry, food processing), which have higher baseline energy needs relative to their scale.

When we recalculate the energy intensity for our sample as a *sales-weighted average*, the aggregate energy share drops significantly to 1.5% (0.015).

This weighted aggregate of 1.5% perfectly aligns with the macroeconomic benchmarks reported

by the U.S. ASM, where the aggregate ratio of purchased fuels and electricity to the total value of shipments historically fluctuates between 1.4% and 2.0% for the overall manufacturing sector.

In conclusion, the behavior of our energy intensity variable confirms the structural integrity of our dataset. It demonstrates that our microdata accurately captures both the high energy dependence of the median firm and the industrial concentration.

A.1 ISO Firms

This appendix provides additional details on the ISO-certified firms used in the analysis. Table 12 reports the source-level breakdown of the 255 unique ISO 14001 and ISO 50001 certificates underlying the count cited in Section 3.1. Most certificates come from the IAF global database (138) and the UNIT national directory (107), with a smaller contribution from the certification body LSQA (10). The 181 firms matched to our administrative panel correspond to roughly 89% of the 203 active certifications officially reported by the IAF for Uruguay in 2018, indicating broad coverage of the certified universe.

Table 12: Summary of ISO-Certified Firms and Data Collection

Category	ISO 14001	ISO 50001	Total Unique Names
Total Identified Certificates	228	27	255^a
<i>Source Breakdown (by certification):</i>			
UNIT (National Institute)	81	26	107
IAF (Global Database)	137	1	138
LSQA (Certification Body)	10	0	10
Total Firms Matched in Database	–	–	181
External Benchmark (IAF Uruguay, 2018)	–	–	203^b


^a Total number of unique firm names identified across all sources. Note that 4 firms were identified as holding both ISO 14001 and ISO 50001 certifications.

^b Official count provided by the International Accreditation Forum (IAF) for active ISO 14001 or ISO 50001 certifications in Uruguay as of 2018.

Figure 4 complements this with two cross-sectional distributions of (log) energy intensity for the year 2016. Panel (a) compares suppliers that transact with at least one ISO-certified buyer to suppliers that do not. Panel (b) compares ISO-certified buyers themselves to non-certified buyers and to the matched control group of non-certified buyers introduced in Section ???. In both panels, the ISO-related distribution lies visibly to the left of its non-ISO counterpart, and the gap survives when restricting to the matched control sample—confirming that the ISO cleanliness premium

documented in the main text is not driven by firm size or network composition.

Figure 4: Energy-Intensity Distributions: ISO vs. Non-ISO Firms and Their Suppliers (2016)



Notes: Both panels report kernel density estimates of the natural logarithm of energy intensity in the 2016 cross-section. Panel (a) plots the distribution among suppliers, separately for suppliers connected to ISO 14001 or ISO 50001 certified buyers (red) and suppliers connected to non-certified buyers (blue). Panel (b) plots the distribution among buyers, comparing ISO-certified buyers, non-certified buyers, and the matched control group of non-certified buyers constructed via one-nearest-neighbor propensity-score matching as described in Section ?? . Energy intensity is defined as the share of fuel and utility expenditures over total sales.

A.1.1 Testable Implication I: Negative Assortative Matching with ISO

Table 13 extends the assortative-matching specification of Testable Implication I by allowing the slope between buyer and supplier energy intensity to differ for ISO-certified buyers. The interaction term $\text{ISO} \times \text{Buyer EI}$ is statistically indistinguishable from zero across most specifications, indicating that the negative assortative matching predicted by the model is a general feature of the network—as the theory implies, since the mechanism operates through the buyer’s energy cost wedge τ_b (equation (24)), not through the specific channel that generates it.

Table 13: Assortative Matching in Energy Intensity, ISO vs. Non-ISO Buyers (2013–2016)

	(1) Weighted Average		(2) Simple Average	
	ISO vs Matched	ISO vs Full	ISO vs Matched	ISO vs Full
ISO $\times$ Buyer EI	-0.034 (0.046)	0.003 (0.021)	0.004 (0.007)	0.013** (0.006)
Employment	0.001 (0.042)	-0.025 (0.016)	0.034 (0.021)	0.085*** (0.006)
Observations	791	79,846	791	79,846
R^2	0.309	0.101	0.309	0.183
Fixed Effects	sector-dept-year	sector-dept-year	sector-dept-year	sector-dept-year

Notes: Regressions of suppliers' energy intensity on that of their buyers, augmented with an interaction between an ISO indicator and the buyer's energy intensity. Columns (1) report the purchase-weighted average of supplier energy intensity; columns (2) report the simple average. "ISO vs Matched" restricts the sample to ISO-certified buyers and the matched control group; "ISO vs Full" uses the full sample of buyers. All specifications include sector-region-year fixed effects. Standard errors clustered at the buyer-sector level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

A.1.2 Testable Implication II: Number of Suppliers of ISO Firms

The model predicts that upgraded buyers activate more suppliers because the marginal gain from each link rises with outsourced dirty spending $X_b^{D,O}$ (equation (13)). ISO 14001 certification provides a natural proxy for environmental upgrading, allowing us to test this prediction directly. Table 14 compares the number of suppliers and energy intensity across ISO-certified buyers, a matched control group of similar size, and the general population of uncertified firms. Consistent with the model, ISO firms are cleaner and maintain on average 15% more suppliers than matched non-ISO firms of comparable size.


Table 14: Buyer Characteristics by Certification and Matching Status

Variable	(1) Non-ISO	(2) Matched Non-ISO	(3) ISO-Certified
Number of Suppliers	57.12 (50.01)	253.21 (180.86)	292.60 (207.94)
Energy Intensity	0.051 (0.094)	0.034 (0.051)	0.027 (0.065)
Observations	29,152 ^a	151	159

Notes: This table compares three groups of buyers: (1) the general population of non-certified firms, (2) a matched control group of non-certified firms, and (3) ISO-certified firms. The matched group is constructed using 1 nearest neighbor propensity score matching based on log sales and log employment. Standard deviations are reported in parentheses. Energy intensity is defined as the share of fuel and utility expenditures over total sales.

^a Observations for the general non-ISO population are derived by excluding the 659 treated and 159 matched control firms from the full cross-sectional sample of 29,152 buyers.

Figure 5 provides further support by comparing the entire distribution of network size across the three groups. The density for ISO-certified buyers is shifted visibly to the right, and crucially, this gap shrinks but persists when comparing against the matched sample, ruling out that the pattern is driven by differences in firm size. The evidence is consistent with the model's prediction that environmentally upgraded firms maintain broader sourcing networks as a consequence of substituting away from in-house energy toward outsourced intermediates.

Figure 5: Kernel Density of Number of Suppliers: ISO vs. Non-ISO Firms

Notes: This figure compares the kernel density of the number of suppliers per buyer across three groups: firms holding ISO 14001 or 50001 environmental certification, uncertified firms, and a matched sample of uncertified firms with comparable sales and employment.

A.2 Institutional Setting: The Architecture of ISO 14001 and ISO 50001

This section provides a detailed overview of the institutional setting and operational mechanics of the ISO 14001 (Environmental Management) and ISO 50001 (Energy Management) standards. These standards do not dictate absolute emission limits; rather, they provide a structured framework for organizations to systematically monitor, control, and improve their environmental and energy performance.

