Product-Level Efficiency and Core Competence in Multi-Product Plants*

Alvaro Garcia-Marin
Universidad de Chile

Nico Voigtländer
UCLA and NBER

First draft: December 2016
This draft: January 2017

Abstract

A growing literature examines trade-related dynamics at the product-level within firms or plants. Product-level efficiency is a key theoretical component, and so is the ranking of products by "core competence." However, data limitations make it difficult to construct product-level efficiency, and productivity patterns across products within plants are largely unexplored. We exploit a uniquely detailed Chilean dataset that allows us to compute several alternative efficiency measures (such as marginal costs, revenue productivity, physical efficiency, and marginal costs), for each product within plants. We present novel stylized facts in three areas. First, on product-level efficiency patterns, we show that productive plants tend to be relatively efficient across the board, not just for their core products. Second, we show that the typically used sales-based product ranks correctly reflect higher physical efficiency (TFPQ); however – seemingly contradictory – marginal costs are higher for top-ranked sales products. We show that this discrepancy is likely driven by product quality and present a stylized model that underlines the importance of the ranking variable. Finally, using the prominent metric of export skewness towards core products, we highlight the importance of using the appropriate ranking variable when testing predictions of flexible manufacturing models. Product ladders based on marginal costs or revenue productivity do not show export skewness, while TFPQ-based rankings do yield skewness towards the most efficient product and thus aggregate efficiency gains from trade.

JEL: D24, F10, F14, L25, L60
Keywords: International Trade, Multi-Product, Productivity, Flexible Manufacturing

*We gratefully acknowledge research support from the UCLA Anderson Center for Global Management.

1 Introduction

A growing literature examines how production within multi-product firms is affected by international trade, and how the optimal response of these firms to competition influences productivity. For example, Bernard, Redding, and Schott (2011), and Mayer, Melitz, and Ottaviano (2014) show that increased competition causes multi-product firms to skew their production towards their best performing ("core") products and to drop less profitable products from their portfolio. Similarly, Eckel and Neary (2010) study cannibalization of own products and diseconomies of scope when firms expand their product lines, moving away from their "core competence." In the theory underlying these studies, product-specific competence depends on the efficiency with which each product is produced. However, data limitations make it difficult to construct product-level efficiency.[1] In fact, productivity patterns across products within plants are largely unexplored. To bypass this limitation, previous studies have constructed the product ladder within firms using total sales (or exports) of each product, rather than efficiency as implied by the theory.

In this paper, we use a uniquely rich dataset to explore product-level efficiency and core competence in Chilean manufacturing. The Chilean data contain information on product-specific inputs. This allows us to estimate markups at the plant-product level, following the method pioneered by De Loecker and Warzynski (2012), which is flexible with respect to the underlying price setting model and the functional form of the production function. Our dataset also includes physical units as well as revenues for each plant-product, allowing us to calculate product prices (unit values). Dividing these by the corresponding markups yields marginal costs at the plant-product level (De Loecker et al., 2016). We also compute physical productivity (TFPQ) and revenue productivity (TFPR) at the plant-product level. We then use this rich set of product-specific performance measures to examine productivity patterns within multi-product plants. In addition, we use our efficiency measures to create product ranks within plants and check to what extent conclusions in the previous literature depend on using sales as the ranking variable.

We document a series of novel stylized facts that can be broadly summarized in three groups. First, we examine productivity patterns within plants, finding that efficient plants tend to be efficient across the board, not merely for their core product. This supports a common feature of flexible manufacturing models: product-level efficiency is driven by a plant-level efficiency draw, in combination with a product-specific term.[2] These specifications imply that firms with

[1] The exception are De Loecker, Goldberg, Khandelwal, and Pavcnik (2016), who study pass-through at the product level; Garcia-Marin and Voigtländer (2013) and Lamorgese, Linarello, and Warzynski (2014), who examine export-related efficiency gains at the product level; and Dhyne, Petrin, Smeets, and Warzynski (2016) who study the effects of import competition on product-level efficiency.

[2] For example, the models of Eckel and Neary (2010) and Mayer et al. (2014), feature a firm-level draw that

relatively high efficiency in their core product (compared to core products of other producers) should also be relatively efficient in lower-ranked products. This is supported by the strong positive correlation between the relative efficiencies of the core product and lower-ranked products. A related finding is that efficiency trends over time are correlated across products within plants. That is, when plants become more efficient at producing one product, the production of other products also becomes more efficient. This common component of efficiency growth has important implications for models with endogenous growth in multi-product plants.

Our second block of results underlines the importance of the variable that is used to rank products. We show that the standard procedure to create product ladders according to sales leads to a seemingly contradictory pattern: top-ranked sales products exhibit higher physical efficiency (TFPQ) but also higher marginal costs (MC). We show that this difference is likely driven by product quality: Unobserved product quality raises marginal costs via higher input prices, but it leaves our measure of TFPQ largely unaffected. Consequently, top-sales products tend to be produced at relatively high efficiency (high TFPQ), but at high marginal costs due to expensive inputs. High TFPQ exerts a downward pressure on MC and prices, while high quality exerts an upward pressure. In sales-based rankings, the latter prevails, so that top products are sold at relatively high prices. Our findings thus support models that emphasize the importance of the quality dimension, but they also point to the importance of using the appropriate efficiency measure when examining gains from reallocation across products within plants.

To rationalize these empirical findings, we built a stylized model, combining Kugler and Verhoogen's (2012) framework of heterogeneous plants and endogenous quality choice with Eckel and Neary's (2010) model of multi-product plants. An important feature of the model is that – contrary to previous contributions – quality capability is distributed independently of physical efficiency at the product level. This yields the feature that physical efficiency is not perfectly correlated with product revenues (which in turn depend on both efficiency and quality). In fact, product rankings based on revenues are 'biased' towards products with higher quality capability draws, which can explain the observed higher MC for core products. On the other hand, product rankings based on physical efficiency (TFPQ) are unaffected by quality, so that core products have lower marginal costs.

Our final block of results examines the extent to which a central finding in models of flexible manufacturing depends on the efficiency measure that is used to rank products. These models (e.g., Mayer et al., 2014) examine how competition across export destinations affects the product mix within-plants. We take advantage of the fact that we observe direct measures of efficiency for each product, and study whether the canonical model holds when actual effi-

serves as the marginal cost of the core product. Increasing distance from the core product then leads to successively higher product-specific marginal cost.

ciency is used to rank products, instead of the typically used ranking by export sales. When using sales-based product ranks, we confirm the skewness of exports sales towards core products in more competitive destination markets. This is also true for TFPQ based product ranks. Thus, exports to competitive markets are indeed skewed towards the most efficiently produced product (as opposed to simply the most prominent one), affirming that this mechanism can lead to aggregate efficiency gains. However, rankings based on marginal cost or revenue productivity (TFPR) do not imply skewness. This underlines the importance of correctly specifying the productivity measure when examining gains from trade.

Our paper relates to a large literature that studies the relationship between international trade and productivity. Papers such as Pavcnik (2002), Bernard, Eaton, Jensen, and Kortum (2003), and Melitz (2003) have examined selection across firms as a driver of productivity increases. Recent contributions on multi-product firms, in turn, have instead focused on the reallocation of resources across products within firms (Bernard et al., 2011; Mayer et al., 2014). Eckel, Iacovone, Javorcik, and Neary (2015) introduce endogenous choice of product quality in the Eckel and Neary (2010) framework. In this context, firms produce more of their core competence products, but these products also have higher margins, providing incentives to invest in their quality. Using Mexican manufacturing data, Eckel et al. (2015) show that firms in differentiated-goods sectors tend to exhibit quality competence, while this is true to a lesser extent for firms operating in homogenous-goods sectors.[3] Papers in this literature have constructed core competence measures based on total sales. The exception are Dhyne et al. (2016), who estimate firm-product efficiency shocks for multi-product plants in Belgium. Since the Belgian data do not include product-level information on inputs, Dhyne et al. (2016) extend the single-product production function methodology to estimate multi-product production functions, defining production possibilities for each firm based on aggregated inputs and outputs. A downside of this framework is that it applies only to given production tuples (e.g., all four-product firms that produce the exact same set of products). In practice, this imposes a severe data restriction, especially in small countries with relatively few firms.[4]

Relative to the existing literature, we make several contributions. First, we provide direct evidence for assumptions that underly prominent models of flexible manufacturing such as Bernard et al. (2011) or Mayer et al. (2014). For example, our finding that productive plants tend to be relatively efficient at all their products supports the setup where a common efficiency draw affects all products within a plant. To the best of our knowledge, we are the first to provide

[3] Quality differences are identified indirectly, through their effect on prices.

[4] To by-pass this issue, the authors aggregate all other outputs produced by the plant in an ad-hoc fashion, using either revenues, or physical output indexes. Our approach, in contrast, does not need to impose any functional form for aggregating product level information, or to assume allocation rules for assigning inputs to outputs, because we infer product-specific input shares directly from the data.

direct evidence that this assumption holds in the data. Second, we document a number of novel stylized facts on the relationship between products' sales rank and their underlying productivity, marginal costs, and markups. We also show how the patterns in the data are related to product quality. Third, our results emphasize the importance of using the appropriate ranking variable (i.e., efficiency measure) when testing predictions of flexible manufacturing models.

The rest of the paper is organized as follows. Section 2 discusses our empirical framework, shedding light on different efficiency measures such as marginal cost, physical productivity, and revenue productivity. We also illustrate the empirical framework to estimate these measures. Section 3 describes our datasets. Section 4 presents our empirical results and novel stylized facts. Section 5 sketches a stylized model that can help to rationalize our empirical findings. Section 6 concludes.

2 Empirical Framework

In this section, we discuss our efficiency measures and explain how we estimate them at the plant-product level. Our first measure of efficiency is *revenue-based* total factor productivity (TFPR) – the standard efficiency measure in the literature that analyzes productivity in the context of international trade. We discuss why this measure may fail to detect productivity differentials. Our second efficiency measure is quantity productivity (TFPQ), and the third, marginal cost. We discuss which potential biases affect the different measures, which is important since we compare efficiency across products within plants.

2.1 Revenue vs. Physical Total Factor Productivity

Productivity is commonly measured in empirical studies as a residual term between total output and the estimated contribution of production factors. Ideally, total output should be computed in terms of physical units of the final good. However, data on physical quantities are generally scarce and have only recently become available for some countries. As a result, the majority of studies use revenue as output variable for measuring productivity. From hereafter, we denote this productivity measure – based on revenues as output variable – by TFPR, to differentiate from its quantity-based counterpart, which we denote by TFPQ.

As shown in previous research, TFPR is a downward biased measure of TFPQ (Foster, Haltiwanger, and Syverson, 2008). The intuition for this result can be illustrated using the definition $TFPR = P \cdot TFPQ$, where P denotes the output price. If more efficient producers charge lower prices, then TFPR will only show a fraction (or in the extreme, none) of the difference in efficiency reflected by TFPQ.[5] For instance, if preferences are CES and there are constant

[5] As we show below, there is an important exception where TFPR fully reflects differences in TFPQ across producers: under constant returns to scale, if input prices are the same for both producers and the two producers charge differential markups in the same proportion as the difference in TFPQ.

returns to scale, then any efficiency difference in TFPQ translates proportionately into difference in prices (since markups are constant with CES demand). Consequently, TFPR will show no differential in efficiency. In empirical studies, the price bias of TFPR is commonly tackled by deflating revenues with industry price indexes. However, *within* industries the bias does not disappear, and cross-sectional differences in TFPR are affected by the difference between individual plants' prices and the corresponding industry price index.

Next, we show that differences in revenue-productivity may actually capture differences in demand-side factors that lead to differential markups. For simplicity, assume for now a Cobb-Douglas production function, where $\gamma = \alpha_L + \alpha_M + \alpha_K$ denotes the degree of returns to scale, with the subscripts L, M, and K denoting labor, material inputs, and capital, respectively. Total and marginal costs are then given by:

$$TC = \left(\frac{Q}{A}\right)^{\frac{1}{\gamma}} \left[\gamma \cdot w_L^{\frac{\alpha_L}{\gamma}} w_M^{\frac{\alpha_M}{\gamma}} w_K^{\frac{\alpha_K}{\gamma}} \left(\frac{1}{\alpha_K^{\alpha_K} \alpha_M^{\alpha_M} \alpha_L^{\alpha_L}}\right)\right] \tag{1}$$

$$MC = \left(\frac{Q^{1-\gamma}}{A}\right)^{\frac{1}{\gamma}} \left[w_L^{\frac{\alpha_L}{\gamma}} w_M^{\frac{\alpha_M}{\gamma}} w_K^{\frac{\alpha_K}{\gamma}} \left(\frac{1}{\alpha_K^{\alpha_K} \alpha_M^{\alpha_M} \alpha_L^{\alpha_L}}\right)\right] \tag{2}$$

where w_i denotes the price of input i and Q is physical output volume. We use the standard notation A for physical efficiency (TFPQ). Assuming that there are approximately constant returns to scale ($\gamma = 1$), it can be shown that (Katayama, Lu, and Tybout, 2009; Garcia-Marin and Voigtländer, 2013):

$$\Delta TFPR = \Delta\mu - \triangle\phi(\mathbf{w}) \tag{3}$$

where we use $\triangle$ to denote percentage (log-point) differences. Equation (3) implies that differential TFPR does not reflect efficiency differences – unless plants with higher TFPQ charge higher markups or face lower input prices. On the other hand, when input price differences are not meaningful, and under constant returns to scale, it can be shown that (3) implies that efficiency differences $\triangle A_{it}$ are fully reflected by differential marginal costs, i.e., $\triangle A_{it} = -\triangle MC_{it}$ (see Garcia-Marin and Voigtländer, 2013, for a detailed discussion).

