

Exporting and Plant-Level Efficiency Gains: It's in the Measure*

Alvaro Garcia-Marin[†]

Nico Voigtländer[‡]

forthcoming in the Journal of Political Economy

Abstract

While there is strong evidence for productivity-driven selection into exporting, the empirical literature has struggled to identify export-related efficiency gains within plants. Previous research typically derived revenue productivity (TFPR), which is downward biased if more efficient producers charge lower prices. Using a census panel of Chilean manufacturing plants, we compute plant-product level marginal cost as an efficiency measure that is not affected by output prices. For export *entrant* products, we find efficiency gains of 15-25%. Because markups remain relatively stable after export entry, most of these gains are passed on to customers in the form of lower prices, and are thus not reflected by TFPR. These results are confirmed when we use tariffs to predict export entry. We also document very similar results in Colombian and Mexican manufacturing plants. In addition, we find sizeable efficiency gains for tariff-induced export expansions of *existing* exporters. Only one quarter of these gains are reflected by TFPR, due to a partial rise in markups. Our results thus imply that within-plant gains from trade are substantially larger than previously documented. Evidence suggests that a complementarity between exporting and investment in technology is an important driver behind these gains.

*We would like to thank the editor, Ali Hortaçsu, two anonymous referees, as well as Roberto Álvarez, Johannes van Biesebroeck, Joaquin Blaum, Lorenzo Caliendo, Donald Davis, Jan De Loecker, Gilles Duranton, Eduardo E. Engel, Amit Khandelwal, Edward Leamer, Steve Redding, Ariell Reshef, Andrés Rodríguez-Clare, Chad Syverson, and Eric Verhoogen, as well as seminar audiences at CEU, Chicago Booth, Copenhagen Business School, Columbia, CREI, Dartmouth, the FREIT Conference in Virginia, KU Leuven, Mannheim, Frankfurt, Pontificia Universidad Católica de Chile, Princeton, RMET Conference, TIGN Conference, UCLA, and Universidad de Chile for helpful comments and suggestions. Claudio Bravo-Ortega and Lucas Navarro kindly shared data to complement our Chilean plant-product panel. Henry Alexander Ballesteros and Oscar Fentanes provided excellent research assistance in Bogotá and Mexico City, respectively. We gratefully acknowledge research support from the UCLA Ziman Center's Rosalinde, Arthur Gilbert Program in Real Estate, Finance and Urban Economics, and CONICYT-Chile through project Fondecyt Iniciación en Investigación 11150432.

[†]Universidad de Chile. Email: algarcia@fen.uchile.cl

[‡]Corresponding author. UCLA, NBER, and CEPR. Email: nico.v@ucla.edu

1 Introduction

A large literature in empirical trade has shown that exporting firms and plants are more productive than their non-exporting counterparts. In principle, this pattern may emerge because exporters have higher productivity to start with, or because they become more efficient after export entry. The former effect – selection across plants – has received strong theoretical and empirical support (c.f. Melitz, 2003; Pavcnik, 2002). On the other hand, evidence for export-related *within*-plant productivity gains is much more sparse, with the majority of empirical studies finding no effects (for recent reviews of the literature see Syverson, 2011; Bernard, Jensen, Redding, and Schott, 2012). In particular, the productivity trajectory of plants or firms typically look flat around the time of export entry, suggesting that producers do not become more efficient after foreign sales begin.¹ This is surprising, given that exporters can learn from international buyers and have access to larger markets to reap the benefits of innovation or investments in productive technology (Bustos, 2011). In other words, there is strong evidence for a complementarity between export expansions and technology upgrading (c.f. Lileeva and Trefler, 2010; Aw, Roberts, and Xu, 2011), and technology upgrading, in turn, should lead to observable efficiency increases. Why has the empirical literature struggled to identify such gains?

In this paper, we use rich Chilean, Colombian, and Mexican data to show that flat productivity profiles after export expansions are an artefact of the measure: previous studies have typically used revenue-based productivity, which is affected by changes in prices. If cost savings due to gains in *quantity* productivity are passed on to buyers in the form of lower prices, then revenue-based productivity will be downward biased (Foster, Haltiwanger, and Syverson, 2008).² Consequently, accounting for pricing behavior (and thus markups) is key when analyzing efficiency trajectories. We show in a simple framework that under a set of non-restrictive assumptions (which hold in our data), marginal costs are directly (inversely) related to quantity productivity, while revenue productivity reflects efficiency gains only if markups rise.