System Boundaries and Significant Aspects A fundamental principle of ISO management standards is that the organization defines the physical and operational boundaries of its management system (EMS or EnMS). Compliance and auditing are strictly tethered to this defined scope. Within this boundary, the standards mandate that organizations identify their environmental aspects and determine which of these have a “Significant Environmental Impact” (SEAs). Similarly, ISO 50001 requires the identification of “Significant Energy Uses” (SEUs).

The determination of significance typically relies on a risk assessment matrix that evaluates both the likelihood of occurrence and the severity of the potential environmental impact. For manufacturing and large-scale commercial entities, in-house electricity consumption and fuel usage are universally classified as significant. Once classified as an SEA or SEU, the firm is required to set measurable environmental objectives and establish rigorous monitoring and measurement processes.

The Auditing and Enforcement Cycle The credibility of these voluntary standards relies heavily on a continuous, multi-year auditing cycle conducted by independent, accredited certification bodies (e.g., UNIT or LSQA in Uruguay). The standard certification cycle lasts three years and is structured to ensure ongoing adherence:

1. **Year 0 (Initial Certification):** The firm undergoes a comprehensive two-stage audit, including an exhaustive documentation review and an on-site facility inspection to verify that the management system is fully operational.
2. **Years 1 and 2 (Surveillance Audits):** The firm faces mandatory annual on-site audits. Auditors sample various components of the EMS to verify continuous improvement and strict adherence to established targets.
3. **Year 3 (Recertification):** A full-scale audit evaluates the EMS’s performance over the entire three-year period to renew the certificate.

During these evaluations, third-party certifiers do not rely on qualitative assurances; they actively verify utility bills, fuel purchasing logs, and emissions inventories against the firm's stated reduction targets and compliance obligations.

Non-Conformities and Certificate Suspension If auditors detect deviations from the standard or a failure to meet established objectives, they issue formal "Non-Conformities" (NCs). A "Minor Non-Conformity" typically addresses localized procedural failures and requires the firm to submit a documented corrective action plan. However, a "Major Non-Conformity"—such as a systemic failure to monitor a significant environmental aspect or a persistent inability to achieve continuous improvement targets—is severely penalized. If a major non-conformity is not resolved within a strict timeframe (often 90 days), it leads to the immediate suspension or withdrawal of the ISO certificate.

Supply Chain Considerations While the primary focus of ISO 14001 has historically been on the firm's direct operations, the standard does incorporate supply chain considerations. Organizations are required to consider environmental impacts associated with the procurement of products and services and, particularly following the 2015 revision, are expected to establish controls for outsourced processes. Firms must communicate their relevant environmental requirements to external providers. However, the core auditing procedures and quantitative reduction targets remain heavily concentrated on the operations and energy consumption that occur within the certified firm's explicitly defined operational boundaries.

B Appendix: Empirical Strategy

B.1 PSM ISO Details

We present the results of the tests performed on the matching:

Table 15: Covariate Balance after Matching

Variable	Mean Treated	Mean Control	% Bias	t-stat	p-value	Var. Ratio (T/C)
Energy Intensity	0.04187	0.04437	-1.2	-1.55	0.121	0.31*
Employment	32.086	31.644	0.3	0.21	0.830	0.70*
Sales	5,000,000	4,600,000	1.0	0.78	0.433	4.53*
Wages	580,000	580,000	-0.0	-0.01	0.991	1.51*
<i>* Variance ratio outside [0.97; 1.03] indicates local variance differences.</i>						
Overall balance statistics:						
Pseudo R^2			0.000			
Mean standardized bias (%)			0.6			
Median standardized bias (%)			0.7			
B			2.2			
R			0.59			
% of variance balanced			100			

Notes: Overall balance statistics follow Rubin's criteria. $B < 25\%$ and R within the $[0.5; 2]$ range indicate that the overall balance is highly satisfactory. * Indicates the variance ratio is outside the strict $[0.97; 1.03]$ interval.


B.2 Robustness: Alternative Thresholds for Clean Buyers

In our baseline specification, we define a “clean” buyer as one whose energy intensity falls within the bottom decile (top 10% cleanest) of its 2-digit sector.

Figure 6 replicates our event-study analysis defining clean buyers as those in the bottom 15% (Panel a) and the bottom 5% (Panel b) of the sectoral energy intensity distribution. As expected, the magnitude of the effect responds to the strictness of the threshold. When expanding the definition to the cleanest 15%, the post-treatment coefficients are slightly attenuated but remain positive and statistically significant for years 3 and 4 after connection. Conversely, when restricting the treatment to the extreme top 5% cleanest buyers, the upward shift in supplier energy intensity more precisely estimated, reinforcing the mechanism that buyers with a very clean energy profile are the ones driving the upstream pollution shifting. Likewise, the 10% threshold continues to yield the most

significant results.

Figure 6: Effect of Connecting to a Clean Top 15% or 5% Buyer




Notes: Event-study coefficients $\hat{\theta}_k$ from equation (34). **Panel (a)** shows the effect for clean buyers, defined as firms in the lowest top 15% of energy intensity within their 2-digit sector, measured in the year prior to the connection. **Panel (b)** shows the effect for clean buyers, defined as firms in the lowest top 5% of energy intensity within their 2-digit sector, measured in the year prior to the connection. For both panels, the dependent variable, omitted period, and confidence intervals are defined as in Figure 1.

B.3 PSM: Clean & Multinationals

Figure 7 confirms these results on the propensity-score matched sample. In Panel (a) for Clean Buyers, the pre-period shows slightly more dispersion than in the ISO specification—consistent with the broader and noisier treatment definition—but the estimates converge to zero by $k = -1$ and rise sharply thereafter. Point estimates reach 10–12 log points by $k = 4$, confirming that the matched sample, produces larger effects. Panel (b) demonstrates the corresponding estimates for connections with Multinational Enterprises (MNE).


Figure 7: Effect of Connecting to a Clean or MNE Buyer: Matched Sample



Notes: Event-study coefficients estimated on the propensity-score matched sample. Each treated supplier is matched to its five nearest neighbors within the same 2-digit sector. **Panel (a)** shows the effect for clean buyers (defined as in Figure 2(a)). **Panel (b)** shows the effect of connecting to a multinational enterprise (MNE). Balance diagnostics are reported in the Appendix.

B.4 Alternative definitions of clean: MNE vs Clean

Figure 8: Connecting to a MNE versus Clean Domestic Buyer



Notes: This figure plots event-study coefficients $\hat{\theta}_k$ following the specification in equation (34). Treatment is defined as forming a first connection with a domestic buyer matched with MNEs using Propensity Score Matching based on energy intensity.

C Appendix: Theoretical Framework

C.1 Proof of Equation (13): Profit Gain from Adding a Supplier

Proof. Let $\mathcal{M}_b^{\mathcal{A}}$ be the current active supplier set of buyer b . Consider adding supplier $s \notin \mathcal{M}_b^{\mathcal{A}}$ with delivered cost $\tilde{r}_{bs}$.