The above discussion shows that ideally, we would like to measure TFPQ directly. Marginal cost is a good alternative measure of physical efficiency if production functions exhibit constant returns to scale, and provided that input price differences are minor.[6] However, there are practical caveats. As we explain in section 2.4, estimating TFPQ may be more demanding from a data perspective and is more likely to be affected by measurement error than marginal cost.

[6] In the presence of increasing returns, marginal costs will tend to overestimate actual efficiency gains. In this case, TFPQ is the preferable efficiency measure, since its estimation allows for flexible returns to scale.

Nevertheless, we compute both TFPQ and marginal costs as alternative efficiency measures.

2.2 Productivity Estimation

To compute productivity we specify a Cobb-Douglas production function with labor (l), capital (k), and materials (m) as inputs. Following De Loecker et al. (2016), we estimate a separate production function for each 2-digit manufacturing sector (s), using the subsample of single product plants.[7] The reason for using single-product plants is that one typically does not observe how inputs are allocated to individual outputs within multi-product plants. For the set of single product plants, no assumption on the allocation of inputs to outputs is needed, and we can estimate the following production function with standard plant-level information:

$$q_{it} = \beta_l^s l_{it} + \beta_k^s k_{it} + \beta_m^s m_{it} + \omega_{it} + \varepsilon_{it} \tag{4}$$

where all lowercase variables are in logs; q_{it} are revenues of single-product plant i in year t, ω_{it} is productivity, k_{it} denotes the capital stock, m_{it} are material inputs, and ε_{it} represents measurement error as well as unanticipated shocks to output. Estimating (4) yields the sector-specific vector of coefficients $\boldsymbol{\beta}^s = \{\beta_l^s, \beta_k^s, \beta_m^s\}$.

When computing TFPR, consistently with the literature, we deflate all nominal variables (revenues, materials, wages) using 4-digit industry specific deflators provided by ENIA. In contrast, when computing TFPQ we use quantities – as opposed to revenues – as output variable, and since we do not observe physical inputs in a consistent way, we implement the correction suggested by De Loecker et al. (2016) to control for the plant-specific variation in input prices.[8]

We estimate (4) following the methodology by Ackerberg, Caves, and Frazer (2015, henceforth ACF), who extend the framework of Olley and Pakes (1996, henceforth OP) and Levinsohn and Petrin (2003, henceforth LP). This methodology controls for the simultaneity bias that arises because input demand and unobserved productivity are positively correlated.[9] The key insight of ACF lies in their identification of the labor elasticity, which they show is in most cases unidentified by the two-step procedure of OP and LP.[10] We modify the canonical ACF proce-

[7]The 2-digit product categories are: Food and Beverages, Textiles, Apparel, Wood, Paper, Chemicals, Plastic, Non-Metallic Manufactures, Basic and Fabricated Metals, and Machinery and Equipment.

[8]This source of bias appears to be less problematic when plant revenues are used as output variable (see De Loecker et al., 2016). Under quality considerations, plants charge higher prices for their outputs and pay more for their inputs (Kugler and Verhoogen, 2012), implying that the input price bias tends to be compensated by the output price variation.

[9]We follow LP in using material inputs to control for the correlation between input levels and unobserved productivity.

[10]The main technical difference is the timing of the choice of labor. While in OP and LP, labor is fully adjustable and chosen in t, ACF assume that labor is chosen at $t-b$ ($0 < b < 1$), after capital is known in $t-1$, but before materials are chosen in t. In this setup, the choice of labor is unaffected by unobserved productivity shocks between $t-b$ and t, but a plant's use of materials now depends on capital, productivity, and labor. In contrast to the OP and LP method, this implies that the coefficients of capital, materials, and labor are all estimated in the second stage.

dure by specifying an endogenous productivity process that can be affected by export status and plant investment. In addition, we include interactions between export status and investment in the productivity process. Thus, the procedure allows exporting to affect current productivity either directly, or through a complementarity with investment in physical capital. This reflects the corrections suggested by De Loecker (2013); if productivity gains from exporting also lead to more investment (and thus a higher capital stock), the standard method would overestimate the capital coefficient in the production function, and thus underestimate productivity (i.e., the residual). Finally, using the set of single-product plants may introduce selection bias because plant switching from single- to multi-product may be correlated with productivity. Following De Loecker et al. (2016), we correct for this source of bias by including the predicted probability of remaining single-product, $\hat{s}_{it}$, in the productivity process as a proxy for the productivity switching threshold.[11] Accordingly, the law of motion for productivity is:

$$\omega_{it} \quad = \quad g(\omega_{it-1}, d^x_{it-1}, d^i_{it-1}, d^x_{it-1} \times d^i_{it-1}, \hat{s}_{it-1}) + \xi_{it} \tag{5}$$

where d^x_{it} is an export dummy, and d^i_{it} is a dummy for periods in which a plant invests in physical capital (following De Loecker, 2013).

In the first stage of the ACF routine, a consistent estimate of expected output $\hat{\phi}_t(\cdot)$ is obtained from the regression

$$q_{it} \quad = \quad \phi_t(l_{it}, k_{it}, m_{it}; \boldsymbol{x}_{it}) + \varepsilon_{it}$$

We use inverse material demand $h_t(\cdot)$ to proxy for unobserved productivity, so that expected output is structurally represented by $\phi_t(\cdot) = \beta^s_l l_{it} + \beta^s_k k_{it} + \beta^s_m m_{it} + h_t(m_{it}, l_{it}, k_{it}, \boldsymbol{x}_{it})$.[12] The vector $\boldsymbol{x}_{it}$ contains other variables that affect material demand (time and product dummies, reflecting aggregate shocks and specific demand components). Next, we use the estimate of expected output together with an initial guess for the coefficient vector $\boldsymbol{\beta}^s$ to compute productivity: for any candidate coefficient vector $\tilde{\boldsymbol{\beta}}^s$, productivity is given by $\omega_{it}(\tilde{\boldsymbol{\beta}}^s) = \hat{\phi}_t - \left(\tilde{\beta}^s_l l_{it} + \tilde{\beta}^s_k k_{it} + \tilde{\beta}^s_m m_{it}\right)$. Finally, we recover the productivity innovation ξ_{it} for the given candidate vector $\tilde{\boldsymbol{\beta}}^s$: following (5), we estimate the productivity process $\omega_{it}(\tilde{\boldsymbol{\beta}}^s)$ nonparametrically as a function of its own lag $\omega_{it-1}(\tilde{\boldsymbol{\beta}}^s)$, prior exporting and investment status, and the plant-specific probability of remaining single-product.[13] The residual is ξ_{it}.

[11] We estimate this probability for single-product plants within each 2-digit sector using a probit model, where the explanatory variables include product fixed effects, labor, capital, material, output price, as well as importing and exporting status.

[12] We approximate the function $\hat{\phi}_t(\cdot)$ with a full second-degree polynomial in capital, labor, and materials.

[13] Following Levinsohn and Petrin (2003), we approximate the law of motion for productivity (the function $g(\cdot)$ stated in (5)) with a polynomial.

The second stage of the ACF routine uses moment conditions on ξ_{it} to iterate over candidate vectors $\tilde{\boldsymbol{\beta}}^s$. In this stage, all coefficients of the production function are identified through GMM using the moment conditions

$$\mathbb{E}\left(\xi_{it}(\boldsymbol{\beta}^s)\mathbf{Z}_{it}\right) = 0 \tag{6}$$

where $\mathbf{Z}_{it}$ is a vector of variables that comprises lags of all the variables in the production function, as well as the current capital stock. These variables are valid instruments – including capital, which is chosen before the productivity innovation is observed. Equation (6) thus says that for the optimal $\boldsymbol{\beta}^s$, the productivity innovation is uncorrelated with the instruments $\mathbf{Z}_{it}$.

Given the estimated coefficients for each product category s (the vector $\boldsymbol{\beta}^s$), TFPR can be calculated both at the plant level and for individual products within plants. For the former, we use the plant-level aggregate labor l_{it}, capital k_{it}, and material inputs m_{it}. We then compute plant-level TFPR, $\hat{\omega}_{it}$:

$$\hat{\omega}_{it} = q_{it} - (\beta_l^s l_{it} + \beta_k^s k_{it} + \beta_m^s m_{it}) \tag{7}$$

where q_{it} are total plant revenues, and the term in parentheses represents the estimated contribution of the production factors to total output in plant i. Note that the estimated production function allows for returns to scale ($\beta_l^s + \beta_k^s + \beta_m^s \neq 1$), so that the residual $\hat{\omega}_{it}$ is not affected by increasing or decreasing returns. When computing *plant*-level TFPR in multi-product plants, we use the vector of coefficients $\boldsymbol{\beta}^s$ that corresponds to the product category s of the predominant product produced by plant i.

Next, we compute our main revenue-based productivity measure – *product*-level TFPR. To perform this step for multi-product plants, the individual inputs need to be assigned to each product j. Here, our sample provides a unique feature: ENIA reports total variable costs (i.e., for labor and materials) TVC_{ijt} for each product j produced by plant i. We can thus derive the following proxy for product-specific material inputs, assuming that total material is used (approximately) in proportion to the variable cost shares:

$$M_{ijt} = s_{ijt}^{TVC} \cdot M_{it} \quad \text{where} \quad s_{ijt}^{TVC} = \frac{TVC_{ijt}}{\sum_j TVC_{ijt}} \tag{8}$$

Taking logs, we obtain m_{ijt}. We use the same calculation to proxy for l_{ijt} and k_{ijt}. Given these values, we can derive plant-product level TFPR, using the vector $\boldsymbol{\beta}^s$ that corresponds to product j:

$$\hat{\omega}_{ijt} = q_{ijt} - (\beta_l^s l_{ijt} + \beta_k^s k_{ijt} + \beta_m^s m_{ijt}) \tag{9}$$

where q_{ijt} are product-specific (log) revenues.

For estimating TFPQ, we modify the first stage of the ACF procedure (given by equation

(6) in the paper), by including a vector of variables to proxy for input prices,[14] and we modify the second stage by adding lags of these variables as instruments to identify the additional parameters. Given the quantity-based estimation of the production function (i.e., the vector of quantity-based elasticities $\boldsymbol{\beta}^s$), we can back out physical productivity TFPQ, using the quantity-equivalent of equation (9). On the output side, physical quantities are directly observed at the plant-product level in the Chilean data. As for inputs, we use deflated plant-level expenditures in the spirit of Foster et al. (2008), and assign these to individual products using the reported expenditure shares from ENIA (as calculated in (8)). With this information, we back out TFPQ at the plant-product-year level.

2.3 Estimating Marginal Cost

To construct a measure of marginal production cost, we follow a two-step process. First, we derive the product-level markup for each plant. Second, we divide plant-product output prices (observed in the data) by the calculated markup to obtain marginal cost.

The methodology for deriving markups follows the production approach proposed by Hall (1986), recently revisited by De Loecker and Warzynski (2012). This approach computes markups without relying on market-level demand information. The main assumptions are that at least one input is fully flexible and that plants minimize costs for each product j. The first order condition of a plant-product's cost minimization problem with respect to the flexible input V can be rearranged to obtain the markup of product j produced by plant i at time t:[15]

$$\underbrace{\mu_{ijt}}_{\text{Markup}} \equiv \frac{P_{ijt}}{MC_{ijt}} = \underbrace{\left(\frac{\partial Q_{ijt}(\cdot)}{\partial V_{ijt}} \frac{V_{ijt}}{Q_{ijt}}\right)}_{\text{Output Elasticity}} \Big/ \underbrace{\left(\frac{P^V_{ijt} \cdot V_{ijt}}{P_{ijt} \cdot Q_{ijt}}\right)}_{\text{Expenditure Share}}, \tag{10}$$

where P (P^V) denotes the price of output Q (input V), and MC is marginal cost. According to equation (10), the markup can be computed by dividing the output elasticity of product j (with respect to the flexible input) by the expenditure share of the flexible input (relative to the sales of product j). Note that under perfect competition, the output elasticity equals the expenditure share, so that the markup is one (i.e., price equals marginal costs).

In our computation of (10) we use materials (M) as the flexible input to compute the output elasticity – based on our estimates of (4) for the quantity version of the production function.[16]

[14] Following De Loecker et al. (2016), we include output prices and plant-product sales relative to the overall sales of the same product, as well as the interaction of these variables with capital and materials.

[15] Note that the derivation of equation (10) essentially considers multi-product plants as a collection of single-product producers, each of whom minimizes costs. This setup does not allow for economies of scope in production. To address this concern, we show below that all our results also hold for single-product plants.

[16] In principle, labor could be used as an alternative. However, in the case of Chile, labor being a flexible input would be a strong assumption due to its regulated labor market. A discussion of the evolution of job security and

We then compute markups based on quantity estimates of the elasticities β^s_m (as opposed to revenue-based estimates). Note that since we use a Cobb-Douglas production function, the output elasticity with respect to material inputs is given by the constant term β^s_m. Consequently, it is absorbed by the product fixed effects (which are implicit in our standardization that constructs the Tornqvist index, see Section 3.2). Thus, potential bias due to mis-measured β^s_m (as described in De Loecker et al., 2016) does not affect our results.

The second component needed in (10) – the expenditure share for material inputs – is directly observed in our data in the case of single-product plants. For multi-product plants, we use the proxy described in equation (8) to obtain the value of material inputs $P^V_{ijt} \cdot V_{ijt} = M_{ijt}$. Since total product-specific revenues $P_{ijt} \cdot Q_{ijt}$ are reported in our data, we can then compute the plant-product specific expenditure shares needed in (10).[17] This procedure yields plant-product-year specific markups μ_{ijt}.

