Production Function Estimation with Multi-Destination Firms

Geoffrey Barrows*, Hélène Ollivier†, Ariell Reshef‡

November 3, 2025

Abstract

We develop a production function estimator for the case when firms endogenously select into multiple destination markets where they compete imperfectly, and when output is denominated only in value. We exploit a novel source of variation to identify the curvature of demand: firm-level *export share of sales*. Ignoring the multi-destination dimension (i.e., exporting) yields biased and inconsistent inference. In contrast, our estimator performs well in Monte Carlo experiments. We estimate in French Manufacturing data increasing total returns to scale, decreasing returns to flexible inputs, demand elasticities between 3.6 and 19.4, and learning-by-exporting effects of 0-4% per year.

Keywords: production function, productivity, returns to scale, non-segmented markets, trade, learning by exporting
JEL Classification: F12, F63, D24

*CNRS, CREST, Ecole Polytechnique. Address: Ecole polytechnique, Departement d'Economie, Route de Saclay, 91128 Palaiseau, France. Email: geoffrey-masters.barrows@polytechnique.edu.

†Paris School of Economics - CNRS, Address: Paris School of Economics, 48 Boulevard Jourdan 74014 Paris, France. Email: helene.ollivier@psemail.eu.

‡Centre d'Économie de la Sorbonne, CNRS, Paris School of Economics, Address: Paris School of Economics, 48 Boulevard Jourdan 74014 Paris, France. Email: ariell.reshef@psemail.eu.

1 Introduction

Production function estimation is a central component of many economic analyses.[1] While early work relied on restrictive assumptions with respect to the evolution of unobserved shocks to supply, remarkable progress has been made in the last 30 years towards specifying flexible conditions under which structural elements of supply can be identified from firm or plant-level data (Olley & Pakes, 1996; Blundell & Bond, 2000; Levinsohn & Petrin, 2003; Wooldridge, 2009; Ackerberg et al., 2015; Gandhi et al., 2020). Nevertheless, significant gaps remain between the conditions assumed by the literature and the real-world datasets confronted by practitioners (De Loecker & Goldberg, 2014; De Loecker & Syverson, 2021).

A fundamental problem that is emphasized in the literature is that since outputs are usually observed in monetary terms—i.e., sales, not physical quantities—unobserved pricing decisions of firms directly influence the outcome variable, sales. In such cases, unobserved shocks to demand can generate bias in the estimation of output elasticities, even if unobserved shocks to supply are adequately controlled for (Klette & Griliches, 1996; Foster et al., 2008; De Loecker & Goldberg, 2014; De Loecker & Syverson, 2021).

When estimating production functions with outputs denominated in value, researchers usually follow one of three approaches: (1) deflate revenues with an industry-wide price deflator and treat the resulting series as if they were quantities, (2) control for an industry-wide aggregate demand shock, or (3) impose constant returns to variable inputs. It is well known that the first approach fails to neutralize omitted variable bias if firms charge different prices (De Loecker & Goldberg, 2014). The second approach yields consistent estimates under parametric restrictions on the inverse demand function, and if theory-consistent price and quantity indices are observed for *all* markets served by firms. Lacking this information, researchers typically include aggregate controls for the domestic market only (Klette & Griliches, 1996; Melitz & Levinsohn, 2006; De Loecker, 2011), which is insufficient to address the bias if firms serve multiple destinations. With the third approach, the elasticity of demand can be recovered from the ratio of variable input costs to revenues, and the elasticity of output to quasi-fixed factors can be estimated via standard GMM techniques (Aw et al., 2011; Almunia et al., 2021); but the assumption of constant returns to variable inputs is hard to defend *a priori*, and constant returns to variable inputs rules out cross-market interactions—a mechanism of interest in the literature.[2]

In this paper, we develop a procedure for estimating structural elements of supply and demand from firm-level data when outputs are only observed in monetary terms, firms serve potentially multiple markets, and returns to variable inputs are potentially non-constant. As in previous work, we substitute for missing price data with assumptions on demand and market structure. But instead of identifying the demand curvature parameter exclusively from time-series variation in aggregate quantity indices, we exploit a novel

[1] Researchers estimate production functions for a wide array of purposes. The structural coefficients can be themselves of interest, as in studies of returns to scale (Caballero & Lyons, 1992) or can be used to estimate markups (Hall, 1986). Alternatively, researchers may aim to control for estimated productivity in order to address omitted variable bias (Almunia et al., 2021), or may estimate productivity as the residual of a production function and then regress this residual on other explanatory variables in order to study the determinants of productivity (Harrigan et al., 2023). For a recent review of the literature, see De Loecker & Syverson (2021).

[2] For example, Almunia et al. (2021) propose a model in which decreasing returns to variable inputs leads firms to substitute sales away from particular markets when those markets experience downturns. Constant returns to variable inputs are inconsistent with such a mechanism. See also Arkolakis et al. (2023) for cross-market complementarities in multinational location choices.

time- and cross-sectional source of variation: the *export share of sales* at the firm level. We show that this readily-observable metric encodes valuable information about the demand conditions faced by firms on all destination markets, given standard assumptions on the profit maximizing behavior of firms. We combine this metric with information derived from the first order necessary conditions of the firm's optimization problem to build model-consistent moment conditions, and thus estimate the curvature of demand, factor output elasticities, and the elasticity of productivity to observed determinants, such as contact with foreign markets.

We start by specifying a data generating process in which heterogeneous firms serve potentially multiple destination markets. Within each destination market, firms face an isoelastic demand curve, with a price elasticity of demand that may vary by destination.[3] Firms face on each market two types of idiosyncratic demand shocks. First, there are idiosyncratic demand shocks that are revealed to the firm prior to making production plans—which we call "*ex ante*" shocks. Second, there are idiosyncratic demand shocks that are realized at the point of sale, and hence are unknown to firms at the time that they choose flexible inputs—which we call "*ex post*" shocks. If not controlled for, the *ex ante* demand shocks generate transmission bias, because they both influence flexible input choices and directly affect revenues. The *ex post* shocks do not generate transmission bias, but rationalize variation in flexible input shares in revenues across firms that use the same quantities of inputs. *Ex post* firm-destination-period shocks also generate variation in prices and markups, both across firms and within firms across destinations.[4]

In each period firms choose simultaneously—conditional on quasi-fixed inputs—the destinations to serve (which we call their "export strategy"), flexible inputs for production, and quantity allocations across destinations to maximize expected profits. While the export strategy is, in principle, the solution to a combinatorial discrete choice problem, in practice our estimation strategy does not require us to solve the optimal export strategy of the firm. As long as the first-order conditions are satisfied conditional on the export strategy, we can characterize firm behavior without imposing additional restrictions on the profit function.[5]

Exploiting nonparametric identification results from Gandhi et al. (2020), we estimate the output elasticities of flexible and quasi-fixed inputs without imposing any assumptions on the production function—in particular, about returns to variable inputs—beyond multiplicative separability between the (unobserved) productivity shifter and the rest of the production function. Estimation proceeds in two steps. First, relying on the first order necessary conditions of the optimization problem conditional on the export strategy and quasi-fixed inputs, we project the expenditure share of flexible inputs (e.g., materials) in revenues on all inputs. This "factor share" regression nonparametrically identifies the "revenue elasticity" of flexible inputs, which combines output elasticities with demand elasticities.

In the second step we use the revenue elasticities of flexible inputs from the first step regression to com-

[3]This is a more general framework than is usually employed in trade models. Researchers often allow that the elasticity of substitution varies by product, but rarely by destination within product (e.g., Broda & Weinstein 2006; Caliendo & Parro 2015; Shapiro 2016; Fontagné et al. 2022). A notable exception is Imbs & Méjean (2017), who identify destination-specific elasticities of substitution from trade flows and price data.

[4]Hence, while we impose structure on demand and market in order to circumvent missing price data, our model is flexible enough to allow for heterogeneous pricing behavior that is potentially important in real-world markets. While our model allows for these features, we cannot study strategic pricing behavior, given the data constraints.

[5]See Arkolakis et al. (2023) for details on how to solve the combinatorial discrete choice problem under additional restrictions on the profit function.

pute the contribution of flexible inputs to domestic revenues. We subtract this contribution from domestic revenues, and then regress the resulting values on predetermined inputs (e.g., capital) and a firm-specific demand proxy constructed from firm-level export revenue shares and the residual from the first stage.[6] Variation in this demand proxy allows us to estimate the curvature of demand without building aggregate demand shifters in each market from industry-wide price indices—an empirical effort plagued by thorny measurement problems that we discuss below—and without even knowing the set of destinations served by firms. The domestic demand elasticity is then used to recover output elasticities and retrieve the production function itself.

In Monte Carlo simulations, we compare the statistical properties of our estimator to (1) the estimator from Aw et al. (2011) that imposes constant returns to variable inputs, (2) the estimator from De Loecker (2011) that allows for different demand elasticities on different markets, (3) the standard practice of deflating firm-level revenues by industry-wide price indices and ignoring demand and price variation across firms, and (4) the estimator that controls for aggregate sales in the domestic market as in Klette & Griliches (1996).[7] Only our multi-destination factor share estimator recovers estimates that are close to the true parameter values in these experiments, whereas the other estimators deviate substantially when the data generating process allows for multiple markets and non-constant returns to variable inputs.

Finally, we use our estimator to study returns to scale, the elasticity of demand, and the effect of exporting on productivity in a panel of French manufacturing firms in 1994–2016. We find price elasticities of demand ranging between -19.4 and -3.6 across industries, which is consistent with estimates from the gravity literature (for example, Shapiro 2016; Fontagné et al. 2022). On the supply side, we estimate that returns to flexible inputs are less than 1, on average. Decreasing returns to flexible inputs imply negative cross-market cost complementarities, as in Almunia et al. (2021). We find overall increasing returns to scale, on average around 1.15. We also find evidence of learning by exporting (LBE) between zero and 4 percent year-on-year. These estimates imply cross-section differences in productivity between exporters and non-exporters of up to 40 percent. The model with no demand correction yields lower returns to scale, consistent with unaddressed transmission bias. The model with a single-market correction delivers unrealistic elasticities of demand (despite the relatively long period) and unrealistic returns to scale. Estimators that apply the control function approach yield extremely low capital elasticities and imprecisely-estimated price elasticity of demand.

Our main contribution is to develop a production function estimator that exploits moment conditions that are consistent with a model in which firms potentially serve multiple destination markets where they face idiosyncratic demand shocks—without relying on quantity data. When quantity data are observed, output elasticities can be estimated without imposing structure on pricing behavior (as in, for example, Roberts et al. 2018, Blum et al. 2023 and Valmari 2023). But in this case, structural assumptions are required on the supply side with respect to how firms apportion inputs across multiple production lines or products. To balance high demands on the data stemming from the existence of multiple production lines,

[6]According to our model, the residual from the first stage regression corresponds to a weighted average of unobserved *ex post* firm-destination-year demand shocks. Combining this residual with firm-level export revenue shares allows us to build a proxy for the weighted average of *ex ante* firm-specific demand shocks.

[7]For the alternatives (3) and (4), we consider both the factor share and the control function approaches.

quantity-based multi-product productivity estimators often rely on restrictive functional form (e.g., Cobb-Douglas) assumption for the production function (Blum et al., 2023; de Roux et al., 2021). Additionally, it may be difficult to compare quantities across firms and products in a meaningful way (De Loecker & Goldberg 2014).[8] Moreover, physical quantity data are only available in rare datasets, and even then only for particular industries, limiting the applicability of quantity-based multi-product estimation techniques.[9] In contrast, we offer an approach that can be implemented in a broad range of differentiated-product markets with information that is widely available.[10]

Beyond the production function estimation literature discussed above, our paper relates to a small literature on cross-market cost complementarities. Berman et al. (2015), Barrows & Ollivier (2021) and Almunia et al. (2021) all estimate the effect of demand shocks in a given market on sales in a different market. If the returns to flexible inputs are decreasing, then more supply to one market increases the cost of serving other markets. Hence, positive *ex ante* demand shocks in one market should lower sales in another market. Neither Berman et al. (2015) nor Barrows & Ollivier (2021) estimate production function parameters, but rather focus on the reduced-form connection between demand shocks in one market and sales in another market. Almunia et al. (2021) specifies a similar model to the one we develop, but adopt a production function estimation that assumes constant marginal costs and myopia relative to *ex post* demand shocks. We show in Appendix A.1 that these assumptions are not consistent with the conditions necessary for cross-market cost complementarities.

Finally, our application is related to the literature on productivity-enhancing effects of exporting (Van Biesebroeck 2005, De Loecker 2007, Wagner 2007, Wagner 2012, De Loecker 2013, Garcia-Marin & Voigtländer 2019, Atkin et al. 2017, Buus et al. 2022). When quantities are observed, then one can study learning by exporting (LBE) without explicitly modeling the export market entry decision and export market itself. However, when outputs are denominated in value it is essential to employ a multi-destination model that includes a correction for demand from multiple markets. Otherwise, an inconsistency arises between studying the impact of exporting and an estimation procedure that allows for only one, domestic, market. In fact, we do not know a paper on LBE—including those cited just above—that addresses this issue.[11]

2 Model

We model the supply and demand for horizontally differentiated varieties that are produced by single-product firms and consumed in multiple destination markets. The model delivers an estimation equation

[8] Quoting directly from De Loecker & Goldberg (2014) referring to quantity data, "the introduction of additional data creates its own challenges; although more data may help alleviate some of the problems discussed above, they are not a panacea" (page 206).

[9] For example, Blum et al. (2023) focus on only 10 3-digit ISIC industries for which there is a standard and uniform measure of physical quantities of output across firms, which yields only 2749 firms for Chilean manufacturing. Dhyne et al. (2022) analyze just firms that produce the same two 6-digit products, thereby excluding roughly 90% of firm-year observations. de Roux et al. (2021) focus only on firms producing rubber and plastic products in the Columbian manufacturing survey (covering 362 firms).

[10] CES demand and monopolistic competition are fairly standard assumptions for most manufacturing industries. In the case of highly concentrated industries, however, additional information may be required to account for oligopolistic pricing behavior.

[11] Van Biesebroeck (2005), De Loecker (2007) and De Loecker (2013) all deflate sales or value added by industry price indices—invariably, domestic price indices—in order to approximate quantities. This leaves firm-level variation in demand shocks a source of transmission bias. In addition, using domestic price indices implies that price conditions faced by exporters are identical in the domestic and foreign markets, which is at odds with the existence of variable trade barriers and different market conditions.

that links observable outputs (revenues) to inputs and observable demand shifters, and also serves as the data generating process for Monte Carlo simulations.

2.1 Demand

There are a fixed number of destination markets indexed by $d \in \{0,1,...,\mathscr{D}\}$, and industries indexed by $i \in \{1,...,\mathscr{I}\}$. Within an industry i, the quantity demanded for a given variety f at time t by consumers in destination market d takes the familiar constant elasticity of substitution (CES) form:

$$X^d_{fit} = \left(\frac{P^d_{fit}}{\Upsilon^d_{it}}\right)^{\frac{1}{\rho^d_i - 1}} B^d_{it} \exp\left(\frac{\varepsilon^d_{fit}}{1-\rho^d_i} + \frac{u^d_{fit}}{1-\rho^d_i}\right) \tag{1}$$

where P^d_{fit} is the price charged for variety f in market d in time t, Υ^d_{it} and B^d_{it} are, respectively, the CES aggregate price and quantity indices, ε^d_{fit} is an *ex ante* variety-specific demand shock (realized prior to production), u^d_{fit} is an *ex post* variety-specific demand shock (realized at the point of sales),[12] and $\eta^d_i = 1/(\rho^d_i - 1) < -1$ is a constant industry-specific price elasticity of demand, which may vary across destinations with the demand curvature parameter ρ^d_i.[13]

Revenues on each destination market d are given by:

$$R^d_{fit} = \left(X^d_{fit}\right)^{\rho^d_i} \frac{\Upsilon^d_{it}}{\left(B^d_{it}\right)^{\rho^d_i - 1}} \exp(\varepsilon^d_{fit} + u^d_{fit}). \tag{2}$$

We will estimate supply and demand parameters industry by industry, and so drop the industry index i from now on.[14]

2.2 Production

Each firm produces a single differentiated variety that may be shipped to any destination market. Hence, we can use f interchangeably to denote varieties or firms. Firms produce outputs using materials M_{ft}, labor L_{ft}, and capital K_{ft}. We assume that materials are flexible inputs, i.e., they are chosen optimally each period, given material input prices W^M_t. Capital is a quasi-fixed input that evolves each period according to the depreciation rate and an endogenous investment choice:

$$K_{ft} = (1-\theta^K)K_{f,t-1} + \iota^K_{ft-1}, \tag{3}$$

where θ^K denotes the rate of depreciation and ι^K_{ft-1} is the investment choice in period $t-1$.

[12] The *ex post* demand shock, u^d_{ft}, also captures firm-destination-time measurement error in revenues.

[13] While the majority of trade models using CES demand and monopolistic competition assumptions assume homogeneity in ρ across destinations (e.g., Broda & Weinstein 2006; Caliendo & Parro 2015; Shapiro 2016; Fontagné et al. 2022), a few papers have considered either a distinction between domestic and foreign ρs (Aw et al., 2011) or full heterogeneity (Imbs & Méjean, 2017). We adopt the most general specification.

[14] Our estimation strategy does not rule out spillover effects from supply and demand shocks in other industries, but these spillovers— if they exist— do not affect the estimation. Hence, we do not specify a top-tier aggregation function in the utility function of consumers.

We can allow for different assumptions on how labor is chosen. For settings with flexible labor markets, we can treat labor as a flexible input—just like materials—with exogenous price W_t^L. But in most settings, labor can be viewed as a dynamic input. For expositional purposes, we treat labor as a quasi-fixed input with evolution:

$$L_{ft} = (1-\theta^L)L_{f,t-1} + \iota^L_{ft-1}, \tag{4}$$

with depreciation rate θ^L and investment ι^L_{ft-1}. We point out below where the procedure needs to be adjusted to allow for alternative assumptions on labor. In particular, our French empirical context, characterized by a dual labor market (with both short, fixed-term contracts and long, indefinite-term contracts), may justify an intermediate approach: labor could be treated as a dynamic input, but not quasi-fixed. The stock of labor may evolve slowly through depreciation and investment, yet there could remain a margin that can be adjusted flexibly within a period t.[15]

Output is a function of a Hicks-neutral productivity shock ω_{ft} and a twice continuously differentiable transformation of variable and quasi-fixed inputs:

$$Q_{ft} = \exp(\omega_{ft})F(m_{ft}, l_{ft}, k_{ft}) \Longleftrightarrow q_{ft} = \omega_{ft} + f(m_{ft}, l_{ft}, k_{ft}), \tag{5}$$

with Q_{ft} (q_{ft}) denoting the quantity in levels (logs) of output produced by firm f in year t, and $F(\cdot)$ ($f(\cdot)$) is an industry-specific function expressed in levels (logs). Productivity ω_{ft} may depend on observable determinants such as export market participation, as in De Loecker (2013). Estimating the elasticity of productivity to these determinants is one objective of this paper.

2.3 Optimization

Firms may ship their variety to any destination market, subject to fixed and variable trade costs. We denote the domestic market by $d = 0$ and assume that researchers observe only firms located in that market, as is typically the case. To serve a given market, firms must pay a firm-destination-time specific fixed cost C^d_{ft} and a destination-time specific ad valorem "iceberg" cost $\tau^d_t \geq 1$.[16] To sell X^d_{ft} units to destination market d, firm f must produce $Q^d_{ft} = \tau^d_t X^d_{ft}$ units. We assume that there are no domestic fixed costs ($C^0_{ft} = 0$), so all firms sell on the domestic market. In contrast, firms may serve any subset of foreign destinations. We denote the set of *foreign* destinations served by firm f in year t as Ω_{ft}, which we refer to as an "export strategy". We remain agnostic about how this set is chosen, as this choice does not affect estimation in any way.[17]

[15]Even though the French dual labor market is known for its rigidity, many French firms adjust their current labor stock to contemporaneous supply and demand shocks, even if not completely (Saint-Paul, 1996; Reshef et al., 2022).

[16]Destination-time specific "iceberg" cost are really origin-destination-time specific costs, but since we assume that researchers observe balance sheet information for a set of firms from a single origin ($d = 0$), we omit the origin subscript, because it is constant across the sample.

[17]Presumably, firms choose the set of destinations they serve to maximize their expected profits in each period, net of the fixed costs of exporting. Solving for the optimal set of destinations to serve is a combinatorial discrete choice problem of high complexity. Under certain conditions on the profit function, this problem can be solved with the algorithm from Arkolakis et al. (2023), for instance. Alternatively, with a relatively small number of potential destinations, this problem can be solved iteratively by comparing all potential strategies and their associated profits. This is what we do in the Monte Carlo experiments.

At each period t, the sum of all units sold to all destination markets must equal total output: $\sum_{d\in 0\cup\Omega_{ft}} Q_{ft}^d = Q_{ft}$. For a general production function, this constraint obliges the firm to simultaneously determine the quantity of output to ship to each destination. Only in the case of constant returns to flexible inputs would it be optimal for the firm to solve for destination-specific quantities separately, i.e. treat the markets as segmented.

For a given export strategy Ω_{ft}, firms choose a vector of destination-specific output shares $\boldsymbol{\chi}_{ft} = (\chi_{ft}^0, \chi_{ft}^1, ..., \chi_{ft}^{\mathscr{D}})$, with $\chi_{ft}^d \equiv Q_{ft}^d / Q_{ft}$, and material inputs M_{ft} to maximize expected profits, given material input prices, quasi-fixed inputs, iceberg trade costs, and market-specific demand conditions. While firms observe *ex ante* variety-destination-specific taste shocks before making their quantity choices, they do not observe *ex post* demand shocks at the time of production. Hence, firms take expectations over *ex post* shocks u_{ft}^d, which are assumed to be i.i.d. with a constant mean u that is known to the firm.[18]

Using (2), we write the optimization problem for a given export strategy Ω_{ft} as

$$\begin{aligned}\max_{m_{ft},\chi_{ft}^0,\{\chi_{ft}^d\}_{d\in\Omega_{ft}}} \quad & \mathbb{E} \quad \Bigg[\Big(\chi_{ft}^0 \exp\left(\omega_{ft}\right) F\left(m_{ft}, l_{ft}, k_{ft}\right)\Big)^{\rho^0} D_t^0 \exp(\varepsilon_{ft}^0)\exp(u_{ft}^0) \\ & + \sum_{d\in\Omega_{ft}} \Big(\chi_{ft}^d \exp\left(\omega_{ft}\right) F\left(m_{ft}, l_{ft}, k_{ft}\right)\Big)^{\rho^d} D_t^d \exp(\varepsilon_{ft}^d)\exp(u_{ft}^d) - W_t^M M_{ft}\Bigg] \\ & + \lambda_{ft}\left(1 - \chi_{ft}^0 - \sum_{d\in\Omega_{ft}} \chi_{ft}^d\right), \end{aligned} \tag{6}$$

where $D_t^d \equiv \Upsilon_t^d \left(B_t^d\right)^{1-\rho^d} \left(\tau_t^d\right)^{-\rho^d}$ is a destination-specific demand shifter for all $d \in \{0 \cup \Omega_{ft}\}$ for the specific industry, and λ_{ft} is the Lagrangian associated to the constraint $\chi_{ft}^0 + \sum_d \chi_{ft}^d = 1$.[19] Assuming monopolistic competition implies that firms take D_t^d as given. We assume that these aggregate demand shifters are observed by firms before making their decisions.

The first-order necessary condition with respect to materials yields

$$\begin{aligned} W_t^M \quad = \quad & \mathbb{E}[\exp(u)]\rho^0 \exp\left(\rho^0\omega_{ft}\right) F\left(m_{ft}, l_{ft}, k_{ft}\right)^{\rho^0-1} \frac{\partial F}{\partial M_{ft}} \left(\chi_{ft}^0\right)^{\rho^0} D_t^0 \exp\left(\varepsilon_{ft}^0\right) \\ & + \sum_{d\in\Omega_{ft}} \mathbb{E}[\exp(u)]\rho^d \exp\left(\rho^d\omega_{ft}\right) F\left(m_{ft}, l_{ft}, k_{ft}\right)^{\rho^d-1} \frac{\partial F}{\partial M_{ft}} \left(\chi_{ft}^d\right)^{\rho^d} D_t^d \exp(\varepsilon_{ft}^d), \end{aligned} \tag{7}$$

since $\mathbb{E}\left[\exp(u_{ft}^d)\right] = \mathbb{E}\left[\exp(u)\right]$, a constant, for all firms and destinations. Rearranging yields the factor share equation

$$\frac{W_t^M M_{ft}}{R_{ft}} \quad = \quad \mathbb{E}[\exp(u)]\rho^0 \sigma_{ft}^M \left[\frac{R_{ft}^0}{R_{ft}} \exp\left(-u_{ft}^0\right) + \sum_{d\in\Omega_{ft}} \frac{\rho^d}{\rho^0}\frac{R_{ft}^d}{R_{ft}} \exp\left(-u_{ft}^d\right)\right], \tag{8}$$

[18] As in Gandhi et al. (2020), *ex post* shocks are necessary to rationalize errors in the factor share equation, which will be defined below. We extend this notion to the multi-market context.

[19] We ignore for simplicity the additional constraints $\chi_{ft}^0 \geq 0, \chi_{ft}^d \geq 0$ for all $d \in \Omega_{ft}$, with the understanding that all these constraints strictly hold. Otherwise, the set Ω_{ft} must be redefined.

where $\sigma^M_{ft} \equiv \frac{\partial F\left(m_{ft},l_{ft},k_{ft}\right)}{\partial M_{ft}} \frac{M_{ft}}{F\left(m_{ft},l_{ft},k_{ft}\right)}$ denotes the output elasticity of materials, which is allowed to vary across firms and over time. This equation serves as the basis for the first step in the estimation procedure.

We denote the term in large brackets in equation (8) by Z_{ft}:

$$Z_{ft} \equiv \frac{R^0_{ft}}{R_{ft}} \exp\left(-u^0_{ft}\right) + \sum_{d\in\Omega_{ft}} \frac{\rho^d}{\rho^0} \frac{R^d_{ft}}{R_{ft}} \exp\left(-u^d_{ft}\right),$$

which, as we can see, is a weighted sum of the inversed *ex post* shocks, where the weights are a combination of destination-specific revenue shares and the ratio of curvature parameters relative to the domestic ρ^0. As we explain below, this term plays an important role in the estimation procedure.

First-order necessary conditions with respect to output shares χ^d_{ft} yield, for each $d \in \{0 \cup \Omega_{ft}\}$,

$$\mathbb{E}[\exp(u)]\rho^d \exp(\rho^d \omega_{ft}) F\left(m_{ft}, l_{ft}, k_{ft}\right)^{\rho^d} (\chi^d_{ft})^{\rho^d - 1} D^d_t \exp(\varepsilon^d_{ft}) = \lambda_{ft}, \tag{9}$$

implying

$$(Q_{ft})^{\rho^d} (\chi^d_{ft})^{\rho^d} = \frac{\rho^0}{\rho^d} \frac{\chi^d_{ft}}{\chi^0_{ft}} (Q_{ft})^{\rho^0} (\chi^0_{ft})^{\rho^0} \frac{D^0_t \exp(\varepsilon^0_{ft})}{D^d_t \exp(\varepsilon^d_{ft})}. \tag{10}$$

Substituting (10) into the expression for total revenues gives

$$R_{ft} = (Q_{ft})^{\rho^0} \left(\chi^0_{ft}\right)^{\rho^0} D^0_t \exp(\varepsilon^0_{ft}) \exp(u^0_{ft}) \left[1 + \frac{\rho^0 \psi^x_{ft}}{\chi^0_{ft} \exp(u^0_{ft})}\right], \tag{11}$$

where $\psi^x_{ft} \equiv \sum_{d\in\Omega_{ft}} \chi^d_{ft} \exp(u^d_{ft})/\rho^d$ is a weighted sum of *ex post* shocks in foreign destinations, with weights equal to the ratio of destination-specific output shares over demand curvature.[20] For non-exporters, this expression of total (domestic) revenues simplifies with $\psi^x_{ft} = 0$ and $\chi^0_{ft} = 1$. From this expression we also observe that when firms export, their domestic output share (χ^0_{ft}) and their domestic revenue share (R^0_{ft}/R_{ft}) are not equal, and the difference depends on relative *ex post* demand shocks and demand curvature for the domestic and foreign markets.