Step 1: Adding a variety reduces the CES price index. Define $\Sigma_0 \equiv \sum_{s' \in \mathcal{M}_b^{\mathcal{A}}} \tilde{r}_{bs'}^{1-\zeta}$, so the initial outsourced price index is $p_b^{M,0} = \Sigma_0^{1/(1-\zeta)}$ and the new price index after adding s is $p_b^{M,1} = (\Sigma_0 + \tilde{r}_{bs}^{1-\zeta})^{1/(1-\zeta)}$. Their ratio is:

$$\frac{p_b^{M,1}}{p_b^{M,0}} = \left(1 + \frac{\tilde{r}_{bs}^{1-\zeta}}{\Sigma_0}\right)^{\frac{1}{1-\zeta}}.$$

The expenditure share that supplier s would receive in the augmented set $\mathcal{M}_b^{\mathcal{A}} \cup \{s\}$ is:

$$\omega_{bs} = \frac{\tilde{r}_{bs}^{1-\zeta}}{\Sigma_0 + \tilde{r}_{bs}^{1-\zeta}}, \quad \text{so} \quad \frac{\tilde{r}_{bs}^{1-\zeta}}{\Sigma_0} = \frac{\omega_{bs}}{1 - \omega_{bs}}.$$

Substituting:

$$\frac{p_b^{M,1}}{p_b^{M,0}} = \left(1 + \frac{\omega_{bs}}{1 - \omega_{bs}}\right)^{\frac{1}{1-\zeta}} = \left(\frac{1}{1 - \omega_{bs}}\right)^{\frac{1}{1-\zeta}} = (1 - \omega_{bs})^{\frac{1}{\zeta-1}}.$$

Since $\zeta > 1$ and $\omega_{bs} \in (0, 1)$, this ratio is strictly less than one: adding a variety always reduces the CES price index. For small ω_{bs} :

$$\Delta \ln p_b^M = \frac{1}{\zeta - 1} \ln(1 - \omega_{bs}) \approx -\frac{\omega_{bs}}{\zeta - 1}.$$

Step 2: From price index to unit cost. By the properties of CES cost functions, the elasticity of each nest's price index with respect to an input price equals that input's cost share:

$$\frac{\partial \ln p_b^D}{\partial \ln p_b^M} = s_b^{D,O}, \quad \frac{\partial \ln c_b}{\partial \ln p_b^D} = s_b^D.$$

By the chain rule:

$$\Delta \ln c_b = s_b^D \cdot s_b^{D,O} \cdot \Delta \ln p_b^M \approx -s_b^D \cdot s_b^{D,O} \cdot \frac{\omega_{bs}}{\zeta - 1}.$$

Step 3: From unit cost to profit. The buyer's profit function is $\Pi_b(c_b) = \max_{q_b} \{p_b(q_b) q_b - c_b q_b\}$.

By the envelope theorem:

$$\frac{d\Pi_b}{dc_b} = -q_b,$$

so every dollar of cost reduction, applied to the optimal output q_b , translates dollar-for-dollar into profit. Therefore:

$$\Delta \Pi_b = q_b \cdot |\Delta c_b| = q_b c_b \cdot |\Delta \ln c_b| = C_b \cdot s_b^D \cdot s_b^{D,O} \cdot \frac{\omega_{bs}}{\zeta - 1},$$

where $C_b \equiv c_b q_b$ is the buyer's total variable cost. Defining the buyer's outsourced dirty cost as $X_b^{D,O} \equiv s_b^D \cdot s_b^{D,O} \cdot C_b$:

$$\Delta \Pi_b(s \mid \mathcal{M}_b^A) \approx \frac{1}{\zeta - 1} \cdot \omega_{bs} \cdot X_b^{D,O}. \quad (35)$$

The link formation condition $\Delta \Pi_b > f$ is therefore:

$$\omega_{bs} \cdot X_b^{D,O} > (\zeta - 1) f. \quad \square \quad (36)$$

□

C.2 Algorithm: Finding the Optimal Supplier Set

The optimal active supplier set always consists of the K cheapest suppliers for some K . This reduces the buyer's problem from choosing among $2^{|\mathcal{M}_b|}$ possible subsets to evaluating just $|\mathcal{M}_b| + 1$ candidates.¹⁰

The algorithm can proceed as follows.

Let $N = |\mathcal{M}_b|$ and label the suppliers in the meeting set so that $\tilde{r}_{b1} \leq \tilde{r}_{b2} \leq \dots \leq \tilde{r}_{bN}$. Define the candidate sets $\mathcal{M}_b^{A^K} = \{1, \dots, K\}$ for $K = 0, 1, \dots, N$, where $\mathcal{M}_b^{A^0} = \emptyset$ means the buyer produces everything in-house.

¹⁰Note that the optimal K^* cannot in general be found by sequentially adding suppliers and stopping when the marginal supplier is unprofitable. The reason is complementarity via the scale effect: adding multiple suppliers jointly may be profitable even when adding each individually is not, because each additional supplier lowers the outsourced price index, which raises total outsourced spending, which in turn makes all other suppliers more valuable.

Step 1: Compute equilibrium objects for each candidate K . For each $K \in \{0, 1, \dots, N\}$, compute:

$$p_b^M(K) = \left[\sum_{k=1}^K \tilde{r}_{bk}^{1-\zeta} \right]^{\frac{1}{1-\zeta}},$$

$$p_b^D(K) = \left[\alpha (\tau_b p_E)^{1-\eta} + (1-\alpha) (p_b^M(K))^{1-\eta} \right]^{\frac{1}{1-\eta}},$$

$$c_b(K) = \frac{1}{z_b} \left[\gamma w^{1-\sigma_b} + (1-\gamma) (p_b^D(K))^{1-\sigma_b} \right]^{\frac{1}{1-\sigma_b}}.$$

For $K = 0$, the buyer does not outsource: set $s_b^{D,O} = 0$ and $p_b^D(0) = [\alpha]^{1/(1-\eta)} \cdot \tau_b p_E$.

Step 2: Compute profits net of fixed costs. The buyer's variable profit under monopolistic competition is proportional to revenue:

$$\Pi_b(K) = \frac{1}{\sigma} R_b(K), \quad R_b(K) \propto c_b(K)^{1-\sigma}.$$

Net profit from activating the top- K suppliers is:

$$\Pi_b^{\text{net}}(K) = \Pi_b(K) - K \cdot f.$$

Step 3: Choose the optimal K . The buyer selects:

$$K_b^* = \arg \max_{K \in \{0, 1, \dots, N\}} \{ \Pi_b^{\text{net}}(K) \}. \quad (37)$$

C.2.1 Extension: Supply-Chain Compliance Policy Quanterfactual

We extend the algorithm to a counterfactual in which upgraded buyers impose a compliance wedge $\tau_s > 1$ on their suppliers' energy use. Under this policy, a supplier serving an upgraded buyer ($\chi_b = 1$) faces an effective energy price of $p_E \cdot \tau_s$ instead of p_E when producing the intermediate for that buyer. This captures policies that extend environmental requirements to the supply chain. Suppliers serving non-upgraded buyers continue to face the market energy price p_E .

Modified supplier cost. When serving buyer b , the supplier's unit cost becomes:

$$r_s(\psi_b, \tau_s^b) = \frac{1}{\tilde{a}_s} \left[w^{1-\rho} + \psi_b (p_E \cdot \tau_s^b)^{1-\rho} \right]^{\frac{1}{1-\rho}}, \quad (38)$$

where $\tau_s^b = \tau_s^{\chi_b}$: the compliance wedge applies only when the buyer is upgraded. The delivered cost is $\tilde{r}_{bs} = r_s(\psi_b, \tau_s^b) \cdot \xi_{bs}$ as before.