Finally, because output prices (unit values) P_{ijt} are also observed at the plant-product-year level, we can derive marginal costs at the same detail, MC_{ijt}. To avoid that extreme values drive our results, we only use observations within the percentiles 2 and 98 of the markup distribution. The remaining markup observations vary between (approximately) 0.4 and 5.6.

2.4 Marginal Cost vs TFPQ

In the following, we briefly discuss the advantages and limitations of marginal cost as compared to quantity productivity (TFPQ) as a measure of efficiency in the context of our study. For now, suppose that the corresponding quantity-based input elasticities $\boldsymbol{\beta}^s$ have been estimated correctly.[18] Then, in order to back out TFPQ by using (7), ideally both output and inputs need to be observed in physical quantities. Output quantities are available in some datasets. But for inputs, this information is typically unavailable. Thus, researchers have adopted the standard practice of using industry-level price indexes to deflate input expenditures (Foster et al., 2008). This approximation may lead to biased TFPQ estimates if input prices or the user cost of capital vary across firms within the same industry. A further complication arises if one aims to compute product-specific TFPQ for multi-product plants, where physical inputs need to be

firing cost in Chile can be found in Montenegro and Pagés (2004).

[17]By using each product's reported variable cost shares to proxy for product-specific material costs, we avoid shortcomings of a prominent earlier approach: since product-specific cost shares were not available in their dataset, Foster et al. (2008) had to assume that plants allocate their inputs proportionately to the share of each product in total *revenues*. This is problematic because differential changes in markups across different products will affect revenue shares even if cost shares are unchanged. De Loecker et al. (2016) address this issue by using an elaborate estimation technique to identify product-specific material costs; this is not necessary in our setting because the uniquely detailed Chilean data allow us to directly compute product-specific material costs from reported data.

[18]To compute TFPQ, the elasticities in the production function (4) must be estimated in quantities. Estimating this vector is challenging in itself: When estimating the production function (4), product-specific output and inputs have to be deflated by proper price indexes. In addition, if input quantities are not available and input expenditure is used instead, the estimation of the production function coefficients is biased (see De Loecker et al., 2016).

assigned to individual products. While our dataset has the unique advantage that plants report the *expenditure* share of each product in total variable costs (which is sufficient to derive the product-specific material expenditure share needed in (10) to compute markups), it does not contain information on how to assign input *quantities* to individual products. Thus, assigning m_{it}, l_{it}, and k_{it} to individual products is prone to error. This is especially true in the case of capital, which is typically not specific to individual output products. In light of these limitations, most studies compute TFPQ at the plant or firm level. An additional complication arises for k_{it} in TFPQ calculations because the capital stock is only available in terms of monetary values and not in physical units.

Contrast this with the computation of markups in (10), still assuming that $\boldsymbol{\beta}^s$ has been correctly estimated. The output elasticity with respect to material inputs is given by β^s_m, and – for single-product plants – the expenditure share for material inputs is readily available in the data. For multi-product plants, we use the approximation with reported variable cost shares in equation (8) to back out plant-product specific input expenditure shares. Thus, plant-product specific markups can be immediately calculated in our Chilean data.[19]

We now turn to the estimation of $\boldsymbol{\beta}^s$, which is challenging and may introduce further error. When using a Cobb-Douglas production function, this issue is less severe for markups than for TFPQ in the context of our analysis. The computation of markups uses only β^s_m from the vector $\boldsymbol{\beta}^s$. Note that measurement error of β^s_m will affect the estimated *level* of markups, but not our analysis across producers of the same product: because we analyze *differences* at the product level, β^s_m is the same across producers and cancels out. In other words, the estimated *differences* in markups in (10) are only driven by the observed material expenditure shares, but not by the estimated output elasticity β^s_m.[20] Contrast this with the computation of TFPQ, which uses all coefficients in $\boldsymbol{\beta}^s$, multiplying each by the corresponding physical input (or deflated input expenditures) in (7). In this case, analyzing differences in TFPQ will not eliminate errors and biases in the level of $\boldsymbol{\beta}^s$.

3 Data

Our primary dataset is a Chilean plant panel for the period 1996-2007, the *Encuesta Nacional Industrial Anual* (Annual National Industrial Survey – ENIA). We combine this dataset with Chilean customs data over the period 2001-2005. A key advantage of the Chilean data is that multi-product plants are required to report product-specific total variable costs. These are

[19]Note that when computing product-level markups for multi-product plants, we only need to proportionately assign the expenditure share of *material* inputs to individual products. This procedure is not needed for labor or capital.

[20]For the same reason, we could in principle use estimates of $\boldsymbol{\beta}^s$ from the *revenue* production function, i.e., the same coefficients used to compute TFPR.

crucial for the calculation of plant-product level markups and marginal costs in multi-product plants, as described in Section 2.3.

Data for ENIA are collected annually by the Chilean *Instituto Nacional de Estadísticas* (National Institute of Statistics – INE). ENIA covers the universe of manufacturing plants with 10 or more workers. It contains detailed information on plant characteristics, such as sales, spending on inputs and raw materials, employment, wages, investment, and export status. ENIA contains information for approximately 5,000 manufacturing plants per year with unique identifiers. Out of these, about 20% are exporters, and roughly 70% of exporters are multi-product plants. Within the latter (i.e., conditional on at least one product being exported), exported goods account for 80% of revenues. Therefore, the majority of production in internationally active multi-product plants is related to exported goods. Finally, approximately two third of the plants in ENIA are small (less than 50 workers), while medium-sized (50-150 workers) and large (more than 150 workers) plants represent 20 and 12 percent, respectively.

In addition to aggregate plant data, ENIA provides rich information for every good produced by each plant, reporting the value of sales, its total variable cost of production, and the number of units produced and sold. Products are defined according to an ENIA-specific classification of products, the *Clasificador Unico de Productos* (CUP). This product category is comparable to the 7-digit ISIC code.[21] The CUP categories identify 2,242 different products in the sample. These products – in combination with each plant producing them – form our main unit of analysis.

Customs data is collected by the Chilean *Servicio Nacional de Aduanas* (National Customs Service) and covers the universe of export transactions over the period 1991-2010. Each export transaction includes an identifier for the exporting firm, the 8-digit Harmonized System category of the product, and the destination country, FOB value, physical volume, and units of each shipment. For the period 2001-2005, we can match this data with ENIA at the plant-product level. For this period, ENIA provides information for the 7-digit Central Product Classification (CPC) code for each product in addition to their CUP. We first use correspondence tables between HS and CPC product categories (provided by the *United Nations Statistical Division*) to consolidate HS-level customs data to the CPC level used by ENIA in 2001-05. Next, we merge the resulting dataset with ENIA at the CPC level. Finally, we collapse the data from the CPC to the CUP level, so as to obtain the same level of disaggregation as the remaining product-level data in ENIA.

Note that the unit of observation in ENIA are plant-products, while in Customs the units are firm-products. However, this does not represent a serious obstacle for matching both datasets.

[21] For example, the wine industry (ISIC 3132) is disaggregated by CUP into 8 different categories, such as "Sparkling wine of fresh grapes," "Cider," "Chicha," and "Mosto."

First, even for multi-plant firms we can match observations at the plant-product level provided that the same exported product is not produced simultaneously in two different plants of the same firm. Second, the vast majority of plants in ENIA (over 97% of the total) are single-plant firms. Thus, the potential for conflict is limited. For the few cases where we cannot establish a unique match between Customs and ENIA, we drop the corresponding observations from the sample.

3.1 Sample Selection and Data Consistency

In order to ensure consistent plant-product categories in our ENIA panel, we follow three steps. First, we exclude plant-product-year observations that have zero values for total employment, demand for raw materials, sales, or product quantities. Second, whenever our analysis involves quantities of production, we have to carefully account for possible changes in the unit of measurement. For example, wine producers change in some instances from "bottles" to "liters." Total revenue is generally unaffected by these changes, but the derived unit values (prices) have to be corrected. This procedure is needed for about 1% of all plant-product observations; it is explained in more detail in Garcia-Marin and Voigtländer (2013). Third, a similar correction is needed because in 2001, ENIA changed the product identifier from CUP to the Central Product Classification (CPC V.1) code. We use a correspondence provided by the Chilean Statistical Institute to match the new product categories to the old ones (see Garcia-Marin and Voigtländer, 2013, for detail). After these adjustments, our sample consists of 118,178 plant-product-year observations.

3.2 Tornqvist Index for Cross-Sectional Comparisons

Productivity measures that are based on units of output – such as physical productivity (TFPQ) or marginal costs – cannot be immediately compared across products, because the output of different products is measured in different units. To tackle this issue, we construct unit-free Tornqvist indexes. This procedure involves two steps. First, for each variable x_{ijt} – defined for product j of plant i in period t – we define its initial normalized value ($\widetilde{x}_{ij0}$) as the log difference of variable x with its average over all plants producing the same product (measured in the same unit of output) in the first period product j is produced by plant i (i.e., $\widetilde{x}_{ij0} = \ln x_{ij0} - (1/I)\sum_{s\in I} \ln x_{sj0}$).[22] Note that our implicit assumption here is that physical units of the same product, measured in the same unit, are comparable. Of course, this ignores possible differences in quality, as we discuss in detail below.

In the second step, once we obtain the initial value for the normalized variable, the levels

[22]Products in ENIA are defined at the 7-digit level. For some products, units of measurement vary. For example, wine may be measured in bottles or boxes. In these cases, we use a separate category for each product-unit. We also trim the data, excluding the top- and bottom 2% within each product category before normalizing. This avoids that the initial levels $\widetilde{x}_{ij0}$ are affected by outliers.

for the remaining periods are computed recursively as:

$$\widetilde{x}_{ijt} = \widetilde{x}_{ij,t-1} + \Delta \ln x_{ijt} \quad (11)$$

where $\Delta \ln x_{ijt} = \ln x_{ijt} - \ln x_{ij,t-1}$. The advantage of this normalization is that, provided that the product is produced by a sufficiently large number of plants, its level in any period can be interpreted as the log-deviation from the average computed over all plants producing the same product.

3.3 Validity of the Sample

Before turning to our empirical results, we check whether our data replicate some well-documented systematic differences between single and multiple-product plants. First, Table 1 reports the prevalence of multi-product plants in our sample. Results suggest that in our sample, multi-product plants are a similarly represented as in the U.S., for which Bernard, Redding, and Schott (2010) provide statistics.[23] Despite the fact that multi-product plants represent less than half of the plants, they account for the majority of output (60 percent). The third row in the table reveals that the average multi-product plant produces 3.7 products. This is also very similar to the number reported for the U.S. by Bernard et al. (2010).

Table 1: Prevalence of plants / firms producing multiple products in Chile / U.S.

	(1) Chile	(2) U.S.
Share of multi-product plants	48.7%	39%
Share of output by multi-product plants	60.0%	87%
Mean products per multi-product plant	3.7%	3.5%

Notes: The table provides statistics for multi-product plants, comparing the US and Chile. Products are defined at the 7-digit level. The third row reports the average number of products produced by a typical multiple-product plant. The numbers for the U.S. (column 2) are for the Census of Manufacturing of 1997, and come from Bernard et al. (2010). The U.S. figures correspond to firms, whereas those for Chile are for plants. However, 97% of all firms are single-plant in Chile, making a comparison viable.

Next, following Bernard et al. (2010), we run the regression

$$\ln(y_{ist}) = \alpha_{st} + \delta \, d_{ist}^{MP} + \varepsilon_{ist} \; , \quad (12)$$

[23] Note that the U.S. statistics from Bernard et al. (2010) are for firms, whereas those for Chile are for plants. However, 97% of all firms in the Chilean data are single-plant, which renders a comparison viable.

where y_{ist} denotes several characteristics of plant i in sector s and period t, d_{ist}^{MP} is a dummy for multi-product plants, and α_{st} denotes sector-year fixed effects.[24] The coefficient δ reports the multi-product premium – the percentage-point difference of the dependent variable between single and multi-product plants. Table 2 reports multi-product plant premia for the Chilean ENIA. Within their respective sectors, multi-product plants are larger both in terms of employment and sales, are more likely to be exporters, but are not more productive (measured by revenue productivity). This is in line with evidence by Bernard et al. (2010) for the United States.

Table 2: Multiple-Product versus Single-Product Firm Characteristics

	(1)	(2)	(3)	(4)
Dependent Variable	log(workers)	log(sales)	Export dummy	ln(TFPR)
Multi-product plant dummy	.325***	.395***	.0411***	.0026
	(.0223)	(.0561)	(.0056)	(.0099)
Sector-Year FE	✓	✓	✓	✓
Observations	53,536	53,536	53,536	53,536
R^2	0.039	0.074	0.018	0.657

Notes: The table reports the percentage-point difference of the dependent variable between multi-product and single-product plants in a panel of approximately 9,600 (4,500 average per year) Chilean plants over the period 1996-2007. All regressions control for sector-year effects at the 2-digit level. Standard errors (in parentheses) are clustered at the sector-year level. Key: *** significant at 1%; ** 5%; * 10%.

4 Empirical Results

In this section we present our empirical results. We begin with results on productivity patterns within multi-product plants. We then turn to product ranks by core competence and establish novel stylized facts on how plant-product performance measure vary along the product ladder. Next, we shed light on the role of product quality by highlighting different patterns for homogenous vs. differentiated products. Finally, we show to what extent the ranking variable that is used to create product ladders within plants affects a prominent mechanism in international trade – efficiency gains due to the skewness of sales towards core products.