Some manipulation of (11) yields

$$R^0_{ft} = (Q_{ft})^{\rho^0} \left(\frac{R^0_{ft}}{R_{ft} Z_{ft}}\right)^{\rho^0} D^0_t \exp(\varepsilon^0_{ft}) \left(\exp(u^0_{ft})\right)^{1-\rho^0}, \tag{12}$$

which holds for both non-exporters and exporters.[21] Equation (12) links domestic revenues, R^0_{ft}, to total

[20]The weights χ^d_{ft}/ρ^d for ψ^x_{ft} do not in general sum to 1 because $\sum_{d\in\Omega_{ft}} \chi^d_{ft} = 1 - \chi^0_{ft} \le 1$ as long as firms sell their products in the domestic market and because ρ^d is included.

[21]For exporters, we take the ratio R^0_{ft}/R_{ft} and solve for χ^0_{ft} as $\chi^0_{ft} = \frac{\rho^0 \psi^x_{ft} \exp(-u^0_{ft})}{(R_{ft}/R^0_{ft} - 1)}$. Substituting χ^0_{ft} into total revenues (11) yields, for exporters,

$$R_{ft} = (Q_{ft})^{\rho^0} \left(R_{ft}/R^0_{ft} - 1\right)^{-\rho^0} \left(\rho^0\right)^{\rho^0} D^0_t \exp(\varepsilon^0_{ft}) \left(\exp(u^0_{ft})\right)^{1-\rho^0} \left(\psi^x_{ft}\right)^{\rho^0} \left(R_{ft}/R^0_{ft}\right),$$

quantity of output, Q_{ft}, the domestic share of revenues, R^0_{ft}/R_{ft}, the weighted sum of the inverse of *ex post* shocks, Z_{ft}, and *ex ante* and *ex post* domestic demand shocks (ε^0_{ft}, u^0_{ft}).

Key to our estimation strategy, neither aggregate nor firm-specific *ex ante* foreign demand conditions intervene directly in equation (12). Instead, the metric $R^0_{ft}/(R_{ft}Z_{ft})$ encodes this information. To see how, consider two firms that produce the same quantity of output and face the same domestic demand conditions. Then the firm with higher foreign demand—either because the firm serves better markets or because the firm sees higher demand shocks—will ship a higher share of output to, and earn a higher share of revenue from, foreign markets (lower value for $R^0_{ft}/(R_{ft}Z_{ft})$). The correlation between this term and domestic revenues identifies the curvature of domestic demand. While the domestic share of revenues is likely endogenous, driven by both unobserved shocks to supply and demand, the elasticity of domestic revenues to $R^0_{ft}/(R_{ft}Z_{ft})$ can be estimated without ever having to observe foreign demand conditions, as we will show.

3 Empirical strategy

We base our estimator on the two-step factor shares approach of Gandhi et al. (2020). Specifically, we extend the "revenue production function" estimator in their Appendix O6-4 to the case where firms serve potentially multiple destination markets, wherein they face heterogeneous demand conditions.[22] Later, we compare the performance of our estimator (denoted "BOR") with alternative estimation strategies, either based on different versions of the same factor shares approach, or on the control function approach.

3.1 Multi-destination estimator: first step

In the first step, we derive the estimation equation from the first-order conditions for material inputs.[23] As implied by the term "factor share approach", output elasticities with respect to material inputs are identified from projecting factor expenditure shares on logs of input use. When outputs are denominated in value, these elasticities include the demand-side parameters.[24]

Conditioning on non-exporters and re-writing equation (8) yields:

$$\ln s^M_{ft} = \ln\left[\mathbb{E}\left[\exp(u)\right]\beta^M_{ft}(m_{ft}, l_{ft}, k_{ft})\right] - u^0_{ft} \tag{13}$$

where we define $s^M_{ft} \equiv M_{ft}W^M_t/R_{ft}$ and where $\beta^M_{ft}(m_{ft}, l_{ft}, k_{ft}) \equiv \rho^0\sigma^M_{ft}$ denotes the output elasticity of

where total revenues appear on both sides of the equation and therefore cancel out, yielding an equation for domestic revenues only. Next, we rewrite Z_{ft} as $Z_{ft} = 1/[\chi^0_{ft}exp(u^0_{ft}) + \rho^0\psi^x_{ft}]$, from which we deduce the expression of ψ^x_{ft} as $\psi^x_{ft} = \frac{R^0_{ft}/R_{ft}-1}{\rho^0 Z_{ft}R_{ft}/R^0_{ft}}$, after substituting for χ^0_{ft}. We then substitute for ψ^x_{ft} in the above equation to obtain (12).

[22] In their main text, Gandhi et al. (2020) treat the case in which outputs are denominated in quantity.

[23] An important decision when taking the factor shares method to the data is whether to classify inputs as flexible or predetermined. It is quite standard in production function applications to treat capital as a quasi-fixed input and to treat materials as a flexible input. The treatment of labor varies by application: sometimes, labor is assumed to be quasi-fixed (as we do), whereas other applications, especially in the developing world, treat labor like a flexible input. In the latter case, researchers must amend our procedure by following the same steps for labor as for materials.

[24] In general, the "revenue elasticities" of flexible inputs are identified from the factor share regressions as long as the markup does not depend on input levels (and as long as the orthogonality condition discussed below is met). But this does not require that markups are fixed. Markups may vary over time and across firms in our model because of *ex post* demand shocks.

materials multiplied by ρ^0, or the "revenue elasticity" of material inputs. It is important to condition on non-exporters for the first step to ensure that the residual (u^0_{ft}) is orthogonal to inputs choices. When the sample includes exporters the residual of the factor share regression depends on a weighted sum of *ex-post* demand shocks. While these *ex-post* demand shocks are by definition orthogonal to input choices, the weighted sum of these *ex-post* shocks is not necessarily orthogonal to input choices, because the weights are influenced both by the shocks and quasi-fixed inputs. It is important to note that conditioning on non-exporters does not introduce selection bias when estimating (13) because u^0_{ft} is unknown to the firm at the time it chooses an export strategy.

Following Gandhi et al. (2020), we approximate $\beta^M_{ft}(\cdot)$ with a complete polynomial function of degree two in all inputs. We estimate $\beta^M_{ft}(\cdot)$ by non-linear least squares (NLLS):

$$\begin{aligned}\min_{\{g^M_j\}} \quad & \sum_f \sum_t \left\{ \ln s^M_{ft} - \ln \left(g^M_0 + g^M_m m_{ft} + g^M_l l_{ft} + g^M_k k_{ft} \right.\right. \\ + \quad & \left.\left. g^M_{mm} m_{ft} m_{ft} + g^M_{ll} l_{ft} l_{ft} + g^M_{kk} k_{ft} k_{ft} + g^M_{mk} m_{ft} k_{ft} + g^M_{ml} m_{ft} l_{ft} + g^M_{lk} l_{ft} k_{ft} \right) \right\}^2 \end{aligned} \tag{14}$$

where all the g^M coefficients include the constant $\mathbb{E}\left[\exp(u)\right]$. To purge this constant from the g^M coefficients, we compute

$$\mathbb{E}\left[\widehat{\exp(u)}\right] = \frac{1}{N} \sum_f \sum_t \exp(-\widehat{u}^0_{ft}), \tag{15}$$

where N is the number of firm-year observations and $\widehat{u}^0_{ft}$ is the residual from (14).[25] We then divide all g^M coefficients by this constant and compute

$$\begin{aligned}\widehat{\beta}^M_{ft}(m_{ft}, l_{ft}, k_{ft}) \quad = \quad & \widehat{g}^M_0 + \widehat{g}^M_m m_{ft} + \widehat{g}^M_l l_{ft} + \widehat{g}^M_k k_{ft} + \widehat{g}^M_{mm} m_{ft} m_{ft} + \widehat{g}^M_{ll} l_{ft} l_{ft} \\ + \quad & \widehat{g}^M_{kk} k_{ft} k_{ft} + \widehat{g}^M_{mk} m_{ft} k_{ft} + \widehat{g}^M_{ml} m_{ft} l_{ft} + \widehat{g}^M_{lk} l_{ft} k_{ft} \end{aligned} \tag{16}$$

where, in a slight abuse of notation, we denote by $\widehat{g}^M_j$ the deflated estimated coefficients from (14). Because $\widehat{g}^M_j$ characterize the production function for all firms, we compute the estimates of $\widehat{\beta}^M_{ft}(m_{ft}, l_{ft}, k_{ft})$ for every firm, including exporters. Consequently, we can estimate Z_{ft} for all firms: $\widehat{Z}_{ft} \equiv (\frac{W^M_t M_{ft}}{R_{ft}}) / (\widehat{\beta}^M_{ft}(\cdot) \mathbb{E}\left[\widehat{\exp(u)}\right])$. We use this residual in the second stage of the estimation.

3.2 Multi-destination estimator: second step

In the second step, we use the information from the first step to recover the rest of the production function. The basic insight from Gandhi et al. (2020) is that the material input elasticity defines a partial differential equation that can be integrated to compute the part of the production function related to material inputs.

[25] Since (15) calls for the use of the exponential function, the estimate of $E\left[\widehat{\exp(u)}\right]$ may be sensitive to the presence of extreme outliers. In the French data, we exclude any firm that ever has a material expenditure share or a labor expenditure share greater than 20 or less than 0.001. This restriction excludes less than 0.1% of the data.

By the fundamental theorem of calculus, we have

$$\int_{\underline{\mathrm{m}}}^{m_{ft}} \beta_{ft}^{M}(m_{ft}, l_{ft} k_{ft}) dm_{ft} = \rho^0 f(m_{ft}, l_{ft}, k_{ft}) + \rho^0 \mathscr{C}^M(l_{ft}, k_{ft}) \tag{17}$$

where $\underline{\mathrm{m}}$ is the minimum possible value for material inputs and $\mathscr{C}^M(\cdot)$ is a constant of integration that depends only on quasi-fixed inputs (here, labor and capital). Therefore, we compute log domestic revenues for all firms net of the contribution of flexible inputs:

$$\widetilde{r}_{ft}^0 \equiv r_{ft}^0 \quad - \quad \int_{\underline{\mathrm{m}}}^{m_{ft}} \widehat{\beta}_{ft}^{M}(m_{ft}, l_{ft} k_{ft}) dm_{ft} \tag{18}$$

and substitute into (12) to derive

$$\begin{aligned} \widetilde{r}_{ft}^0 &= b_l l_{ft} + b_k k_{ft} + b_{ll} l_{ft} l_{ft} + b_{kk} k_{ft} k_{ft} + b_{lk} l_{ft} k_{ft} \\ &+ \ln D_t^0 + \rho^0 \ln\left(\frac{R_{ft}^0}{R_{ft}\widehat{Z}_{ft}}\right) + \varepsilon_{ft}^0 + \left(1-\rho^0\right) u_{ft}^0 + \rho^0 \omega_{ft} \end{aligned} \tag{19}$$

where the term $-\rho^0 \mathscr{C}^M(l_{ft}, k_{ft})$ is approximated by a complete polynomial function of degree two in quasi-fixed factors—i.e., first line of the right hand side of (19).

In equation (19), ω_{ft}, ε_{ft}^0, and u_{ft}^0 are all endogenous to $\ln[R_{ft}^0/(R_{ft}\widehat{Z}_{ft}])$, both through the endogenous choice of destinations in the export strategy and through realized sales in the domestic market. Additionally, if ω_{ft}, ε_{ft}^0, and u_{ft}^0 are persistent, then they correlate with all quasi-fixed inputs through the investment rule. As in many production function estimation routines, we exploit assumptions on the evolution of these unobserved shocks for identification. In particular, we assume that productivity shocks ω_{ft} and domestic demand shocks ε_{ft}^0 evolve according to first order Markov processes with the same persistence parameter h.[26] In addition, we allow ω_{ft} to depend on a lagged firm decision, for example, an export participation indicator $e_{f,t-1}$ (as in De Loecker 2013). Combining all unobservables into a composite shock, $\nu_{ft} \equiv \varepsilon_{ft}^0 + \left(1-\rho^0\right) u_{ft}^0 + \rho^0 \omega_{ft}$, we then have

$$\nu_{ft} = h\nu_{f,t-1} + \rho^0 \mu e_{f,t-1} + \xi_{ft}, \tag{20}$$

where $\omega_{ft} = h\omega_{f,t-1} + \mu e_{f,t-1} + \widetilde{\omega}_{ft}$ and $\varepsilon_{ft}^0 = h\varepsilon_{f,t-1}^0 + \widetilde{\varepsilon}_{ft}^0$, with $\widetilde{\omega}_{ft}$ and $\widetilde{\varepsilon}_{ft}^0$ indicating i.i.d. innovations to productivity and domestic demand shock. The parameter μ indicates the effect of "learning by exporting" (LBE), i.e., the potential gains in productivity from selling to foreign markets. Note that we make no assumptions on the evolution of ε_{ft}^d for $d > 0$, i.e., on any other market that is not the domestic one.[27] The term $\xi_{ft} \equiv \widetilde{\varepsilon}_{ft}^0 + \rho^0 \widetilde{\omega}_{ft} + (1-\rho^0) u_{ft}^0 - h(1-\rho^0) u_{ft-1}^0$ is an MA(1) error term (due to the presence of u_{ft-1}^0).

[26] Allowing for non-unit-root persistence in both productivity and domestic demand shocks entails that both ω_{ft} and ε_{ft}^0 must follow the same linear AR(1), as explained by Melitz & Levinsohn (2006); De Loecker (2011). Alternatively, one could impose that ε_{ft}^0 is i.i.d., conditional on observables, in which case one could allow for a more flexible first-order Markov process in ω_{ft}.

[27] The LBE coefficient μ is mediated by the curvature of domestic demand ρ^0. The reason is that (20) is identified by domestic revenues. An increase in (log) output increases domestic revenues by ρ^0. LBE increases output by μ, so domestic revenues by $\rho^0 \times \mu$. Had we been using revenues from a different market to identify LBE, say, market $d > 0$, then the coefficient to $e_{f,t-1}$ would have been $\rho^d \times \mu$.

Substituting ν_{ft} into (19) yields

$$\tilde{r}^0_{ft} = b_l l_{ft} + b_k k_{ft} + b_{ll} l_{ft} l_{ft} + b_{kk} k_{ft} k_{ft} + b_{lk} l_{ft} k_{ft} + \rho^0 \ln\left(\frac{R^0_{ft}}{R_{ft}\widehat{Z}_{ft}}\right) + \alpha_t + \nu_{ft} \tag{21}$$

where α_t is a time fixed effect that absorbs the domestic demand shifter, $\ln D^0_t$.[28]

For any candidate vector $(\rho^{0^*}, b^*_l, b^*_k, b^*_{ll}, b^*_{kk}, b^*_{lk})$, we can compute

$$\widehat{\nu_{ft} + \alpha_t} \quad = \quad \tilde{r}^0_{ft} - \rho^{0^*} \ln\left(\frac{R^0_{ft}}{R_{ft}\widehat{Z}_{ft}}\right) - b^*_l l_{ft} - b^*_k k_{ft} - b^*_{ll} l_{ft} l_{ft} - b^*_{kk} k_{ft} k_{ft} - b^*_{lk} l_{ft} k_{ft}, \tag{22}$$

and then regress $\widehat{\nu_{ft} + \alpha_t}$ on $\nu_{f,t-1} \widehat{+ \alpha}_{t-1}$, the past exporting decision $e_{f,t-1}$ and time fixed effects, and compute the residual $\widehat{\xi}_{ft}\left(\rho^{0^*}, b^*_l, b^*_k, b^*_{ll}, b^*_{kk}, b^*_{lk}\right)$. We then build the following moment conditions:

$$E\left\{\widehat{\xi}_{ft}\left(\rho^{0^*}, b^*_l, b^*_k, b^*_{ll}, b^*_{kk}, b^*_{lk}\right)\begin{pmatrix}\ln\left(\frac{R^0_{f,t-2}}{R_{f,t-2}\widehat{Z}_{f,t-2}}\right)\\ l_{ft}\\ k_{ft}\\ ll_{ft}\\ kk_{ft}\\ lk_{ft}\end{pmatrix}\right\} = 0 \tag{23}$$

and minimize deviations from these moments by GMM.

At the true parameter values, $\widehat{\xi}_{ft}$ is orthogonal to all quasi-fixed inputs in period t. This follows from the fact that $\widehat{\xi}_{ft}$ contains only period t innovations to productivity $\tilde{\omega}_{ft}$ and domestic demand $\tilde{\varepsilon}^0_{ft}$, and *ex post* domestic demand shocks u^0_{ft} and $u^0_{f,t-1}$, none of which influence the investment decision in period $t-1$. However, even at the true parameter values, $\widehat{\xi}_{ft}$ correlates with $\ln\left(\frac{R^0_{ft}}{R_{ft}\widehat{Z}_{ft}}\right)$ and $\ln\left(\frac{R^0_{f,t-1}}{R_{f,t-1}\widehat{Z}_{f,t-1}}\right)$ through the endogenous set of destinations and through sales on the domestic market.[29] Thus, we use $\ln\left(\frac{R^0_{f,t-2}}{R_{f,t-2}\widehat{Z}_{f,t-2}}\right)$, which is orthogonal to $\widehat{\xi}_{ft}$, to ensure that the moment condition hold.

Even though this moment condition holds, one may wonder whether it is relevant for identifying ρ^0 in (23). We offer two explanations for the correlation between $\ln\left(\frac{R^0_{f,t-2}}{R_{f,t-2}\widehat{Z}_{f,t-2}}\right)$ and $\ln\left(\frac{R^0_{ft}}{R_{ft}\widehat{Z}_{ft}}\right)$, conditional on quasi-fixed inputs, time fixed effects, and $\nu_{f,t-1} \widehat{+ \alpha}_{t-1}$. First, persistence in the *ex ante* foreign demand shocks ε^d_{ft}, for $d \neq 0$, yields correlation between $\ln\left(\frac{R^0_{f,t-2}}{R_{f,t-2}\widehat{Z}_{f,t-2}}\right)$ and $\ln\left(\frac{R^0_{ft}}{R_{ft}\widehat{Z}_{ft}}\right)$. To see this, consider two firms with the same evolution of ν_{ft} (which includes only domestic demand shocks ε^0_{ft}) and quasi-fixed inputs, serving the same set of destinations Ω_{ft}. Suppose that the first firm has persistently higher draws for ε^d_{ft} for some $d > 0$ than the second firm. The former firm will tend to earn a higher share of revenues from

[28] By absorbing the (industry-specific) time-varying demand shifter into a fixed effect, we exploit different variation than estimators that rely on time series variation in a domestic aggregate quantity index, such as Klette & Griliches (1996).

[29] Recall that in order to build $\ln\left(\frac{R^0_{f,t-1}}{R_{f,t-1}\widehat{Z}_{f,t-1}}\right)$ we use realized domestic sales, which are directly affected by $u^0_{f,t-1}$.

the export market than the latter, and hence will tend to have lower $\ln\left(\frac{R^0_{ft}}{R_{ft}\widehat{Z}_{ft}}\right)$ in all periods.

The second mechanism that generates correlation between $\ln\left(\frac{R^0_{f,t-2}}{R_{f,t-2}\widehat{Z}_{f,t-2}}\right)$ and $\ln\left(\frac{R^0_{ft}}{R_{ft}\widehat{Z}_{ft}}\right)$ relies on persistent firm-specific fixed costs of market entry. To see this, consider two firms with the same evolution of ν_{ft} and quasi-fixed inputs, but different fixed costs of reaching different markets. In this case, the two firms will likely serve different markets. If these fixed costs are persistent, then the two firms will be exposed to different aggregate shocks. Suppose that the first firm has lower fixed costs compared to the second firm for serving a particular large foreign market. Then the former firm will tend to earn more from exporting than the latter, all else equal, and thus will tend to have a lower $\ln\left(\frac{R^0_{ft}}{R_{ft}\widehat{Z}_{ft}}\right)$ in all periods.[30] While both mechanisms give rise to a "relevant" moment condition, we remain agnostic about their relative importance.

Finally, we compute the output elasticity for each quasi-fixed input:[31]

$$\widehat{\sigma}^K_{ft} = \frac{1}{\widehat{\rho}^0}\left(\frac{\partial \widehat{r}^0_{ft}}{\partial k_{ft}} + \frac{\partial}{\partial k_{ft}}\left[\int \widehat{\beta}^M_{ft}(m_{ft}, l_{ft}k_{ft})dm_{ft}\right]\right) \quad (24)$$

$$\widehat{\sigma}^L_{ft} = \frac{1}{\widehat{\rho}^0}\left(\frac{\partial \widehat{r}^0_{ft}}{\partial l_{ft}} + \frac{\partial}{\partial l_{ft}}\left[\int \widehat{\beta}^M_{ft}(m_{ft}, l_{ft}k_{ft})dm_{ft}\right]\right), \quad (25)$$

and for flexible inputs:

$$\widehat{\sigma}^M_{ft} = \widehat{\beta}^M_{ft}(m_{ft}, l_{ft}k_{ft})/\widehat{\rho}^0. \quad (26)$$

We compute returns to scale as the sum of output elasticities from variable and quasi-fixed inputs, and LBE as the point estimate on the lagged export indicator from the regression of the Markov process, all deflated by $\widehat{\rho}^0$. As a result, for our estimation procedure, knowledge of just the domestic elasticity of demand is sufficient to estimate factor output elasticities, returns to scale, and elasticities of productivity to observables, such as learning-by-exporting. Since the second step uses estimated objects from the first step, we bootstrap the entire two-step procedure to compute standard errors. The bootstrap procedure samples firms rather than individual observations, which is akin to clustering standard errors by firm.

[30] Alternatively, we could exploit a shift-share instrument instead of $\ln\left(\frac{R^0_{f,t-2}}{R_{f,t-2}\widehat{Z}_{f,t-2}}\right)$ in (23), where the weights would be pre-period market shares for each firm and the shocks would reflect industry-destination-period demand B_t. However, this would require knowledge of the entire destination network of each firm and measures of aggregate demand. We prefer to use $\ln\left(\frac{R^0_{f,t-2}}{R_{f,t-2}\widehat{Z}_{f,t-2}}\right)$ as the instrument because it requires only knowledge of the domestic share in revenues and it allows for persistence in the ε^d_{ft} draws to contribute to the relevance of the instrument.

[31] Specifically, assuming a second degree polynomial function for both the first step and the second step yields

$$\widehat{\sigma}^K_{ft} = \left(\widehat{b}_k + 2\widehat{b}_{kk}k_{ft} + \widehat{b}_{kl}l_{ft} + \widehat{g}^M_k m_{ft} + 2\widehat{g}^M_{kk}m_{ft}k_{ft} + \widehat{g}^M_{lk}l_{ft}m_{ft} + \frac{1}{2}\widehat{g}^M_{mk}m_{ft}m_{ft}\right)/\widehat{\rho}^0$$

$$\widehat{\sigma}^L_{ft} = \left(\widehat{b}_l + 2\widehat{b}_{ll}l_{ft} + \widehat{b}_{kl}k_{ft} + \widehat{g}^M_l m_{ft} + 2\widehat{g}^M_{ll}m_{ft}l_{ft} + \widehat{g}^M_{lk}k_{ft}m_{ft} + \frac{1}{2}\widehat{g}^M_{ml}m_{ft}m_{ft}\right)/\widehat{\rho}^0.$$

4 Existing Estimators

In this section we present existing approaches to estimating production functions when output is denominated in value and discuss the biases that arise from ignoring the multi-destination dimension. The reader uniquely interested in the performance of our multi-destination estimator can skip this section and proceed directly to Sections 5.

4.1 Factor share method with no demand correction

When estimating production function parameters with data denominated in value, researchers typically deflate firm-level revenues by the domestic price deflator and treat the resulting series as if they were quantities. Given our model, this would convert firm-level revenues into firm-level quantities only if firms serve only the domestic market, where they face perfect competition $\rho^0 = 1$. In this case, the conditions discussed in the main text of Gandhi et al. (2020) would be met, and thus their factor shares estimator could be applied. By contrast, when goods are not perfect substitutes ($\rho^0 < 1$), or when firms sell to multiple markets, deflating revenues by the domestic price index and implementing the estimation procedure from Gandhi et al. (2020) will bias the estimator. We present the main steps involved when implementing this estimator and discuss the source of the bias.

The factor shares estimator from Gandhi et al. (2020) follows two steps. The first step proceeds exactly as we describe in section 3.1, except that the entire sample of firms is used to estimate the factor share equation. In the second step, the contribution from materials and the residual are computed exactly as described in section 3.2. Both terms are subtracted off from deflated revenues, and the resulting value, $\widetilde{r}_{ft}^{NoD}$ (indexed by NoD for "no demand correction"), is projected on polynomial functions of the quasi-fixed inputs

$$\widetilde{r}_{ft}^{NoD} \quad = \quad b_l^{NoD} l_{ft} + b_k^{NoD} k_{ft} + b_{ll}^{NoD} l_{ft} l_{ft} + b_{kk}^{NoD} k_{ft} k_{ft} + b_{lk}^{NoD} l_{ft} k_{ft} + \nu_{ft}^{NoD}, \tag{27}$$

where ν_{ft}^{NoD} collects all unobserved shocks. For a candidate vector of parameters, the structural residual ξ_{ft}^{NoD} can be computed, moments evaluated, and parameters chosen to minimize the usual GMM objective function. Compared with the moments described in (23), only the moment combining the structural residual with $\ln\left(\frac{R_{f,t-2}^0}{R_{f,t-2}\widehat{Z}_{f,t-2}}\right)$ is omitted.

If the true data generating process (DGP) coincides with the model from Section 2, but the conditions ensuring that firms sell at the same price are not met, then the moment conditions for the "no demand correction" estimator fail at two junctures. First, if firms serve multiple destination markets with heterogeneous demand curvatures (ρ^ds), and all firm-year observations are included in the first-stage NLLS, then correlation between Z_{ft} and inputs leads to a violation of the moment condition. Our procedure circumvents this problem by conditioning on non-exporters in the first stage.

Second, even if the contribution of materials could be accurately estimated and subtracted off from deflated revenues, the endogenous choice of destination-specific quantity shares leads to a violation the moment conditions in the second step GMM. This can be seen from (11), where, after moving the contribution of materials to the left-hand side, both χ_{ft}^0 and ψ_{ft}^x directly affect revenues, and are included in the residual ξ_{ft}^{NoD} in (20). Since destination-specific output shares depend on quasi-fixed input levels, failure to control for χ_{ft}^0 and ψ_{ft}^x implies a violation of the second step moment conditions.[32]

The violation causes biases in ways that are hard to assess and likely depend on parameter values. For example, χ_{ft}^0 depends negatively on quasi-fixed inputs, since higher quasi-fixed input levels lead to lower marginal costs, higher marginal revenues, higher likelihood of exporting to any given destination, and hence higher export share. But it is not clear whether leaving χ_{ft}^0 for the error term leads to upward or downward bias in the b coefficients, because the residual in (11) may depend either positively or negatively on χ_{ft}^0.[33]

4.2 Factor share method with a single-market correction

The few papers that explicitly address the value-versus-quantity distinction in the context of production function estimation typically implement some versions of Klette & Griliches (1996)'s procedure that controls for the domestic aggregate CES quantity index (e.g., De Loecker 2011 and Grieco et al. 2016).[34] We present first the version using the factor share approach that follows from Appendix O6-4 in Gandhi et al. (2020), and then consider the version adapted from the control function approach.