The prefix property still holds. The compliance wedge τ_s^b is buyer-specific, not supplier-specific: all suppliers serving the same buyer b face the same effective energy price $p_E \cdot \tau_s^b$. The supplier's unit cost $r_s(\psi_b, \tau_s^b)$ is therefore common across all suppliers of a given buyer. The ranking of suppliers by delivered cost $\tilde{r}_{bs}$ is still determined by the supplier's productivity $\tilde{a}_s$ and the bilateral friction ξ_{bs} , exactly as in the baseline model. Proposition 1 applies without modification: the optimal active set consists of the K cheapest met suppliers for some K .

Modified algorithm. The algorithm of Section C.2 requires only one change: in Step 1, replace p_E with $p_E \cdot \tau_s^b$ in the supplier's unit cost when computing the price indices for buyer b . Specifically, for each candidate K :

$$\begin{aligned} r_s(\psi_b, \tau_s^b) &= \frac{1}{\tilde{a}_s} \left[w^{1-\rho} + \psi_b (p_E \cdot \tau_s^b)^{1-\rho} \right]^{\frac{1}{1-\rho}}, \\ \tilde{r}_{bs} &= r_s(\psi_b, \tau_s^b) \cdot \xi_{bs}, \\ p_b^M(K) &= \left[\sum_{k=1}^K \tilde{r}_{bk}^{1-\zeta} \right]^{\frac{1}{1-\zeta}}. \end{aligned}$$

All subsequent steps—computing $p_b^D(K)$, $c_b(K)$, $\Pi_b^{\text{net}}(K)$, and selecting K_b^* —are identical to the baseline algorithm.

Supplier's energy intensity under compliance. The supplier's energy intensity when serving buyer b under the compliance wedge is:

$$\text{EI}_s(b) = \frac{\psi_b (p_E \cdot \tau_s^b)^{1-\rho}}{w^{1-\rho} + \psi_b (p_E \cdot \tau_s^b)^{1-\rho}} \cdot \frac{p_E}{p_E \cdot \tau_s^b}, \quad (39)$$

where the first fraction is the energy expenditure share at the supplier's effective price ($p_E \cdot \tau_s^b$), and the second factor converts from the supplier's cost-based energy expenditure to physical energy

valued at the market price p_E . With $\rho < 1$ (complements), the compliance wedge lowers energy intensity: the supplier substitutes away from expensive energy. With $\rho > 1$ (substitutes), the effect depends on the relative magnitude of the substitution and price effects.

C.3 Micro-Foundation for the Demand-Composition Parameter ψ_b

This appendix derives the reduced-form parameterization $\psi_b = \psi_0 \cdot \tau_b^\nu$ used in the main text from two primitives: the buyer's optimal specification choice and supplier learning.

$$D_b = \left[\alpha^{\frac{1}{\eta}} e_b^{\frac{\eta-1}{\eta}} + (1 - \alpha)^{\frac{1}{\eta}} \psi_b M_b^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}$$

The buyer's benefit from higher ψ_b . With $\psi_b \cdot M_b$ in the dirty nest (6), each unit of the outsourced composite delivers ψ_b units of effective dirty service. The effective outsourced price per unit of dirty service is $q_b = r_s(\psi_b)/\psi_b$, which is decreasing in ψ_b when $\rho > 1$: dirtier intermediates are both cheaper to produce and more useful per unit. Absent any friction, the buyer would demand the dirtiest possible specification.

Reorganization costs. Changing the specification from the industry default ψ_0 is costly. The default reflects established product designs, existing supplier capabilities, and decades of accumulated know-how about how buyer and supplier production lines interface. Deviating from it requires the buyer to redesign product requirements. To fix ideas, consider a garment company that sources fabric from a textile mill. The default specification ψ_0 reflects a standard dyeing process—moderately energy-intensive, using conventional dyes at standard temperatures. If the garment company wants a dirtier specification (higher ψ_b)—say, a heavier dye process that uses more energy but produces more durable coloring—it must change its cutting patterns, adjust sewing procedures to handle the stiffer fabric.

We model these frictions as a convex cost in units of labor:

$$C^{\text{reorg}}(\psi_b) = \frac{\kappa}{2} \left(\frac{\psi_b - \psi_0}{\psi_0} \right)^2, \quad (40)$$

where $\kappa > 0$ governs the magnitude of reorganization frictions.

The buyer's problem. An upgraded buyer ($\chi_b = 1$) chooses ψ_b to maximize operating profit net of reorganization costs:

$$\max_{\psi_b} \frac{R_b(\psi_b, \tau_b)}{\sigma} - \frac{\kappa}{2} \left(\frac{\psi_b - \psi_0}{\psi_0} \right)^2. \quad (41)$$

Revenue $R_b \propto c_b(\psi_b, \tau_b)^{1-\sigma}$ depends on ψ_b through the cost chain $\psi_b \rightarrow q_b \rightarrow p_b^D \rightarrow c_b$: a higher ψ_b lowers the effective outsourced price, which lowers the buyer's unit cost and raises revenue. The buyer trades off this revenue gain against the reorganization cost of changing the specification.

Non-upgraded buyers ($\chi_b = 0$) use the industry default ψ_0 . They face no environmental wedge and therefore no additional incentive to change their specification.

First-order condition. The FOC equates the marginal revenue gain to the marginal reorganization cost:

$$\underbrace{\frac{(\sigma - 1) R_b F(\psi_b^*)}{\sigma \psi_b^*}}_{\text{marginal revenue gain from dirtier specification}} = \underbrace{\frac{\kappa(\psi_b^* - \psi_0)}{\psi_0^2}}_{\text{marginal reorganization cost}}, \quad (42)$$

where $F(\psi_b) \equiv -\partial \ln c_b / \partial \ln \psi_b > 0$ is the cost-saving elasticity:

$$F(\psi_b) = s_b^D \cdot s_b^{D,O} \cdot \underbrace{\left(1 + \frac{\text{EI}_s(\psi_b)}{\rho - 1} \right)}_{\Phi(\psi_b) > 1 \text{ when } \rho > 1}. \quad (43)$$

The left side of equation (42) is positive: dirtier outsourcing always saves cost when $\rho > 1$. The right side is zero at $\psi_b = \psi_0$ and increasing. Therefore $\psi_b^* > \psi_0$: the upgraded buyer always demands a dirtier specification than the default, limited by reorganization costs.

Comparative static. When τ_b rises, the outsourced share $s_b^{D,O}$ increases: the buyer channels more of its dirty spending through suppliers. This raises the cost-saving elasticity F in equation (43), shifting the left side of the FOC (42) upward at the initial ψ_b^* . To restore equality, ψ_b^* must increase: the buyer demands a dirtier specification.

Proposition 2. *The optimal specification satisfies $d\psi_b^*/d\tau_b > 0$: upgraded buyers demand dirtier intermediates when the environmental wedge increases.*

We define the outsourcing elasticity as:

$$\nu_d \equiv \frac{\ln(\psi_b^*(\tau)/\psi_0)}{\ln \tau} > 0, \quad (44)$$

so that $\psi_b^*(\tau) = \psi_0 \cdot \tau^{\nu_d}$ by construction.