4.1 Product- and Plant-Specific Efficiency

A common assumption in models of flexible manufacturing such as Bernard et al. (2011) and Mayer et al. (2014) is that producers draw a firm-specific productivity component that effects all products. On top of this, there is a product-specific efficiency component (or a demand component, which is typically isomorphic to efficiency). As a result of this setup, these models

[24] We control for sector-year effects at the 2-digit level.

feature selection both across firms within industries, and also across products within firms. So far, data limitation have prevented a direct test of the fundamental assumption that the efficiency with which individual products are produced should be positively related to the efficiency of the firm overall. Instead, the literature has largely focused on testing the *predictions* of the selection models, such as the skewness of exports in top- vs. lower-ranked products. We get back to these patterns below in Section 4.5. Here, we directly examine the relationship between product- and plant-level efficiency.

Table 3 presents our results, using log TFPQ of the top-ranked (core) product as dependent variable. The sample includes all plants that produce at least 5 products in any given sample year.[25] Product ranks are computed based on product-specific TFPQ, which is made comparable across products using the Tornqvist index described in Section 3.2. In columns 1-4, we compute the rank in each sample period, potentially allowing products to switch ranks within plants over time. Before describing our results, note that all regressions include product-year fixed effects. Thus, we compare the efficiency of a given plant's top-ranked product to the efficiency in production of the *same* product by all other plants that also produce this product. We refer to this as "relative efficiency."

Column 1 in Table 3 shows that there is a strong positive correlation between the relative efficiency in producing the top product and the relative efficiency of producing the second-ranked product. The coefficients can be interpreted as elasticities, so that a doubling in the relative efficiency of producing the second product is associated with a 67% increase in the efficiency of the top product. In columns 2 and 3, we confirm a strong correlation also for the 3rd and 4th ranked products, respectively. As one should expect, the magnitude of the coefficient declines as we move to lower-ranked products, but it remains both statistically and economically highly significant. In column 4, we compute the average relative efficiency for all non-top products (i.e., below rank 1) produced by a plant. Again, we find a strong positive correlation with efficiency of the top-product. Doubling the average relative efficiency of all other products is associated with a 56% increase in TFPQ of the top-product.

In columns 5-8, we change the way in which we rank products. We now keep the rank from the first year of the sample constant over the entire sample period. This excludes the possibility that products change their ranks within plants. We obtain coefficients that are very similar to those in columns 1-4, suggesting that rank switches are unlikely to affect our results. In sum, the strong correlation of product efficiencies establishes our first stylized fact:

Stylized Fact 1. Plants with high relative efficiency in their core products tend to be relatively efficient also in other products.

[25] If a plant produces fewer than 5 years in some years, but not in others, it is dropped from our sample during the years in which it produces fewer than 5 products.

Table 3: Co-movement of TFPQ across products (by TFPQ rank)

Dep. Var.: TFPQ of the best-performing product	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log(TFPQ) Top 2nd	.671***	—	—	—	.517***	—	—	—
	(.0207)				(.0352)			
log(TFPQ) Top 3rd	—	.529***	—	—	—	.468***	—	—
		(.0231)				(.0342)		
log(TFPQ) Top 4th	—	—	.374***	—	—	—	.379***	—
			(.0223)				(.0368)	
Avg. Log(TFPQ) Rest	—	—	—	.559***	—	—	—	.510***
				(.0224)				(.0329)
Plant FE	No	No	No	No	No	No	No	No
Product-year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Reference period for rank‡	Current	Current	Current	Current	First	First	First	First
Observations	2,305	2,246	2,141	2,305	1,545	1,429	1,227	1,848
R-squared	.645	.495	.337	.474	.395	.377	.294	.310

Notes: The table regresses plant-product physical productivity (TFPQ) of the best performing product of the plant on TFPQ of the second, third and fourth top products (columns 1-3, and 5-7 respectively), and against the average log TFPQ of all products below rank 1 produced by the plant (columns 4 and 8). The sample includes all plants that produce at least 5 products. Within-plant product rankings are computed in terms of normalized product TFPQ (based on the Torqvist index described in section 3.2). Standard errors (clustered at the plant level) are in parenthesis. Key: *** significant at 1%; ** 5%; * 10%.
‡ This describes the sample year in which we rank products. In columns 1-4, we compute the rank in each sample period, potentially allowing products to switch ranks within plants over time. In columns 5-8, we keep the rank from the first year constant over the entire sample period.

Next, in Table 4, we introduce plant fixed effects. This analysis exploits only within-plant variation over time, thus exploring the co-movement of product-specific efficiency within plants. In other words, we examine whether there is a tendency for the top-ranked product's efficiency to rise when the production of other products in the same plant becomes more efficient (or vice-versa).[26] We again obtain statistically highly significant coefficients. As one would expect, the magnitude of the coefficients is smaller than in Table 3, because plant fixed effects absorb the plant-level efficiency component that affects all products. Nevertheless, the coefficients are economically meaningful, suggesting that a doubling in the average efficiency of all non-top products is associated with a 25% increase in efficiency of the top product. The strong positive coefficients in Table 4 suggest that there is a significant tendency of efficiency co-movement between products produced by the same plant. This is our second stylized fact:

Stylized Fact 2. Efficiency tends to co-move across products within plants.

[26] In Table 4, we only present results for stable product ranks. In fact, one reason to rank products in the first sample year and keep this ranking constant over time is that otherwise – with product rank switches – it would be impossible to examine efficiency over time of products with a given rank.

This stylized fact goes beyond the standard framework of models of flexible multi-product manufacturing, where firms receive a given efficiency draw that does not change over time, and where re-allocation of resources across plants and products drives efficiency gains. Stylized fact 2, in contrast, focuses only on efficiency trends *within* plants.[27] Our results thus imply that potential extensions that introduce innovation into models of flexible manufacturing need to allow for co-movement of efficiency gains across products. This can be achieved either by focusing on innovation in the plant-level efficiency component, or by introducing spillovers from innovation in one product to other products produced by the same plant.

Table 4: Within-Plants Comovement of TFPQ across products

	(1)	(2)	(3)	(4)
log(TFPQ) Top 2nd	.264***	—	—	—
	(.0569)			
log(TFPQ) Top 3rd	—	.266***	—	—
		(.0580)		
log(TFPQ) Top 4th	—	—	.327***	—
			(.0601)	
Avg. Log(TFPQ) Rest	—	—	—	.248***
				(.0509)
Plant FE	Yes	Yes	Yes	Yes
Product-year FE	Yes	Yes	Yes	Yes
Reference period for rank‡	First	First	First	First
Observations	1,545	1,429	1,227	1,848
R-sq	.929	.928	.929	.924

Notes: The table regresses plant-product physical productivity (TFPQ) of the best performing product of the plant on TFPQ of the second, third and fourth top products (columns 1–3, respectively), and against the average log TFPQ of all products below rank 1 produced by the plant (column 4). The sample includes all plants that produce at least 5 products. Within-plant product rankings are computed in terms of normalized product TFPQ (based on the Torqvist index described in section 3.2). Standard errors (clustered at the plant level) are in parenthesis. Key: *** significant at 1%; ** 5%; * 10%.
‡ In determining product ranks, we keep the product's efficiency rank from the first year constant over the entire sample period.

4.2 Core Competence and Plant-Product Performance

In the following, we examine various product-specific performance measures within multi-product plants. In the first part of our analysis, we follow the standard procedure of ranking products within plants by their sales revenues. Later, we present results for alternative product

[27] Contributions such as Bustos (2011) and Garcia-Marin and Voigtländer (2013) suggest that these within-plant efficiency gains can be substantial after export entry, and Amiti and Konings (2007) and De Loecker et al. (2016) show similar within-plant efficiency gains when trade provides access to new or cheaper imported inputs.

ranks based on product-level efficiency (TFPQ). We regress different outcome measures y_{ijt} for products j produced by plant i in year t on product rank dummies R^r_{ijt}, with $r = \{1, ..., 4\}$:

$$y_{ijt} = \sum_{r=1}^{4} \beta_r R^r_{ijt} + \delta_{it} + \gamma_j + \varepsilon_{ijt} \tag{13}$$

For consistency across the different specifications, we standardize all dependent variables using the Tornqvist index, i.e., we also standardize those variables that can be compared in their raw form (such as sales revenues).[28] Due to the standardization, all regressions implicitly account for product fixed effects. The regressions also include plant-year fixed effects δ_{it}, so that we only exploit variation across products within plants. Finally, ε_{ijt} denotes the error term. The excluded category in regression (13) comprises all products with rank 5 or higher. Consequently, coefficients β_r are to be interpreted as percentage increase in outcome y when going from products with rank below 5 to product rank r.

Table 5 presents our results. Column 1 merely serves illustrative purposes, showing by how much sales increase when going to higher-ranked products. Top-ranked products account for more than three times higher revenues than products ranked 5th or below. Columns 2 and 3 split the difference in revenues into differences in quantities and prices, respectively. Quantity sold accounts for the largest part of the sales differences along the product ladder. Sales prices also increase with product rank, but this is less pronounced: top-ranked products are sold at about 21% higher prices than products ranked 5th or below (column 3). The fact that core products are sold at higher prices is in line with quality-based models of flexible manufacturing such as Eckel et al. (2015). The findings in columns 1-3 thus replicate previous findings. Next, we move towards results that are new to the literature, because product-specific efficiency measures have not been available. Column 4 shows that there are no significant differences in TFPR across products within plants. This is our third stylized fact:

Stylized Fact 3. Within plants, revenue productivity (TFPR) is fairly uniform across product ranks.

If we take simple models of misallocation such as Hsieh and Klenow (2009) to the product level *within* plants, uniform TFPR means that managers efficiently allocate resources to the individual products. However, this is not astonishing, given that most of the frictions that are typically discussed in the literature (such as access to finance) apply at the plant level, and thus equally to different products *within* plants.

[28] We also use the standardized sales revenues when ranking products within plants. Results are almost identical when we rank products by their raw sales instead. Also, for outcome variables that can be compared across products in their raw form (sales revenues, TFPR, and markups) results are very similar when we use the non-standardized variables instead.

Table 5: Core Competence by Sales Rank and Plant-Product Performance

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dep. Var.:	log(Sales)	log(Volume)	log(Price)	log(TFPR)	log(TFPQ)	log(MC)	log(Markup)
Top product	3.245***	3.106***	.211***	.0124	1.945***	.183***	.00401
	(.0303)	(.0408)	(.0267)	(.00907)	(.0393)	(.0268)	(.00623)
Top 2nd	2.455***	2.424***	.139***	.0154*	1.544***	.106***	.00324
	(.0259)	(.0370)	(.0259)	(.00862)	(.0363)	(.0265)	(.00581)
Top 3rd	1.804***	1.788***	.0863***	.0189**	1.144***	.0504*	.00887
	(.0242)	(.0371)	(.0266)	(.00807)	(.0375)	(.0269)	(.00605)
Top 4th	1.180***	1.188***	.0659**	.0111	.712***	.0319	.00814
	(.0229)	(.0384)	(.0269)	(.00867)	(.0390)	(.0272)	(.00632)
Plant-year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry-year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	14,304	14,304	14,523	14,304	14,304	14,304	14,304
R-sq	.834	.726	.475	.629	.603	.499	.771

Notes: The table regresses each column variable against categorical variables for the top, second, third and fourth best performing product of the plant. Within-plant product rankings are computed in terms of normalized product TFPQ (based on the Torqvist index described in section 3.2). We update the rank in each sample period, potentially allowing products to switch ranks within plants over time. The sample includes all plants that produce at least 5 products. Standard errors (clustered at the plant-year level) are in parenthesis. Key: *** significant at 1%; ** 5%; * 10%.

In column 5 of Table 5 we compare physical efficiency (TFPQ) across product ranks. There are two important features that distinguish TFPQ from TFPR (see Section 2). First, TFPQ is computed based on physical quantities and thus not affected by differences in output prices. Second, TFPQ is estimated using physical input quantities (based on detailed plant-specific input price indexes). Thus, TFPQ is not affected by differences in input costs (e.g., due to different quality of inputs). Given that output prices are higher for top products (column 3), we would expect a tendency for TFPR to be higher. On the other hand, if top products are produced at higher quality, then TFPR would tend to be lower than TFPQ. We find a strong trend of TFPQ to *increase* in product rank. This is our fourth stylized fact:

Stylized Fact 4. Within plants, physical efficiency (TFPQ) is significantly larger for core products.

The fact that TFPQ differences (column 5) are substantially larger than those in TFPR (column 4) suggests that top products are produced at higher input costs. Consequently, core products are likely produced at higher quality. Marginal costs (MC) offer a way to check whether this interpretation is true. The main difference between marginal costs and TFPQ is that the former are affected by input prices, while the latter is not.[29] Thus, if input prices are significantly higher for core products, then two counter-acting forces are at play: on the one hand, higher

[29] An additional difference is that marginal costs are affected by increasing (or decreasing) returns to scale, while

efficiency (TFPQ) for core products would imply lower MC. On the other hand, more expensive inputs would raise MC. Thus, a quality driven increase in input prices would imply that marginal costs fall less strongly for core products, as compared to the corresponding increase in TFPQ of top products. In fact, if the quality-driven increase in input prices is strong, marginal costs may even *increase* for core products, turning the the TFPQ pattern around. This is indeed the case in our data. Column 6 in Table 5 shows that marginal costs are significantly higher for top-ranked products. This is our fifth stylized fact:

Stylized Fact 5. In a sales-based product ranking, marginal costs are higher for core products within plants.