The first step NLLS estimation is exactly the same as in section 4.1. The second step also proceeds exactly as in section 4.1, except that an additional regressor denoted by $\ln B_t^{proxy}$ is included on the right hand side of the deflated revenues equation:

$$\widetilde{r}_{ft}^{SMC} = b_l^{SMC} l_{ft} + b_k^{SMC} k_{ft} + b_{ll}^{SMC} l_{ft} l_{ft} + b_{kk}^{SMC} k_{ft} k_{ft} + b_{lk}^{SMC} l_{ft} k_{ft} + \left(1-\rho^0\right) \ln B_t^{proxy} + v_{ft}^{SMC}, \tag{28}$$

where variables and parameters are indexed by SMC for "single-market correction". Construction of $\widehat{\xi}_{ft}^{SMC}$ proceeds as in section 4.1, accounting for the extra regressor, moment conditions are constructed, and parameters chosen to minimize the objective function.

This additional regressor, B_t^{proxy}, captures the aggregate domestic demand conditions. If firms serve only one destination market—the domestic one—and if researchers could compute the empirical price index and B_t^{proxy} in a theory-consistent way (see Appendix A.2 for details), then we would have $B_t^{proxy} = B_t/\Upsilon_0$, where Υ_0 captures the price index normalization. In this case, the constructed residual $\widehat{\xi}_{ft}^{SMC}$ would be orthogonal to quasi-fixed inputs in period t because $\xi_{ft}^{SMC} \equiv \widetilde{\varepsilon}_{ft} + \rho\widetilde{\omega}_{ft}$ at the true parameter values. Indeed, when all firms serve a single market, the *ex post* demand shock u_{ft} is identified by the factor share regression in the first step, and hence does not appear in the second step. Moreover, the aggregate demand shifter $\ln B_t^{proxy}$ is orthogonal to $\widehat{\xi}_{ft}^{SMC}$ by assumption. Hence, the parameter ρ^0 would be identified by time series variation in industry-wide demand aggregates, and this estimation procedure would also identify all output elasticities.

However, when firms select endogenously into multiple destination markets, several sources of bias

[32] In addition, χ_{ft}^0 cannot follow the same AR(1) as ε_{ft}^0 since χ_{ft}^0 depends directly on ε_{ft}^0. Thus, χ_{ft}^0 cannot be absorbed into v_{ft}.

[33] If there is only one quasi-fixed input that enters linearly in (27), and the residual depends positively on χ_{ft}^0, then the bias is clearly negative. With multiple quasi-fixed inputs and higher order terms and interactions, it is not clear that omitting χ_{ft}^0 leads to downward bias in all estimated b terms.

[34] While De Loecker (2011) uses a control function method, Grieco et al. (2016) make strong structural assumptions (i.e., constant returns to scale, CES production function) that simplify their estimator at the cost of generality.

arise. First, unless χ^0_{ft} follows exactly the same AR(1) process as ε^0_{ft} and ω_{ft}—which is not possible, given the model—then $\widehat{\xi}^{SMC}_{ft}$ includes χ^0_{ft}. Since χ^0_{ft} depends on quasi-fixed inputs, this implies a violation of the moment conditions, which results in a biased estimator.

Second, the demand shifter B^{proxy}_t is likely measured with error. Gandhi et al. (2020) cite De Loecker (2011) for how to construct B^{proxy}_t from the data as a weighted sum of deflated *total* revenues of domestic firms. We show in Appendix A.2 that if the price deflator is constructed in a theory-consistent way, then the domestic quantity index B^0_t can be computed up to a normalization from price deflators and *total domestic absorption*, i.e. total domestic sales of domestic firms plus total imports from foreign firms. Incorrectly building B^{proxy}_t from total revenues of domestic firms (whether using weights or not) results in a discrepancy with B^0_t, that corresponds to the trade deficit, which would be relegated to the error term, multiplied by $1-\rho^0$. The trade deficit may be positively or negatively correlated with B^{proxy}_t, depending on whether local demand shocks or foreign supply shocks dominate. This, too, leads to a violation of the moment conditions and a biased estimator.

In light of these concerns, a tempting strategy would be to estimate the factor shares approach with a single-market correction for a set of non-exporters. For non-exporters, the first stage factor share regression identifies the revenue elasticity of flexible inputs as well as u^0_{ft}, the residual from this regression. Additionally, for all non-exporting firms, $\chi^0_{ft} = 1$ and $Z_{ft} = \exp(-u^0_{ft})$. In this case, equation (21) becomes

$$\begin{aligned} \widetilde{r}^{SMC}_{ft} &= b^{SMC}_l l_{ft} + b^{SMC}_k k_{ft} + b^{SMC}_{ll} l_{ft} l_{ft} + b^{SMC}_{kk} k_{ft} k_{ft} + b^{SMC}_{lk} l_{ft} k_{ft} \\ &+ \left(1-\rho^0\right) \ln B^{proxy}_t + \varepsilon^0_{ft} + \rho^0 \omega_{ft}, \end{aligned} \tag{29}$$

where $\varepsilon^0_{ft} + \rho^0 \omega_{ft}$ combine as an error term that follows an AR(1) process. However, bias persists for two reasons. First, B^{proxy}_t is still likely measured with error. Second, sample selection bias violates the orthogonality conditions in the second step: the residual from the AR(1) process does not have a zero mean conditional on quasi-fixed inputs. If higher levels of quasi-fixed inputs are associated with a greater probability to export—for instance due to increasing returns to scale—then the conditional mean of the residual will be negatively correlated with them because the sample never admits exporters. The direction of the bias will vary by type of cross-market complementarities and by the degree of returns to variable inputs.

Whether estimating the single-market demand-side correction model in the full sample or in a subsample of non-exporters, two additional problems arise. First, identification in this model relies on time-series variation in aggregate demand, which may not be sufficient in short panels. Second, given our application—the productivity effects of learning by exporting—using a method with a single-market correction entails an additional conceptual issue: there is no exporting in a single market model. There is a logical inconsistency in relying on an estimator that assumes only one, domestic destination to study the effect of serving different, foreign markets.

4.3 Control function versions of the standard approaches

The two preceding estimation models could also be estimated by the control function method, wherein unobserved productivity shocks are proxied by a model-inversion procedure. There are many ways to implement the control function approach. Here we follow the gross output control function estimator from Gandhi et al. (2020). Using B_t^{proxy} to identify the demand curvature parameter yields an estimator that is very similar to the homogeneous-ρ version of De Loecker (2011) or to the original demand-side correction model from Klette & Griliches (1996).

To fix ideas, we first invert the material demand function to substitute for the unobserved shocks to supply and demand. This substitution generates the estimation equation

$$r_{ft}^{CF} \quad = \quad \Phi(m_{ft}, l_{ft}, k_{ft}) + \delta_t + \vartheta_{ft}^{CF} \tag{30}$$

where $\Phi(m_{ft}, l_{ft}, k_{ft})$ is an unknown function that combines $f(\cdot)$ and the inverted material demand function and δ_t is a time fixed effect that controls for material input prices. When controlling for a single-market demand shifter, we simply include B_t^{proxy} in $\Phi(\cdot)$. We approximate $\Phi(\cdot)$ with a second-degree polynomial function, estimate by OLS, and extract the residual $\widehat{\vartheta}_{ft}^{CF}$.

Next, we subtract $\widehat{\vartheta}_{ft}^{CF}$ from deflated revenues. In the case that the demand-side correction is ignored we obtain

$$\widetilde{r}_{ft}^{CF-NoD} = f(m_{ft}, l_{ft}, k_{ft}) + \nu_{ft}. \tag{31}$$

For the case in which a single-market demand correction is employed we obtain

$$\widetilde{r}_{ft}^{CF-SMC} = \rho^0 f(m_{ft}, l_{ft}, k_{ft}) + \left(1 - \rho^0\right) \ln B_t^{proxy} + \nu_{ft}, \tag{32}$$

where ν_{ft} collects unobserved shocks to supply and demand. We adopt a complete polynomial function of degree 2 to approximate $f(\cdot)$ and assume that ν_{ft} follows an AR(1) process. For any candidate vector of parameters, we can compute $\widehat{\nu}_{ft}$, then regress $\widehat{\nu}_{ft}$ on $\widehat{\nu}_{f,t-1}$ and $e_{f,t-1}$ (the indicator for exporting) and compute the residual $\widehat{\xi}_{ft}^{CF}$. We build moment conditions by multiplying $\widehat{\xi}_{ft}^{CF}$ by the levels of all quasi-fixed inputs, lags of flexible inputs, along with the appropriate interaction and square terms, as well as $\ln B_t^{proxy}$ in the case of the single-market correction model. We then minimize the objective function to estimate parameters.

Beyond the model misspecification issues mentioned above, there are two problems highlighted by the literature that suggest that neither of these control function estimators are likely to generate consistent estimates of model parameters. First, the control function method relies on time series variation in material input prices for identification. Gandhi et al. (2020) show that the control function approach generates inconsistent estimates when material price variation is low, even if the moment conditions hold. This is because low material price variation renders its lag to be a weak instrument for contemporaneous materials demand, conditional on productivity and capital. We replicate this finding in Appendix B in our Monte Carlo simulations (extending the Monte Carlo experiments from Gandhi et al. (2020) to the case of imperfect competition) for a single-market version of the model.

Second, even when the sample size goes to infinity, the GMM objective function admits multiple solutions in the standard control function framework, as demonstrated by Ackerberg et al. (2023). Ackerberg et al. (2023) argue that choosing among these candidate solutions is not as simple as just choosing the parameter combination that yields the lowest objective function value. This is because there are, in fact, multiple parameter vectors for which the moment conditions are satisfied and the objective equals zero—among them the OLS parameter vector. Hence, the optimization problem is under identified.[35] Despite these issues, we estimate control function versions of the no-demand-correction and the single-market correction estimators both in Monte Carlo experiments and in the French data, for comparison.

4.4 Multi-destination estimator with constant returns to flexible inputs

Researchers in the international trade literature often estimate the elasticity of demand from the ratio of revenues to total variable costs (Aw et al. 2011, Almunia et al. 2021). We can derive the estimating equations for this strategy from our general model by imposing particular parametric restrictions.

To illustrate, we consider the estimation strategy from Aw et al. (2011). Starting from the model in section 2, if we impose (i) the same demand curvature for all foreign destinations: $\rho^d = \rho^x$ for all $d = 1, \ldots \mathscr{D}$, (ii) a Cobb-Douglas production function in materials and capital, (iii) constant returns to materials $\sigma^M_{ft} = 1$, and (iv) the absence of *ex post* demand shocks: $u^d_{ft} = 0$ for all d, then the factor-share equation (8) simplifies to

$$W^M_t M_{ft} \quad = \quad \rho^0 R^0_{ft} + \rho^x R^x_{ft}, \tag{33}$$

with $R^x_{ft} \equiv \sum_{d \in \Omega_{ft}} R^d_{ft}$. Adding a regression error term, equation (33) can be estimated by OLS to recover ρ^0 and ρ^x, which is precisely what Aw et al. (2011) do.

Next, Aw et al. (2011) estimate returns to capital and LBE. To do so, they derive an expression for R^0_{ft} which depends on just K_{ft}, ω_{ft}, W^M_t, and the domestic aggregate demand shifter D^0_t. The key to identification is the assumption of constant returns to materials, under which materials drops out of the expression. To derive this equation from our general model, we impose no *ex post* shocks and $\rho^d = \rho^x$ for all $d > 1$ and combine (10) with (8) to obtain

$$W^M_t M_{ft} \quad = \quad \sigma^M_{ft} \rho^0 \left(Q_{ft}\right)^{\rho^0} \left(\chi^0_{ft}\right)^{\rho^0} D^0_t \exp\left(\varepsilon^0_{ft}\right) \left[1 + \frac{1-\chi^0_{ft}}{\chi^0_{ft}}\right], \tag{34}$$

which implies $\chi^0_{ft} = \rho^0 R^0_{ft} / (W^M_t M_{ft})$ given the assumption $\sigma^M_{ft} = 1$. Next, we substitute this expression into the definition of domestic revenues using the Cobb-Douglas assumption with $\sigma^M_{ft} = \gamma^M_{ft} = 1$:

$$R^0_{ft} \quad = \quad \left(M_{ft} K^{\gamma^K}_{ft} \exp\left(\omega_{ft}\right)\right)^{\rho^0} \left(\frac{\rho^0 R^0_{ft}}{M_{ft} W^M_t}\right)^{\rho^0} D^0_t \exp\left(\varepsilon^0_{ft}\right), \tag{35}$$

where γ^j indicates the output elasticity of factor $j \in \{M, K\}$, and where materials M_{ft}—present both in Q_{ft}

[35] It is well known that nonlinear estimation like GMM can be sensitive to initial values as well as searching algorithms (Knittel et al., 2014). As shown by Ackerberg et al. (2023), the problem with the control function method is more severe than mere numerical challenges.

and in χ_{ft}^{0}—cancel out. Rearranging yields

$$R_{ft}^{0} = \left(K_{ft}\right)^{\frac{\gamma^{K}\rho^{0}}{1-\rho^{0}}}\left(W_{t}^{M}\right)^{\frac{-\rho^{0}}{1-\rho^{0}}}\left(\rho^{0}\right)^{\frac{\rho^{0}}{1-\rho^{0}}}\left(\exp\left(\omega_{ft}\right)\right)^{\frac{\rho^{0}}{1-\rho^{0}}}\left(D_{t}^{0}\exp\left(\varepsilon_{ft}^{0}\right)\right)^{\frac{1}{1-\rho^{0}}}, \tag{36}$$

which can be log-linearized. Aw et al. (2011) proceed to identify γ^{K} and LBE via GMM from this equation, exploiting the AR(1) assumption in the usual way.

This illustrates that the estimation strategy from Aw et al. (2011) (i.e., estimating (33) and then (36)) is a special case of our multi-destination estimator, subject to restrictions (i)-(iv) just above. If these restrictions do not hold, we would not expect this procedure to deliver consistent estimates of model parameters, but we implement it in Monte Carlo experiments and in the French data, for comparison.

4.5 Estimator with multiple aggregate quantity indices

De Loecker (2011) provides an alternative extension of Klette & Griliches (1996)'s model and (similarly to us) proposes an estimation strategy for the case in which outputs are observed only in value. There is just a single destination (Europe) and a single industry (Textile & Apparel) in De Loecker (2011), but there are multiple product markets (e.g., yarns, clothing, interior fabrics, etc.)—which De Loecker (2011) calls "segments"—so the model has a similar structure to the multi-destination model we propose. Hence, *modulo* some adjustments to translate the multi-segment dimension into a multi-destination dimension, and some additional restrictions on the production function and the demand shocks, we could also use the approach from De Loecker (2011) to estimate returns to scale, curvature of demand, and LBE in our setting.

However, there are two significant advantages to the procedure we propose. First, to implement De Loecker (2011)'s procedure, we would need to build analogues to the aggregate CES quantity indices on all destination markets. This would require computing theory-consistent price indices for every industry-destination cell in the world, which is, to our knowledge, not possible with currently available data. Second, when firms serve multiple markets, the estimation equation from De Loecker (2011) only holds as a first-order approximation. If higher order terms of this approximation are important in this nonlinear setting, then the values estimated by this procedure may be far from the truth, even if the data generating process coincides with the structural model.

To see this last point we re-write the relevant equations from De Loecker (2011) in our notation. We start from our general multi-destination model in Section 2, and impose (i) $\varepsilon_{ft}^{d} = \varepsilon_{ft}$ for all $d \in \{0 \cup \Omega_{ft}\}$, (ii) $u_{ft}^{d} = u_{ft}$ for all $d \in \{0 \cup \Omega_{ft}\}$, and (iii) a Cobb-Douglas structure for $F(\cdot)$, as in De Loecker (2011). Total revenues can then be written as

$$R_{ft} = \sum_{d\in\{0\}\cup\Omega_{ft}} \left(\exp\left(\omega_{ft}\right) M_{ft}^{\gamma^{M}} L_{ft}^{\gamma^{L}} K_{ft}^{\gamma^{K}}\right)^{\rho^{d}} \left(\chi_{ft}^{d}\right)^{\rho^{d}} D_{t}^{d} \exp\left(\varepsilon_{ft}\right) \exp\left(u_{ft}\right), \tag{37}$$

where γ^{j} indicates the output elasticity of factor $j \in \{M, L, K\}$.

De Loecker (2011) estimates the following equation (corresponding to his equation (8), written in our notation) via the control function approach:

$$\widetilde{r}_{ft} = \beta^M m_{ft} + \beta^L l_{ft} + \beta^K k_{ft} + \beta_N N_f + \sum_{d \in \{0\} \cup \Omega_{ft}} s_f^d \left(1 - \rho^d\right) \ln B_t^d + \omega_{ft} + \varepsilon_{ft} + u_{ft}, \tag{38}$$

where $\widetilde{r}_{ft}$ denotes deflated log revenues, β^j indicates the revenue elasticity of factor $j \in \{M, L, K\}$, and s_f^ds is a time-invariant measure of the importance of market d for firm f (e.g., base-year market shares). The term N_f denotes the number of products in the context of De Loecker (2011), but could represent the number of destinations in our multi-destination setting.

Careful examination shows that (38) does not follow from (37). With heterogeneous values of ρ^d, inputs do not pull out of the sum in (37), nor does ω_{ft}. Hence, the estimation equation in De Loecker (2011) is, strictly speaking, inconsistent with our model, even after imposing the restrictions specified above.

Nevertheless, by taking a first-order Taylor series approximation of (37) around a base-year equilibrium, it is possible to recover an estimation equation that is very close to (38) (see Appendix A.3 for details).[36] In this case, nontrivial movements in the quantity shares (χ_{ft}^d) over time could introduce bias. Furthermore, identification of the parameter ρ^d is complicated by all the measurement issues for B_t^d discussed above, as well as the concerns regarding short panels. By exploiting a novel source of identifying variation—the export revenues share—we circumvent both the approximation bias and the measurement issues related to B_t^d.

5 Monte Carlo simulations

In this section, we study the finite sample properties of the multi-destination estimator presented in Section 3, and the alternative estimators presented in Section 4 using Monte Carlo simulations. We use the multi-destination model from Section 2 as the data generating process.

Specifically, we impose for the data generating process that firms produce with a Cobb-Douglas production function with one flexible input, materials (M), and one quasi-fixed input, capital (K):

$$Q_{ft} = \exp\left(\omega_{ft}\right) M_{ft}^{\gamma^M} K_{ft}^{\gamma^K}, \tag{39}$$

with $\gamma^M = 0.8$ and $\gamma^K = 0.3$. We draw initial capital stocks $K_{f1} \sim U(1, 201)$, initial productivity shocks $\omega_{f1} \sim N(0, 0.0009)$, and initial *ex ante* demand shocks $\varepsilon_{f1}^d \sim N(0, 0.0009)$. We let ω and *ex ante* demand shocks for the domestic market ($d = 0$) update according to an AR(1) process with persistence parameter $h = 0.8$ and where the innovations follow a normal distribution: $\widetilde{\omega}_{ft} \sim N(0, 0.0009)$ and $\widetilde{\varepsilon}_{ft}^0 \sim N(0, 0.0009)$. Foreign *ex ante* demand shocks are unconstrained in their evolution. We draw *ex post* demand shocks using $u_{ft}^d \sim N(0, 0.0009)$. We set $\rho^0 = 0.8$ and all foreign $\rho^d = 0.6$ for $d \neq 0$.

We simulate 200 samples of a single industry with 2,000 firms over 10 periods. In order to keep the com-

[36] De Loecker (2011) never derives this equation as an approximation, nor even describes it as such. Appendix A.3 presents our own derivation of this approximation. Note that the first-order approximation holds for a fixed set of destinations over time. If the export strategy Ω_{ft} changes over time, then the equation holds as an approximation, though not necessarily a "first-order" approximation.

putational burden manageable, we allow for only four destination markets. Destination-specific industry-wide expenditures and quantity indices are drawn randomly each period, along with homogeneous (across firms) material input prices. There are no fixed costs of serving the domestic market, whereas fixed costs of entry to foreign markets are drawn from a log normal distribution with mean 6 and standard deviation 0.6, which induces heterogeneous participation in the export market.

We simulate the model period by period. In the first period, we solve for the set of destinations that maximizes expected profits for each firm. From these values, we determine which firms are active on the export market. We then update firm productivity for period 2, which includes the LBE effect, setting $\mu = 0.05$. We also update the capital stock using a fixed investment rule, which corresponds to $\iota_{ft}^{K} = \exp[0.8*(\rho^0*\omega_{ft}+\varepsilon_{ft}^0)]K_{ft}^{0.2}$, and a depreciation rate of 10 percent.[37] Given ω_{f2} and K_{f2}, we then solve the combinatorial problem for each firm in period 2. We again determine which firms are active on the export market in period 2, and update firm productivity accordingly. We continue in this fashion until the final period. In each period, given the set of export strategies, we solve for optimal material input choices and destination specific quantity shares according to (6). Finally, given u_{ft}^{d} draws, we compute optimal prices and revenues.

For each simulated sample, we implement (1) our multi-destination factor shares procedure denoted "BOR", along with procedures for (2) the estimator from Aw et al. (2011) that imposes constant returns to variable inputs (denoted "ARX"), (3) the estimator from De Loecker (2011) (denoted by "DL") that allows for different CES parameters on different markets, (4) both factor shares and control function versions of the standard practice of deflating firm-level revenues by industry-wide price indices and ignoring demand-side corrections (denoted by "NoD"), and (5) both factor shares and control function versions of the estimator that controls for aggregate sales in the domestic market as in Klette & Griliches (1996) (denoted by "SMC"). In each case, we assume researchers observe R_{ft}, M_{ft}, K_{ft}, W_t^M, B_t, Υ_t and domestic revenue shares R_{ft}^0/R_{ft}.[38] We use a common practitioners' approach to choosing initial conditions for the nonlinear optimizations based on OLS.[39]

[37] This investment rule is similar to the rule used to generate simulation data in Grieco et al. (2016). We could alternatively solve for investment as the solution to a dynamic problem, but this would not change anything about the first order necessary conditions with respect to materials and quantity shares, which are the only conditions we need to generate our estimation equation.

[38] To implement the procedure from De Loecker (2011) we also assume that the researcher observes all s_{ft}^{d}s.

[39] For the factor shares model, we set the initial conditions for the first-step NLLS estimation for M_{ft} based on an OLS estimation of the regression

$$\ln(W_{ft}^{m}M_{ft}/R_{ft}) = g_0^m + g_m^m m_{ft} + g_k^m k_{ft} + g_{mm}^m m_{ft}m_{ft} + g_{kk}^m k_{ft}k_{ft} + g_{mk}^m m_{ft}k_{ft} + \vartheta_{ft},$$

where ϑ_{ft} is a regression residual. For our multi-destination procedure, we condition the OLS regression on non-exporters.

For the second step GMM in our multi-destination estimation procedure, we set initial conditions based on an OLS estimation of the regression

$$\widetilde{r}_{ft}^{0} = g_k k_{ft} + g_{kk} k k_{ft} + g_D \ln\left[R_{ft}^0/(R_{ft}\widehat{Z}_{ft})\right] + \delta_t + \vartheta_{ft}'.$$

For the factor shares model that makes no correction for demand or controls for the aggregate domestic CES quantity index, we adjust the dependent variable and the demand proxy accordingly.

For the control function approach, we set initial conditions for the second-step GMM based on an OLS estimation of the regression

$$\widetilde{r}_{ft}^{CF} = g_0^{CF} + g_B^{CF}\ln B_t^0 + g_m^{CF} m_{ft} + g_k^{CF} k_{ft} + g_{mm}^{CF} m_{ft}m_{ft} + g_{kk}^{CF} k_{ft}k_{ft} + g_{mk}^{CF} m_{ft}k_{ft} + \vartheta_{ft}^{CF},$$

Table 1 reports the median and standard deviation of parameter estimates across 200 simulated data samples by estimator.[40] The first row of the table reports the true value of the parameters across all samples. The second row presents the results from our multi-destination factor shares procedure ("BOR"). Subsequent rows present results from alternative estimators, grouped by whether they rely on factor shares or control function approaches.

As expected, we find in the second row that our BOR procedure recovers median estimates of population parameters that are very close to the true underlying model parameters used to simulate the data. Median estimates for all parameter estimates are within 1 percentage point of the true parameter values, and standard deviations of the distribution of estimates are relatively tight. Investigating the distributions visually (see Figure C.1) reveals that the distributions are approximately centered on the true parameter values. Our multi-destination estimator hence performs well even in relatively small samples.

In contrast, we find that all other estimators are biased. In the third row of Table 1, we present results from the estimator of Aw et al. (2011), which imposes constant returns to variable inputs ($\sigma^M = 1$) and sets $u^d_{ft} = 0$. Both restrictions are inconsistent with the assumed DGP, so it is unsurprising that this estimator fails to recover the model parameters. Indeed, we find that the median estimates of σ^K, ρ^0, and μ are 220% higher, 20% lower, and 352% higher than the true underlying values, respectively. For each parameter, almost the entire distribution of estimates lies to one side of the true value (see Figure C.2). Having omitted materials from the second-step GMM estimation, the effect of materials loads onto the estimated revenue elasticity of capital, leading to the reported overestimate. Additionally, underestimating ρ^0 further inflates the estimated capital elasticity, as well as the estimate on LBE. These results indicate that the procedure from Aw et al. (2011) could lead to severely biased estimates when returns to variable inputs are non-constant, or when there are *ex post* demand shocks.

In rows 4 and 5, we report the results from the factor shares versions of the single-market demand correction estimator (SMC) and the no-demand-correction estimator (NoD), respectively. Given that firms can potentially serve multiple destination markets, these procedures are biased and yield inaccurate estimates.

The SMC estimator adapted to the factor shares approach tends to underestimate ρ^0 (by 10% for the median estimate), and thus overestimate both returns to materials and capital (by 7% and 46%, respectively). The learning by exporting effect is also underestimated with this procedure, with a median estimate 33% lower than the true value. By allowing $\rho^0 \neq 1$, the SMC factor shares estimator gets closer to the true parameter values than the NoD factor shares estimator (except for the capital elasticity). But by ignoring the multi-destination dimension of the problem, transmission bias leads to an underestimate of ρ^0, which leads to overestimates of factor returns.

The NoD estimator tends to underestimate the median returns to materials by 22% and the learning by exporting effect by 54%. This underestimation is in part due to the fact that the NoD implicitly imposes $\rho^0 = 1$, and thus, neglects to deflate parameter values (while the true value of ρ^0 is less than 1). The median returns to capital, however, is close to the true parameter value.

where $\tilde{r}^{CF}_{ft}$ represents log revenues net of the residual from the control function first step, and ϑ^{CF}_{ft} is a regression residual.

[40] Results using average values instead of medians are qualitatively similar, though some extreme outliers can produce non-trivial deviations in the average values from medians, so we focus on medians. Estimators using the control function approach suffer the most from this issue.

Table 1: Monte Carlo Experiments

	σ^M	σ^K	ρ^0	μ
True Value	0.8000	0.3000	0.8000	0.0500
Factor Shares Estimators				
Multi-Market Demand Correction (BOR)	0.7934 (0.0975)	0.3003 (0.0247)	0.8066 (0.1212)	0.0492 (0.0139)
Multi-Market, Constant Returns to Flexible Inputs (ARX)	-	0.9595 (0.0334)	0.6377 (0.0006)	0.2264 (0.0153)
Single Market Demand Correction (SMC)	0.8583 (0.2174)	0.4373 (0.1587)	0.7210 (0.1163)	0.0335 (0.0129)
No Demand Correction (NoD)	0.6198 (0.0048)	0.3172 (0.0736)	-	0.0228 (0.0069)
Control Function Estimators				
Multi-Market Demand Correction (DL)	-0.1126 (18.11)	0.5345 (13.60)	0.3146 (3.866)	0.0451 (1.212)
Single Market Demand Correction (CF-SMC)	0.8391 (16.287)	0.3753 (0.9096)	0.5520 (0.6889)	0.0301 (0.3935)
No Demand Correction (CF-NoD)	0.8360 (0.5606)	0.1372 (0.6045)	-	0.0154 (0.1433)

Notes: Table reports median and standard deviation of parameter estimates across 200 simulated data samples, by estimator. Data generating process includes 4 destination markets, 2,000 firms, and 10 time periods. True parameter values are reported in the first row. Additionally, all foreign $\rho^d = 0.6$ for $d \neq 0$.