C.4 Step-by-Step Derivation of the Comparative Static

From the buyer's problem (42) to the sign of $d\psi^*/d\tau$ (Appendix C.3)

From the buyer's problem to the FOC. The buyer chooses the specification to maximize operating profit net of reorganization costs:

$$\max_{\psi} \frac{R_b(\psi, \tau_b)}{\sigma} - \frac{\kappa}{2} \left(\frac{\psi - \psi_a}{\psi_a} \right)^2. \quad (42)$$

Differentiating the revenue term in logarithmic form,

$$\frac{1}{\sigma} \frac{\partial R_b}{\partial \psi} = \frac{R_b}{\sigma} \frac{\partial \ln R_b}{\partial \psi} = \frac{R_b}{\sigma} \cdot \frac{1}{\psi} \cdot \frac{\partial \ln R_b}{\partial \ln \psi} = \frac{(\sigma - 1)}{\sigma} R_b \frac{F(\psi, \tau_b)}{\psi},$$

where the last step uses $R_b \propto c_b^{1-\sigma}$ and the definition $F \equiv -\partial \ln c_b / \partial \ln \psi$. The FOC is therefore

$$\underbrace{\frac{(\sigma - 1)}{\sigma} R_b(\psi^*, \tau_b) \frac{F(\psi^*, \tau_b)}{\psi^*}}_{\text{MB}(\psi^*, \tau_b)} = \underbrace{\kappa \frac{\psi^* - \psi_a}{\psi_a^2}}_{\text{MC}(\psi^*)}, \quad F = s_b^D s_b^{D,O} \Phi(\psi). \quad (43)$$

Note two features for later: MB carries the revenue *level* R_b as a factor, and MC contains no revenue at all (κ is a fixed cost in units of labor).

The implicit function theorem. Write the FOC as $G(\psi, \tau) \equiv \text{MB}(\psi, \tau) - \text{MC}(\psi) = 0$. Totally differentiating,

$$\frac{d\psi^*}{d\tau} = - \frac{\partial G / \partial \tau}{\partial G / \partial \psi} = - \frac{\partial \text{MB} / \partial \tau}{\partial (\text{MB} - \text{MC}) / \partial \psi}. \quad (\text{IFT})$$

The denominator is negative by the second-order condition (at a maximum, MB crosses MC from above). Hence

$$\text{sign}\left(\frac{d\psi^*}{d\tau}\right) = \text{sign}\left(\frac{\partial \text{MB}}{\partial \tau}\right) = \text{sign}\left(\frac{\partial \ln \text{MB}}{\partial \ln \tau}\right),$$

the last equality because $\text{MB} > 0$ at an interior optimum. Since MB is a product of factors,

$$\ln \text{MB} = \ln \frac{\sigma-1}{\sigma} + \ln R_b + \ln s_b^D + \ln s_b^{D,O} + \ln \Phi(\psi) - \ln \psi,$$

so, at fixed ψ ,

$$\frac{\partial \ln \text{MB}}{\partial \ln \tau} = \underbrace{\frac{\partial \ln R_b}{\partial \ln \tau}}_{\text{Term 1}} + \underbrace{\frac{\partial \ln s_b^D}{\partial \ln \tau}}_{\text{Term 2}} + \underbrace{\frac{\partial \ln s_b^{D,O}}{\partial \ln \tau}}_{\text{Term 3}} + \underbrace{\frac{\partial \ln \Phi}{\partial \ln \tau}}_{\text{Term 4}}. \quad (\text{D0})$$

We now compute the four terms of (D0), one by one, in this order.

Objects used throughout. The wedge enters the model in exactly one place: the buyer's in-house energy price τp_E in the middle nest. The two price indices are

$$p_b^D = \left[\alpha (\tau p_E)^{1-\eta} + (1-\alpha) q_b^{1-\eta} \right]^{\frac{1}{1-\eta}}, \quad c_b = \frac{1}{z_b} \left[\gamma w^{1-\sigma_b} + (1-\gamma) (p_b^D)^{1-\sigma_b} \right]^{\frac{1}{1-\sigma_b}},$$

with $R_b = A c_b^{1-\sigma}$ (A independent of ψ and τ), and q_b the effective price of the outsourced composite, which does not depend on τ because suppliers transact at the market energy price. By Shephard's lemma, the CES expenditure shares are

$$s_b^{D,I} = \frac{\alpha (\tau p_E)^{1-\eta}}{\alpha (\tau p_E)^{1-\eta} + (1-\alpha) q_b^{1-\eta}}, \quad s_b^{D,O} = \frac{(1-\alpha) q_b^{1-\eta}}{\alpha (\tau p_E)^{1-\eta} + (1-\alpha) q_b^{1-\eta}},$$

$$s_b^D = \frac{(1-\gamma) (p_b^D)^{1-\sigma_b}}{\gamma w^{1-\sigma_b} + (1-\gamma) (p_b^D)^{1-\sigma_b}}.$$

Term 1 of (D0): the revenue derivative $\partial \ln R_b / \partial \ln \tau$.

Since $\ln R_b = \ln A + (1-\sigma) \ln c_b$, we need how τ moves the unit cost. The chain is $\tau \rightarrow p_b^D \rightarrow c_b$, in two links.

Link (a): from τ to the dirty-nest price. Take logs of the price index:

$$\ln p_b^D = \frac{1}{1-\eta} \ln \left[\alpha (\tau p_E)^{1-\eta} + (1-\alpha) q_b^{1-\eta} \right].$$

Inside the bracket only the in-house term contains τ , and it is proportional to $\tau^{1-\eta}$, so

$$\frac{\partial [\alpha (\tau p_E)^{1-\eta}]}{\partial \ln \tau} = (1-\eta) \alpha (\tau p_E)^{1-\eta}, \quad \frac{\partial [(1-\alpha) q_b^{1-\eta}]}{\partial \ln \tau} = 0.$$

Therefore

$$\frac{\partial \ln p_b^D}{\partial \ln \tau} = \frac{1}{1-\eta} \cdot \frac{(1-\eta) \alpha (\tau p_E)^{1-\eta}}{\alpha (\tau p_E)^{1-\eta} + (1-\alpha) q_b^{1-\eta}} = s_b^{D,I} : \quad (\text{D1a})$$

the $(1-\eta)$ from the exponent cancels against the $(1-\eta)$ from the derivative, leaving the in-house share.

Link (b): from the dirty-nest price to unit cost. The same computation one layer up. From $\ln c_b = -\ln z_b + \frac{1}{1-\sigma_b} \ln [\gamma w^{1-\sigma_b} + (1-\gamma) (p_b^D)^{1-\sigma_b}]$, only the second term in the bracket contains p_b^D , proportional to $(p_b^D)^{1-\sigma_b}$, so

$$\frac{\partial \ln c_b}{\partial \ln p_b^D} = \frac{1}{1-\sigma_b} \cdot \frac{(1-\sigma_b)(1-\gamma) (p_b^D)^{1-\sigma_b}}{\gamma w^{1-\sigma_b} + (1-\gamma) (p_b^D)^{1-\sigma_b}} = s_b^D. \quad (\text{D1b})$$

Combine. Chaining (D1a) and (D1b) and multiplying by $(1-\sigma)$:

$$\frac{\partial \ln R_b}{\partial \ln \tau} = (1-\sigma) \cdot \frac{\partial \ln c_b}{\partial \ln p_b^D} \cdot \frac{\partial \ln p_b^D}{\partial \ln \tau} = -(\sigma-1) s_b^D s_b^{D,I} < 0. \quad (\text{D1})$$

Intuition: the wedge raises the buyer's cost and shrinks its sales; since MB scales with the revenue level while the reorganization cost does not, this pushes the whole marginal-benefit schedule down.