In the context of this stylized fact, the sales-based product ranking plays an important role, as we discuss in detail below. Stylized facts 4 and 5 constitute an important – previously undocumented – pattern: While core products (in terms of sales) exhibit significantly higher physical efficiency in production, they are produced at *higher* marginal costs. This apparent contradiction can be reconciled by quality: as has been previously argued by Eckel et al. (2015) and Antoniades (2015), product quality is a potentially powerful dimension that can help to explain empirical patterns in multi-product firms. If the higher input cost due to product quality raises marginal cost, then this explains the reversed patterns of TFPQ vs. MC. However, it is crucial that product quality affects marginal costs, i.e., that the quality-related cost is not fixed. Since previous models of product quality in flexible manufacturing have typically featured a fixed investment in product quality, they cannot explain the pattern in our data. We discuss this in more detail in Section 4.3

Finally, we turn to the behavior of markups along the product ladder. Column 7 in Table 5 shows that there is no significant difference in markups across products. This is our sixth stylized fact:

Stylized Fact 6. In a sales-based product ranking, markups do not vary systematically across product ranks within plants.

The absence of markup differences along the product ladder can be rationalized in the context of the above discussion. Core products (when ranked by sales) tend to have higher marginal costs. Higher MC are associated with lower markups in demand systems that allow for flexible markups, such as Melitz and Ottaviano (2008). On the other hand, quality upgrading is particularly pronounced for products with more scope for product differentiation, i.e., in industries with long quality ladders (Khandelwal, 2010). Since we find evidence that core products are produced at relatively high quality, this should lead to higher markups. In sum, the non-results

TFPQ is not. However, we show in Garcia-Marin and Voigtländer (2013) that this is unlikely to affect TFPQ vs MC in the Chilean data, because returns to scale are close to one.

for markups may emerge because of the tendencies towards lower markups (due to higher MC) and higher markups (due to higher quality) of core products cancel each other.

Note that in this context, the product ranking by sales is important. Typically, models of flexible manufacturing feature (unobserved) product-specific efficiency, which the theory uses to rank products. In the data, however, products are ranked by sales revenues. This is valid as long as product efficiency ranks map one-to-one into product sales ranks. Our results suggest that this is not necessarily the case. Below, in Section 4.4 we rank products by TFPQ and discuss the arising differences in results.

4.3 Product Differentiation and the Role of Quality in Production

Models that introduce quality in flexible manufacturing derive their basic insight from a mechanism whereby producers have higher incentives to invest in quality of more efficient (core) products. There is an original draw of plant-level marginal costs. Then, for the core product, marginal cost corresponds to the plant-level MC draw, and it increases successively with products' distance from the core. In other words, product-specific MC is lowest for the core product. Thus, the core product offers higher profit margins, which in turn makes it easier to recover fixed costs that are incurred when raising product quality. The incentives to invest in product quality are particularly strong when there is a high scope for product differentiation (Eckel et al., 2015). We explore this dimension in the following, after discussing an important discrepancy between existing theories and our findings.

The investment in product quality is typically modeled as a fixed cost. This implies that on top of the original product-specific marginal cost, there is an *average* cost of investment in quality, which declines with volume produced. Together, the two cost component yield the "full marginal cost" (Eckel et al., 2015). This raises the question which exact cost components are captured by our different efficiency measures. First, recall from our discussion above that TFPQ is not affected by higher input prices due to higher product quality. Thus, TFPQ is close to the "original marginal cost" of products, before quality adjustments.[30] Second, note that marginal costs, by construction, do not reflect fixed cost of improving product quality. Thus, if the standard setup with fixed cost of quality was correct, we should not observe an increase in marginal costs as product quality grows. However, the fact that we do see significantly higher MC for core products (while TFPQ is also higher) suggests that higher quality drives up *marginal* costs – and not only average total costs, as assumed by existing theories. Consequently, to match the patterns in the data, future models of flexible manufacturing with a quality dimension should feature increasing marginal costs as quality rises.

[30] This holds as long as higher product quality does not slow down the production process. For example, in the case of high-quality rug production, more time is dedicated for to each rug (Atkin, Khandelwal, and Osman, 2014). In this context, TFPQ would be lower for high-quality products. If, in turn, higher quality is mostly associated with more expensive inputs that are otherwise processed similarly, TFPQ will be unaffected by quality.

We have argued that the different patterns of core product for the performance measures TFPQ vs. MC reflect higher product quality of core products. In the following, we explore this point further, using the insight that investment in quality is particularly profitable in more differentiated industries. In Table 6 we repeat the regressions from Table 5, but now estimating separate coefficients for β_r in (13) for plants operating in industries with homogenous products (Panel A) vs. differentiated products (Panel B).[31] For sales revenues and volume, the results in Table 6 are very similar for homogenous and differentiated products. In contrast, the pattern for prices differs substantially: in the homogenous category, there is only a small difference in output prices for the top product as compared to lower ranked products. In the differentiated category, on the other hand, prices are significantly higher for core products. This is in line with quality playing a more important role in differentiated products, confirming the results in Eckel et al. (2015).

Next, we turn to our productivity measures. TFPR shows essentially no differences across product ranks in either of the two subsets (column 4). The pattern for TFPQ is very similar in both subsets: there is a strong increase in TFPQ as we move up the product ranks towards core products. Since TFPQ is unlikely to be confounded by quality differences, finding very similar patterns for homogenous vs. differentiated products makes sense. In contrast, results differ substantially for MC in the two subsets: for homogenous products, there is no apparent difference by product rank, while for differentiated products, MC increase strongly for core products. This result is in line with rising MC reflecting increasing product quality. The following stylized fact summarizes the results for product differentiation:

Stylized Fact 7. The pattern of systematically higher TFPQ for core products (Stylized fact 4) holds equally for homogenous and differentiated products. In contrast, the pattern of increasing MC for core products (Stylized fact 5) holds only in the subset of differentiated products.

This stylized fact underlines the important differences in efficiency measures. It also supports our interpretation that TFPQ is unlikely to be affected by quality differences, while rising MC – with simultaneously increasing TFPQ – reflect the costs of higher product quality.

[31] We define the degree of differentiation at the plant level based on the liberal classification in Rauch (1999). For this, we use concordances between SITC (used by Rauch) and ISIC codes of the main product (used by the Chilean ENIA). This yields a plant-level classification into homogenous and differentiated. "Homogeneous" is for product categories that according to Rauch (1999) are "traded on organized exchanges" or are "referenced priced"; "differentiated" is based on Rauch's "differentiated" category.

Table 6: Core Competence by Sales Rank: Sample Splits by Product Differentiation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dep. Var.:	log(sales)	log(volume)	log(Price)	log(TFPR)	log(TFPQ)	log(MC)	log(Markup)
			Panel A: Homogeneous Products				
Top product	3.396***	3.166***	.0904*	.00927	1.993***	.0648	-.0010
	(.0443)	(.0653)	(.0523)	(.0138)	(.0701)	(.0503)	(.0104)
Top 2nd	2.549***	2.471***	.0628	.00569	1.647***	.0224	-.0078
	(.0365)	(.0595)	(.0502)	(.0135)	(.0657)	(.0493)	(.0100)
Top 3rd	1.826***	1.753***	.00278	.00303	1.252***	-.0342	-.0031
	(.0348)	(.0607)	(.0517)	(.0119)	(.0671)	(.0501)	(.0097)
Top 4th	1.154***	1.075***	.0292	.0100	.734***	-.0208	.0116
	(.0330)	(.0663)	(.0527)	(.0124)	(.0716)	(.0507)	(.0108)
			Panel B: Differentiated Products				
Top product	3.121***	3.052***	.307***	.0140	1.912***	.275***	.0152**
	(.0409)	(.0511)	(.0253)	(.0120)	(.0429)	(.0264)	(.0076)
Top 2nd	2.382***	2.383***	.187***	.0225**	1.466***	.167***	.0115*
	(.0359)	(.0467)	(.0275)	(.0111)	(.0393)	(.0273)	(.0068)
Top 3rd	1.792***	1.815***	.146***	.0313***	1.062***	.112***	.0179**
	(.0329)	(.0457)	(.0271)	(.0109)	(.0414)	(.0276)	(.0077)
Top 4th	1.207***	1.279***	.0863***	.0113	.702***	.0661**	.00447
	(.0313)	(.0443)	(.0282)	(.0119)	(.0420)	(.0284)	(.0075)
Plant-year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Product FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	14,304	14,304	14,304	14,304	14,304	14,304	14,304
R-squared	.835	.726	.472	.629	.603	.501	.771

Notes: The table regresses each column variable against categorical variables for the top, second, third and fourth best performing product of the plant, interacted with a dummy for homogenous products. Within-plant product rankings are computed in terms of normalized product sales (based on the Torqvist index described in section 3.2). We update the rank in each sample period, potentially allowing products to switch ranks within plants over time. The sample includes all plants that produce at least 5 products. We define degree of differentiation at the plant level based on the liberal classification in Rauch (1999). For this, we use concordances between SITC (used by Rauch) and ISIC codes of the main product (used by the Chilean ENIA). This yields a plant-level classification into homogenous and differentiated. "Homogeneous" is for product categories that according to Rauch (1999) are "traded on organized exchanges" or are "referenced priced"; "differentiated" is based on Rauch's "differentiated" category. Standard errors (clustered at the plant-year level) are in parenthesis. Key: *** significant at 1%; ** 5%; * 10%.

4.4 TFPQ-Based Product Ranking

Most models of flexible manufacturing rank products in terms of physical efficiency. However, in empirical tests of these models' prediction, products are typically ranked based on sales. Given our product-specific efficiency measures, we can make progress on this front. First, we need to pick the 'right' efficiency measure for our ranking. Our results above suggest that TFPQ is the most appropriate variable to capture physical efficiency. In the following, we thus use TFPQ to rank products and examine whether this affects the results that we obtained above for the commonly used sales-based rankings.

Table 7 replicates the results from Table 5, estimating regression (13) for TFPQ-based product ranks. Already the first column suggests that sales-based ranks do not map one-to-one into TFPQ-based ranks. While sales revenues increase significantly with higher product rank, the top product is now only sold 217% more than products ranked fifth or lower. In contrast, this number was 324% for the sales-based ranking in Table 5. While the pattern of physical units sold (column 2) is similar to our results above, product prices (column 3) show a striking difference: In our sales-based ranking, prices were higher for top-ranked products. For the TFPQ-based ranking, the opposite is true – prices are *lower* for top-ranked products. TFPR again shows no significant differences across product ranks (column 4), while TFPQ is increasing by construction (column 5). Next, column 6 shows that the pattern for marginal costs is also reversed, as compared to the sales-based ranking: top products are produced at lower marginal costs. Finally, markups are slightly higher for top-ranked products, but the magnitude of the coefficient is small. Our next stylized fact summarizes the most striking difference between sales- and TFPQ-based product rankings.

Stylized Fact 8. When products are ranked by their physical efficiency, the pattern for prices and marginal costs (Stylized fact 5) is reversed: Prices and marginal costs are lower for products with higher TFPQ.

This finding implies that prominent results in the literature depend crucially on the variable that is used to rank products. For example, Eckel et al. (2015) find that core products (ranked by sales) are sold at higher prices, and they interpret this as evidence for higher product quality. The model of Eckel et al. (2015) actually classifies core products based on their physical efficiency (marginal costs, not accounting for quality). This is most closely reflected by TFPQ. We replicate the finding by Eckel et al. (2015) in the sales based ranking. But when we use TFPQ as a ranking variable (i.e., the ranking that is more closely reflecting their model), we obtain the opposite results.

We can rationalize the striking difference in results in the context of our discussion of the alternative efficiency measures. Recall that TFPQ is largely unaffected by product quality, so that the corresponding ranking reflects largely physical efficiency (and thus lower marginal

Table 7: Plant-product outcomes by TFPQ-based product rank

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dep. Var.:	log(Sales)	log(Volume)	log(Price)	log(TFPR)	log(TFPQ)	log(MC)	log(Markup)
Top product	2.172***	3.172***	-.479***	-.00753	3.142***	-.498***	.0183***
	(.0417)	(.0434)	(.0293)	(.00859)	(.0332)	(.0288)	(.00631)
Top 2nd	1.664***	2.427***	-.334***	-.00695	2.414***	-.341***	.0129**
	(.0381)	(.0397)	(.0278)	(.00873)	(.0296)	(.0270)	(.00580)
Top 3rd	1.177***	1.798***	-.211***	-.00173	1.793***	-.214***	.00827
	(.0387)	(.0385)	(.0262)	(.00806)	(.0271)	(.0256)	(.00590)
Top 4th	.689***	1.067***	-.118***	.000486	1.129***	-.127***	.00872
	(.0397)	(.0383)	(.0262)	(.00822)	(.0260)	(.0258)	(.00619)
Plant-year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry-year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	14,304	14,304	14,304	14,304	14,304	14,304	14,304
R-sq	.629	.720	.489	.628	.802	.519	.771

Notes: The table regresses each column variable against categorical variables for the top, second, third and fourth best performing product of the plant. Within-plant product ranking are computed in terms of normalized product TFPQ (see section 3.2 for details). We compute the rank in each sample period, potentially allowing products to switch ranks within plants over time. The sample includes all plants that produce at least 5 products. Standard errors (clustered at the plant-year level) are in parenthesis. Key: *** significant at 1%; ** 5%; * 10%.

costs, which translate into lower prices). In contrast, total sales are also (positively) affected by product quality. Thus, ranking products based on sales lifts high-quality products to the top ranks, implying higher marginal costs and higher prices. Consequently, empirical studies have to carefully choose the ranking variable, according to the mechanism that they are seeking to examine.