In the lower panel of Table 1 we present estimates from three different versions of control function estimators. The first control function estimator (DL) is the procedure adapted from De Loecker (2011) for multiple destinations and presented in Section 4.5. Even though we assume that quantity indices B^ds can be computed from the data and that destination-specific revenue shares are observed (which is typically not the case in the data), the 6th row of Table 1 shows that this procedure is biased. The median estimates for σ^K, ρ^0, and μ are biased up by 78%, down by 61%, and down by 10% respectively, and the median estimate for σ^M has the wrong sign. For each parameter, the distribution of estimated coefficients is centered well away from the true values (see Figure C.3). Furthermore, in several replications, the estimated parameters are orders of magnitude away from the true values, generating extremely large standard deviations. Extreme outliers arise when the estimate of ρ^0 approaches zero because all other parameters are deflated by $\widehat{\rho}^0$.

In rows 7 and 8 of Table 1. we present results from control function versions of the single-market demand correction (SMC) and the no-demand-correction (NoD) estimators, respectively. With a DGP in which firms charge different prices and can potentially serve multiple destination markets, these procedures are biased, and indeed the estimates are off mark. For the SMC estimator the median estimate of ρ^0 is 31% lower than the true value. Median returns to materials and capital are thus overestimated by 5% and 25%, while the median LBE effect is 40% lower than the true one. In the NoD estimator the median estimates for σ^M, σ^K, and μ are up by 5%, down by 58%, and down by 71%, respectively.[41]

[41] Given the poor performance of the control function estimators, one might ask whether there exist conditions under they are unbiased. Gandhi et al. (2020) show, in a single-market homogeneous-good model, that with a long panel (roughly 50 years) and

In summary, given the assumed data generating process and available data, only the multi-destination factor shares BOR estimator recovers estimates of the output elasticities, the curvature of the domestic demand function, and LBE that are close to the true values in the Monte Carlo experiments. Alternative estimators can be severely biased, leading to incorrect inference.

6 Application to French manufacturing

In this section we describe our data sources, report descriptive statistics, and then report estimates of returns to scale, the elasticity of demand, the output elasticities of inputs, and learning-by-exporting effects for French manufacturing firms. We report results for our multi-market estimator, the estimator from Aw et al. (2011) that imposes constant returns to variable inputs, the single-market correction estimator, and the standard estimator that makes no correction for demand—all based on the factor share approach. For the main results we assume that labor is predetermined in the previous period, like capital. In Appendices E.2 and E.3 we report additional results that assume a partial adjustment process for labor and results using the control function method, respectively. While the former set of results serve as a robustness check for our main results, the latter are just for comparison, as we expect this estimator to be biased, as demonstrated in the previous section.

6.1 Data and descriptive statistics

We use administrative data sources to build a quasi-exhaustive panel of the universe of French manufacturing firms in 1994–2016. Most of the data comes from firm balance sheets from the FICUS and FARE datasets, which originate in firms' tax declarations. We use total revenues, material expenditures, employment, and book-value of capital stocks. We obtain information on firms' exports from the French Customs. It is straightforward to merge the customs data to FICUS/FARE because they use the same firm-level SIREN identifier. We deflate expenditures on materials by industry-level input price indices that we obtain from the EU KLEMS dataset. We build firm-level capital stocks using the methodology of Bonleu et al. (2013) and Cette et al. (2015). Appendix D provides further details on the data sets and explains how we construct firm-level capital stocks.

We report descriptive statistics in Table 2. The skewed firm size distribution is apparent from the difference between means and medians, for example, in revenue and employment. This feature is common in many manufacturing datasets. The high percentages of exporting firms is typical of European economies, who trade intensively within Europe. On average, 25% of firms in our data export at least once, but this varies considerably across industries, with a low of 6.6% in "Food, beverage, tobacco", and a high of 71% in "Chemical products".

substantial time series variation in material input prices (roughly 10 times that observed in the Chilean manufacturing data), the gross output control function yields estimates approximately centered on the true values in Monte Carlo experiments. We replicate this result in Appendix B using a single-market differentiated-good model. A key finding is that the control function SMC estimator performs well in finite samples only under the conditions specified by Gandhi et al. (2020) and under the additional condition that the nonlinear optimization in the second stage starts from the *true* underlying structural parameters. If the latter condition does not hold and the nonlinear optimization starts from the OLS values—as would be the natural initial conditions when estimating in the real data—the estimator is biased, as we can see in Appendix B.

In Table 3 we report descriptive statistics for the export intensity among firms that export. Within exporting firms, the export share also varies considerably, both across industries and across firms within industries. While the median exporter obtains 4.2% of revenue from exporting, the 90^{th} percentile firm obtains almost 40% of revenue from foreign markets.

Tables 2 and 3 make two important points. First, the fact that many firms export in addition to serving the domestic market implies that estimation methods that assume that all sales are on a single, domestic market ignore important information. In particular, controlling for just the *domestic* aggregate quantity index does not control for aggregate demand conditions faced by exporting firms. Second, variation in the extensive exporting margin and the high variation in export intensity among exporting firms jointly indicate that there is sufficient variation to identify ρ^0 in our setting, coming from both the cross section of firms and firm-level changes over time. This is in contrast to methods that assume only one market, where only time series variation in an aggregate demand shifter identifies the demand curvature.

Table 2: Sample Descriptive Statistics

No.	Industry		Revenue (mn euros)	Labor (employment)	Materials (mn euros)	Capital (mn euros)	Obs.	No. firms	No. exporters	Percent exporters
1	Autos and transport equipment	Mean	39.93	164.05	26.61	22.95	49982	5480	2771	50.6
		Median	1.04	10.00	0.41	0.22				
2	Chemical products	Mean	52.95	98.04	30.15	28.07	51296	4889	3471	71.0
		Median	2.38	14.00	0.95	0.56				
3	Computer, electronics	Mean	11.34	61.48	5.17	4.93	51583	5641	3127	55.4
		Median	0.83	8.00	0.27	0.13				
4	Electrical equipment	Mean	13.07	70.48	6.97	5.11	42285	4601	2333	50.7
		Median	0.97	9.00	0.35	0.13				
5	Food, beverage, tobacco	Mean	3.11	12.73	1.82	1.21	880040	112476	7498	6.7
		Median	0.24	3.75	0.08	0.10				
6	Machinery and equipment	Mean	3.82	22.86	1.74	1.03	318464	34436	11184	32.5
		Median	0.56	5.00	0.18	0.09				
7	Basic metal and fabricated metal	Mean	4.32	26.66	1.89	2.18	353028	33765	12909	38.2
		Median	0.80	9.00	0.16	0.23				
8	Other manufacturing	Mean	1.55	12.19	0.61	0.55	246572	30639	6550	21.4
		Median	0.22	3.00	0.05	0.07				
9	Rubber and plastic	Mean	6.93	39.27	3.06	4.14	161883	16000	6960	43.5
		Median	0.89	8.00	0.30	0.27				
10	Textiles, wearing apparel	Mean	3.32	24.44	1.41	0.98	147807	21319	9066	42.5
		Median	0.49	6.00	0.14	0.08				
11	Wood, paper products	Mean	2.76	17.38	1.21	1.62	294072	32232	9771	30.3
		Median	0.47	5.00	0.11	0.14				
	Total	Mean	5.44	25.17	2.88	2.56	2597012	301478	75640	25.1
		Median	0.38	5.00	0.11	0.11				

Notes. The table reports descriptive statistics for the estimation sample, where capital is the book value reported by the firm and materials are expenditures. Exporters are defined as firms that exported at least once during the sample. Source: FICUS/FARE datasets and French Customs.

Table 3: Percent exports in revenue for exporters

No.	Industry	Mean	p5	p10	p50	p90	p95
1	Autos and transport equipment	14.6	0.2	0.4	5.8	43.1	55.9
2	Chemical products	22.5	0.2	0.6	11.3	64.6	77.5
3	Computer, electronics	19.6	0.2	0.5	8.4	58.2	74.1
4	Electrical equipment	15.5	0.2	0.5	6.1	46.7	60.8
5	Food, beverage, tobacco	10.1	0.1	0.1	2.8	31.2	48.5
6	Machinery and equipment	12.4	0.1	0.3	4.1	38.5	57.2
7	Basic metal and fabricated metal	10.6	0.1	0.3	3.5	31.7	47.8
8	Other manufacturing	12.8	0.2	0.5	4.9	38.1	52.7
9	Rubber and plastic	11.5	0.1	0.2	3.7	36.0	52.2
10	Textiles, wearing apparel	17.8	0.4	0.8	9.5	48.4	62.2
11	Wood, paper products	7.8	0.1	0.1	1.6	23.8	41.9
	Total	12.8	0.1	0.3	4.3	39.8	56.7

Notes. The table reports the distribution of the percent of exports in revenue for exporters in the estimation sample. Percent exports in revenue for exporters is computed for firms and years in which exports are positive. Source: FICUS/FARE datasets and French Customs.

6.2 Main results

We present our main results graphically in Figures 1 and 2, which facilitates comparison of estimated parameters across estimators. The figures show estimated total returns to scale, variable returns to scale, the domestic price elasticity of demand, and the learning by exporting parameter. The corresponding detailed estimates are provided in Tables E.1–E.5, where we also report the persistence parameter h, the demand curvature parameter ρ^0, and the long run effect of exporting $\mu/(1-h)$, as well as bootstrapped standard errors for all estimates.[42]

Total Returns to Scale. We present estimates of total returns to scale (RTS) in the top left panels of Figures 1 and 2. Our multi-market estimator, denoted "BOR" (depicted by blue circles) yields estimated total returns to scale that range from 1.05 (Electrical equipment) to 1.24 (Wood, paper products), and for one industry up to 1.35 (Food, beverage, tobacco). Bars represent 95% bootstrapped confidence intervals. The mean (median) estimate across the 11 industries is 1.15 (1.13), which in line with estimates in Antweiler & Trefler (2002), and is substantially higher than corresponding OLS estimates (Figure E.5).

In contrast, the estimator that makes no demand correction, depicted in Figure 1 with red triangles, yields estimated total returns to scale slightly below 1 for most industries. The mean and median of the industry-specific average returns to scale are both equal to 0.96, with tight confidence intervals. These estimates are close to what researchers tend to find with this approach. For example, after deflating revenues by industry-wide price indices and interpreting them as quantities, Gandhi et al. (2020) find in their factor shares approach average returns to scale in Colombia between 0.99 and 1.06, and between 1.04 and 1.15 for Chile, which is close to what we find using the control function method and OLS (see Figure E.5). Comparing

[42]While our main specification treat labor as a quasi-fixed input, which is often assumed in the French context (Harrigan et al., 2023), we investigate the sensitivity of the results to the possibility that firms partially adjust labor to contemporaneous shocks. Appendix E.2 shows that this does not materially alter the main results.

these estimates with the ones obtained with the BOR estimator indicates that the constant returns to scale estimated by the no demand correction model mask returns that are actually increasing (in our notation, $\sigma^M_{ft} + \sigma^L_{ft} + \sigma^k_{ft} > 1$), as hypothesized by Klette & Griliches (1996).[43] The reason is simple: estimators that fail to adjust for the demand curvature report "revenue elasticities" that are, in our notation, a factor $1/\rho^0$ smaller than output elasticities.

Next, in Figure 2, we compare our BOR estimates (blue circles) with the single market correction (black squares) and the multi-market estimator with constant returns to flexible inputs (orange crosses). Because the scale of the y-axis must be expanded to accommodate the larger magnitude of the estimates from these two alternatives, we split the presentation of the results into two figures. If the data-generating process corresponds to the multi-destination model from Section 2, the SMC estimator only partially addresses the transmission bias from missing output prices. As shown in the top left panel of Figure 2, the single-market correction estimator from Section 4.2 delivers highly dispersed estimates across industries. We find that the average returns to scale range from -3 (Rubbers and plastics) to 7.3 (Chemicals, omitted from the figure to ease visibility of other results). We find only three industries with plausible estimates for returns to scale: Auto and transportation (1.09), Communication electronics (1.29), and Electrical equipment (1.34). Estimated average returns to scale are implausibly high or implausibly low for all other industries.[44] The large range of estimates for returns to scale can be explained by the range of estimates of ρ, which are used to "deflate" the revenue elasticities. With the SMC estimator, we estimate a range for ρ from -0.61 to 1.49 across industries. When ρ is estimated to be close to zero, then the returns to scale become very large in absolute value, as they do for Chemicals. When the estimate of ρ is negative, this leads to negative estimates of returns to scale (Rubbers and plastics and Other manufacturing). These extreme estimates of ρ are not merely due to estimation uncertainty; the estimates are quite precise (see Table E.3 for bootstrapped standard errors).

The wide range of SMC estimates of returns to scale is notable because we find a much smaller range of estimated returns to scale in the Monte Carlo experiments when using this estimator. A likely explanation for the difference is that in the Monte Carlo simulations we assume the researcher observes the true domestic quantity index B^0_t. In the application, we follow the common practice to approximate B^0_t using all of the firms' revenues (not distinguishing exports and domestic sales, nor adding imports) and use price indices built from producer prices in the domestic market (as discussed above in Section 4.2). These discrepancies highlight an additional, practical advantage for our multi-market estimator: it does not require making such data compromises, as it does not require building B^0_t, nor does it require to deflate firm revenues.

Given the theoretical drawbacks of applying the SMC estimator to a sample of multi-destination firms, we can alternatively apply the SMC estimator to a sample of never exporters, as discussed in Section 4.2.

[43] Klette & Griliches (1996) argue that ignoring unobserved firm-specific prices would tend to lead to a downward bias in estimated returns to scale, *ceteris paribus*. As we discuss in section 4.1, there are, in fact, several forces that bias the estimator with no demand correction, and the overall sign cannot be determined in general. Nevertheless, the evidence in the upper left panel of Figure 1 is consistent with the central hypothesis from Klette & Griliches (1996). For comparison, estimates of average returns to scale using the control function approach with no demand correction are all close to 1 (except for Wood), with very low estimates of the capital output elasticity; see Appendix E.3.

[44] The high and low estimates of average returns to scale is not just a matter of outlier observations either. Medians within the industry are very close to the means.

Figure 1: Factor Share Estimates by Industry, BOR vs. No Demand Correction

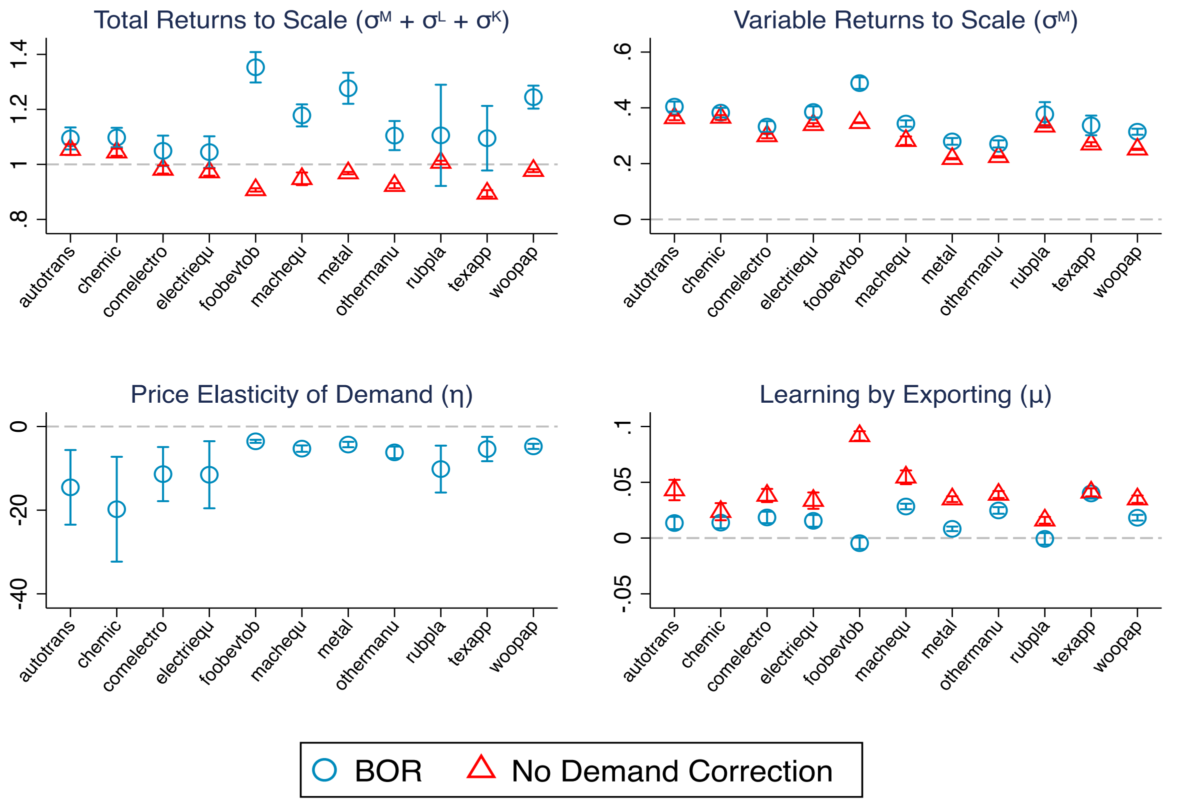

Notes. The figure reports factor share estimates of average Total Returns to Scale, Variable Returns to Scale, the Price Elasticity of Demand on the domestic market $\eta^0 = 1/\left(\rho^0 - 1\right)$, and the LBE learning by exporting (LBE) parameter μ by industry and estimator. Bars represent 95% bootstrap confidence intervals. Detailed estimates are reported in Tables E.1 and E.2.

In this case, measurement error in B_t^{proxy} and selection bias could still lead to biased estimates. Table E.4 reports results for the sample of non exporters, where we find very similar results to the main SMC estimates when we do not drop exporters (see Table E.3). Overall, the wide dispersion in the SMC estimates can be explained by measurement error in B_t^{proxy}. This is compounded by the fact that identification in the SMC estimator rests solely on time-series variation in aggregate consumption, with at most 21 observations per industry in our sample. Limited variation or imprecise measurement of this key variable naturally leads to unstable estimates.[45]

Figure 2 also shows that estimates of returns to scale using the estimator of Aw et al. (2011) ($\times$'s) are implausibly high, ranging from 2.27 to 4.21. This is driven not only by the output elasticity for materials, which is fixed to unity by construction, but also by the output elasticity of predetermined labor being estimated above 1 for all industries but one (see Table E.5).

In our preferred specification (BOR, blue circles) average returns to scale are greater than 1 for all industries. This indicates that there are efficiency gains from size embedded in the technology used by firms, regardless of how total factor productivity evolves. Our estimates imply that returns to scale are increasing

[45] Because the measurement error is non-classical, the resulting bias is not necessarily toward zero.

Figure 2: Factor Share Estimates by Industry, BOR vs. Alternative Estimators

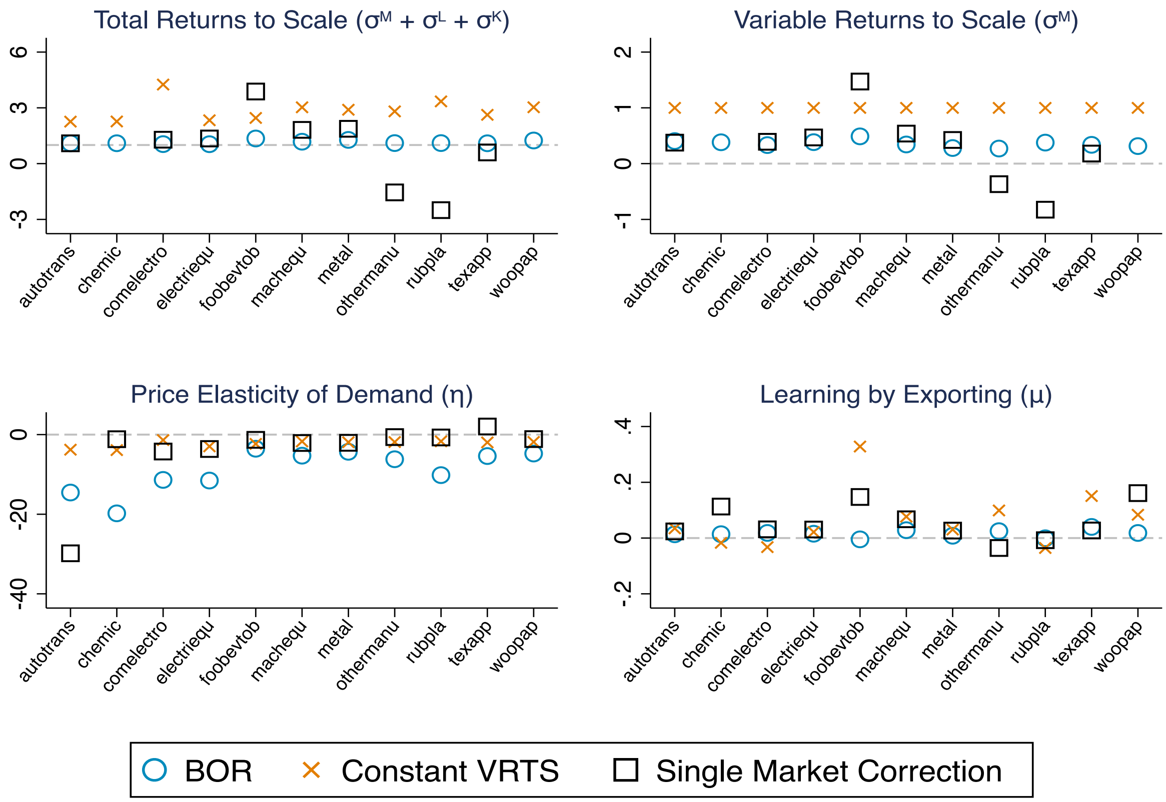

Notes. The figure reports factor share estimates of average Total Returns to Scale, Variable Returns to Scale, the Price Elasticity of Demand on the domestic market $\eta^0 = 1/\left(\rho^0 - 1\right)$, and the LBE learning by exporting (LBE) parameter μ by industry and estimator. Confidence intervals have been suppressed for ease of viewing. Detailed estimates are reported in Table E.3. Results for Chemicals and Wood and Paper for Single Market Correction in top row have been suppressed for ease of viewing.

for virtually all firm-year observations, not just on average (see Figure E.4).[46] From a welfare perspective, increasing returns imply a cost to diversification that weighs against love of variety, as hypothesized by Krugman (1979). In addition, increasing returns imply larger business cycle fluctuations, and may provide a rationale for targeted interventions during downturns.

Returns to flexible inputs. In our main specification, only materials are treated as flexible inputs, hence the variable returns to scale (VRTS) are just the output elasticity with respect to materials. The top right panel of Figure 1 reports the average VRTS by industry and by estimator. The mean (median) of the average estimates across industries is 0.34 (0.35) with our multi-market estimator and 0.30 (0.30) with the no demand correction estimator. Both estimators share the same first step. But the first step without demand correction identifies the revenue elasticity with respect to materials, which must be divided by ρ^0 to recover the output elasticity—as the BOR multi-market correction does. Since $\rho^0 < 1$ in our estimates, the implied output

[46] Several mechanisms could explain this phenomenon. The simplest explanation for increasing returns to scale is the presence of fixed costs of operation. Alternatively, complementarities within the firm could generate increasing returns at any point along the firm-size distribution. For example, externalities across workers could lead to increasing returns (e.g., learning by doing), as in, for example Kellogg (2011) and Hjort (2014).

elasticity of materials is larger under our BOR approach.

Returns to flexible inputs well below 1 imply negative cross-market complementarities in the short-to-medium run. For example, a positive demand shock in one market leads to more sales to that market, an increase in marginal costs, lower sales to other markets and a reduction in the likelihood of selling to other markets. This is consistent with findings in Almunia et al. (2021), who argue that the massive negative demand shock in Spain during the financial crisis caused an increase in exporting, presumably due to a reduction in scale and in marginal costs.[47]

The top right panel of Figure 2 reports the implicit variable returns to scale from the estimator described in section 4.4 (with $\sigma^M = 1$) and the estimated VRTS from the single-market correction method. While the former is likely too high, the latter approximately coincides with our preferred BOR estimates for 6 industries out of 11, with implausibly large or even negative estimates in the remaining industries. The SMC estimates range from -0.83 (Rubber and plastic) to 2.54 (Chemical products, omitted from the Figure), with an average and median around 0.45. As noted above, this variation is mostly due to variation in the estimates of ρ.

Elasticity of Demand. In the bottom left panel of Figures 1 and 2, we report estimates of the domestic price elasticity of demand $\eta^0 = 1/(\rho^0 - 1)$. With our multi-market estimator, we estimate a range of demand elasticities from -19.42 (Chemicals) to -3.56 (Food, beverage and tobacco), with no outlier estimates. The mean (median) estimate across industries is -9.9 (-5.8), which is comparable to the mean and median estimates typically obtained in gravity regressions (e.g., Shapiro 2016; Fontagné et al. 2022).[48] Bootstrapped standard errors are reported in Table E.1. As ρ^0 approaches 1 (e.g., in electrical equipment), standard errors for η^0 become extremely large, making it more informative to examine the standard errors for ρ^0. The bootstrapped confidence intervals for ρ^0 are narrow, and we can reject $\rho^0 = 1$ in all industries, supporting our key assumption that manufacturing products are differentiated.

We compare these BOR estimates with the ones resulting from the SMC estimator and the estimator with constant returns to variable inputs in Figure 2. With the SMC estimator, most of the demand elasticities fall within the range -4.23 to -0.62, which is significantly lower in magnitude than our BOR estimates. Additionally, with the single market correction, we estimate a positive demand elasticity for Textile and Apparel ($\eta = 2.03$), and a very high (in magnitude) demand elasticity for Autos and transport equipment (-29.8).

The demand elasticities obtained by the Aw et al. (2011) approach (constant VRTS) are all much smaller (in absolute value) than those estimated with BOR. The domestic demand elasticities range from -1.37 to -3.89, while the demand elasticities on foreign markets ($\eta^x = 1/(\rho^x - 1)$) range from -1.68 to -2.57—both with a median of about -2. They are much lower in magnitude than what is typically found in the literature.[49]

[47] Almunia et al. (2021) perform production function and productivity estimation, but when developing their estimator they ignore cross-market complementarities. We explain how their estimator differs fundamentally from ours in Appendix A.1.

[48] For example, using data on trade flows and trade costs, Shapiro (2016) estimates an average trade elasticity across industries of -8.16, which translates into a price elasticity of demand of -9.16. This estimate can be interpreted as the demand elasticity under the assumption of an Armington trade model.

[49] Estimates of the demand elasticity using the control function approach with no demand correction or the single market correction are very noisy; see Appendix E.3.

Learning by Exporting. Finally, in the bottom right panels of Figures 1 and 2, we report estimates of LBE effects by industry and estimator. With our multi-market estimator, we estimate LBE effects in the range of -0.005 (Food, beverage and tobacco) to 0.040 (Textile and apparel). We can reject the null hypothesis of zero LBE for all but two industries with very low estimated LBE effects (Food, beverage and tobacco and Rubbers and plastics). The average (median) estimate of LBE across all eleven industries is 0.017 (0.018). These estimates are lower than the estimates with the no demand correction estimator (average = 0.04 and median = 0.038), the SMC estimator (average = 0.035 and median = 0.030), and the estimator with constant VRTS (average = 0.067 and median = 0.035, with Food, beverage, tobacco an outlier at 0.34). The comparison suggests that estimated LBE effects are biased upward in the other estimators, possibly because at least the NoD and the SMC estimators mistakenly attribute the effect of foreign demand shocks to LBE.

Overall, we find robust evidence of significant LBE effects in French manufacturing data. Estimates using our multi-market procedure imply as much as 40% long run cross-sectional differences in productivity between exporters and non-exporters (Table E.1, last column).[50] Compared to previous work, these effects are smaller than those estimated via RCT with Egyptian firms (Atkin et al., 2017) and via structural approaches with Chilean, Colombian, and Mexican firms, e.g., (Garcia-Marin & Voigtländer, 2019), but larger than estimates from Danish firms using a quasi-natural experiment (Buus et al., 2022).[51]

7 Conclusion

Production function estimation is central to many economic analyses, yet the conditions assumed in theory rarely match those faced by applied researchers (De Loecker & Goldberg, 2014; De Loecker & Syverson, 2021). In particular, most datasets report output only in values, not in quantities. In addition, many firms serve multiple destination markets wherein they face heterogeneous demand conditions, which is inconsistent with models that control for a single market demand shifter. These discrepancies matter for estimation, even when researchers are not interested in exporting or the effect of exporting *per se*.