Term 2 of (D0): the dirty-share derivative $\partial \ln s_b^D / \partial \ln \tau$.

Take logs of the share:

$$\ln s_b^D = \ln [(1-\gamma) (p_b^D)^{1-\sigma_b}] - \ln [\gamma w^{1-\sigma_b} + (1-\gamma) (p_b^D)^{1-\sigma_b}].$$

Differentiate each piece with respect to $\ln p_b^D$: the first is proportional to $(p_b^D)^{1-\sigma_b}$, so its derivative

is $1 - \sigma_b$; the second is the log of the cost bracket, whose derivative we already computed in (D1b), namely $(1 - \sigma_b) s_b^D$. Subtracting,

$$\frac{\partial \ln s_b^D}{\partial \ln p_b^D} = (1 - \sigma_b) - (1 - \sigma_b) s_b^D = (1 - \sigma_b)(1 - s_b^D),$$

and multiplying by $\partial \ln p_b^D / \partial \ln \tau = s_b^{D,I}$ from (D1a):

$$\frac{\partial \ln s_b^D}{\partial \ln \tau} = (1 - \sigma_b)(1 - s_b^D) s_b^{D,I}. \quad (\text{D2})$$

Intuition: the wedge makes the dirty nest more expensive. If $\sigma_b < 1$ (clean and dirty inputs are complements), a pricier dirty nest *gains* expenditure share and the term is positive; if $\sigma_b > 1$ it is negative.

Term 3 of (D0): the outsourced-share derivative $\partial \ln s_b^{D,O} / \partial \ln \tau$.

Take logs of the share:

$$\ln s_b^{D,O} = \ln[(1 - \alpha) q_b^{1-\eta}] - \ln[\alpha (\tau p_E)^{1-\eta} + (1 - \alpha) q_b^{1-\eta}].$$

The first piece contains no τ at all, so its derivative is zero. The second piece is the log of the dirty-nest bracket, whose derivative with respect to $\ln \tau$ we already computed inside (D1a): it equals $(1 - \eta) s_b^{D,I}$. Subtracting,

$$\frac{\partial \ln s_b^{D,O}}{\partial \ln \tau} = 0 - (1 - \eta) s_b^{D,I} = (\eta - 1) s_b^{D,I} > 0 \text{ for } \eta > 1. \quad (\text{D3})$$

Intuition: the wedge does not touch the numerator of $s_b^{D,O}$; it only moves the in-house component of the denominator, $\alpha (\tau p_E)^{1-\eta}$. With $\eta > 1$, costlier in-house energy shrinks that component, so the outsourced share rises: dirty spending is pushed toward suppliers. This is the outsourcing force.

Term 4 of (D0): the supplier-side derivative $\partial \ln \Phi / \partial \ln \tau$.

Φ depends only on ψ , w , and the market energy price p_E — the wedge applies to the buyer's

in-house energy, not to suppliers:

$$\Phi(\psi) = 1 + \frac{1}{\rho - 1} \cdot \frac{\psi p_E^{1-\rho}}{w^{1-\rho} + \psi p_E^{1-\rho}} \implies \frac{\partial \ln \Phi}{\partial \ln \tau} = 0. \quad (\text{D4})$$

Back to D0

Substituting (D1)–(D4) into (D0), every term carries the common factor $s_b^{D,I}$ — natural, since the wedge bites only through the in-house energy price:

$$\begin{aligned} \frac{\partial \ln \text{MB}}{\partial \ln \tau} &= -(\sigma - 1) s_b^D s_b^{D,I} + (1 - \sigma_b)(1 - s_b^D) s_b^{D,I} + (\eta - 1) s_b^{D,I} + 0 \\ &= s_b^{D,I} \left[\underbrace{-(\sigma - 1) s_b^D}_{\text{revenue}} + \underbrace{(\eta - 1)}_{\text{outsourcing}} + \underbrace{(1 - \sigma_b)(1 - s_b^D)}_{\text{dirty share}} \right]. \end{aligned} \quad (\text{D5})$$

The sign of the bracket — and hence of $d\psi^*/d\tau$ — is theoretically ambiguous, because it collects two forces of opposite sign. The *outsourcing force*, $(\eta - 1) > 0$, pushes toward a dirtier specification: costlier in-house energy channels dirty spending through suppliers, and a larger outsourced share raises the return to dirt. The *revenue force*, $-(\sigma - 1) s_b^D < 0$, pushes toward a cleaner one: the wedge raises the buyer's cost and shrinks its sales, and since MB scales with the revenue level while the reorganization cost does not, the entire marginal-benefit schedule shifts down. (The dirty-share term reinforces one side or the other depending on whether σ_b is below or above one, and is second order.) The buyer therefore demands a dirtier specification when the outsourcing margin η is strong relative to the demand elasticity σ ,

$$\frac{d\psi^*}{d\tau} \geq 0 \iff (\eta - 1) + (1 - \sigma_b)(1 - s_b^D) \geq (\sigma - 1) s_b^D,$$

and a cleaner one otherwise. The theory does not restrict which case obtains: the sign of the specification response is an empirical question.

Supplier learning. The buyer's specification ψ_b^* determines what intermediate the supplier must produce. But how the supplier produces it can improve through exposure to the buyer's environmental management system. Through audits, training, process documentation, and regular monitoring, the supplier learns energy-saving techniques—waste heat recovery, fuel switching, process optimization—that reduce the energy required to produce the demanded intermediate without

changing the specification itself. The supplier produces the same product but burns less fuel. The effective energy weight in the supplier's CES is:

$$\psi_b^{\text{eff}} = \psi_b^* \cdot \tau_b^{-\nu_l \cdot \chi_b}, \quad (45)$$

where $\nu_l \geq 0$ is the learning elasticity. Learning is energy-specific: with $\rho > 1$, lower ψ_b^{eff} means the supplier substitutes toward labor and away from energy, raising unit cost. The supplier becomes cleaner but more expensive—not more productive overall.

Derivation of the reduced form. Substituting $\psi_b^* = \psi_0 \cdot \tau_b^{\nu_d}$ from (44) into (45):

$$\psi_b^{\text{eff}} = \psi_0 \cdot \tau_b^{\nu_d} \cdot \tau_b^{-\nu_l \chi_b}. \quad (46)$$

For non-upgraded buyers ($\chi_b = 0$, $\tau_b = 1$): $\psi_b^{\text{eff}} = \psi_0$. For upgraded buyers ($\chi_b = 1$, $\tau_b = \tau$):

$$\psi_b^{\text{eff}} = \psi_0 \cdot \tau^{\nu_d - \nu_l} = \psi_0 \cdot \tau^{\nu}, \quad (47)$$

where:

$$\boxed{\nu = \nu_d - \nu_l}. \quad (48)$$

Proposition 3 (Sign of ν). *The demand-composition elasticity $\nu = \nu_d - \nu_l$ is genuinely ambiguous:*

- (i) $\nu > 0$: the buyer demands a sufficiently dirtier intermediate that, even after the supplier learns clean practices, the net specification exceeds the default. Suppliers get dirtier.
- (ii) $\nu = 0$: exact offset.
- (iii) $\nu < 0$: the supplier learns enough from the clean buyer to more than compensate for the dirtier specification. Suppliers get cleaner.