4.5 Exporting and Skewness of Sales

One of the main implications of models with flexible manufacturing and variable markups (e.g., Mayer et al., 2014) relates to the impact of competition across export destinations on the within-plant product mix. In these models, in response to a more competitive market, plants lower their markups in all products and concentrate their sales in their best products. In practical terms, the literature tests this prediction by analyzing whether top products account for a larger share of exports in more competitive (tougher) markets – approximated by market size in terms of gross domestic product.[32]

In this section, we take advantage of the fact that we observe direct measures of efficiency for each product, and study whether the canonical model holds when actual efficiency is used to rank products, instead of ranking products by export sales. Following Mayer et al. (2014), we

[32] In the model, the parameter that measures competition is the number of firms/varieties in each destination market. Tougher markets are characterized by a larger number of varieties; in them each firm has a lower residual demand, and thus reduces its price-cost markups.

estimate the relationship between export sales skewness between the top and the second-ranked product for each export market. Skewness is defined as $s_{ict} = \ln(y_{ict}^{r=1}/y_{ict}^{r=2})$, where $y_{ict}^{r=1}$ and $y_{ict}^{r=2}$ are export sales of the top and second-ranked product, respectively, of plant i exporting to destination country c in year t. Note that thus, the top and second-ranked export product are determined for each destination country separately.[33] We check whether the variable that is used to rank products r affects our results. We run the following regression:

$$s_{ict} = \beta \ln GDP_{ct} + \gamma X_c + \delta_{it} + \varepsilon_{ict} \tag{14}$$

We use a set of controls for geographical distance (distance between Santiago – the capital city of Chile – and the capital city from the destination country), location (whether the importing country shares a border with Chile), and similarity to Chile (whether the importing country shares the same official language as Chile, and whether it was a Spanish colony). All regressions include plant-year fixed effects δ_{it}.

Table 8 presents our results. First, in column 1 we use export data for the universe of Chilean exporters between 2001 and 2005. Here, we intend to verify whether the Chilean export data displays the same behavior as the sample used by Mayer et al. (2014) for the case of France. The positive and significant coefficient on log GDP suggests that in response to increased competition, Chilean exporters also tend to concentrate their sales in their most important product(s). In column 2 we repeat the analysis, but this time restricting the sample to the set of plant-products in the Chilean ENIA that we can confidently match to customs data. Despite the considerably smaller sample size, we find the same pattern for this restricted sample. Thus implies that the Mayer et al. (2014) mechanism also holds in our matched sample. We also confirm the main pattern in column 3, where we use total sales (both in the domestic and export markets, as reported in ENIA) to rank products.[34]

Next, we focus on our main efficiency measures. We find no relation between export sales skewness and the degree of competition in the export market, when using TFPR or marginal cost (columns 4 and 6, respectively). In contrast, when using TFPQ to define the ranking of best performing products (column 5), we find a positive and significant coefficient on log GDP.

[33] In principle, it would be possible to use an alternative global export sales ranking to construct a different skewness measure, e.g., define the main products in terms of their sales to *all* countries. However, as Mayer et al. (2014) show with customs level French data for 2003, the correlation between local and global rankings is relatively high – above 60 percent – and very stable across samples. In addition, when we use the global ranking, our sample shrinks considerably. Most exported products are sold to only a few export destinations, which makes it impossible to apply the ranking of the top two products to many countries, simply because the top two products are not sold to all countries.

[34] Note that in this case, the top product for a given destination market is defined as the one with highest domestic sales that is exported to this destination. The same applies to columns 4-6, but with the ranking based on TFPR, TFPQ, and MC, respectively.

Table 8: Export sales skewness across destinations

Dep.Var.: Skewness of destination-specific export sales between top- and second-ranked product

Products ranked	Export sales		Sales	TFPR	TFPQ	MC
according to:	(1)	(2)	(3)	(4)	(5)	(6)
log(RGDP)	.0391***	.111***	.128**	.0275	.174**	.0845
	(.0082)	(.0407)	(.0557)	(.0703)	(.0827)	(.0790)
log(Distance)	.0008	-.202	-.286	-.444**	-.776***	-.327
	(.0212)	(.131)	(.203)	(.218)	(.270)	(.284)
Colony	-.0565	.189	.519	.525	-.518	.550
	(.0634)	(.358)	(.448)	(.552)	(.540)	(.608)
Border with Chile	-.128***	-.176	-.217	-.0538	-.230	.275
	(.0325)	(.172)	(.288)	(.338)	(.356)	(.347)
Common Official Language	-.0507	-.0581	-.162	-.431	-.277	.322
	(.0348)	(.181)	(.251)	(.291)	(.350)	(.268)
Sample:	Customs Only	Customs & ENIA	Customs & ENIA	Customs & ENIA	Customs & ENIA	Customs & ENIA
Plant-year FE	Yes	Yes	Yes	Yes	Yes	Yes
N	45,107	1,952	1,816	1,509	1,075	1,256
R-sq	.545	.470	.480	.557	.577	.538

Notes: The table reports the relation between export sales skewness and market size – measured in terms of log real GDP – across destinations, by different ranking variables. The dependent variable in all regressions is the logarithmic difference between export sales of the top and the second-ranked product for each destination. In columns 1, the export sales rank is computed in each destination market with information from customs data. Columns 2 repeats the exercise, but only for plants that are reported by both customs and ENIA. Columns 3-6 rank the top two products in terms of total sales, TFPR, TFPQ and marginal cost, respectively. In all columns, the top product in each destination is defined conditionally on the product being exported to that destination (i.e., the top product in the destination may not coincide with the top product globally). The same is true for the second-ranked product. All regressions control for plant-year fixed effects. Standard errors (in parentheses) are clustered at the plant-year level. Key: *** significant at 1%; ** 5%; * 10%.

These results lead to our final stylized fact:

Stylized Fact 9. The finding that skewness (higher sales of core products) in more competitive markets depends crucially on which variable is used to rank products. TFPR or MC as ranking variables do not imply skewness. The standard result of skewness is only obtained when ranking products by sales (as is common) or in terms of physical productivity (TFPQ).

Why do product rankings based on TFPR or marginal cost fail to show a higher skewness of exports in more competitive markets? As we discussed before, marginal cost reflects not only efficiency, but also differences in product quality. Thus, a higher ranked product in terms of marginal cost (lower value), may be related to either higher production efficiency, or lower product quality. Does, marginal costs reflect two opposing forces, so that the resulting ranking is not meaningful for efficiency-based mechanisms. Similarly, TFPR does not fully reflect efficiency differences when plants pass part of their efficiency advantage on to customers in the form of lower prices. Thus, TFPR and marginal cost-based rankings are likely to be less informative than the TFPQ-based counterpart for defining the best performing products of the plant.

Stylized fact 9 is important because it confirms the underlying mechanism behind the within-plant productivity enhancing effect of competition as in Bernard et al. (2011) or Mayer et al. (2014): in response to increased competition, plants skew their sales towards their more efficient performing products. As a consequence, plant-level productivity increases. So far, only indirect evidence has been provided for this mechanism, since previous empirical studies have relied on sales-based rankings (which may be affected by many other factors that are unrelated to production efficiency) or on structural simulation. We provide direct evidence, by showing that firms skew their exports more towards products with higher physical efficiency in more competitive countries.

5 A Stylized Multi-Product Plant Model with Quality Choice

In this section we present a stylized model that can reconcile the empirical patterns documented above. In particular, the model can help to explain why the results for prices depend crucially on whether we rank products by sales revenues or by physical productivity. The model combines Kugler and Verhoogen's (2012) framework of heterogeneous plants and endogenous quality choice with Eckel and Neary's (2010) model of multi-product plants. Since our main goal is to understand factors behind price dispersion within plants in a given year, we solve the model in partial equilibrium. We focus on a single monopolistically competitive industry producing a differentiated final good. A caveat of this model is that it features constant markups. However, this is not a major constraint because the data show only minor markup differences across products within plant-years (Stylized fact 6).

5.1 Preferences

The representative consumer derives utility from the consumption of a continuum of differentiated varieties, each of which may have different degrees of product quality:

$$U = \left[\int_{\omega \in \Omega} [q(\omega) x(\omega)]^{\frac{\sigma-1}{\sigma}} \, d\omega \right]^{\frac{\sigma}{\sigma-1}} \tag{15}$$

where ω is an index for varieties from the product space Ω; σ is the elasticity of substitution between varieties; $x(\omega)$ is the quantity of each variety ω consumed; and $q(\omega)$ is product quality, interpreted here in broad terms as product appeal.

Cost minimization leads to the usual CES demand equation, augmented by product quality:

$$x(\omega) = X q(\omega)^{\sigma-1} \left(\frac{p(\omega)}{P} \right)^{-\sigma} \tag{16}$$

where $p(\omega)$ is the price of variety ω, P is the aggregate (quality-adjusted) price index, and X is aggregate (quality-adjusted) consumption.

5.2 Production

Our model features a continuum of plants i, each of them producing several products that we index by ij, i.e., product j produced by plant i. We assume that each plant produces unique varieties from the product space Ω. Consequently, demand for a given variety $x(\omega)$ given by (16) has its equivalent x_{ij} on the production side. As in Melitz (2003), multi-product plants are heterogenous in their overall (plant-level) efficiency, given by λ_i. We assume that λ is drawn from a fixed distribution $G(\lambda)$ with support $[\underline{\lambda}, \infty]$, and that this capability draw is known to the plant only after it enters the market.[35]

In addition, plants vary in the efficiency of the individual varieties they produce. This product-specific efficiency term requires a more elaborate explanation. Production of each variety ij is carried out using an intermediate input, supplied by a perfectly competitive market. Plants can purchase inputs of different quality. Perfect competition in the input producing sector allows us to assume that one unit of input with quality c is sold at cost c. Following Eckel and Neary (2010) and Mayer et al. (2014), we assume that plants can produce any number of varieties, but each additional variety entails a higher marginal cost. There is a fixed production cost f that the plant has to incur for each product that is produced (Bernard et al., 2010). Each plant has a "core competence," corresponding to the product the plant produces most efficiently. To rank products within plants, we adopt the notation $m_{ij} \in \mathbb{N}_0$, representing increasing distance

[35] To get a productivity draw, plants have to pay a fixed cost as in Melitz (2003). Since our model is in partial equilibrium, we take the existence of plants as given and thus abstract from explicitly modeling entry decisions at the plant level. Nevertheless, we do model the decision on how many products each plant produces.

from the core competence. For the core product, $m_{ij} = 0$, and for periphery products, $m_{ij} \geq 1$. This establishes a product ladder, where products m_{ij} are produced with marginal cost

$$MC(m_{ij}, \lambda_i) = \frac{c_{ij}}{\phi^{m_{ij}} \lambda_i} \equiv \frac{c_{ij}}{\widetilde{\lambda}_{ij}} \tag{17}$$

with $\phi \in (0, 1)$, and where c_{ij} is the quality (and cost) of the intermediate input used to produce variety ij. Note that product-specific efficiency $\widetilde{\lambda}_{ij} \equiv \phi^{m_{ij}} \lambda_i$ is increasing in λ_i and decreasing in product rank m_{ij}. In words, more productive plants (higher λ_i) tend to be more efficient at any given ladder step m_{ij}. However, within plants, products further from the core (higher m_{ij}) are produced at higher marginal cost (for any given input quality c_{ij}). Next, we derive the optimal choice of input and output quality, c_{ij} and q_{ij}.

In modeling the relationship between efficiency and product quality, we adapt the first variant of the firm-level model described in Kugler and Verhoogen (2012) to the context of multi-product plants. This involves two assumptions. First, quality upgrading involves no fixed-cost. Second, in the production of product quality, product-specific efficiency $\widetilde{\lambda}_{ij}$ and input quality c_{ij} are complements according to the following CES function:

$$q_{ij}(m_{ij}) \;=\; \left[\frac{1}{2}\left(\widetilde{\lambda}_{ij}^{b}\right)^{\theta} + \frac{1}{2}\left(c_{ij}^{2}\right)^{\theta}\right]^{\frac{1}{\theta}} \tag{18}$$

where $\theta < 0$ measures the degree of complementarity between input quality and plant-product capability $\widetilde{\lambda}_{ij}$ in the production of product quality, and b is a parameter that reflects each plant's ability for translating higher plant-product capability into higher product quality. Consequently, higher b increases a plant's incentives to improve product quality. We assume that b is plant-product specific, i.e., $b \equiv b_{ij}$, and it is randomly drawn – before plants decide their product range – from a fixed distribution $F(b)$ with support $[\underline{b}, \infty]$. Importantly, we assume that b_{ij} is independently drawn for each product ij, so that it is independent of the product rank m_{ij}. Thus, our setup allows for different plants that produce similar products to have different product-specific quality capability. For example, among two furniture plants, one may have an advantage at producing high-quality chairs, and the other, at producing high-quality beds. In addition, since $b \equiv b_{ij}$ is drawn at random, production may be particularly quality-sensitive for lower-ranked products (higher m_{ij}). Nevertheless, the complementarity in (18) introduces a *tendency* for products closer to the core to be produced at higher quality (since these have higher $\widetilde{\lambda}_{ij}$). This will be crucial for our results.