In this paper, we show how to estimate output elasticities, the domestic price elasticity of demand, the elasticity of productivity to observable determinants, and productivity itself when firms serve multiple destination markets and when outputs are denominated only in monetary terms. We show that existing production function estimators that either do not address the missing price issue or do so using restrictive assumptions (e.g., single market or constant returns to variable inputs) yield biased inference in this case. By contrast, our estimator recovers estimates that are close to the true parameters values in Monte Carlo simulations.

Our estimator is no more difficult to implement than existing methods and requires only one additional piece of information: firms' export shares. This provides a practical advantage, since we exploit this novel

[50]The long run, cross-sectional difference is computed as $\mu/(1-h)$, where μ is the effect of exporting today on productivity tomorrow and h is the persistence parameter in the AR(1) process for productivity.

[51]The other notable comparisons in the literature is De Loecker (2013), who estimates LBE effects in the range of 0.017 to 0.066 across Slovenian manufacturing industries. Since De Loecker (2013) does not correct for demand conditions, these estimates are best compared to the results from our no-demand-correction estimator, which we find in a similar range of 0.016 to 0.092.

source of firm-level variation to identify the demand curvature faced by firms, instead of relying solely on time-series variation in aggregate demand shifters. Moreover, we do not need to construct measures of these aggregate demand shifters, which would be very demanding if required for all destination-industry-years. Finally, our estimator does not rely on functional form assumptions for the production function, and although it assumes destination-industry-specific constant elasticities of demand, it allows for firm-destination-year variation in prices and markups.

We compare the estimates obtained with our estimator to the ones from existing approaches using balance sheet information for the universe of French manufacturing firms. In the French data, we estimate demand elasticities between -19.42 and -3.56, which are in a range that is consistent with much of the literature. We estimate average returns to scale ranging from 1.05 to 1.24 with one industry at 1.35, and average returns to flexible inputs uniformly below 1. The latter result implies cross-market complementarities: additional sales for a given market raises the cost of serving all other markets in the short run. Alternative approaches yield estimates of returns to scale or demand curvature that deviate from our preferred estimates, sometimes very substantially. We also estimate learning-by-exporting effects ranging from 0 to 4% per year, which imply cross-sectional differences in productivity between exporters and non-exporters up to 40%. Hence, our approach allows us to study a determinant of productivity such as learning by exporting using a productivity estimation framework that is consistent with heterogeneous firms' export decisions and pricing to market.

References

Ackerberg, D. A., Caves, K., & Frazer, G. (2015). Identification properties of recent production function estimators. *Econometrica*, 83(6), 2411–2451.

Ackerberg, D. A., Frazer, G., il Kim, K., Luo, Y., & Su, Y. (2023). Under-identification of structural models based on timing and information set assumptions. *Journal of Econometrics*, 237(1), 105463.

Almunia, M., Antràs, P., Lopez-Rodriguez, D., & Morales, E. (2021). Venting Out: Exports during a Domestic Slump. *American Economic Review*, 111(11), 3611–62.

Antweiler, W. & Trefler, D. (2002). Increasing returns and all that: a view from trade. *American Economic Review*, 92(1), 93–119.

Arkolakis, C., Eckert, F., & Shi, R. (2023). *Combinatorial Discrete Choice: A Quantitative Model of Multinational Location Decisions*. Working Paper 31877, National Bureau of Economic Research.

Atkin, D., Khandelwal, A. K., & Osman, A. (2017). Exporting and firm performance: Evidence from a randomized experiment. *The Quarterly Journal of Economics*, 132(2), 551–615.

Aw, B. Y., Roberts, M. J., & Xu, D. Y. (2011). R&D investment, exporting, and productivity dynamics. *American Economic Review*, 101(4), 1312–44.

Barrows, G. & Ollivier, H. (2021). Foreign demand, developing country exports, and CO2 emissions: Firm-level evidence from India. *Journal of Development Economics*, 149, 102587.

Berman, N., Berthou, A., & Héricourt, J. (2015). Export dynamics and sales at home. *Journal of International Economics*, 96(2), 298–310.

Blum, B. S., Claro, S., Horstmann, I., & Rivers, D. A. (2023). The ABCs of firm heterogeneity when firms sort into markets: The case of exporters. *Journal of Political Economy*, Just Accepted.

Blundell, R. & Bond, S. (2000). GMM estimation with persistent panel data: an application to production functions. *Econometric reviews*, 19(3), 321–340.

Bonleu, A., Cette, G., & Horny, G. (2013). Capital utilization and retirement. *Applied Economics*, 45(24), 3483–3494.

Broda, C. & Weinstein, D. E. (2006). Globalization and the gains from variety. *The Quarterly Journal of Economics*, 121(2), 541–585.

Buus, M. T., Munch, J. R., Rodrigue, J., & Schaur, G. (2022). Do export support programs affect prices, quality, markups and marginal costs? Evidence from a natural policy experiment. *Review of Economics and Statistics*, (pp. 1–45).

Caballero, R. J. & Lyons, R. K. (1992). External effects in US procyclical productivity. *Journal of Monetary Economics*, 29(2), 209–225.

Caliendo, L. & Parro, F. (2015). Estimates of the trade and welfare effects of nafta. *The Review of Economic Studies*, 82(1), 1–44. NBER Working Paper No. 18508 (2012):contentReference[oaicite:1]index=1.

Cette, G., Dromel, N., Lecat, R., & Paret, A.-C. (2015). Production factor returns: the role of factor utilization. *Review of Economics and Statistics*, 97(1), 134–143.

De Loecker, J. (2007). Do exports generate higher productivity? evidence from slovenia. *Journal of International Economics*, 73(1), 69–98.

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

De Loecker, J. (2013). Detecting learning by exporting. *American Economic Journal: Microeconomics*, 5(3), 1–21.

De Loecker, J. & Goldberg, P. K. (2014). Firm performance in a global market. *Annu. Rev. Econ.*, 6(1), 201–227.

De Loecker, J. & Syverson, C. (2021). An Industrial Organization Perspective on Productivity. *Becker Friedman Institute Working Paper No. 2021-109*.

de Roux, N., Eslava, M., Franco, S., & Verhoogen, E. (2021). Estimating Production Functions in Differentiated-Product Industries with Quantity Information and External Instruments. *National Bureau of Economic Research Working Paper*, 28323.

Dhyne, E., Petrin, A., Smeets, V., & Warzynski, F. (2022). Theory for extending single-product production function estimation to multi-product settings. *National Bureau of Economic Research Working Paper*, 30784.

Feenstra, R. C. (1994). New product varieties and the measurement of international prices. *American Economic Review*, (pp. 157–177).

Fontagné, L., Guimbard, H., & Orefice, G. (2022). Tariff-based product-level trade elasticities. *Journal of International Economics*, 137, 103593.

Foster, L., Haltiwanger, J., & Syverson, C. (2008). Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *American Economic Review*, 98(1), 394–425.

Gandhi, A., Navarro, S., & Rivers, D. A. (2020). On the identification of gross output production functions. *Journal of Political Economy*, 128(8), 2973–3016.

Garcia-Marin, A. & Voigtländer, N. (2019). Exporting and plant-level efficiency gains: It's in the measure. *Journal of Political Economy*, 127(4), 1777–1825.

Grieco, P. L., Li, S., & Zhang, H. (2016). Production function estimation with unobserved input price dispersion. *International Economic Review*, 57(2), 665–690.

Hall, R. E. (1986). Market structure and macroeconomic fluctuations. *Brookings papers on economic activity*, 1986(2), 285–338.

Harrigan, J., Reshef, A., & Toubal, F. (2023). Techies, Trade and Skill-Biased Productivity. NBER Working Paper No. 31341.

Hjort, J. (2014). Ethnic divisions and production in firms. *The Quarterly Journal of Economics*, 129(4), 1899–1946.

Imbs, J. & Méjean, I. (2017). Trade elasticities. *Review of International Economics*, 25(2), 383–402.

Kellogg, R. (2011). Learning by drilling: Interfirm learning and relationship persistence in the Texas oilpatch. *The Quarterly Journal of Economics*, 126(4), 1961–2004.

Klette, T. J. & Griliches, Z. (1996). The inconsistency of common scale estimators when output prices are unobserved and endogenous. *Journal of Applied Econometrics*, 11(4), 343–361.

Knittel, C. R., Metaxoglou, K., et al. (2014). Estimation of Random-Coefficient Demand Models: Two Empiricists' Perspective. *The Review of Economics and Statistics*, 96(1), 34–59.

Krugman, P. R. (1979). Increasing returns, monopolistic competition, and international trade. *Journal of International Economics*, 9(4), 469–479.

Levinsohn, J. & Petrin, A. (2003). Estimating production functions using inputs to control for unobservables. *The review of economic studies*, 70(2), 317–341.

Melitz, M. & Levinsohn, J. A. (2006). Productivity in a differentiated products market equilibrium. *Mimeo, Princeton University*.

Nelson, C. R. & Startz, R. (1990). Some further results on the exact small sample properties of the instrumental variable estimator. *Econometrica*, 58(4), 967.

Olley, S. & Pakes, A. (1996). Market share, market value and innovation in a panel of British manufacturing firms. *Econometrica*, 64(6), 1263–1297.

Redding, S. J. & Weinstein, D. E. (2020). Measuring aggregate price indices with taste shocks: Theory and evidence for CES preferences. *The Quarterly Journal of Economics*, 135(1), 503–560.

Reshef, A., Saint-Paul, G., & Toubal, F. (2022). Labor Market Flexibility and Firms' Resilience. *Mimeo, Paris School of Economics*.

Roberts, M. J., Yi Xu, D., Fan, X., & Zhang, S. (2018). The role of firm factors in demand, cost, and export market selection for Chinese footwear producers. *The Review of Economic Studies*, 85(4), 2429–2461.

Saint-Paul, G. (1996). *Dual labor markets*. MIT Press Cambridge.

Sato, K. (1976). The ideal log-change index number. *The Review of Economics and Statistics*, (pp. 223–228).

Shapiro, J. S. (2016). Trade costs, CO2, and the environment. *American Economic Journal: Economic Policy*, 8(4), 220–254.

Valmari, N. (2023). Estimating production functions of multiproduct firms. *The Review of Economic Studies*, 90(6), 3315–3342.

Van Biesebroeck, J. (2005). Exporting raises productivity in sub-saharan african manufacturing firms. *Journal of International economics*, 67(2), 373–391.

Vartia, Y. O. (1976). Ideal log-change index numbers. *Scandinavian Journal of Statistics*, (pp. 121–126).

Wagner, J. (2007). Exports and productivity: A survey of the evidence from firm-level data. *World economy*, 30(1), 60–82.

Wagner, J. (2012). International trade and firm performance: A survey of empirical studies since 2006. *Review of World Economics*, 148, 235–267.

Wooldridge, J. M. (2009). On estimating firm-level production functions using proxy variables to control for unobservables. *Economics letters*, 104(3), 112–114.

Appendix

A Other Production Function Estimators

A.1 Comparison to Almunia et al. (2021)

In recent work, Almunia et al. (2021) also specify a model with multiple destinations, monopolistic competition, and quasi-fixed capital. They also estimate production function parameters and an elasticity of demand. There are four key differences between the production function estimation in Almunia et al. (2021) and ours. First, Almunia et al. (2021) assume that firms are myopic with respect to *ex post* destination specific demand shocks ($E\left[e^{u_{ft}}\right]=1$). Second, as we explain below, Almunia et al. (2021) implicitly assume constant returns to flexible inputs when estimating the elasticity of demand, which is inconsistent with the thrust of their main findings. Third, it seems that Almunia et al. (2021) do not consider the multi-destination dimension when estimating firm-level productivity (see Appendix F in their paper). Finally, Almunia et al. (2021) assume that the elasticity of demand is identical across all destinations (within an industry), whereas one of our contributions is to allow for variation in elasticities across markets.

In what follows, we analyze their estimator in the context of a version of our model where we impose the same elasticity of demand across destinations, as they do. We demonstrate that their estimator is not consistent with variable (increasing or decreasing) returns to variable inputs.

Almunia et al. (2021) claim (see Appendix F.1 of Almunia et al. (2021)) that the assumption of monopolistic competition implies (in their notation) that

$$R_{it}-C_{it}^{v}=\frac{1}{\sigma}R_{it}, \tag{A.1}$$

where R_{it} is total revenues of firm i at time t, C_{it}^{v} denotes the total variable cost of firm i at time t, and σ is the elasticity of substitution (and demand). Equivalently, we can define in our notation the total variable cost of firm f in time t,

$$Cost_{ft}\left(Q_{ft}\right)=M_{ft}W_{t}^{M}, \tag{A.2}$$

We can then re-write (A.1) in our notation

$$R_{ft}-Cost_{ft}\left(Q_{ft}\right)=\left(1-\rho\right)R_{ft}. \tag{A.3}$$

Based on assumption (A.1), and substituting for C_{t}^{v} with expenditures on flexible inputs, Almunia et al. (2021) derive the following moment condition (in their notation)

$$E\left[\ln\left(\frac{\sigma-1}{\sigma}\right)+r_{it}^{obs}-\ln\left(P_{it}^{M}M_{it}+w_{it}L_{it}\right)\right]=0, \tag{A.4}$$

or in our notation

$$E\left[\ln\rho + r_{ft} - \ln\left(W_t^M M_{ft}\right)\right] = 0. \tag{A.5}$$

If the assumption in (A.1) were to hold, then indeed the moment condition (A.5) could be exploited to identify the curvature of demand (ρ, in our notation). But as we show below, (A.1) requires both (in our notation) $E\left[e^{u_{ft}}\right] = 1$ and constant marginal costs. This implies that firms are myopic with respect to *ex post* destination specific demand shocks, and that variable returns to scale are unitary.

We can rewrite the optimization problem of a firm from section 2 for a fixed set of destinations (fixed "export strategy") using the variable cost function (A.2):

$$\max_{\boldsymbol{\chi}_{ft}, Q_{ft}} \quad \mathscr{L} \quad = \quad E\left[Q_{ft}^{\rho} \sum_{d\in\Omega_{ft}} \left(\chi_{ft}^{d}\right)^{\rho} D_t^d \exp\left(\varepsilon_{ft}^d + u_{ft}^d\right)\right] - Cost_{ft}\left(Q_{ft}\right) + \lambda_{ft}\left(1 - \sum_{d\in\Omega_{ft}} \chi_{ft}^d\right) \tag{A.6}$$

which leads to first order condition for Q_{ft}

$$\rho\left(Q_{ft}\right)^{\rho-1}\left[\sum_{d\in\Omega_{ft}}\left(D_t^d \exp\left(\varepsilon_{ft}^d\right)\right)^{\frac{1}{1-\rho}}\right]^{1-\rho} E\left[\exp\left(u\right)\right] = \frac{\partial Cost_{ft}\left(Q_{ft}\right)}{\partial Q_{ft}}, \tag{A.7}$$

Multiplying both sides by Q_{ft}, we have

$$\rho\left(Q_{ft}\right)^{\rho}\left[\sum_{d\in\Omega_{ft}}\left(D_t^d \exp\left(\varepsilon_{ft}^d\right)\right)^{\frac{1}{1-\rho}}\right]^{1-\rho} E\left[\exp\left(u\right)\right] = \frac{\partial Cost_{ft}\left(Q_{ft}\right)}{\partial Q_{ft}} Q_{ft}, \tag{A.8}$$

and substituting with total revenues

$$\rho E\left[\exp\left(u\right)\right] R_{ft} \psi_{ft}^{-1} = \frac{\partial Cost_{ft}\left(Q_{ft}\right)}{\partial Q_{ft}} Q_{ft}, \tag{A.9}$$

Now if we set $E\left[\exp\left(u\right)\right] = 1$ and we assume $\frac{\partial Cost_{ft}\left(Q_{ft}\right)}{\partial Q_{ft}} Q_{ft} = Cost_{ft}\left(Q_{ft}\right)$, we get the moment condition (A.5). But with non-constant marginal cost, $\frac{\partial Cost_{ft}\left(Q_{ft}\right)}{\partial Q_{ft}} Q_{ft} \neq Cost_{ft}\left(Q_{ft}\right)$. In particular, with decreasing returns to flexible inputs, which is the necessary condition for cross-market complementarities, $\frac{\partial Cost_{ft}\left(Q_{ft}\right)}{\partial Q_{ft}} Q_{ft} \neq Cost_{ft}\left(Q_{ft}\right)$.

A.2 Building market quantity proxy from price indices

In a single-market estimation model, the demand-side parameter is identified from time series variation in the industry-wide CES demand index. In this section, we discuss how to construct this index.

Essentially, the quantity index can be recovered from expenditure data and industry-wide price deflators.

Assuming just a single market (hence dropping the d superscript), and using (1) and (2), we can write

$$B_t^{\rho} \quad = \quad \sum_{f\in\Theta_t} \exp\left(\varepsilon_{ft} + u_{ft}\right) X_{ft}^{\rho} = \sum_{f\in\Theta_t} \frac{R_{ft}}{\Upsilon_t} B_t^{\rho-1} \tag{A.10}$$

This implies

$$B_t \quad = \quad \sum_{f\in\Theta_t} \frac{R_{ft}}{\Upsilon_t} \tag{A.11}$$

Hence, if we observe the true CES price index in levels, we can construct the CES quantity index from aggregate deflated revenues.[52]

But of course, the true CES price index is not observed in levels. First, price indices are almost always reported relative to some base year normalization. This implies

$$B_t = \sum_{f\in\Theta_t} \frac{R_{ft}}{\Upsilon_t} = \Upsilon_0 \underbrace{\sum_{f\in\Theta_t} \frac{R_{ft}}{\Lambda_t}}_{\equiv B_t^{proxy}} \tag{A.12}$$

where Λ_t is the empirical analogue to the true CES price index normalized to base-year $t = 0$, and Υ_0 is the unobserved base-year normalization.

Second, the CES price index is a theoretical construct that depends on structural parameters. How does this theoretical object correspond to Λ_t? Sato (1976) and Vartia (1976) prove for a symmetric CES with no entry and exit, there exists a set of weights wt_{ft} such that

$$\ln \frac{\Upsilon_t}{\Upsilon_0} \quad = \quad \sum_{f\in\Theta_t} wt_{ft} \ln\left(\frac{p_{ft}}{p_{fi0}}\right) \tag{A.13}$$

I.e., the log change in the true CES price index is a weighted average of the log change in the prices of individual firms. Sato (1976) and Vartia (1976) give the analytical expression for these weights, which ends up being very close to a simple chain weight. Feenstra (1994) extends to the case of entry and exit. Redding & Weinstein (2020) extends to the asymmetric CES (which corresponds to our demand system (1)). If we assume that Λ_t is computed using Weinstein-Redding weights, then (A.12) holds.

A.3 Multi-destination estimator with control function

In this section, we explain the steps that are required to recover an estimation equation that resembles the one developed by De Loecker (2011). In particular, we show how (38) can be rationalized as a first-order Taylor approximation to a function relating destination-specific revenues to total firm-level revenues, with some amendments. De Loecker (2011) does not derive this approximation, nor does he describe this equation as an approximation; nevertheless, this equation can be motivated by an approximation, as we will now show.

[52] In the only work we are aware of that explains how to construct the CES quantity index, De Loecker (2011) computes the weighted average of deflated revenues (see De Loecker (2011) equation B.1.9 in the appendix), though – as we show in (A.11) – theory indicates the gross sum is called for.

Starting from the model in Section 2, we impose: i) $\varepsilon_{ft}^{d}=\varepsilon_{ft}$ for all $d\in\{0\cup\Omega_{ft}\}$, (ii) $u_{ft}^{d}=u_{ft}$ for all $d\in\{0\cup\Omega_{ft}\}$, and (iii) a Cobb-Douglas structure for $F(\cdot)$, as in De Loecker (2011). Then from (2) and (5) we have

$$r_{ft}^{d}\equiv\ln R_{ft}^{d}=\rho^{d}\omega_{ft}+\rho^{d}\gamma^{M}m_{ft}+\rho^{d}\gamma^{L}\ell_{ft}+\rho^{d}\gamma^{K}k_{ft}+\rho^{d}\ln\chi_{ft}^{d}+\ln D_{t}^{d}+\varepsilon_{ft}+u_{ft}, \tag{A.14}$$

Define $\bar{\Omega}_{f}$ the export strategy of firm f in some baseline year. Then by definition, we have

$$\begin{aligned}\ln R_{ft} &= \ln\left(\sum_{d\in\{0\}\cup\Omega_{ft}}\exp(r_{ft}^{d})\right)\\ &= \ln\left(\sum_{d\in\{0\}\cup\left(\bar{\Omega}_{f}\cap\Omega_{ft}\right)}\exp(r_{ft}^{d})+\sum_{d\in\cup\Omega_{ft}\setminus\bar{\Omega}_{f}}\exp(r_{ft}^{d})\right)\\ &= \underbrace{\ln\left(\sum_{d\in\{0\}\cup\left(\bar{\Omega}_{f}\cap\Omega_{ft}\right)}\exp(r_{ft}^{d})\right)}_{\equiv G(\mathbf{r_{ft}})}+\underbrace{\ln\left(1+\frac{\sum_{d\in\left(\Omega_{ft}\setminus\bar{\Omega}_{f}\right)}\exp\left(r_{ft}^{d}\right)}{\sum_{d\in\{0\}\cup\left(\bar{\Omega}_{f}\cap\Omega_{ft}\right)}\exp\left(r_{ft}^{d}\right)}\right)}_{\equiv E_{ft}}\end{aligned} \tag{A.15}$$

where $\mathbf{r_{ft}}$ is the vector $\{r_{ft}^{d}\},d\in\{0\}\cup\left(\bar{\Omega}_{f}\cap\Omega_{ft}\right)$, i.e. the destination-specific log revenues for the set of destinations common between the export strategies $\bar{\Omega}_{f}$ and Ω_{ft}. The term $G(\mathbf{r_{ft}})$ represents the component of log revenues for firm f in year t from destinations in this set. The term E_{ft} indicates the component of log revenues from destinations in Ω_{ft}, but not in $\bar{\Omega}_{f}$.

For a fixed export strategy, $\Omega_{ft}=\bar{\Omega}_{f}$ for all f and t, and $E_{ft}=0$ for all f and t. In this case, a first-order Taylor approximation of $G(\mathbf{r_{ft}})$ around the baseline $\bar{\mathbf{r}}_{\mathbf{f}}$ will generate the estimation equation. If the export strategy varies from the baseline, then E_{ft} introduces potentially higher-order bias. Hence, the approximation still holds, but not necessarily to a first order.

Taking a first-order Taylor approximation of $G(\mathbf{r_{ft}})$ around the baseline $\bar{\mathbf{r}}_{\mathbf{f}}$ we have

$$\begin{aligned}G(\mathbf{r_{ft}}) &= G(\bar{\mathbf{r}}_{\mathbf{f}})+\sum_{d\in\{0\}\cup\left(\bar{\Omega}_{f}\cap\Omega_{ft}\right)}\frac{\partial G}{\partial\vartheta_{ft}^{d}}\Big|_{\vartheta_{ft}^{d}=\bar{r}_{f}^{d}}\left(r_{ft}^{d}-\bar{r}_{f}^{d}\right)\\ &= G(\bar{\mathbf{r}}_{\mathbf{f}})+\sum_{d\in\{0\}\cup\left(\bar{\Omega}_{f}\cap\Omega_{ft}\right)}\bar{s}_{f}^{d}\left(r_{ft}^{d}-\bar{r}_{f}^{d}\right)\\ &= \sum_{d\in\{0\}\cup\left(\bar{\Omega}_{f}\cap\Omega_{ft}\right)}\bar{s}_{f}^{d}r_{ft}^{d}+\underbrace{G(\bar{\mathbf{r}}_{\mathbf{f}})-\sum_{d\in\{0\}\cup\left(\bar{\Omega}_{f}\cap\Omega_{ft}\right)}\bar{s}_{f}^{d}\bar{r}_{f}^{d}}_{\equiv Cons_{f}}\end{aligned} \tag{A.16}$$

where $\bar{s}_{f}^{d}\equiv\frac{\bar{R}_{f}^{d}}{\sum_{z\in\{0\}\cup\left(\bar{\Omega}_{f}\cap\Omega_{ft}\right)}\bar{R}_{f}^{z}}$ is the revenues share of destination d in the common set $\{0\}\cup\left(\bar{\Omega}_{f}\cap\Omega_{ft}\right)$, and $Cons_{f}$ is a firm-specific constant that does not depend on t.

Hence, to an approximation,

$$\ln R_{ft} \approx \sum_{d \in \{0\} \cup \left(\bar{\Omega}_f \cap \Omega_{ft}\right)} \bar{s}_f^d r_{ft}^d + Cons_f + E_{ft} \tag{A.17}$$

A natural choice for baseline $\bar{\mathbf{r}}_\mathbf{f}$ would be initial period outcomes $\bar{\mathbf{r}}_\mathbf{f} = \mathbf{r}_{\mathbf{f0}}$. Substituting in (A.14) and deflating the log of total revenues by the weighted average of price indices in each destination, we have

$$\begin{aligned}
\ln R_{ft} - \sum_{d \in \{0\} \cup \left(\bar{\Omega}_f \cap \Omega_{ft}\right)} s_{f0}^d \ln \Gamma_t^d &= \sum_{d \in \{0\} \cup \left(\bar{\Omega}_f \cap \Omega_{ft}\right)} s_{f0}^d \rho^d \gamma^M m_{ft} + \sum_{d \in \{0\} \cup \left(\bar{\Omega}_f \cap \Omega_{ft}\right)} s_{f0}^d \rho^d \gamma^L \ell_{ft} \\
&+ \sum_{d \in \{0\} \cup \left(\bar{\Omega}_f \cap \Omega_{ft}\right)} s_{f0}^d \rho^d \gamma^K k_{ft} \sum_{d \in \{0\} \cup \left(\bar{\Omega}_f \cap \Omega_{ft}\right)} s_{f0}^d (1 - \rho^d) \ln B_t^d \\
&+ \sum_{d \in \{0\} \cup \left(\bar{\Omega}_f \cap \Omega_{ft}\right)} s_{f0}^d \rho^d \ln \chi_{ft}^d - \sum_{d \in \{0\} \cup \left(\bar{\Omega}_f \cap \Omega_{ft}\right)} s_{f0}^d \rho^d \ln \tau_t^d \\
&+ Cons_f + E_{ft} + \left(\sum_{d \in \{0\} \cup \left(\bar{\Omega}_f \cap \Omega_{ft}\right)} s_{f0}^d \rho^d \right) \omega_{ft} + \varepsilon_{ft} + u_{ft}.
\end{aligned}$$

If we denote $\beta^j \equiv \sum_{d \in \{0\} \cup \left(\bar{\Omega}_f \cap \Omega_{ft}\right)} s_{f0}^d \rho^d \gamma^j$ and $\beta^\omega \equiv \sum_{d \in \{0\} \cup \left(\bar{\Omega}_f \cap \Omega_{ft}\right)} s_{f0}^d \rho^d$, this expression simplifies to

$$\begin{aligned}
\widetilde{r}_{ft} &= \beta^M m_{ft} + \beta^L l_{ft} + \beta^K k_{ft} + \sum_{d \in \{0\} \cup \left(\bar{\Omega}_f \cap \Omega_{ft}\right)} s_{f0}^d \left(1 - \rho^d\right) \ln B_t^d + \sum_{d \in \{0\} \cup \left(\bar{\Omega}_f \cap \Omega_{ft}\right)} s_{f0}^d \rho^d \ln \chi_{ft}^d \\
&- \sum_{d \in \{0\} \cup \left(\bar{\Omega}_f \cap \Omega_{ft}\right)} s_{f0}^d \rho^d \ln \tau_t^d + Cons_f + E_{ft} + \beta^\omega \omega_{ft} + \varepsilon_{ft} + u_{ft}.
\end{aligned} \tag{A.18}$$

Comparing (A.18) with (38) in the text, we find several difficulties and discrepancies. First, deflating revenues by the weighted average of price indices in each destination requires a lot of information. The isomorphism between a multi-segment model and a multi-destination model does not imply similar data requirements. In practice, De Loecker (2011) uses data on Belgium prices only, which ignores the fact that Belgian firms export.