Identification. The event-study coefficient recovers $\nu = \nu_d - \nu_l$, not the individual components. This is sufficient: the structural equation, counterfactual analysis, and leakage calculations all depend only on ν . The reorganization cost parameter κ and the learning elasticity ν_l are not separately identified and are not needed for any computation. Our empirical finding $\hat{\beta}_1 > 0$ implies $\nu > 0$: the outsourcing channel dominates whatever learning takes place.

D Appendix: Quantitative Model

D.1 Solution Algorithm

This appendix gives the full solution algorithm summarized in Section 4.2. The equilibrium is a triple—an energy price p_E , an upgrading profile $\{\chi_b\}$, and an active-supplier count K_b for each buyer—together with the prices and allocations they imply. We solve it with the nested procedure in Algorithm 2, which calls the supplier-set subroutine in Algorithm 3.

Algorithm 2 Full equilibrium solver

```

1: Input: parameters; draws  $\{z_b\}, \{\tilde{a}_s\}, \{\mathcal{M}_b\}, \{\xi_{bs}\}, \{f_b^{\text{up}}\}$ ; tolerances  $\text{tol}_E, \text{tol}_\chi$ 
2: bracket the energy price in  $[p_E^-, \bar{p}_E]$ 
3: repeat ▷ outer: energy market
4:    $p_E \leftarrow$  safeguarded secant/Illinois update within the bracket (bisection fallback)
5:   initialize  $\{\chi_b\}, P, Y$ 
6:   repeat ▷ middle: upgrading fixed point
7:     for each buyer  $b$  do
8:       for  $\chi \in \{0, 1\}$  do
9:          $(K_b(\chi), \Pi_b^{\text{op}}(\chi)) \leftarrow \text{ACTIVATE\_SUPPLIERS}(b, \chi, p_E, P, Y)$ 
10:      end for
11:       $\Delta\pi_b \leftarrow \Pi_b^{\text{op}}(1) - \Pi_b^{\text{op}}(0); \chi_b \leftarrow \mathbf{1}\{\Delta\pi_b > f_b^{\text{up}}\}$ 
12:    end for
13:    recompute  $P$  from  $\{p_b\}$  and  $Y$  from the budget constraint
14:  until  $\max_b |\chi_b^{\text{new}} - \chi_b^{\text{old}}| = 0$  and  $|\Delta P|, |\Delta Y| < \text{tol}_\chi$ 
15:   $Z(p_E) \leftarrow$  aggregate energy demand  $- \bar{E} p_E^e$ 
16: until  $|Z(p_E)| < \text{tol}_E$ 
17: Output:  $p_E^*, \{\chi_b\}, S^*$ ; prices and allocations; welfare  $W = U - \text{SCC} \cdot E^{\text{total}}$ 

```

Algorithm 3 ACTIVATESUPPLIERS(b, χ, p_E, P, Y)

- 1: set $(\tau_b, \delta_b) = (\tau, \delta)$ if $\chi = 1$, else $(1, 1)$
 - 2: for each met supplier $s \in \mathcal{M}_b$, compute the delivered cost $\tilde{r}_{bs} = r_s(\psi_b) \xi_{bs}$ from (10) and (11)
 - 3: sort the meeting set so that $\tilde{r}_{b(1)} \leq \dots \leq \tilde{r}_{b(n_b)}$
 - 4: **for** $K = 0, 1, \dots, n_b$ **do** ▷ evaluate nested prefixes
 - 5: compute $p_b^M(K), p_b^D(K), c_b(K)$ from (17), (19), (21)
 - 6: $\Pi_b(K) \leftarrow \frac{1}{\sigma} \delta_b^{\sigma-1} (p_b/P)^{1-\sigma} Y$, with $p_b = \frac{\sigma}{\sigma-1} c_b(K)$
 - 7: $\Pi_b^{\text{net}}(K) \leftarrow \Pi_b(K) - K f_b$
 - 8: **end for**
 - 9: $K_b^* \leftarrow \arg \max_K \Pi_b^{\text{net}}(K)$
 - 10: **return** K_b^* and $\Pi_b^{\text{op}} = \Pi_b^{\text{net}}(K_b^*)$
-

The prefix property—that the optimal active set is the cost-ranked prefix of the meeting set—is established in Appendix C.2, which also shows why the per-supplier threshold is not a valid stopping rule, since link profitability is supermodular in the active set.

Convergence and implementation. The outer loop converges when the energy excess-demand function falls within tolerance across a shrinking bracket, with the secant/Illinois step accelerated by adaptive damping and guarded by bisection; the middle loop converges when the upgrading profile is unchanged across iterations and the aggregates are stationary. The supplier-set subroutine is the only loop over individual links, and because the prefix indices are computed in closed form and the economy is built once per equilibrium rather than re-drawn at each trial price, a full solve completes in a fraction of a second, which makes the repeated solves required by the simulated method of moments feasible.

D.2 Calibration and Moment Construction

This appendix details the construction of the calibration targets in Section 4.3.

Buyer cost shares (γ, α). Buyer variable cost is exhausted by three inputs: clean labor (the wage bill), in-house energy (firm energy expenditure), and the outsourced composite M_b . With $\text{VC} = \text{WageBill} + \text{EnergyExp} + M_b$, we set $\gamma = \text{WageBill}/\text{VC}$ and $\alpha = \text{EnergyExp}/(\text{EnergyExp} + M_b)$,

normalizing baseline price indices to one so that the CES weights equal the observed expenditure shares. We measure M_b as buyers' intermediate-input purchases through the firm-to-firm network, net of three carve-outs: energy and utility purchases (which belong in in-house energy and would otherwise be double-counted), capital goods (the model has no capital), and imported intermediates (the modeled suppliers are domestic). Within these carve-outs, M_b comprises all non-labor, non-energy input purchases.

Baseline supplier energy intensity (ψ_0). We recover ψ_0 by inverting the supplier energy-intensity equation (25) on links serving non-upgraded buyers, for which $\tau_b = 1$ and hence $\psi_b = \psi_0$:

$$\psi_0 = \frac{\text{EI}_s^0}{1 - \text{EI}_s^0} \left(\frac{w}{p_E} \right)^{1-\rho}.$$

In the model the supplier uses only labor and energy and prices at marginal cost, so its energy intensity is an energy share of value added—energy over labor plus energy. The empirical moment must be built on the same denominator: an energy-over-sales ratio understates the value-added share and biases ψ_0 downward, which is why the value in Table 6 is provisional pending the value-added rebuild. The baseline moment is the transaction-value-weighted mean supplier energy intensity over non-upgraded links, consistent with the model's value-weighted aggregation of supplier energy demand.