5.3 Profit maximization

Plants maximize each product's profits over their choices of output price p_{ij} and input quality c_{ij}, taking as given each product's demand:[36]

$$\pi(p_{ij}, c_{ij}; \widetilde{\lambda}_{ij}) = \left(p_{ij} - \frac{c_{ij}}{\widetilde{\lambda}_{ij}}\right) x_{ij} - f \tag{19}$$

where x_{ij} is given by product demand as given by (16). Optimization yields the optimal levels:

$$c^*_{ij}(\lambda_i, m_{ij}, b_{ij}) = \left(\widetilde{\lambda}_{ij}\right)^{b_{ij}/2} \tag{20}$$

$$q^*_{ij}(\lambda_i, m_{ij}, b_{ij}) = \left(\widetilde{\lambda}_{ij}\right)^{b_{ij}} \tag{21}$$

$$p^*_{ij}(\lambda_i, m_{ij}, b_{ij}) = \left(\frac{\sigma}{\sigma - 1}\right)\frac{c^*_{ij}}{\widetilde{\lambda}_{ij}} = \left(\frac{\sigma}{\sigma - 1}\right)\left(\widetilde{\lambda}_{ij}\right)^{\frac{b_{ij}}{2} - 1} \tag{22}$$

$$r^*_{ij}(\lambda_i, m_{ij}, b_{ij}) = X \cdot P^{\sigma}\left(\frac{\sigma - 1}{\sigma}\right)^{\sigma - 1}\left(\widetilde{\lambda}_{ij}\right)^{(\sigma-1)\left(\frac{b_{ij}}{2}+1\right)} \tag{23}$$

where c^*_{ij} represents optimal input quality chosen by plant i for its product j, q^*_{ij} is optimal output quality, p^*_{ij} is the optimal output price, and r^*_{ij} is (maximized) revenue.

5.4 Decision to Produce Products

Plants receive their efficiency draw λ_i together with quality capability draws b_{ij} for a sufficiently large number of (potentially produced) products ij, which are ranked by $m_{ij} = \{0, 1, 2, ...\}$. Plants decide to produce a product ij if its variable profits exceed the (annualized) fixed cost f. Note that with CES demand (and thus constant markup $\sigma/(\sigma - 1)$) and marginal costs being constant in quantity x_{ij}, variable profits for a product are given by $\frac{1}{\sigma}r_{ij}$. Thus, total profits associated with product ij follow directly from (23): $\pi_{ij} = \frac{1}{\sigma}r_{ij} - f$, which depends on plant-level efficiency λ_i, product rank m_{ij}, and quality capability b_{ij}. Thus, the zero-profit condition associated with product ij is:

$$\pi_{ij}\left(\lambda_i, m_{ij}, b_{ij}\right) = \frac{1}{\sigma} \cdot r^*\left(\lambda_i, m_{ij}, b_{ij}\right) - f = 0 \tag{24}$$

where λ_i and m_{ij} can be combined to the product-specific efficiency term $\widetilde{\lambda}_{ij} \equiv \phi^{m_{ij}}\lambda_i$, which is reflected in (23). Thus, the model features two sources of heterogeneity ($\widetilde{\lambda}_{ij}$ and b_{ij}). However, for the purpose of the simulation exercise we present below, we only need to know if the

[36] We assume as in Mayer et al. (2014) that there is no within- or across-plant-product cannibalization (i.e., product demand is not directly affected by individual products produced by the same plant, or by products produced by other plants.

combined capability component $\gamma_{ij} \equiv \left(\widetilde{\lambda}_{ij}\right)^{(\sigma-1)\left(\frac{b_{ij}}{2}+1\right)}$ is above or below a threshold at which product-specific profits are zero. This threshold component follows from (23) and (24), and it is defined by the following expression:

$$\overline{\gamma}_{ij} \;=\; \frac{f}{\kappa_1} \tag{25}$$

where $\kappa_1 \equiv (\sigma-1)^{\sigma-1}\sigma^{-\sigma}XP^{\sigma}$, which is constant across plants. Plants decide to produce products for which $\gamma_{ij} = (\phi^{m_{ij}}\lambda_i)^{(\sigma-1)\left(\frac{b_{ij}}{2}+1\right)} \geq \overline{\gamma}_{ij}$. Note that the steps on the product ladder that are observed in the data do note necessarily reflect the theoretically possible ladder steps $m_{ij} = \{0, 1, 2, ...\}$ for a given plant i. For example, if the product with $m_{ij} = 1$ has a very low draw b_{ij}, it may not be produced, while (some) subsequent products $m_{ij} > 1$ may be produced. While the remaining (i.e., produced) products are still clearly ranked by their efficiency $\widetilde{\lambda}_{ij}$, the differences in $\widetilde{\lambda}_{ij}$ across produced products may vary substantially, depending on a plant's draws of b_{ij}.

5.5 Model Predictions

The stylized model delivers several insights. We discuss these following the order of the stylized facts documented above.[37] Both stylized facts 1 and 2 follow directly from the model setup. Product-specific efficiency is given by $\widetilde{\lambda}_{ij} \equiv \phi^{m_{ij}}\lambda_i$. For the core product, this equals λ_i. Consequently, for plants with high core efficiency λ_i, other products $m_{ij} \geq 1$ also tend to be produced more efficiently (stylized fact 1), and efficiency tends to co-move across products provided that innovation (or shocks) affect the plant-level λ_i (stylized fact 2).

Revenue-based product ranking

We now turn to stylized facts 4-7, all of which are derived from ranking products based on sales revenues.[38] Following (23), product-specific revenues are proportional to $\widetilde{\lambda}_{ij}^{(\sigma-1)\left(\frac{b_{ij}}{2}+1\right)}$, where product-specific efficiency $\widetilde{\lambda}_{ij}$ reflects TFPQ (since it is independent of product quality). Thus, products with higher $\widetilde{\lambda}_{ij}$ or higher b_{ij} will be ranked closer to the core. Note that b_{ij} affects the product rank – a crucial difference to the purely TFPQ-based rankings that we discuss below. At the same time, since $b_{ij} \geq 0$ is distributed independently of the product rank m_{ij}, products with high revenues will also – at least on average – tend to have high $\widetilde{\lambda}_{ij}$. In words, for revenue based rankings, TFPQ will tend to be larger for core products (stylized fact 4). Qualitatively, this pattern holds irrespective of the average capability for quality differentiation, i.e., irrespective of the mean of b_{ij} (first part of stylized fact 7). In particular, it holds even in the extreme case

[37] Our model yields also a prediction for product variety: Plants with higher core efficiency λ_i tend to produce a larger product portfolio. Since this prediction follows also from previous models, we discuss it in the appendix.

[38] We discuss stylized fact 3 below.

when there is no potential for quality differentiation ($b_{ij} = 0$).

Next, marginal costs are given by (17), so that (20) implies $MC_{ij} = \widetilde{\lambda}_{ij}^{\frac{b_{ij}}{2}-1}$. Since products with higher $\widetilde{\lambda}_{ij}$ or higher b_{ij} are closer to the core, these products' marginal costs will also tend to be higher (stylized fact 5). Note that for relatively small b_{ij} (low degree of quality differentiation), the revenue-based ranking is mostly driven by differences in physical efficiency $\widetilde{\lambda}_{ij}$. At the same time, for products with low b_{ij}, the exponent in MC_{ij} is close to zero or even becomes negative. Thus, the relationship between MC_{ij} and (revenue-based) product rank is ambiguous for (relatively) undifferentiated products. In fact, for completely homogenous products ($b_{ij} = 0$), the relationship is negative, and for somewhat differentiated products with $b_{ij} \approx 2$, the relationship is flat (second part of stylized fact 7). Nevertheless, since the revenue-based ranking is 'biased' towards high-b_{ij} products, core products are more likely to have particularly high b_{ij} draws and thus also high MC_{ij}. Consequently, core products can exhibit systematically higher marginal costs in revenue-based rankings even if b_{ij} is relatively small on average. We show this in the simulation below. Finally, stylized fact 6 on relatively constant markups along the product ladder follows by construction in our model, due to CES demand.

Efficiency-based product ranking

We now turn to the predicted patterns when products are ranked by their physical efficiency $\widetilde{\lambda}_{ij}$ within plants (stylized fact 8). Recall from our discussion above that $MC_{ij} = \widetilde{\lambda}_{ij}^{\frac{b_{ij}}{2}-1}$. This reflects the two-sided effect that higher product-specific efficiency $\widetilde{\lambda}_{ij}$ has on marginal costs: on the one hand, according to (17) a more efficient product is produced with lower marginal cost; but on the other hand, plants choose a higher quality for more efficient products, which raises input cost as given by (20). As long as $b_{ij} < 2$ (i.e., for plant-products with relatively low quality capability), the latter effect dominates so that marginal costs (and thus prices, given the constant markup) are *lower* for more efficient products. Consequently, if the parameter b is relatively small on average, then our model predicts that marginal costs and prices are on average lower for core products, in line with stylized fact 8. A low average b can have two reasons: our sample may be dominated by relatively undifferentiated products, or b may be low even for differentiated products. To shed more light on this angle, we replicate the results from Table 7 for the subsamples of homogenous vs. differentiated products.[39] Table A.1 in the appendix shows that marginal costs and prices are significantly smaller for core products in the subsample of 'homogenous' products (Panel A). However, this pattern prevails (albeit, as expected, somewhat weaker) in the subsample with differentiated products (Panel B). This suggests that the parameter b is relatively small on average even for differentiated products. We

[39] We define the degree of product differentiation at the plant level based on the liberal classification in Rauch (1999). "Homogeneous" is for product categories that are "traded on organized exchanges" or are "referenced priced"; "differentiated" is based on Rauch's "differentiated" category.

use this below in the calibration to choose a conservative parametrization such that b has a mean of 2.

Finally, we turn to stylized fact 3 on TFPR, which is given by the product of p_{ij} and TFPQ. Note that this stylized fact holds for both revenue- and efficiency-based rankings in Tables 5 and 7. Here, we discuss it in the context of the latter. Using (22), we obtain $TFPR_{ij} = p_{ij}\widetilde{\lambda}_{ij} = \widetilde{\lambda}_{ij}^{\frac{b_{ij}}{2}}$. Consequently, TFPR should be increasing in physical efficiency, unless b_{ij} is very small on average. In revenue based rankings, where b_{ij} also affects the ranking, core products should show an even clearer pattern of higher TFPR. A possible explanation why we do not find this in the data is that we follow the most commonly used methodology to compute TFPR, using sector-level deflators for input prices (as opposed to plant-level deflators, as we do for TFPQ). If core products are of higher quality, they will also tend to have higher input costs – which TFPQ empirically takes into account, but not TFPR. Consequently, the common strategy for estimating TFPR may fail to correctly identify the (revenue-based) productivity advantage of core products.

5.6 Main Predications and Simulation

In the theoretical description above we argued that our stylized model can replicate the stylized facts found in the data if i) quality capability is drawn independently of the product-specific efficiency rank m_{ij}, and ii) if quality capability b is relatively small on average. In the following, we calibrate and simulate the model to further support this argument.

We present a simulation of the model based on equations (20)–(23) and show how our simple framework can account for the divergent results that we obtain for prices when we rank products within plants by sales revenues versus physical productivity.[40] Following the recent trade literature, we specify a Pareto distribution for core efficiency λ. For the particular values of the shape parameter k of the core efficiency distribution, and for the size of the step of the efficiency ladder (ω), we follow the values presented in Mayer et al. (2014). In particular, we set $k = 3.25$, and $\omega = 0.96$. The former parameter is in the middle of the range of shape parameters considered by Mayer et al. (2014). Regarding ω, we set its value to be consistent with $k = 3.25$ according to the strategy described in Mayer et al. (2014).[41] We set the elasticity of substitution $\sigma = 5$, which implies an average markup of 1.25. As in the case of core efficiencies, we specify

[40] Before proceeding, a word of caution is due. The simulation exercise provides evidence that the model can potentially fit the puzzling patterns documented in the data; however, it is not calibrated to *fit* the data. That is, we do not estimate or calibrate the model to particular moments in the data. Instead, we guide the choice of parameters based on estimates available in the related literature whenever possible, and choose values that generate product distributions close to the empirical distribution in Table 1 for the remaining parameters that are not readily available from previous contributions.

[41] Mayer et al. (2014) show that when core efficiencies are distributed Pareto, there is a linear relationship between log-revenues of each variety and its associated ladder step m, with the slope $\vartheta \equiv k \ln \omega$. We take their estimates of $\hat{\vartheta} = -0.13$ from French export data and $k = 3.2$ as given, to recover ω.

a Pareto distribution for b over $[1, +\infty)$, with shape parameter equal to 2,[42] and impose zero correlation between λ and b. Finally, we choose the combined capability threshold $\gamma^*_{ij}(m_{ij})$ so that 48.7% of all simulated plants produce multiple products. This matches the share of multi-product plants in ENIA reported in Table 1.

The simulation algorithm involves the following steps: (i) Draw N plant-specific core efficiency components λ (where N is the number of plants drawn); (ii) Draw $N \times M$ plant-product specific b_{ij} (with M the maximum number of products that we allow each plant to have); (iii) Given (λ, b_{ij}), check for each product whether $\gamma_{ij}(m_{ij})$ is greater than $\gamma^*_{ij}(m_{ij})$; if not, drop the product form the sample, and (iv) Compute variables of interest, such as plant-product price and revenues. For the particular exercise we present in this section, we set $N = 10,000$ and $M = 30$.

Figure 1 shows the simulation results. The left panel ranks products within plants by revenues as given by (23), while the right panel ranks products by plant-product-specific efficiency, as given by $\tilde{\lambda}(m) = \omega^m \lambda$.[43] The figure shows the same pattern as found in the Chilean manufacturing data: prices are higher for core products when ranked by revenues, but they are lower for core products when ranked by physical productivity. The same reversal holds for marginal costs, which we do not separately show since MC are proportional to prices, given the constant markup. The intuition for the reversal is as discussed in the theoretical section above: revenue-based product ranks are 'biased' towards products with high b draws, which are also associated with higher MC and prices. Thus, the simulation also underlines the importance of the ranking variable: product ladders based on efficiency vs. revenues yield radically different results for prices and marginal costs.