Second, there are 4 terms in (A.18) that are missing in (38)— the terms associated with χ_{ft}^ds and τ_t^d, and the terms related to the approximation ($Cons_f$ and E_{ft}). For the first term, De Loecker (2011) assumes a fixed product mix over time within firms to proxy for revenues shares with the number of products N_f. In our setting, this proxy would be equivalent to the number of destinations. But most firms change their set of destinations over time. Assuming all destination-specific revenues shares s_{ft}^ds are observed, we could use that information to proxy for χ_{ft}^ds, or we could proxy it by the maximum number of destinations. The second term that contains trade costs is specific to our setting. The constant is not time-varying, but it is firm specific, hence it could be absorbed by firm fixed effects in the Markov process (which could introduce Nickell bias). But these fixed effects are not present in (38), thereby relegating this constant to the error term. As a result, the baseline period export share could generate an omitted variable problem in the specification proposed by De Loecker (2011), and the moment conditions may not hold. Finally, our derivation indicates that the extensive margin for destinations (E_{ft}) should be controlled for.

B Single Market Monte Carlo Simulations

In this appendix, we test the performance of the factor shares single-market estimator from Section 4.2 and a single-market version of the control function method from Section 4.3 when the underlying data generating process features just a single market. We thus simulate data as described in Section 2 assuming there is just a single market. With these simulated data, we estimate output elasticities and the curvature of the demand function. These simulations extend the Monte Carlo experiments from Gandhi et al. (2020) and Ackerberg et al. (2023) to the case of heterogeneous products with missing output price data, and highlight the advantages of the factor share method over the control function approach.

For the data generating process, we impose that firms produce with a Cobb-Douglas production function with one flexible input, materials (M), and one quasi-fixed input, capital (K):

$$Q_{ft} = \exp\left(\omega_{ft}\right) M^{\gamma^M} K^{\gamma^K} \tag{B.1}$$

with $\gamma^M = 0.8$ and $\gamma^K = 0.3$. Capital updates each period according to the law of motion: $K_{ft} = 0.9 K_{f,t-1} + \iota_{f,t-1}$, where $\iota_{ft} = \exp(0.8\rho\,\omega_{ft} + 0.8\varepsilon_{ft})\,(K_{ft})^{0.2}$. We fix $\rho = 0.8$. While we build the data according to these restrictions for simplicity, we obviously need not impose any functional form in the estimation.

Within each replication we draw total expenditures and quantity series, and homogeneous (across firms) material input prices. At the firm level, we draw initial capital stocks $K_{f,1} \sim U(1,201)$, initial productivity shocks $\omega_{f,1} \sim N(0,0.01)$, and initial *ex ante* demand shocks $\varepsilon_{f,1} \sim N(0,0.0009)$. We let ω and ε update according to the same AR(1) process described in the main text, with $h = 0.8$ and where $\widetilde{\omega}_{ft} \sim N(0,0.01)$ and $\widetilde{\varepsilon_{ft}} \sim N(0,0.0009)$. We draw *ex post* demand shocks $u_{ft} \sim N(0,0.0009)$. Firm-period quantities, revenues and inputs M_{ft} are determined given productivity, capital, materials prices and aggregate demand. We simulate 100 samples of a single industry with 500 firms over 50 periods.[53]

We estimate in each sample of simulated data the factor shares approach and the control function approach assuming researchers observe R_{ft}, M_{ft}, K_{ft}, W_t^M, B_t and Υ_t. For the factor shares model, we set initial conditions for the first-step NLLS estimation for M based on an OLS estimation of the regression

$$\ln\left[\frac{W_{ft}^m M_{ft}}{R_{ft}}\right] = g_0^m + g_m^m m_{ft} + g_k^m k_{ft} + g_{mm}^m m_{ft} m_{ft} + g_{kk}^m k_{ft} k_{ft} + g_{mk}^m m_{ft} k_{ft} + \vartheta_{ft},$$

where ϑ_{ft} is a regression residual. For the second step GMM, we set initial conditions based on an OLS estimation of the regression

$$\widetilde{r}_{ft} = g_k k_{ft} + g_{kk} k k_{ft} + g_D B_t + \vartheta'_{ft},$$

where ϑ'_{ft} is a regression residual and B_t is the true CES quantity index.

[53]In the single-market case, we follow closely the experiments presented in Gandhi et al. (2020). Gandhi et al. (2020) posit a long panel (50 periods) in order to give the control function a reasonable chance to identify the structural parameters. Since the output elasticity for materials is identified purely from time-series variation in the material input price, there is little chance that the control function identifies structural parameters in panels of only 10-15 years (the type of duration one usually observes in balance sheet datasets). In the multi-market simulations below, we can entertain much shorter panels.

For the control function approach, we set initial conditions for the second-step GMM based on an OLS estimation of the regression

$$\widetilde{r}_{ft}^{CF} \quad = \quad g_0 + g_D B_t + g_m^m m_{ft} + g_k^m k_{ft} + g_{mm}^m m_{ft} m_{ft} + g_{kk}^m k_{ft} k_{ft} + g_{mk}^m m_{ft} k_{ft} + \vartheta_{ft}'',$$

where $\widetilde{r}_{ft}^{CF}$ represents log revenues net of the residual from the control function first step, and ϑ_{ft}'' is a regression residual. We denote the resulting control function estimates as "CF ls", since the GMM starts from the OLS point estimates. We also start the control function estimation from the true parameter values and refer to the resulting estimates as "CF tr".

In Figure B.1 we present the distribution of estimates of ρ and the average material and capital output elasticities across the 100 samples by estimator, along with the average ("av") and median ("md") of the distributions. In the top row we present the case of high input price variation. True material and capital output elasticities are constant across firms and over time ($\sigma^M = \gamma^M = 0.8$ and $\sigma^K = \gamma^K = 0.3$), and are depicted with vertical black lines.[54]

The distribution of the estimates from the factor shares approach is depicted in solid blue. For each empirical object ρ, σ^K , σ^M, the distribution of estimates appears to be centered on the true values. Averages and medians of the distributions are identical with the truth out to at least two decimal places. Similarly, the distribution of the control function estimates *taking true values as the initial conditions* (black solid line) also appears to be centered on the truth, with averages and medians of distributions identical to the truth out to at least two decimal places. The distribution of the solid blue line is clearly narrower than the distribution of the solid black line, indicating that the factor shares approach is more efficient.

In contrast, when the second step of the control function method starts the optimization algorithm from the OLS values (black dashed line), the distributions of estimates are clearly biased. Estimates of ρ and σ^M tend to center around 1, and estimates of σ^K center around 0. These values coincide roughly with the median results from a naïve OLS estimate of the production function (depicted with a vertical red line).

Also in Figure B.1, we present in the bottom row the distribution of estimates for the case of low input price variation. The factor shares method still recovers unbiased estimates of structural parameters. However, the estimates from the control function method are biased even when the GMM optimization starts from the true parameters. As explained by Nelson & Startz (1990), instrumental variables estimators are biased towards OLS in finite samples with weak instruments. We can see this pattern from the medians of the distributions in solid black. The distributions are wide, and outliers severely distort the means, but the medians indicate that the estimates of σ^M are biased up and the estimates of σ^K are biased down, as they are in OLS. This is the same pattern found in Monte Carlo simulations by Gandhi et al. (2020) for the single-market case in which quantities are observed.

Comparing the solid black line to the dashed black line, it is clear that the control function is sensitive to initial starting conditions. We explore this sensitivity further by plotting the distribution of parameter estimates from the control function method when varying systematically the initial values of the second step of the GMM procedure below. We run the same Monte Carlo simulations as in the main text. We

[54]Estimated material and capital output elasticities vary both due to sampling error and with the level of capital because we allow for higher order terms in capital and interactions between inputs in the estimation process (see equations (24) and (26)).

Figure B.1: Parameter Estimates in the Single-Market Simulations

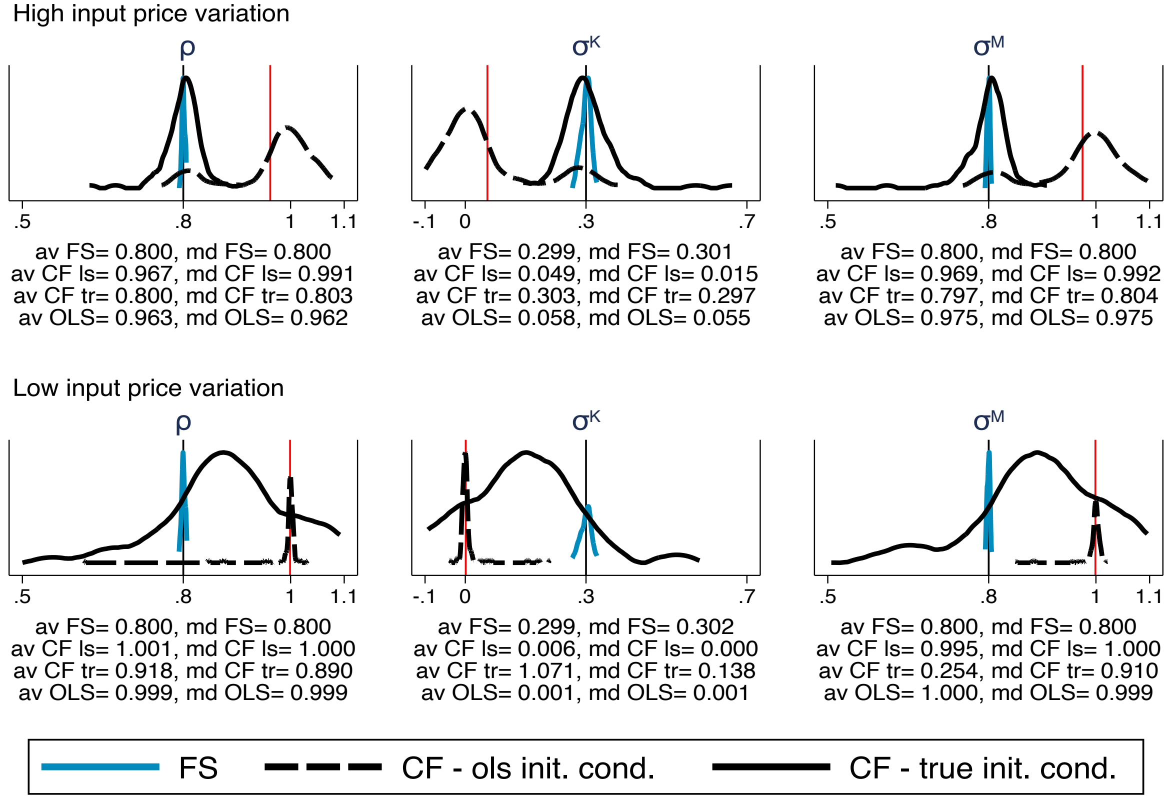

Notes. The figure reports the distribution of averages of estimates across 100 Monte Carlo samples. Top (bottom) row presents results for high (low) input price variation. True parameter values are depicted as vertical lines. Averages ("av") and medians ("md") of distributions for each estimator are reported below each subfigure. "FS" indicates the factor shares method. "CF - ls" indicates control function method where the GMM optimization starts from the OLS values. "CF - tr" indicates control function method where the GMM optimization starts from the true model parameters. Red line indicates the median estimate based on an OLS regression.

vary initial conditions for $\beta_0^D \in \{0, 0.05, 0.1, 0.15, 0.2, 0.25\}$, $\beta_0^M \in \{0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8\}$ and $\beta_0^K \in \{0, 0.1, 0.2, 0.3, 0.4\}$. We compute the GMM solution starting from every combination defined by these three sets. Recall that the true parameter values are $\beta^D = 1 - \rho = 0.2$, $\beta^M = \rho\sigma^M = 0.64$ and $\beta^K = \rho\sigma^K = 0.24$.

In Figure B.2 (B.3), we present results when the data is simulated with a high (low) degree of time series variation in material input prices. We display the distribution of GMM solutions for each parameter for two different Monte Carlo samples, one in each row. Results starting from the OLS estimates are depicted with a vertical red dashed line. Results from all other starting values are depicted in solid bars in blue.

In Figure B.2, we see that the estimates based on the OLS initial values coincide with the results in Figure B.1: starting from the OLS values, the GMM solution tends towards $\rho = 1$, $\sigma^M = 1$, and $\sigma^K = 0$ (red dashed line). In Figure B.2, we also see that when the GMM starts from other initial conditions there is a mass point of convergence around the same values, although we also see other mass points. This bunching pattern is consistent with Ackerberg et al. (2023).

In Figure B.3, we see that the estimates converges more often towards the OLS results. That is, the distribution is bunched more tightly around the OLS estimates than in Figure B.2. This is what we would expect, as with low input price variation, the GMM suffers from weak instruments, which tends to bias the estimates towards the OLS result.

Figure B.2: Different Starting Values for Control Function Estimation, High Material Price Variation

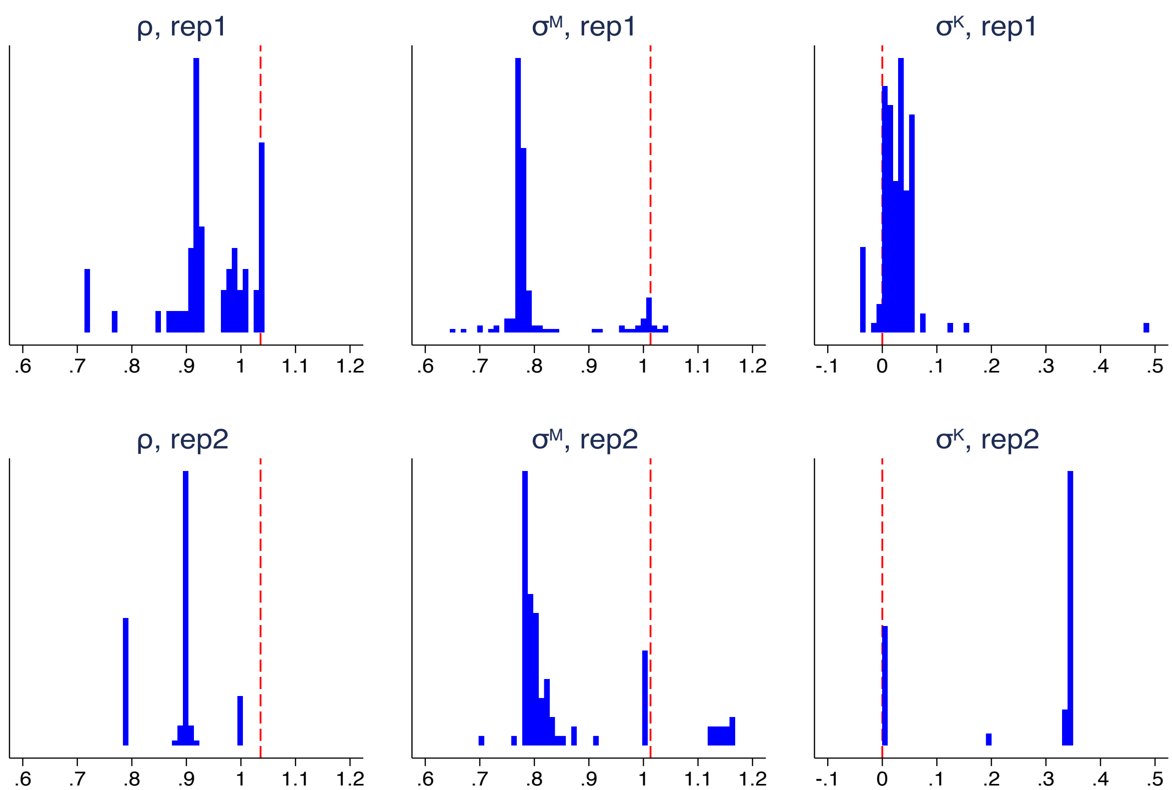

Notes. The figure reports the distribution of estimates of ρ and averages of estimates of factor output elasticities resulting from the control function method taking different parameter vectors as starting values for the second-step GMM procedure. Each row reports results for a single Monte Carlo sample. Data is generated based on the single-market scenario described in Section B, with high input price variation. The true parameter values are $\rho = 0.8$, $\sigma^M = 0.8$ and $\sigma^K = 0.3$.

Figure B.3: Different Starting Values for Control Function Estimation, Low Material Price Variation

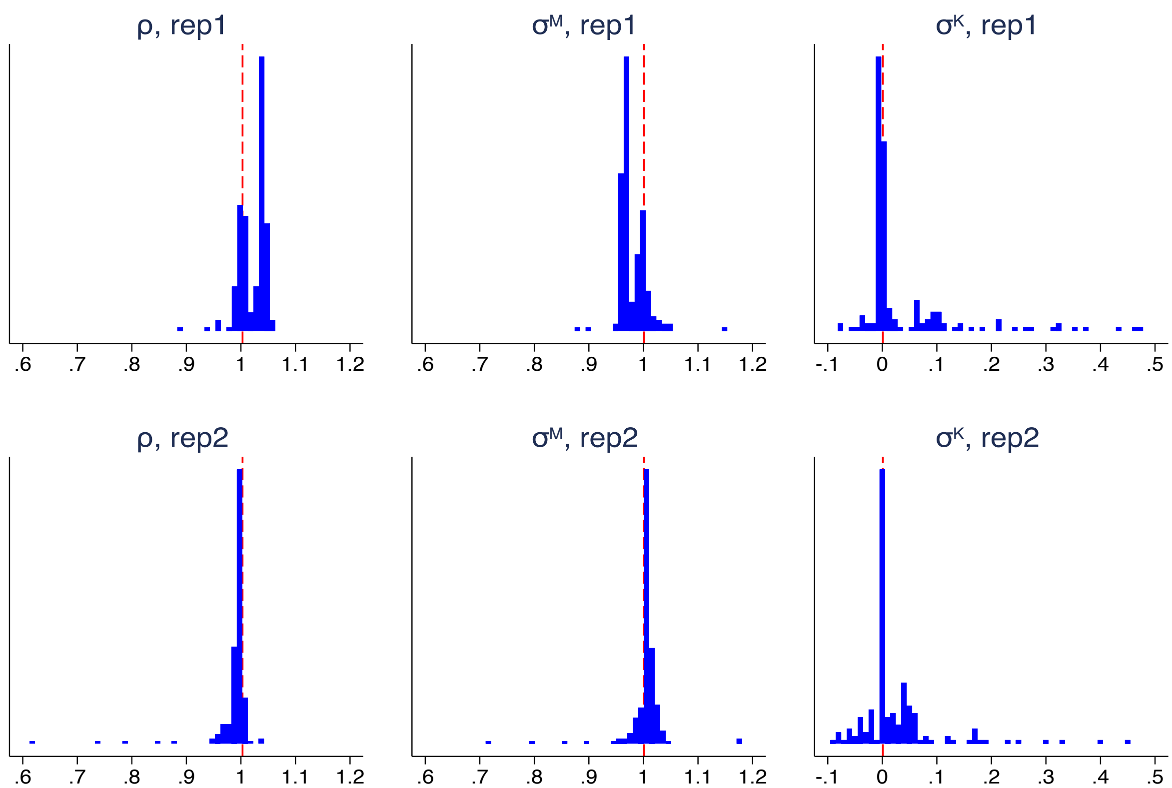

Notes. The figure reports the distribution of estimates of ρ and averages of estimates of factor output elasticities resulting from the control function method taking different parameter vectors as starting values for the second-step GMM procedure. Each row reports results for a single Monte Carlo sample. Data is generated based on the single-market scenario described in Section B, with low input price variation. The true parameter values are $\rho = 0.8$, $\sigma^M = 0.8$ and $\sigma^K = 0.3$.

C Additional Monte Carlo Simulation Results

Figure C.1: Distribution of Parameter Estimates, Multi-Destination Factor Shares Approach

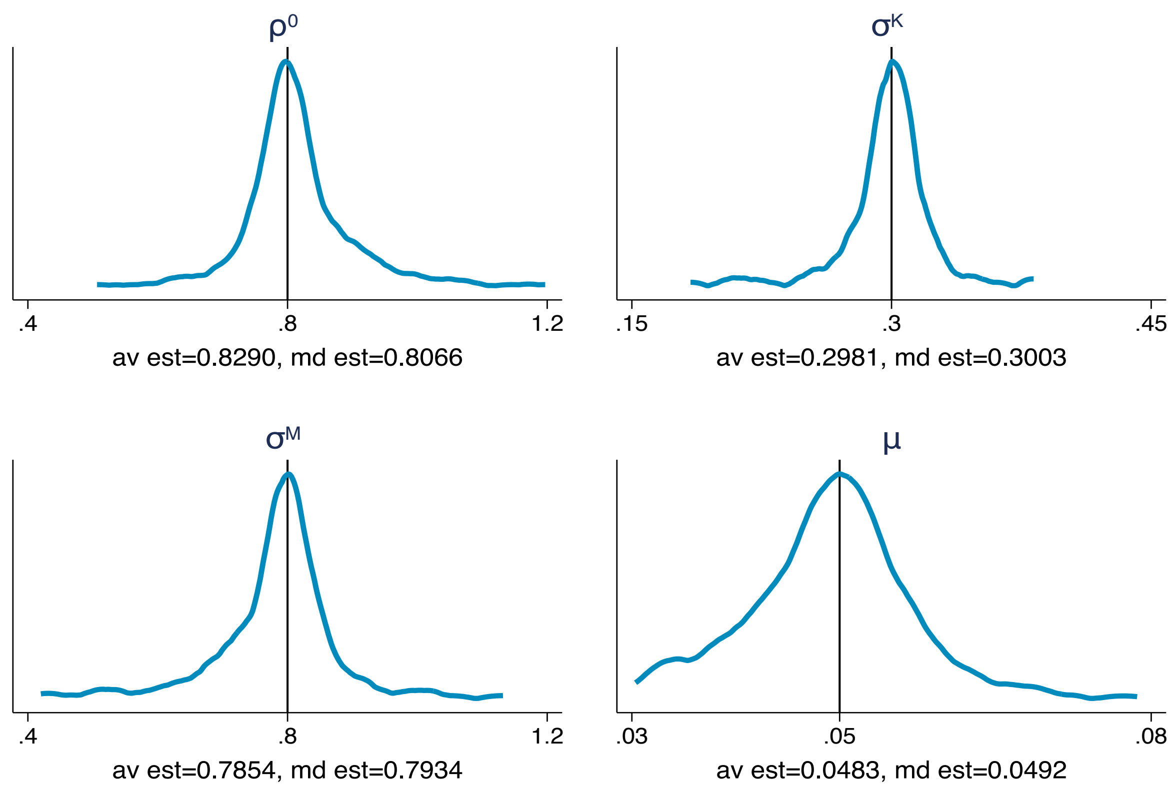

Notes. The figure reports distributions of point estimates across 200 simulations estimated by our multi-destination factor shares approach. Means ("av est") and medians ("md est") of each distribution is reported below each subfigure.

Figure C.2: Distribution of Parameter Estimates, Alternative Factor Shares Approach

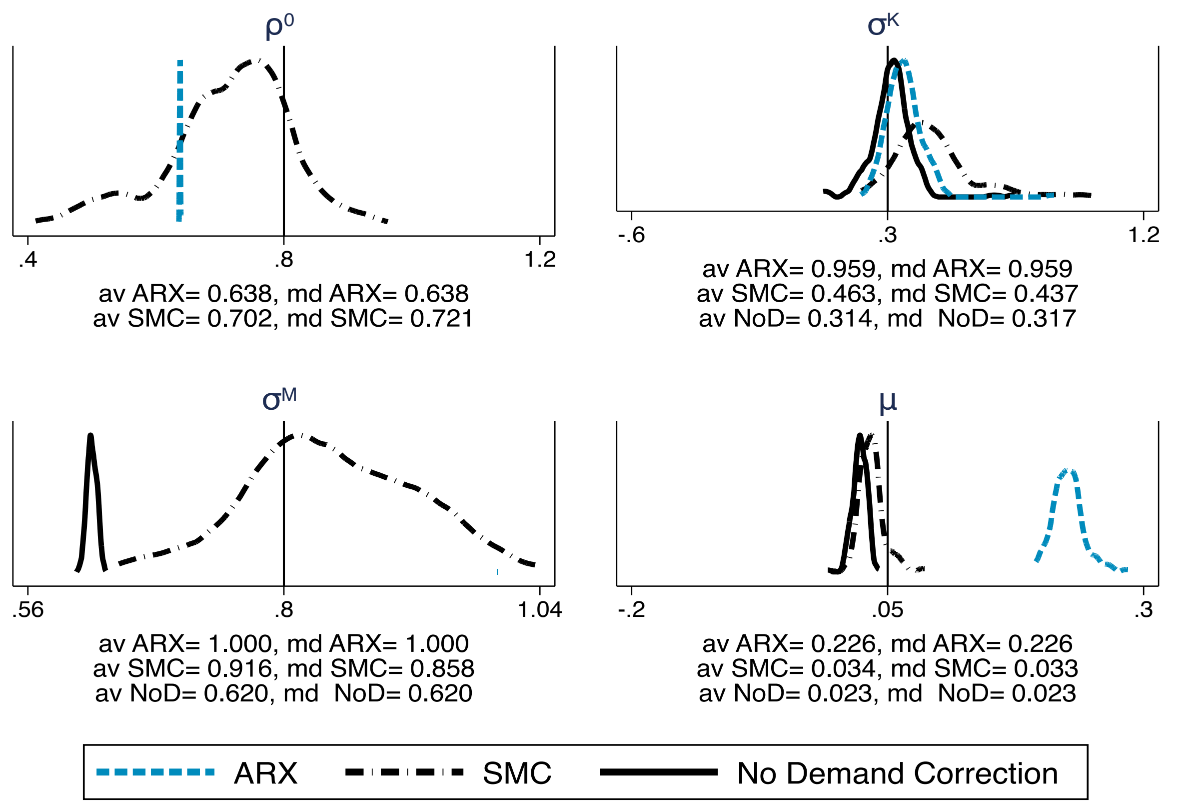

Notes. The figure reports distributions of point estimates across 200 simulations estimated by alternative factor shares approaches. Means (“av est”) and medians (“md est”) of each distribution is reported below each subfigure. “ARX” refers to the estimator from Aw et al. (2011) that imposes constant returns to variable inputs. “SMC” refers to the single market correction model (appendix of Gandhi et al. (2020)) “NoD” refers to the estimator in which revenues are deflates by an industry wide deflator and the resulting series is treated as if were a quantities series (main text of Gandhi et al. (2020)).