Environmental wedge (τ). Upgrading raises the effective price of a buyer's in-house energy to τp_E . Because the in-house-energy share within the dirty composite is $s_b^{D,H} = \alpha(\tau_b p_E / p_b^D)^{1-\eta}$, the ratio of this share between matched upgraded and non-upgraded buyers isolates $\tau^{1-\eta}$, giving the closed form (31),

$$\tau \approx \left(\frac{\bar{s}_{\text{clean}}^{D,H}}{\bar{s}_{\text{non}}^{D,H}} \right)^{\frac{1}{1-\eta}},$$

which we evaluate on propensity-score-matched buyers at the calibrated η .

Demand-composition elasticity (ν). The structural equation (29) maps the supplier energy-intensity response to connecting with an upgraded buyer into ν . Inverting the event-study coefficient $\hat{\beta}_1$ gives equation (32),

$$\hat{\nu} = \frac{\hat{\beta}_1}{\bar{\omega}^{\text{clean}} (1 - \bar{\text{EI}}_s) \ln \hat{\tau}},$$

where $\bar{\omega}^{\text{clean}}$ is the average revenue share of upgraded buyers in supplier sales, $(1 - \bar{\text{EI}}_s)$ is the elasticity of supplier energy intensity with respect to the energy weight, and $\hat{\tau}$ is the wedge above.

The baseline energy intensity here uses the same value-added denominator as ψ_0 . Because all three factors are positive, the sign of $\hat{\beta}_1$ determines the sign of ν .

Scale and energy ($L, \bar{E}$). L is total employment in the calibration sample, in thousands of workers. Observed energy expenditure is converted to numeraire units through the sample wage bill, and $\bar{E}$ is chosen so that, in the baseline equilibrium, $p_E^0 E_0$ equals $L \cdot V_E^{\text{data}} / W^{\text{data}}$, the energy account expressed in model units. The scale $\bar{E}$ is then held fixed in counterfactuals, where energy use, price, and expenditure are outcomes.

SMM moments. The certification moments are the cross-sectional ISO sales premium (a buyer-level regression of log sales on certification status, controlling for sector, region, and year), the share of buyers certified, and the standard deviation of log sales among certified buyers, where a buyer is certified if it holds ISO 14001 or ISO 50001. The dispersion of log sales within the certified set disciplines the dispersion of certification costs f_σ : when costs are tightly concentrated, certification is a near-deterministic function of firm size and the certified set occupies a narrow size band; when costs are dispersed, certified firms span a wider range of the size distribution. The network moments are network density (active links over $N_B \times N_S$), the degree–size slope (the coefficient of log buyer degree on log buyer sales), the residual spread of degree at a given size, and the mean within-buyer Herfindahl of supplier spending shares conditional on the supplier count. The productivity-tail moments are the top-decile shares of buyer and supplier sales.

Buyer–supplier split and robustness. The bipartite split is by two-digit sector, as in Section 4.1, and is used to fix N_B , N_S , and the meeting rate λ , and to define the buyer universe over which certification incidence is measured. All other moments are computed on the full in-sample universe. As a robustness check, we also report a fully sector-restricted moment set, under which the headline result is unchanged ($\nu > 0$, with $\nu \approx 0.86$).

D.3 Identification and Inference

Identification of the estimated parameters. Figure 9 plots the SMM criterion as each estimated parameter moves $\pm 30\%$ around its estimate (the productivity tails are truncated below near $\zeta - 1$, where supplier revenue ceases to have a finite mean), holding the remaining eight parameters at their estimated values. Every panel displays an interior minimum at the estimate (dashed line). The criterion rises steeply away from the estimate for the network block (λ , f_{link} , $f_{\text{link},\sigma}$), the productivity tails (θ_z , θ_a), and the demand premium δ ; the plateau to the left of the δ panel is the region $\delta \leq 1$ in which no buyer certifies. The friction-dispersion parameter θ_ξ is identified asymmetrically: the

criterion rejects lower values sharply but rises only gently above the estimate, which is reflected in its comparatively large standard error.




Figure 9: Identification: the SMM criterion around the estimate, one parameter at a time

Standard errors. The model is exactly identified (nine parameters, nine moments), so the SMM asymptotic variance reduces to

$$\widehat{\text{Var}}(\hat{\theta}) = \left(1 + \frac{1}{S}\right) G^{-1} \widehat{\Omega} (G^{-1})',$$

where G is the numerical Jacobian of the simulated moments at the estimate (two-sided differences under common random numbers, with step-size robustness checks), $S = 16$ is the number of

simulation replications, and $\hat{\Omega}$ is the sampling variance of the data moments, estimated by a nonparametric bootstrap of the calibration sample ($B = 500$; firms resampled with replacement within each moment’s sampling frame, clustering by firm in the pooled panel underlying the ISO sales premium). The certification-block moments are measured with sampling errors of 9–11% of their values—the certified set is small—which translates into the wider standard errors on $(\delta, f_\mu, f_\sigma)$ in Table 7; the network and tail moments are measured within 2%. A conservative variant in the sense of [Cocci and Plagborg-Møller \(2025\)](#), which uses only the bootstrapped marginal variances and takes the worst case over the unknown cross-moment correlations, roughly doubles the reported standard errors without changing any qualitative conclusion.

Sensitivity to moments. Table 16 reports the sensitivity matrix of [Andrews et al. \(2017\)](#), which in the exactly identified case reduces to (minus) the inverse moment Jacobian. We express it in elasticity units and in the data-moment reading: entry (j, i) is the percent change in the estimate of parameter j induced by a one-percent increase in the observed value of moment i . The matrix confirms the identification discussion of Section 4.3: each parameter loads primarily on its designated moment, with the economically expected signs—for example, a higher observed supplier top-decile share calls for a fatter supplier tail (a negative θ_a entry), and a higher observed buyer top-decile share for a fatter buyer tail—together with the expected cross-effects within the certification and network blocks. The large sensitivities of θ_ξ mirror its flat identification region in Figure 9.

Table 16: Sensitivity of the estimates to the moments ([Andrews et al., 2017](#)): percent change in each estimate per one-percent increase in each data moment

	Density	Slope	Spread	HHI	Top-10 (b)	Top-10 (s)	Premium	ISO share	S.d. log ISO
δ	0.01	0.04	−0.01	0.00	0.15	−0.43	0.16	0.01	0.03
f_μ	0.00	−0.01	−0.05	−0.04	2.95	−0.94	0.31	0.43	−2.25
f_σ	−0.06	−0.21	0.11	0.04	−1.44	2.10	−0.80	0.16	1.39
λ	0.89	0.27	1.72	1.31	1.91	0.94	−0.01	−0.01	0.07
f_{link}	−1.18	1.27	−0.28	−0.17	−4.51	−1.86	−0.02	−0.02	0.10
$f_{\text{link},\sigma}$	−0.01	−0.99	2.06	0.18	0.44	−2.13	−0.03	−0.01	0.06
θ_ξ	0.35	3.98	−6.99	−4.11	−7.24	8.12	0.21	0.10	−0.44
θ_z	−0.02	0.03	−0.02	−0.02	−1.06	−0.16	0.12	0.01	−0.03
θ_a	−0.02	0.03	−0.02	−0.02	0.08	−1.23	0.01	0.00	0.01