6 Conclusion

Product-level efficiency is a key theoretical component in a growing literature that examines trade-related dynamics within firms or plants. So far, data limitations have made it difficult to construct product-level efficiency, and productivity patterns across products within plants were largely unexplored. We exploited a uniquely detailed Chilean dataset to compute several alternative efficiency measures at the product level within plants. We have established numerous novel stylized facts in three areas. First, on product-level efficiency patterns, we showed that productive plants tend to be relatively efficient across the board, not just for their core products. This provides direct evidence for a central assumption that underlies prominent models of flexible manufacturing such as Bernard et al. (2011) or Mayer et al. (2014). There, a common efficiency draw affects all products within a plant. Our finding that productive plants tend to

[42] The main results presented below are unchanged in a reasonable neighborhood for this value.

[43] To be consistent with results in Tables 5 and 7, we control in both panels of Figure 1 for plant fixed effects. Thus, the figure displays variation across products *within* plants.

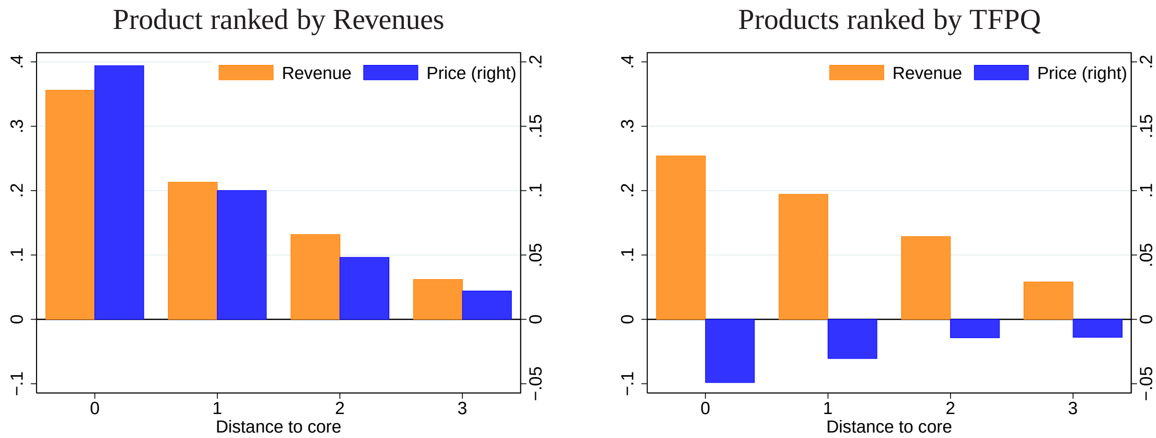

Figure 1: Simulated Revenue and Price Across Products within Plants

Notes: The figure shows the simulated average revenue and price across products within plants, resulting from simulating the model in section 5. Product $m = 0$ corresponds to the core product defined according to physical productivity (left) and product revenue (right). We only consider plants producing 5 or more products. The parameters underlying the simulation are listed in Section 5.6. Both panels controls for plant fixed effects and simulate the model for a single year.

be relatively efficient at all their products supports this setup. Second, we have shown that the typically used sales-based product ranks correctly reflect higher physical efficiency (TFPQ); however – seemingly contradictory – marginal costs are higher for top-ranked sales products. We showed that this discrepancy is likely driven by product quality. Our results thus emphasize the importance of using the appropriate ranking variable when testing predictions of flexible manufacturing models. The same is true for results that involve the prominent metric of export skewness towards core products. Product ladders based on marginal costs or revenue productivity do not show export skewness, while TFPQ-based rankings do yield skewness towards the most efficient product and thus aggregate efficiency gains from trade.

References

Ackerberg, D. A., K. Caves, and G. Frazer (2015). Identification Properties of Recent Production Function Estimators. *Econometrica 83*(6), 2411–2451.

Amiti, M. and J. Konings (2007). Trade Liberalization, Intermediate Inputs, and Productivity: Evidence from Indonesia. *American Economic Review 97*(5), 1611–1638.

Antoniades, A. (2015). Heterogeneous Firms, Quality, and Trade. *Journal of International Economics 95*(2), 263–273.

Atkin, D., A. K. Khandelwal, and A. Osman (2014). Exporting and Firm Performance: Evidence from a Randomized Trial. Working Paper.

Bernard, A. B., J. Eaton, J. B. Jensen, and S. Kortum (2003). Plants and Productivity in International Trade. *American Economic Review 93*(4), 1268–1290.

Bernard, A. B., S. Redding, and P. K. Schott (2011). Multi-Product Firms and Trade Liberalization. *Quarterly Journal of Economics 126*(3), 1271–1318.

Bernard, A. B., S. J. Redding, and P. K. Schott (2010). Multiple-Product Firms and Product Switching. *American Economic Review 100*(1), 70–97.

Bustos, P. (2011). Trade Liberalization, Exports, and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinian Firms. *American Economic Review 101*(1), 304–40.

De Loecker, J. (2013). A Note on Detecting Learning by Exporting. *American Economic Journal: Macroeconomics*. forthcoming.

De Loecker, J., P. K. Goldberg, A. K. Khandelwal, and N. Pavcnik (2016). Prices, Markups and Trade Reform. *Econometrica 84*(2), 445–510.

De Loecker, J. and F. Warzynski (2012). Markups and Firm-Level Export Status. *American Economic Review 102*(6), 2437–2471.

Dhyne, E., A. Petrin, V. Smeets, and F. Warzynski (2016). Multi-Product Firms, Import Competition, and the Evolution of Firm-product Technical Efficiencies. Working Paper.

Eckel, C., L. Iacovone, B. Javorcik, and J. P. Neary (2015). Multi-Product Firms at Home and Away: Cost- versus Quality-Based Competence. *Journal of International Economics 95*(2), 216–232.

Eckel, C. and J. P. Neary (2010). Multi-Product Firms and Flexible Manufacturing in the Global Economy. *Review of Economic Studies 77*(1), 188–217.

Foster, L., J. Haltiwanger, and C. Syverson (2008). Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability? *American Economic Review 98*(1), 394–425.

Garcia-Marin, A. and N. Voigtländer (2013). Exporting and Plant-Level Efficiency Gains: It's in the Measure. NBER working paper #19033.

Hall, R. E. (1986). Market Structure and Macroeconomic Fluctuations. *Brookings Papers on Economic Activity 17*(2), 285–338.

Hsieh, C.-T. and P. J. Klenow (2009). Misallocation and Manufacturing TFP in China and India. *Quarterly Journal of Economics 124*(4), 1403–1448.

Katayama, H., S. Lu, and J. R. Tybout (2009). Firm-Level Productivity Studies: Illusions and a Solution. *International Journal of Industrial Organization 27*(3), 403–413.

Khandelwal, A. (2010). The Long and Short (of) Quality Ladders. *Review of Economic Studies 77*(4), 1450–1476.

Kugler, M. and E. Verhoogen (2012). Prices, Plant Size, and Product Quality. *Review of Economic Studies 79*(1), 307–339.

Lamorgese, A. R., A. Linarello, and F. Warzynski (2014). Free Trade Agreements and Firm-Product Markups in Chilean Manufacturing. Working Paper, School of Economics and Management, University of Aarhus.

Levinsohn, J. and A. Petrin (2003). Estimating Production Functions Using Inputs to Control for Unobservables. *Review of Economic Studies 70*(2), 317–341.

Mayer, T., M. J. Melitz, and G. I. P. Ottaviano (2014). Market Size, Competition, and the Product Mix of Exporters. *American Economic Review 104*(2), 494–536.

Melitz, M. J. (2003). The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. *Econometrica 71*(6), 1695–1725.

Melitz, M. J. and G. I. P. Ottaviano (2008). Market Size, Trade, and Productivity. *Review of Economic Studies 75*(1), 295–316.

Montenegro, C. E. and C. Pagés (2004). Who Benefits from Labor Market Regulations? Chile, 1960-1998. In *Law and Employment: Lessons from Latin America and the Caribbean*, NBER Chapters, pp. 401–434. National Bureau of Economic Research, Inc.

Olley, G. S. and A. Pakes (1996). The Dynamics of Productivity in the Telecommunications Equipment Industry. *Econometrica 64*(6), 1263–1297.

Pavcnik, N. (2002). Trade Liberalization, Exit and Productivity Improvements: Evidence from Chilean Plants. *Review of Economic Studies 69*, 245–76.

Rauch, J. E. (1999). Networks versus Markets in International Trade. *Journal of International Economics 48*(1), 7–35.

APPENDIX

Model Prediction on Product Variety

Our model yields a prediction for product variety produced by plants: Plants with higher core efficiency λ_i tend to produce a larger product portfolio. This prediction is not novel. It also follows from the models by Eckel and Neary (2010), Bernard et al. (2010), Bernard et al. (2011), and Mayer et al. (2014). In the context of our model, the prediction is obtained as follows. For *given* λ_i and m_{ij}, the zero-profit condition (24) implicitly defines a threshold $\overline{b}_{ij}(\lambda_i, m_{ij})$ such that for quality capability draws $b_{ij} \geq \overline{b}_{ij}$, the plant chooses to produce the product. Note that (24) implies that a plant-product with lower efficiency $\widetilde{\lambda}_{ij}$ (either due to low λ_i or due to high m_{ij}) requires a higher b_{ij} to be produced profitably.[1] Thus, varieties closer to the core (low m_{ij}) are produced even if b_{ij} is relatively low. In contrast, periphery varieties require increasingly higher b_{ij} to be produced profitably. Since b_{ij} is distributed independently of m_{ij}, the model predicts that – on average – plants with relatively high core efficiency λ_i tend to produce a wider range of varieties than plants with relatively low core efficiency. Note that this proposition holds on average. However, the model allows for heterogeneity: plants with low λ_i may still produce a rich set of products if they have particularly high draws b_{ij} for many products ij.

Additional Empirical Results

[1] This can be shown by solving (24) for $\overline{b}_{ij}$ as a function of $\widetilde{\lambda}_{ij} \equiv \phi^{m_{ij}} \lambda_i$:

$$\overline{b}_{ij}(\widetilde{\lambda}_{ij}) = 2\left[\frac{1}{(\sigma-1)} \frac{[\log(f) - \log((\sigma-1)^{\sigma-1}\sigma^{-\sigma} X P^{\sigma})]}{\log(\widetilde{\lambda}_{ij})} - 1\right]$$

Table A.1: Core Competence by TFPQ Rank: Sample Splits by Product Differentiation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dep. Var.:	log(sales)	log(volume)	log(Price)	log(TFPR)	log(TFPQ)	log(MC)	log(Markup)
			Panel A: Homogeneous Products				
Top product	1.962***	3.384***	-.666***	-.0260**	3.735***	-.727***	.0116
	(.0677)	(.0728)	(.0535)	(.0129)	(.0524)	(.0543)	(.0101)
Top 2nd	1.480***	2.565***	-.460***	-.0300**	2.878***	-.503***	.00384
	(.0619)	(.0668)	(.0490)	(.0118)	(.0478)	(.0507)	(.00950)
Top 3rd	.930***	1.866***	-.296***	-.0162	2.128***	-.337***	.00408
	(.0652)	(.0648)	(.0479)	(.0119)	(.0447)	(.0487)	(.00959)
Top 4th	.527***	1.057***	-.119**	-.0105	1.341***	-.138***	.00119
	(.0658)	(.0634)	(.0482)	(.0132)	(.0432)	(.0479)	(.00984)
			Panel B: Differentiated Products				
Top product	2.323***	3.008***	-.302***	.00579	2.695***	-.322***	.0229***
	(.0520)	(.0522)	(.0267)	(.0115)	(.0374)	(.0281)	(.00801)
Top 2nd	1.793***	2.324***	-.222***	.0103	2.076***	-.221***	.0198***
	(.0472)	(.0480)	(.0280)	(.0125)	(.0334)	(.0271)	(.00720)
Top 3rd	1.358***	1.754***	-.131***	.00822	1.563***	-.128***	.0109
	(.0461)	(.0465)	(.0248)	(.0109)	(.0305)	(.0253)	(.00741)
Top 4th	.800***	1.087***	-.122***	.00761	1.000***	-.132***	.0142*
	(.0486)	(.0471)	(.0262)	(.0103)	(.0292)	(.0271)	(.00790)
Plant-year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Product FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	14,304	14,304	14,304	14,304	14,304	14,304	14,304
R-squared	.630	.721	.496	.629	.811	.524	.771

Notes: The table regresses each column variable against categorical variables for the top, second, third and fourth best performing product of the plant, interacted with a dummy for homogenous products. Within-plant product rankings are computed in terms of normalized product sales (based on the Torqvist index described in section 3.2). We update the rank in each sample period, potentially allowing products to switch ranks within plants over time. The sample includes all plants that produce at least 5 products. We define degree of differentiation at the plant level based on the liberal classification in Rauch (1999). For this, we use concordances between SITC (used by Rauch) and ISIC codes of the main product (used by the Chilean ENIA). This yields a plant-level classification into homogenous and differentiated. "Homogeneous" is for product categories that according to Rauch (1999) are "traded on organized exchanges" or are "referenced priced"; "differentiated" is based on Rauch's "differentiated" category. Standard errors (clustered at the plant-year level) are in parenthesis. Key: *** significant at 1%; ** 5%; * 10%.