Figure C.3: Distribution of Parameter Estimates, Control Function Methods

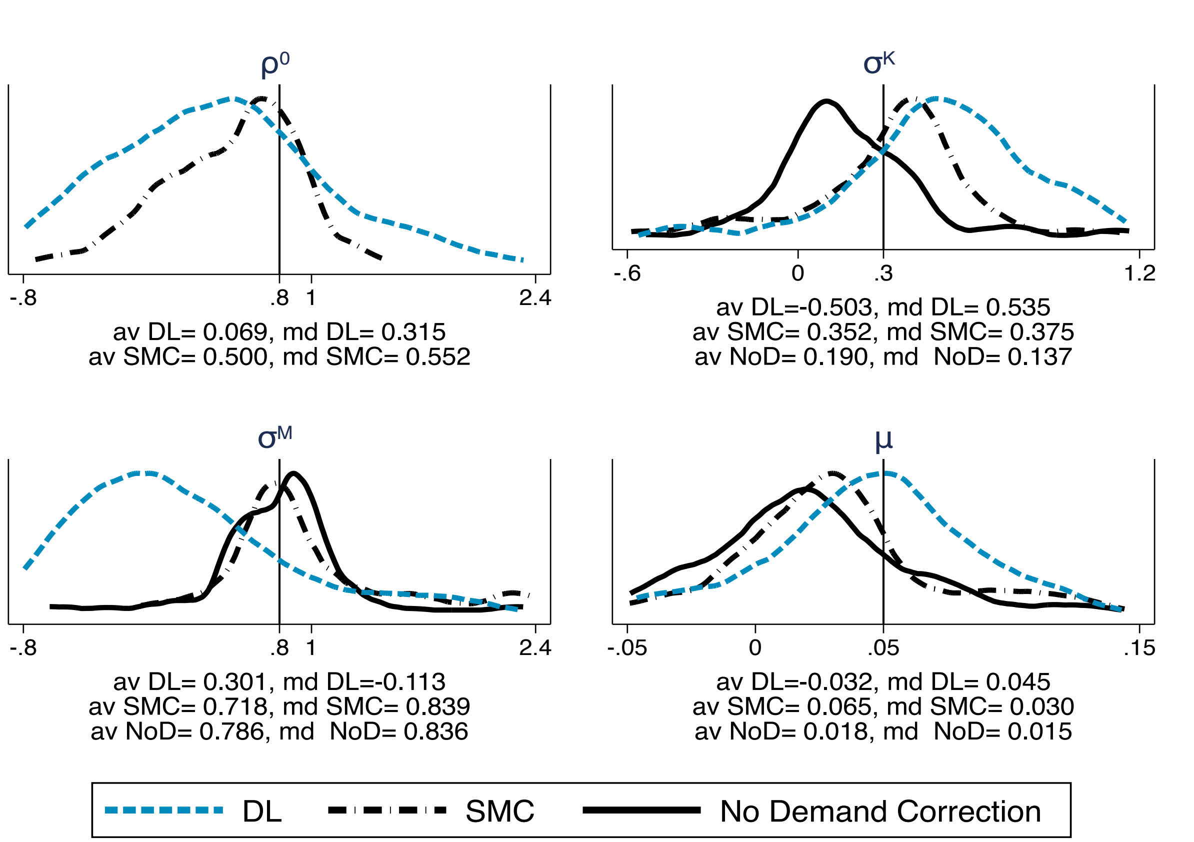

Notes. The figure reports distributions of point estimates across 200 simulations. Means ("av est") and medians ("md est") of each distribution is reported below each subfigure. "DL" refers to the multi-destination control function estimator of De Loecker (2011). "SMC" refers to the single market correction model, estimated by control function. "NoD" refers to the estimator in which revenues are deflates by an industry wide deflator and the resulting series is treated as if were a quantities series (control function approach described in the main text of Gandhi et al. (2020)).

D Data

Firm-level balance sheet information is reported in the FICUS (*Fichier complet unifié de SUSE*) and FARE (*Fichier Approché des Résultats ESANE*) datasets, which cover the periods 1994–2007 and 2008–2016, respectively. These data originate in tax declarations of all firms in France, and are collected by the French National Institute of Statistics and Economic Studies, INSEE. We use total revenue, expenditure on materials, employment and the book value of capital.

We construct capital stocks following the methodology proposed by Bonleu et al. (2013) and Cette et al. (2015). We start with the book value of capital. Since the stocks are recorded at historical cost, i.e., the value at the time of entry into the firm i's balance sheet, an adjustment has to be made to move from stocks valued at historic cost ($K^{BV}_{i,s,t}$) to stocks valued at current prices ($K_{i,s,t}$). We deflate K^{BV} by an industry-specific price index (sourced from INSEE) that assumes that the price of capital is equal to the sectoral price of investment T years before the date when the first book value was available, where T is the corrected average age of capital, hence $p^K_{s,t+1} = p^I_{s,t-T}$. The average age of capital is computed using the share of depreciated capital, $DK^{BV}_{i,s,t}$ in the capital stock at historical cost:

$$T = \frac{DK^{BV}_{i,s,t}}{K^{BV}_{i,s,t}} \times \widetilde{A}$$

where

$$\widetilde{A} = median_{i \in S} \left(\frac{K^{BV}_{i,s,t}}{\Delta DK^{BV}_{i,s,t}} \right)$$

where S the set of firms in a sector. We use the median value $\widetilde{A}$ to reduce the volatility in the data, as investments within firms are discrete events.

Data on firms' exports are from the French Customs. For each observation, we know the value of exports of the firm. We use the firm-level SIREN identifier to match the trade data to FICUS/FARE. This match is not perfect. The imperfect match is because there are SIRENs in the trade data for which there is no corresponding SIREN in FICUS/FARE. This may lead to measurement error: for some firms, we will observe zero exports even when true exports are positive. This is not a big concern because most of the missing values are in the oil refining industry, which we drop from our sample.

E Additional Results from French Manufacturing

In this section, we first report the detailed results that are presented in the paper graphically and then report two additional sets of results. The first allows for dynamic, partial adjustment for labor and serves as a robustness check for the main results. The second set of results applies the control function method, which is just for comparison, as we expect finite sample bias and weak moments problems.

E.1 Results with pre-determined labor

Table E.1: Estimates using Multi-Market Estimator, Predetermined Labor Input

Industry	σ^M	σ^L	σ^K	RTS	ρ^0	$\eta^0 = \frac{1}{(\rho^0-1)}$	μ	h	$\mu/(1-h)$
Autos and transport equipment	0.403	0.524	0.169	1.096	0.932	-14.685	0.013	0.816	0.073
	(0.009)	(0.016)	(0.010)	(0.020)	(0.017)	(4.554)	(0.003)	(0.015)	(0.014)
Chemical products	0.386	0.537	0.176	1.098	0.949	-19.423	0.014	0.877	0.110
	(0.009)	(0.019)	(0.010)	(0.018)	(0.013)	(6.395)	(0.003)	(0.013)	(0.017)
Computer, electronics	0.332	0.575	0.138	1.045	0.915	-11.772	0.018	0.845	0.119
	(0.010)	(0.019)	(0.010)	(0.028)	(0.023)	(3.302)	(0.003)	(0.008)	(0.018)
Electrical equipment	0.386	0.521	0.140	1.047	0.911	-11.245	0.016	0.839	0.100
	(0.010)	(0.018)	(0.008)	(0.029)	(0.023)	(4.091)	(0.003)	(0.010)	(0.016)
Food, beverage, tobacco	0.488	0.536	0.326	1.350	0.719	-3.563	-0.005	0.783	-0.021
	(0.011)	(0.013)	(0.007)	(0.028)	(0.015)	(0.188)	(0.003)	(0.006)	(0.012)
Machinery and equipment	0.344	0.686	0.147	1.177	0.812	-5.331	0.028	0.796	0.136
	(0.006)	(0.013)	(0.004)	(0.020)	(0.014)	(0.387)	(0.001)	(0.006)	(0.005)
Basic metal and fabricated metal	0.278	0.763	0.230	1.271	0.771	-4.369	0.008	0.846	0.054
	(0.006)	(0.018)	(0.007)	(0.029)	(0.017)	(0.334)	(0.001)	(0.004)	(0.007)
Other manufacturing	0.270	0.616	0.218	1.104	0.839	-6.205	0.024	0.852	0.166
	(0.007)	(0.013)	(0.009)	(0.027)	(0.022)	(0.710)	(0.002)	(0.006)	(0.008)
Rubber and plastic	0.376	0.529	0.199	1.104	0.903	-10.328	-0.000	0.870	-0.003
	(0.022)	(0.044)	(0.029)	(0.094)	(0.045)	(2.860)	(0.003)	(0.009)	(0.019)
Textiles, wearing apparel	0.338	0.535	0.228	1.102	0.810	-5.270	0.040	0.896	0.380
	(0.018)	(0.030)	(0.015)	(0.060)	(0.043)	(1.487)	(0.002)	(0.004)	(0.016)
Wood, paper products	0.314	0.732	0.197	1.244	0.790	-4.755	0.018	0.837	0.112
	(0.006)	(0.013)	(0.005)	(0.021)	(0.014)	(0.313)	(0.001)	(0.006)	(0.006)

Notes. The table reports estimates based on the multi-market estimator, treating labor as predetermined (like capital): average output elasticities σ^j for materials input ($j = M$), labor ($j = L$) and capital ($j = M$), overall returns to scale (RTS), the demand elasticity on the domestic market $\eta^0 = 1/(\rho^0 - 1)$, the coefficient to learning by exporting μ, the persistence parameter in the controlled Markov h, and the long run effect of exporting $\mu/(1-h)$. Bootstrap standard errors clustered by firm in parentheses.

Figure E.4: Returns to Scale by Industry

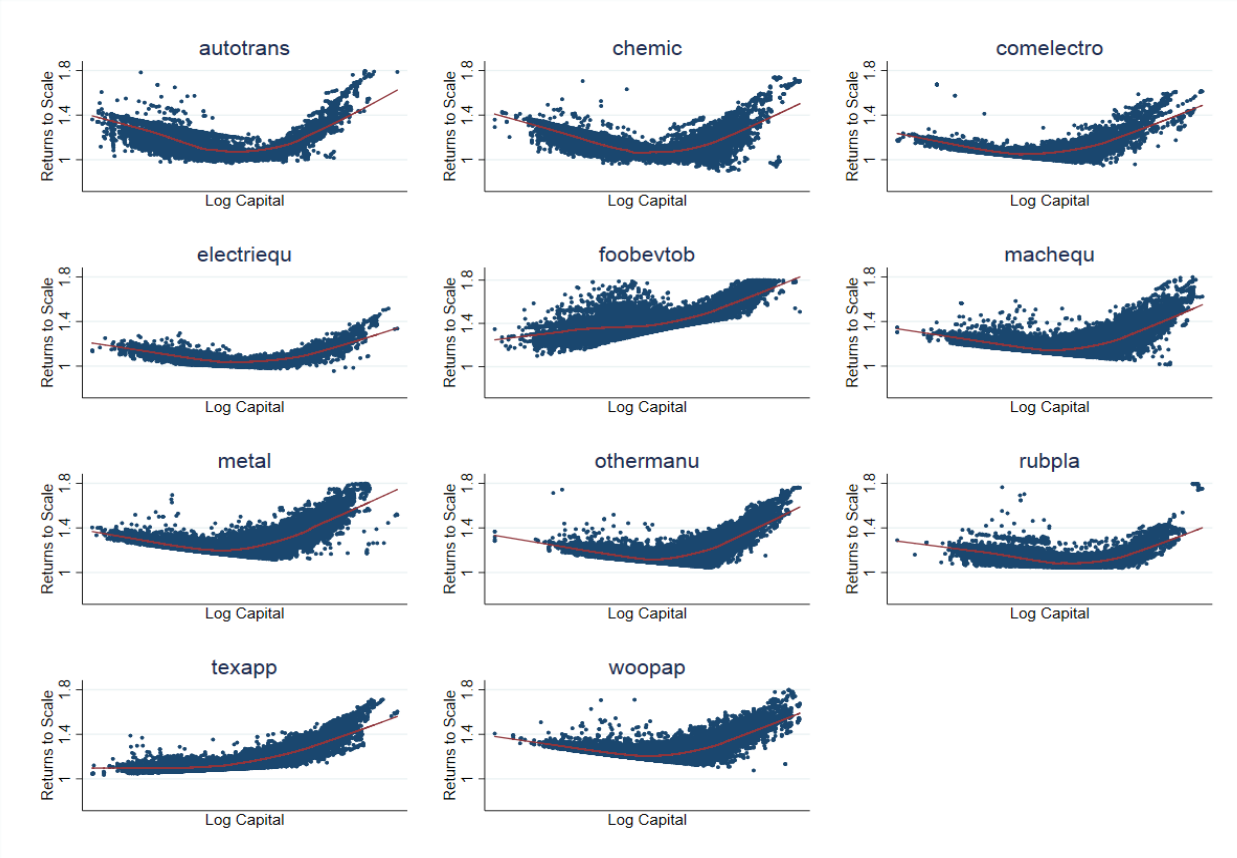

Notes. This figure presents estimated returns to scale via the multi-market factor shares estimator by firm-year against log of capital. These are built using the same specification in Table E.1, where labor is assumed to be pre-determined.

Table E.2: Estimates using no Demand Correction, Predetermined Labor Input

Industry	σ^M	σ^L	σ^K	RTS	ρ	$\eta = \frac{1}{(\rho-1)}$	μ	h	$\mu/(1-h)$
Autos and transport equipment	0.364	0.546	0.145	1.054	-	-	0.044	0.930	0.622
	(0.004)	(0.012)	(0.007)	(0.010)			(0.005)	(0.013)	(0.065)
Chemical products	0.366	0.514	0.164	1.044	-	-	0.024	0.968	0.748
	(0.004)	(0.012)	(0.009)	(0.007)			(0.004)	(0.005)	(0.072)
Computer, electronics	0.299	0.534	0.148	0.981	-	-	0.038	0.939	0.619
	(0.004)	(0.010)	(0.007)	(0.008)			(0.003)	(0.006)	(0.054)
Electrical equipment	0.339	0.491	0.145	0.974	-	-	0.034	0.952	0.717
	(0.003)	(0.009)	(0.007)	(0.007)			(0.004)	(0.005)	(0.054)
Food, beverage, tobacco	0.346	0.401	0.160	0.908	-	-	0.092	0.922	1.185
	(0.001)	(0.003)	(0.002)	(0.003)			(0.002)	(0.002)	(0.027)
Machinery and equipment	0.283	0.558	0.109	0.950	-	-	0.054	0.881	0.450
	(0.008)	(0.007)	(0.003)	(0.012)			(0.003)	(0.010)	(0.039)
Basic metal and fabricated metal	0.217	0.596	0.155	0.968	-	-	0.035	0.930	0.503
	(0.001)	(0.003)	(0.003)	(0.002)			(0.001)	(0.002)	(0.013)
Other manufacturing	0.224	0.530	0.167	0.922	-	-	0.039	0.943	0.688
	(0.001)	(0.004)	(0.003)	(0.005)			(0.002)	(0.002)	(0.023)
Rubber and plastic	0.333	0.494	0.180	1.007	-	-	0.016	0.971	0.543
	(0.002)	(0.006)	(0.004)	(0.003)			(0.002)	(0.002)	(0.034)
Textiles, wearing apparel	0.269	0.456	0.170	0.894	-	-	0.040	0.951	0.823
	(0.004)	(0.006)	(0.004)	(0.006)			(0.002)	(0.002)	(0.035)
Wood, paper products	0.251	0.580	0.146	0.977	-	-	0.035	0.950	0.695
	(0.001)	(0.004)	(0.002)	(0.003)			(0.002)	(0.002)	(0.019)

Notes. The table reports estimates without correcting for demand at all, using the factor share approach, treating labor as predetermined (like capital): average output elasticities σ^j for materials input ($j = M$), labor ($j = L$) and capital ($j = M$), overall returns to scale (RTS), the demand elasticity $\eta = 1/(\rho - 1)$, the coefficient to learning by exporting μ, the persistence parameter in the controlled Markov h, and the long run effect of exporting $\mu/(1-h)$. Bootstrap standard errors clustered by firm in parentheses.

Table E.3: Estimates using Single-Market Estimator (SMC), Predetermined Labor Input

Industry	σ^M	σ^L	σ^K	RTS	ρ	$\eta = \frac{1}{(\rho-1)}$	μ	h	$\mu/(1-h)$
Autos and transport equipment	0.376	0.559	0.159	1.094	0.967	-29.864	0.024	0.792	0.117
	(0.005)	(0.013)	(0.007)	(0.011)	(0.005)	(4.671)	(0.003)	(0.020)	(0.012)
Chemical products	2.535	3.557	1.162	7.254	0.144	-1.169	0.103	0.874	0.813
	(0.613)	(0.888)	(0.284)	(1.768)	(0.024)	(0.032)	(0.032)	(0.014)	(0.245)
Computer, electronics	0.392	0.711	0.184	1.287	0.764	-4.231	0.032	0.838	0.194
	(0.083)	(0.167)	(0.046)	(0.295)	(0.085)	(1.467)	(0.010)	(0.010)	(0.074)
Electrical equipment	0.464	0.693	0.184	1.340	0.730	-3.707	0.031	0.833	0.186
	(0.021)	(0.034)	(0.010)	(0.059)	(0.031)	(0.422)	(0.004)	(0.012)	(0.021)
Food, beverage, tobacco	0.902	1.034	0.440	2.376	0.384	-1.623	0.090	0.849	0.598
	(0.047)	(0.058)	(0.023)	(0.126)	(0.007)	(0.013)	(0.008)	(0.005)	(0.039)
Machinery and equipment	0.504	0.997	0.197	1.699	0.561	-2.278	0.061	0.794	0.297
	(2.534)	(5.491)	(1.165)	(9.190)	(0.109)	(0.257)	(0.339)	(0.023)	(2.946)
Basic metal and fabricated metal	0.430	1.114	0.356	1.899	0.506	-2.024	0.027	0.851	0.181
	(0.012)	(0.030)	(0.013)	(0.054)	(0.015)	(0.064)	(0.001)	(0.006)	(0.010)
Other manufacturing	-0.366	-0.853	-0.319	-1.538	-0.614	-0.620	-0.035	0.876	-0.285
	(0.064)	(0.152)	(0.053)	(0.269)	(0.099)	(0.038)	(0.008)	(0.006)	(0.051)
Rubber and plastic	-0.833	-1.223	-0.470	-2.525	-0.400	-0.714	-0.008	0.872	-0.064
	(0.098)	(0.149)	(0.051)	(0.296)	(0.047)	(0.024)	(0.002)	(0.010)	(0.016)
Textiles, wearing apparel	0.180	0.309	0.112	0.602	1.494	2.026	0.027	0.895	0.255
	(0.005)	(0.007)	(0.003)	(0.013)	(0.028)	(0.106)	(0.002)	(0.007)	(0.011)
Wood, paper products	1.709	3.905	1.045	6.659	0.147	-1.173	0.117	0.849	0.775
	(154.128)	(344.430)	(95.748)	(594.305)	(0.053)	(0.063)	(9.705)	(0.006)	(66.938)

Notes. The table reports estimates based on the single-market estimator *à la* Klette & Griliches (1996), using the factor share approach, treating labor as predetermined (like capital): average output elasticities σ^j for materials input ($j = M$), labor ($j = L$) and capital ($j = M$), overall returns to scale (RTS), the demand elasticity $\eta = 1/(\rho - 1)$, the coefficient to learning by exporting μ, the persistence parameter in the controlled Markov h, and the long run effect of exporting $\mu/(1-h)$. Bootstrap standard errors clustered by firm in parentheses.

Table E.4: Estimates using Single-Market Estimator (SMC) on Sample of Non-Exporters, Predetermined Labor Input

Industry	σ^M	σ^L	σ^K	RTS	ρ	$\eta = \frac{1}{(\rho-1)}$	μ	h	$\mu/(1-h)$
Autos and transport equipment	0.376	0.570	0.169	1.115	0.966	-29.242	-	0.799	-
	(0.005)	(0.013)	(0.007)	(0.012)	(0.005)	(4.653)		(0.019)	
Chemical products	2.837	4.060	1.339	8.236	0.129	-1.148	-	0.874	-
	(0.877)	(1.306)	(0.407)	(2.575)	(0.025)	(0.032)		(0.014)	
Computer, electronics	0.462	0.855	0.237	1.554	0.647	-2.833	-	0.844	-
	(0.268)	(0.551)	(0.164)	(0.983)	(0.093)	(0.692)		(0.009)	
Electrical equipment	0.475	0.720	0.207	1.402	0.714	-3.493	-	0.837	-
	(0.023)	(0.038)	(0.011)	(0.066)	(0.032)	(0.378)		(0.011)	
Food, beverage, tobacco	0.879	1.015	0.448	2.342	0.394	-1.651	-	0.856	-
	(0.046)	(0.057)	(0.023)	(0.124)	(0.008)	(0.013)		(0.005)	
Machinery and equipment	0.518	1.043	0.222	1.783	0.546	-2.204	-	0.804	-
	(0.757)	(1.667)	(0.400)	(2.823)	(0.117)	(0.254)		(0.022)	
Basic metal and fabricated metal	0.437	1.143	0.376	1.955	0.498	-1.992	-	0.853	-
	(0.013)	(0.031)	(0.014)	(0.055)	(0.015)	(0.061)		(0.006)	
Other manufacturing	-0.319	-0.764	-0.288	-1.371	-0.704	-0.587	-	0.881	-
	(0.047)	(0.114)	(0.039)	(0.200)	(0.097)	(0.033)		(0.005)	
Rubber and plastic	-0.831	-1.228	-0.472	-2.531	-0.401	-0.714	-	0.872	-
	(0.097)	(0.149)	(0.051)	(0.295)	(0.046)	(0.023)		(0.010)	
Textiles, wearing apparel	0.177	0.316	0.127	0.619	1.524	1.907	-	0.913	-
	(0.005)	(0.007)	(0.004)	(0.013)	(0.032)	(0.105)		(0.005)	
Wood, paper products	2.193	5.067	1.404	8.664	0.115	-1.129	-	0.853	-
	(109.253)	(252.483)	(70.471)	(432.207)	(0.053)	(0.059)		(0.006)	

Notes. The table reports estimates based on the single-market estimator *à la* Klette & Griliches (1996), using the factor share approach, treating labor as predetermined (like capital), where we restrict the sample to non-exporting firms: average output elasticities σ^j for materials input ($j = M$), labor ($j = L$) and capital ($j = M$), overall returns to scale (RTS), the demand elasticity $\eta = 1/(\rho - 1)$, the coefficient to learning by exporting μ, the persistence parameter in the controlled Markov h, and the long run effect of exporting $\mu/(1-h)$. Bootstrap standard errors clustered by firm in parentheses.

Table E.5: Estimates using the approach of Aw et al. (2011), Predetermined Labor Input

Industry	σ^M	σ^L	σ^K	RTS	ρ^0	$\eta^0 = \frac{1}{(\rho^0-1)}$	ρ^x	$\eta^x = \frac{1}{(\rho^x-1)}$	μ	h	$\mu/(1-h)$
Autos and transport equipment	1.000	1.011	0.248	2.259	0.735	-3.779	0.610	-2.562	0.035	0.813	0.188
		(1.427)	(0.339)	(1.763)	(0.001)	(0.295)	(0.002)	(0.071)	(0.059)	(0.011)	(0.319)
Chemical products	1.000	0.925	0.341	2.266	0.743	-3.888	0.405	-1.680	-0.016	0.868	-0.123
		(0.105)	(0.042)	(0.136)	(0.001)	(0.187)	(0.003)	(0.023)	(0.008)	(0.007)	(0.057)
Computer, electronics	1.000	2.614	0.594	4.208	0.271	-1.373	0.486	-1.946	-0.032	0.826	-0.185
		(0.493)	(0.115)	(0.592)	(0.003)	(0.009)	(0.002)	(0.028)	(0.024)	(0.009)	(0.143)
Electrical equipment	1.000	1.048	0.292	2.340	0.656	-2.906	0.494	-1.974	0.025	0.841	0.160
		(0.113)	(0.040)	(0.146)	(0.002)	(0.125)	(0.002)	(0.032)	(0.009)	(0.010)	(0.060)
Food, beverage, tobacco	1.000	0.959	0.552	2.511	0.549	-2.218	0.610	-2.567	0.341	0.793	1.648
		(0.051)	(0.030)	(0.081)	(0.000)	(0.007)	(0.001)	(0.028)	(0.020)	(0.003)	(0.092)
Machinery and equipment	1.000	1.223	0.248	2.471	0.573	-2.340	0.425	-1.739	0.055	0.789	0.263
		(0.631)	(0.126)	(0.757)	(0.001)	(0.025)	(0.001)	(0.007)	(0.030)	(0.005)	(0.143)
Basic metal and fabricated metal	1.000	1.489	0.407	2.896	0.477	-1.911	0.480	-1.923	0.030	0.818	0.164
		(0.387)	(0.108)	(0.494)	(0.001)	(0.011)	(0.001)	(0.007)	(0.010)	(0.003)	(0.052)
Other manufacturing	1.000	1.288	0.524	2.813	0.462	-1.859	0.465	-1.869	0.096	0.815	0.521
		(0.062)	(0.031)	(0.090)	(0.000)	(0.006)	(0.001)	(0.007)	(0.009)	(0.005)	(0.046)
Rubber and plastic	1.000	1.486	0.922	3.408	0.393	-1.648	0.457	-1.842	-0.035	0.857	-0.245
		(0.159)	(0.098)	(0.255)	(0.001)	(0.004)	(0.001)	(0.009)	(0.008)	(0.005)	(0.054)
Textiles, wearing apparel	1.000	1.107	0.518	2.625	0.489	-1.955	0.417	-1.717	0.150	0.866	1.119
		(0.090)	(0.047)	(0.133)	(0.001)	(0.008)	(0.000)	(0.004)	(0.015)	(0.004)	(0.106)
Wood, paper products	1.000	1.507	0.524	3.032	0.454	-1.832	0.569	-2.321	0.084	0.807	0.434
		(0.149)	(0.053)	(0.201)	(0.000)	(0.005)	(0.001)	(0.021)	(0.010)	(0.005)	(0.050)

Notes. The table reports estimates based on the estimator of Aw et al. (2011), using the factor share approach, treating labor as predetermined (like capital): average output elasticities σ^j for materials input ($j = M$), labor ($j = L$) and capital ($j = M$), overall returns to scale (RTS), the domestic demand elasticity $\eta^0 = 1/(\rho^0 - 1)$, the demand elasticity on foreign markets $\eta^x = 1/(\rho^x - 1)$, the coefficient to learning by exporting μ, the persistence parameter in the controlled Markov h, and the long run effect of exporting $\mu/(1-h)$. Bootstrap standard errors clustered by firm in parentheses, except for η^0, ρ^0, η^x, ρ^x, where standard errors are clustered by firm but need not be bootstrapped.

E.2 Results allowing partial adjustment for labor

We present results that allow labor to partially adjust to contemporaneous productivity and demand shocks. This permits entertaining the possibility that firms have the ability to flexibly adjust part of employment, while another part is pre-determined within the period. This is particularly interesting in the context of the French dual labor market, which features both short-term fixed employment contracts and long term indefinite duration contracts. Even though the French dual labor market is known for its rigidity, it is certainly possible that French firms adjust the current labor stock to contemporaneous supply and demand shocks, even if not completely (Saint-Paul, 1996; Reshef et al., 2022). To allow for this possibility, we need only adjust the factor shares second step moment condition (23) to replace all contemporaneous labor measures with lagged measures.

We report detailed results for the four models estimated above (multi-market, no demand correction, single market, and single market with no exporters) in Tables E.6–E.9. The results for the multi-market estimator are quite similar to our main specification, in which we treat labor as quasi-fixed. We estimate slightly higher returns to scale and lower elasticities of demand for the multi-market estimator, and a slightly larger range of values for LBE (-0.017 to 0.042). Estimates based on the single-market correction estimator still vary wildly by industry, and the estimates of LBE still appear biased up in the two misspecified estimators.

Table E.6: Estimates using Multi-Market Estimator, Partially Adjusting Labor Input

Industry	σ^M	σ^L	σ^K	RTS	ρ^0	$\eta^0 = \frac{1}{(\rho^0-1)}$	μ	h	$\mu/(1-h)$
Autos and transport equipment	0.401	0.613	0.117	1.132	0.936	-15.649	0.004	0.787	0.017
	(0.008)	(0.021)	(0.011)	(0.021)	(0.015)	(4.233)	(0.003)	(0.012)	(0.016)
Chemical products	0.393	0.720	0.096	1.209	0.930	-14.335	-0.000	0.829	-0.001
	(0.008)	(0.030)	(0.013)	(0.022)	(0.011)	(2.581)	(0.004)	(0.010)	(0.021)
Computer, electronics	0.336	0.698	0.078	1.113	0.903	-10.348	0.008	0.815	0.046
	(0.010)	(0.034)	(0.009)	(0.036)	(0.024)	(2.562)	(0.004)	(0.008)	(0.019)
Electrical equipment	0.390	0.641	0.079	1.110	0.903	-10.273	0.011	0.803	0.054
	(0.010)	(0.027)	(0.011)	(0.032)	(0.022)	(2.838)	(0.004)	(0.008)	(0.018)
Food, beverage, tobacco	0.487	0.817	0.192	1.497	0.720	-3.566	0.003	0.715	0.012
	(0.010)	(0.017)	(0.005)	(0.031)	(0.015)	(0.186)	(0.003)	(0.004)	(0.009)
Machinery and equipment	0.352	0.842	0.091	1.285	0.793	-4.839	0.017	0.766	0.071
	(0.006)	(0.019)	(0.003)	(0.024)	(0.014)	(0.320)	(0.002)	(0.005)	(0.007)
Basic metal and fabricated metal	0.286	0.955	0.130	1.371	0.751	-4.009	0.003	0.805	0.015
	(0.006)	(0.025)	(0.004)	(0.030)	(0.015)	(0.247)	(0.001)	(0.003)	(0.007)
Other manufacturing	0.292	0.942	0.134	1.367	0.777	-4.494	0.008	0.774	0.035
	(0.008)	(0.038)	(0.005)	(0.046)	(0.022)	(0.362)	(0.003)	(0.004)	(0.011)
Rubber and plastic	0.392	0.744	0.115	1.251	0.865	-7.434	-0.017	0.816	-0.091
	(0.019)	(0.076)	(0.009)	(0.100)	(0.038)	(1.168)	(0.006)	(0.007)	(0.032)
Textiles, wearing apparel	0.375	0.762	0.176	1.312	0.732	-3.727	0.042	0.850	0.281
	(0.017)	(0.059)	(0.012)	(0.073)	(0.034)	(0.513)	(0.003)	(0.005)	(0.019)
Wood, paper products	0.326	0.976	0.103	1.404	0.762	-4.201	0.011	0.780	0.048
	(0.005)	(0.020)	(0.004)	(0.024)	(0.012)	(0.217)	(0.002)	(0.005)	(0.007)

Notes. The table reports estimates based on the multi-market estimator, allowing labor to partially adjust to current period shocks: average output elasticities σ^j for materials input ($j = M$), labor ($j = L$) and capital ($j = M$), overall returns to scale (RTS), the demand elasticity on the domestic market $\eta^0 = 1/(\rho^0 - 1)$, the coefficient to learning by exporting μ, the persistence parameter in the controlled Markov h, and the long run effect of exporting $\mu/(1-h)$. Bootstrap standard errors clustered by firm in parentheses.

Table E.7: Estimates using Single-Market Estimator (SMC), Partially Adjusting Labor Input

Industry	σ^M	σ^L	σ^K	RTS	ρ	$\eta = \frac{1}{(\rho-1)}$	μ	h	$\mu/(1-h)$
Autos and transport equipment	0.378	0.632	0.127	1.136	0.963	-26.900	0.012	0.770	0.053
	(0.005)	(0.018)	(0.010)	(0.014)	(0.006)	(4.826)	(0.003)	(0.015)	(0.012)
Chemical products	17.830	33.893	4.219	55.942	0.021	-1.021	-0.174	0.836	-1.060
	(102.891)	(188.297)	(27.161)	(318.169)	(0.179)	(0.276)	(0.712)	(0.013)	(4.304)
Computer, electronics	0.418	0.866	0.149	1.434	0.715	-3.510	0.021	0.811	0.110
	(0.066)	(0.156)	(0.031)	(0.250)	(0.073)	(0.877)	(0.006)	(0.009)	(0.033)
Electrical equipment	0.475	0.828	0.132	1.435	0.713	-3.483	0.023	0.799	0.112
	(0.025)	(0.059)	(0.017)	(0.077)	(0.034)	(0.398)	(0.005)	(0.011)	(0.025)
Food, beverage, tobacco	0.499	0.967	0.090	1.556	0.695	-3.275	0.058	0.775	0.257
	(0.021)	(0.035)	(0.009)	(0.061)	(0.018)	(0.093)	(0.004)	(0.003)	(0.016)
Machinery and equipment	0.560	1.231	0.174	1.965	0.505	-2.021	0.056	0.777	0.251
	(8.411)	(21.179)	(2.992)	(32.582)	(0.106)	(0.214)	(1.006)	(0.016)	(6.144)
Basic metal and fabricated metal	0.344	0.878	0.266	1.488	0.632	-2.719	0.023	0.809	0.122
	(0.009)	(0.032)	(0.011)	(0.042)	(0.017)	(0.144)	(0.002)	(0.005)	(0.010)
Other manufacturing	2.216	6.735	1.421	10.372	0.101	-1.113	0.234	0.804	1.198
	(32.744)	(101.633)	(21.390)	(155.761)	(0.058)	(0.074)	(3.031)	(0.005)	(16.075)
Rubber and plastic	-0.840	-1.614	-0.308	-2.762	-0.397	-0.716	0.019	0.828	0.113
	(0.098)	(0.175)	(0.041)	(0.310)	(0.045)	(0.023)	(0.003)	(0.009)	(0.017)
Textiles, wearing apparel	0.169	0.289	0.114	0.572	1.594	1.685	0.028	0.854	0.191
	(0.004)	(0.020)	(0.010)	(0.015)	(0.031)	(0.094)	(0.002)	(0.009)	(0.025)
Wood, paper products	1.975	5.545	0.776	8.296	0.127	-1.146	0.097	0.810	0.509
	(13.364)	(38.381)	(5.059)	(56.802)	(0.046)	(0.055)	(0.667)	(0.004)	(3.555)

Notes. The table reports estimates based on the single-market estimator *à la* Klette & Griliches (1996), using the factor share approach, allowing labor to partially adjust to current period shocks: average output elasticities σ^j for materials input ($j = M$), labor ($j = L$) and capital ($j = M$), overall returns to scale (RTS), the demand elasticity $\eta = 1/(\rho - 1)$, the coefficient to learning by exporting μ, the persistence parameter in the controlled Markov h, and the long run effect of exporting $\mu/(1-h)$. Bootstrap standard errors clustered by firm in parentheses.

Table E.8: Estimates using Single-Market Estimator (SMC) on Sample of Non-Exporters, Partially Adjusting Labor Input

Industry	σ^M	σ^L	σ^K	RTS	ρ	$\eta = \frac{1}{(\rho-1)}$	μ	h	$\mu/(1-h)$
Autos and transport equipment	0.378	0.645	0.127	1.150	0.962	-26.525	-	0.773	-
	(0.005)	(0.018)	(0.010)	(0.014)	(0.006)	(4.816)		(0.015)	
Chemical products	16.111	30.526	3.828	50.466	0.023	-1.023	-	0.836	-
	(133.255)	(256.665)	(31.451)	(421.089)	(0.235)	(0.090)		(0.013)	
Computer, electronics	0.456	0.974	0.164	1.593	0.655	-2.901	-	0.814	-
	(0.103)	(0.249)	(0.047)	(0.396)	(0.077)	(0.622)		(0.009)	
Electrical equipment	0.483	0.867	0.137	1.488	0.701	-3.344	-	0.801	-
	(0.026)	(0.062)	(0.017)	(0.081)	(0.034)	(0.372)		(0.011)	
Food, beverage, tobacco	0.483	0.956	0.089	1.528	0.718	-3.542	-	0.777	-
	(0.020)	(0.035)	(0.010)	(0.060)	(0.020)	(0.111)		(0.003)	
Machinery and equipment	0.585	1.328	0.188	2.101	0.484	-1.937	-	0.783	-
	(0.928)	(2.424)	(0.341)	(3.692)	(0.109)	(0.203)		(0.014)	
Basic metal and fabricated metal	0.348	0.909	0.272	1.529	0.625	-2.664	-	0.810	-
	(0.009)	(0.032)	(0.011)	(0.043)	(0.016)	(0.133)		(0.004)	
Other manufacturing	2.941	9.297	1.848	14.087	0.076	-1.083	-	0.805	-
	(8.198)	(25.930)	(5.197)	(39.321)	(0.062)	(0.073)		(0.005)	
Rubber and plastic	-0.850	-1.610	-0.314	-2.773	-0.392	-0.718	-	0.829	-
	(0.100)	(0.179)	(0.041)	(0.317)	(0.045)	(0.023)		(0.009)	
Textiles, wearing apparel	-3.370	-3.241	-3.427	-10.038	-0.080	-0.926	-	0.919	-
	(0.010)	(0.029)	(0.030)	(0.018)	(0.096)	(2.212)		(0.014)	
Wood, paper products	2.325	6.622	0.923	9.869	0.108	-1.121	-	0.811	-
	(12.392)	(35.354)	(4.846)	(52.588)	(0.045)	(0.053)		(0.004)	

Notes. The table reports estimates based on the single-market estimator *à la* Klette & Griliches (1996), using the factor share approach, allowing labor to partially adjust to current period shocks, where we restrict the sample to non-exporting firms: average output elasticities σ^j for materials input ($j = M$), labor ($j = L$) and capital ($j = M$), overall returns to scale (RTS), the demand elasticity $\eta = 1/(\rho - 1)$, the coefficient to learning by exporting μ, the persistence parameter in the controlled Markov h, and the long run effect of exporting $\mu/(1-h)$. Bootstrap standard errors clustered by firm in parentheses.

Table E.9: Estimates using no Demand Correction, Partially Adjusting Labor Input

Industry	σ^M	σ^L	σ^K	RTS	ρ	$\eta = \frac{1}{(\rho-1)}$	μ	h	$\mu/(1-h)$
Autos and transport equipment	0.364	0.607	0.117	1.087	-	-	0.045	0.942	0.783
	(0.004)	(0.016)	(0.009)	(0.012)			(0.004)	(0.008)	(0.078)
Chemical products	0.366	0.809	0.030	1.205	-	-	0.013	0.866	0.095
	(0.004)	(0.471)	(0.243)	(0.229)			(0.133)	(0.030)	(1.994)
Computer, electronics	0.299	0.607	0.118	1.024	-	-	0.044	0.902	0.451
	(0.004)	(0.022)	(0.011)	(0.014)			(0.004)	(0.019)	(0.071)
Electrical equipment	0.339	0.581	0.104	1.024	-	-	0.054	0.920	0.668
	(0.003)	(0.038)	(0.020)	(0.020)			(0.008)	(0.031)	(0.154)
Food, beverage, tobacco	0.346	0.635	0.070	1.051	-	-	0.142	0.860	1.017
	(0.001)	(0.008)	(0.004)	(0.006)			(0.003)	(0.004)	(0.039)
Machinery and equipment	0.283	0.624	0.086	0.993	-	-	0.055	0.858	0.391
	(0.008)	(0.008)	(0.003)	(0.007)			(0.004)	(0.008)	(0.037)
Basic metal and fabricated metal	0.217	0.629	0.128	0.975	-	-	0.041	0.895	0.395
	(0.001)	(0.009)	(0.004)	(0.006)			(0.001)	(0.006)	(0.018)
Other manufacturing	0.224	0.680	0.130	1.035	-	-	0.053	0.922	0.682
	(0.001)	(0.012)	(0.005)	(0.009)			(0.002)	(0.003)	(0.037)
Rubber and plastic	0.333	0.684	0.099	1.117	-	-	0.026	0.944	0.464
	(0.002)	(0.020)	(0.010)	(0.011)			(0.002)	(0.005)	(0.040)
Textiles, wearing apparel	0.269	0.438	0.191	0.898	-	-	0.045	0.948	0.868
	(0.004)	(0.029)	(0.011)	(0.021)			(0.003)	(0.004)	(0.051)
Wood, paper products	0.251	0.708	0.093	1.052	-	-	0.046	0.909	0.510
	(0.001)	(0.011)	(0.005)	(0.007)			(0.001)	(0.005)	(0.025)

Notes. The table reports estimates without correcting for demand at all, using the factor share approach, allowing labor to partially adjust to current period shocks: average output elasticities σ^j for materials input ($j = M$), labor ($j = L$) and capital ($j = M$), overall returns to scale (RTS), the demand elasticity $\eta = 1/(\rho - 1)$, the coefficient to learning by exporting μ, the persistence parameter in the controlled Markov h, and the long run effect of exporting $\mu/(1-h)$. Bootstrap standard errors clustered by firm in parentheses.

E.3 Results using the control function method in French manufacturing data

We report here estimates of production functions, demand parameters and controlled Markov processes across different models using the control function method. We consider the model without correction for demand (Table E.10), the single market model with all firms (Table E.11) and the single market model in a sample without exporters (Table E.12). In doing so we cannot apply a quasi-non-parametric approach as we did above when using the factor shares approach. Instead, we must make a slightly stronger assumption on the structure of the production function. Since the data clearly reject a Cobb-Douglas production function, we apply a translog, which is a second order approximation.

We use OLS estimates for setting the initial values for the GMM search, which is a common practice. This procedure is prone to the Ackerberg et al. (2023) critique, whereby the GMM search tends not to move away from the OLS point estimates. However, this does not restrict the results to be similar across different models. Indeed, the results are distinct across models.

Table E.10: Estimates using No Demand Correction, Control Function Estimator

Industry	σ^M	σ^L	σ^K	RTS	ρ	$\eta = \frac{1}{(\rho-1)}$	μ	h	$\mu/(1-h)$
Autos and transport equipment	0.621	0.287	0.090	0.999	-	-	0.023	0.912	0.264
	(0.051)	(0.079)	(0.021)	(0.013)			(0.020)	(0.037)	(0.203)
Chemical products	0.589	0.291	0.120	1.000	-	-	0.023	0.899	0.224
	(0.037)	(0.097)	(0.030)	(0.036)			(0.017)	(0.043)	(0.385)
Computer, electronics	0.279	0.748	-0.019	1.008	-	-	0.125	0.884	1.083
	(0.200)	(0.247)	(0.039)	(0.024)			(0.042)	(0.055)	(0.353)
Electrical equipment	0.455	0.514	0.017	0.986	-	-	0.129	0.871	1.000
	(0.186)	(0.276)	(0.055)	(0.051)			(0.070)	(0.042)	(0.563)
Food, beverage, tobacco	0.612	0.262	0.095	0.969	-	-	0.026	0.947	0.497
	(0.004)	(0.008)	(0.004)	(0.006)			(0.002)	(0.002)	(0.054)
Machinery and equipment	0.169	0.790	0.036	0.995	-	-	0.052	0.857	0.364
	(0.405)	(0.459)	(0.034)	(0.026)			(0.037)	(0.050)	(0.235)
Basic metal and fabricated metal	0.359	0.495	0.121	0.976	-	-	0.014	0.956	0.317
	(0.013)	(0.020)	(0.005)	(0.005)			(0.002)	(0.002)	(0.034)
Other manufacturing	0.337	0.559	0.097	0.993	-	-	0.024	0.931	0.354
	(0.058)	(0.116)	(0.021)	(0.038)			(0.013)	(0.032)	(0.095)
Rubber and plastic	0.572	0.251	0.097	0.921	-	-	0.032	0.924	0.415
	(0.164)	(0.425)	(0.039)	(0.243)			(0.027)	(0.018)	(0.233)
Textiles, wearing apparel	0.539	0.294	0.115	0.948	-	-	0.005	0.967	0.152
	(0.039)	(0.063)	(0.009)	(0.027)			(0.009)	(0.016)	(0.164)
Wood, paper products	-3.074	6.450	-0.250	3.126	-	-	0.305	0.828	1.769
	(0.719)	(1.244)	(0.069)	(0.458)			(0.050)	(0.053)	(0.299)

Notes. The table reports estimates without correcting for demand at all, using the control function estimator: average output elasticities σ^j for materials input ($j = M$), labor ($j = L$) and capital ($j = M$), overall returns to scale (RTS), the demand elasticity $\eta = 1/(\rho - 1)$, the coefficient to learning by exporting μ, the persistence parameter in the controlled Markov h, and the long run effect of exporting $\mu/(1-h)$. Bootstrap standard errors clustered by firm in parentheses.

Table E.11: Estimates using Single-Market Model (SMC), Control Function Estimator

Industry	σ^M	σ^L	σ^K	RTS	ρ	$\eta = \frac{1}{(\rho-1)}$	μ	h	$\mu/(1-h)$
Autos and transport equipment	0.132	0.993	0.045	1.170	0.934	-15.244	0.009	0.830	0.051
	(0.408)	(0.606)	(0.059)	(0.154)	(0.063)	(313.136)	(0.015)	(0.038)	(0.090)
Chemical products	1.402	0.485	0.328	2.214	0.440	-1.785	0.022	0.856	0.154
	(0.982)	(0.750)	(0.217)	(1.852)	(0.161)	(188.785)	(0.048)	(0.043)	(0.275)
Computer, electronics	0.490	0.459	0.084	1.033	0.969	-31.930	0.001	0.758	0.003
	(0.017)	(0.095)	(0.018)	(0.085)	(0.080)	(190.105)	(0.003)	(0.035)	(0.011)
Electrical equipment	0.535	0.395	0.081	1.010	0.988	-81.723	0.001	0.860	0.007
	(0.028)	(0.157)	(0.014)	(0.130)	(0.090)	(335.190)	(0.004)	(0.022)	(0.028)
Food, beverage, tobacco	-0.466	0.774	0.049	0.356	2.857	0.538	-0.017	0.750	-0.069
	(2.536)	(0.530)	(0.349)	(3.144)	(0.346)	(1.176)	(0.017)	(0.038)	(0.069)
Machinery and equipment	1.182	0.918	0.175	2.275	0.435	-1.769	-0.014	0.805	-0.070
	(0.119)	(1.103)	(0.029)	(1.178)	(0.056)	(0.106)	(0.050)	(0.013)	(0.205)
Basic metal and fabricated metal	0.666	0.049	0.221	0.936	0.908	-10.857	-0.004	0.931	-0.053
	(0.172)	(1.411)	(0.065)	(1.307)	(0.378)	(4.692)	(0.006)	(0.007)	(0.104)
Other manufacturing	1.726	2.310	0.454	4.489	0.218	-1.279	-0.002	0.783	-0.011
	(1.463)	(3.805)	(0.316)	(5.476)	(0.923)	(1.876)	(0.038)	(0.009)	(0.180)
Rubber and plastic	-1.051	-1.586	-0.183	-2.819	-0.371	-0.729	-0.002	0.806	-0.011
	(4.924)	(1.077)	(0.976)	(6.799)	(0.161)	(0.099)	(0.226)	(0.031)	(1.772)
Textiles, wearing apparel	0.280	0.334	0.050	0.664	1.583	1.715	-0.006	0.817	-0.035
	(0.115)	(0.134)	(0.040)	(0.031)	(0.258)	(1.983)	(0.013)	(0.014)	(0.071)
Wood, paper products	0.329	-0.919	0.017	-0.573	-3.291	-0.233	0.003	0.819	0.016
	(0.351)	(0.911)	(0.093)	(1.334)	(1.130)	(77.877)	(0.008)	(0.025)	(0.060)

Notes. The table reports estimates based on the single-market model *à la* Klette & Griliches (1996), using the control function estimator: average output elasticities σ^j for materials input ($j = M$), labor ($j = L$) and capital ($j = M$), overall returns to scale (RTS), the demand elasticity $\eta = 1/(\rho - 1)$, the coefficient to learning by exporting μ, the persistence parameter in the controlled Markov h, and the long run effect of exporting $\mu/(1-h)$. Bootstrap standard errors clustered by firm in parentheses.

Table E.12: Estimates using Single-Market Model (SMC), on Sample of Non-Exporters, Control Function Estimator

Industry	σ^M	σ^L	σ^K	RTS	ρ	$\eta = \frac{1}{(\rho-1)}$	μ	h	$\mu/(1-h)$
Autos and transport equipment	0.530	0.359	0.035	0.924	1.097	10.303	-	0.738	-
	(0.346)	(0.519)	(0.056)	(0.134)	(0.053)	(118.600)		(0.030)	
Chemical products	1.131	-0.084	0.352	1.399	0.644	-2.807	-	0.855	-
	(0.604)	(0.793)	(0.100)	(1.379)	(0.167)	(0.891)		(0.042)	
Computer, electronics	0.491	0.458	0.084	1.033	0.970	-32.860	-	0.756	-
	(0.016)	(0.136)	(0.020)	(0.122)	(0.373)	(72.781)		(0.034)	
Electrical equipment	0.556	0.381	0.089	1.026	0.974	-37.956	-	0.844	-
	(0.040)	(0.130)	(0.023)	(0.102)	(0.107)	(586.978)		(0.025)	
Food, beverage, tobacco	21.106	3.332	5.291	29.729	0.030	-1.031	-	0.820	-
	(3.943)	(0.532)	(0.839)	(5.140)	(0.249)	(1.034)		(0.038)	
Machinery and equipment	1.171	1.056	0.179	2.407	0.410	-1.694	-	0.801	-
	(0.275)	(3.081)	(0.074)	(3.271)	(0.069)	(0.129)		(0.015)	
Basic metal and fabricated metal	0.665	0.033	0.217	0.915	0.919	-12.381	-	0.930	-
	(0.113)	(0.631)	(0.058)	(0.598)	(0.286)	(4.688)		(0.007)	
Other manufacturing	2.024	1.874	0.655	4.552	0.208	-1.263	-	0.796	-
	(2.362)	(5.769)	(0.404)	(8.461)	(2.158)	(0.862)		(0.009)	
Rubber and plastic	-1.298	-1.505	-0.227	-3.030	-0.335	-0.749	-	0.802	-
	(141.529)	(18.905)	(26.296)	(186.708)	(0.190)	(0.157)		(0.032)	
Textiles, wearing apparel	0.167	0.416	0.022	0.604	1.891	1.123	-	0.820	-
	(0.153)	(0.163)	(0.049)	(0.049)	(0.268)	(8.025)		(0.013)	
Wood, paper products	0.377	-1.298	0.011	-0.911	-1.790	-0.358	-	0.820	-
	(2.039)	(7.559)	(0.719)	(10.314)	(1.373)	(52.194)		(0.019)	

Notes. The table reports estimates based on the single-market model *à la* Klette & Griliches (1996), using the control function estimator, where we restrict the sample to non-exporting firms: average output elasticities σ^j for materials input ($j = M$), labor ($j = L$) and capital ($j = M$), overall returns to scale (RTS), the demand elasticity $\eta = 1/(\rho - 1)$, the coefficient to learning by exporting μ, the persistence parameter in the controlled Markov h, and the long run effect of exporting $\mu/(1-h)$. Bootstrap standard errors clustered by firm in parentheses.

In Figure E.5, we plot with black triangles average returns to scale (RTS) and output elasticities for materials, labor, and capital estimated by industry via the control function method with no demand correction, and with red circles the same estimates obtained from an OLS regression. Estimated average returns to scale are very near unity for 9 out of 10 industries, and output elasticities vary from 0.4 to 0.7 for materials and from 0.2 to 0.4 for labor for the most part, outside of two outlier industries (Rubbers and Plastics and Communication Electronics). Industry-specific median estimates are very similar (not reported). These results are quite similar to control function estimates from other panel datasets.

In Figure E.5, we also plot with black Xs the average factor share in revenue by industry. For materials and labor, this value is the average expenditure share in revenue across the sample. For capital, we compute the factor share as 1 minus the expenditure share for materials and labor, what we term the "residual capital share". If returns to scale were really constant, and firm-specific price deviations were exogenous, or null, then factor shares should sum to one, so the computation of the capital share as the residual would be valid. But in Figure E.5, we find that estimated capital elasticities are only about one quarter to one third of the residual capital share computed from the data. In contrast, estimated material elasticities are about 50% higher than materials expenditure share in revenue. Under the assumptions of perfect competition and constant returns to scale—both necessary assumptions to rationalize the findings in Figure E.5—it appears that the estimated capital elasticity is too low and the estimated material elasticity is too high.

To gain further insight into this discrepancy, it is useful to examine the results from OLS estimates, plotted with red circles in Figure E.5. Just as with the control function, the OLS estimates indicate that returns to scale are roughly unity, and capital elasticities are only about a quarter to a third of the computed residual factor share. These results are somewhat surprising, because the control function is specifically designed to address omitted variable bias in the OLS. If the OLS suffers from omitted variable bias, and the control function estimator corrects for it, then we would expect a wedge between the OLS and the control function estimates. Rather, they seem to almost perfectly coincide. The result indicates that either (i) both the OLS and the control function estimator are consistent, or (ii) the control function estimator fails to neutralize the transmission bias that leads to bias in OLS.

From Tables E.11–E.12, we find that the models that apply some correction for demand (single market and single-market while excluding exporters) yield some really high or really low returns to scale and output elasticities, or some very high or very low demand curvatures, or both. The single market correction model yields erratic estimates of both returns to scale and demand elasticities, regardless of whether we exclude exporters. The model with no correction for demand yields mostly plausible estimates of returns to scale (except for Wood and paper products), but with quite low estimates of returns to capital and high returns to materials—a telltale sign of transmission bias. We conclude that the control function approach delivers, in practice, a poor estimator of the production function and demand parameters.

Figure E.5: Control Function with No Demand Correction and OLS Estimates in the French Data

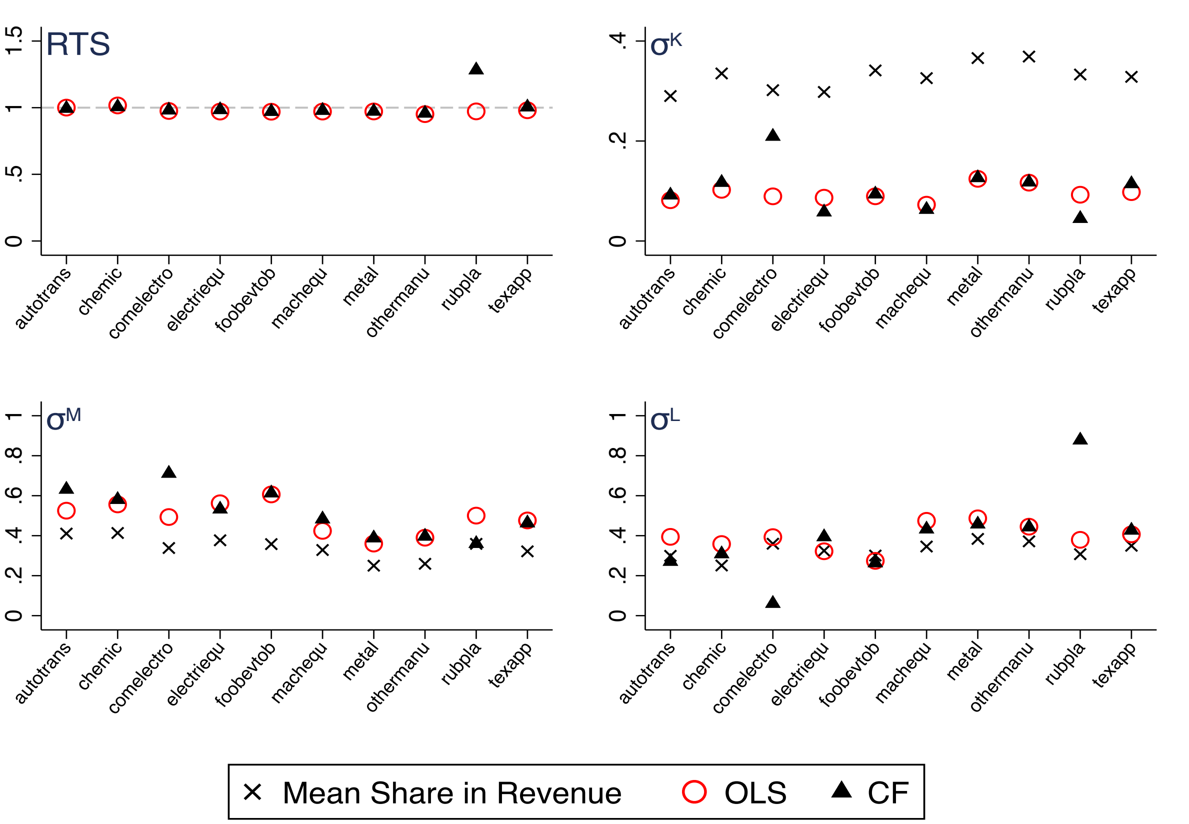

Notes. The figure reports average estimated returns to scale (RTS) and output elasticities to capital (σ^K), materials (σ^M), and labor (σ^L). Black Xs plot the average input expenditure in revenues for materials and labor and the average factor share residual for capital.