We begin by using our main dataset – an unusually rich panel of Chilean manufacturing plants between 1996 and 2007 – to analyze the trajectories of marginal cost, markups, and prices around

¹Early contributions that find strong evidence for selection, but none for within-firm efficiency gains, include Clerides, Lach, and Tybout (1998) who use data for Colombian, Mexican, and Moroccan producers, and Bernard and Jensen (1999) who use U.S. data. Most later studies have confirmed this pattern. Among the few studies that document within-plant productivity gains are De Loecker (2007) and Lileeva and Trefler (2010). Further reviews of this ample literature are provided by Wagner (2007, 2012).

²Recent evidence suggests that this downward bias also affects the link between trade and productivity. Smeets and Warzynski (2013) construct a firm level price index to deflate revenue productivity and show that this correction yields larger international trade premia in a panel of Danish manufacturers. Eslava, Haltiwanger, Kugler, and Kugler (2013) use a similar methodology to show that trade-induced reallocation effects across firms are also stronger for price-adjusted productivity.

export entry and export expansions. To derive markups at the plant-product level, we apply the method pioneered by De Loecker and Warzynski (2012), in combination with the uniquely detailed reporting of product-specific input cost shares by Chilean multi-product plants. In addition, our dataset comprises 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, Goldberg, Khandelwal, and Pavcnik, 2016). This procedure is flexible with respect to the underlying price setting model and the functional form of the production function. Importantly, by disentangling the individual components, we directly observe the extent to which efficiency gains (lower marginal costs) are translated into higher revenue productivity (by raising markups), or passed on to customers (by reducing prices). To compare our results with the typically used efficiency measure, we also compute revenue productivity (TFPR) at the plant-product level.

Figure 1 presents our main results – within plant-product trajectories for export entrants in Chile. Time on the horizontal axis is normalized so that zero represents the export entry year. The left panel confirms that, in line with most of the previous literature, the trajectory of TFPR is flat around export entry. The right panel disentangles this pattern and shows that (i) marginal costs within plant-products drop by approximately 15-25% during the first three years after export entry; (ii) prices fall by a similar magnitude as marginal costs; (iii) markups do not change significantly during the first years following export entry. Our findings suggest that export entrants do experience efficiency gains, but that these are passed on to their customers. In other words, constant markups and falling prices explain why revenue productivity is flat around export entry.

Our results for export entrants are very similar when we use propensity score matching to construct a control group of plant-products that had an a-priori comparable likelihood of entering the export market. In addition, we show that we obtain very similar results when (i) computing physical efficiency (TFPQ, which requires stronger assumptions than marginal costs at the plant-product level, as discussed in Section 2.5), and (ii) when using reported average variable costs at the plant-product level. This suggests that our findings are not an artefact of the methodology used to calculate marginal costs; in fact, the computed marginal costs are strongly correlated with the reported average variable costs. We also discuss that our results are unlikely to be confounded by changes in product quality.³ We then exploit falling tariffs on Chilean products in destination

³The bias that may result from changes in quality works against finding efficiency gains with our methodology: exported goods from developing countries are typically of higher quality than their domestically sold counterparts (c.f. Verhoogen, 2008) and use more expensive inputs in production (Kugler and Verhoogen, 2012). Thus, exporting should *raise* marginal costs. This is confirmed by Atkin, Khandelwal, and Osman (2014) who observe that quality upgrading of Egyptian rug exporters is accompanied by higher input prices. Using Mexican data, Iacovone and Javorcik (2012) provide evidence for quality upgrading right before, but not after, export entry.




Figure 1: Trajectories for Export Entrants in Chile

Notes: Data are from the Chilean Annual Industrial Survey (ENIA) for the period 1996-2007. The figure shows the trajectories for our main outcome variables before and after export entry; period $t = 0$ corresponds to the export entry year. The left panel shows the trajectory for revenue productivity (TFPR); the right panel, for marginal cost, price, and markup. All results are at the plant-product level. A plant-product is defined as an entrant if it is the first product exported by a plant and is sold domestically for at least one period before entry into the export market (see Section 3.2). Coefficient estimates are reported in Table 1. The lines and whiskers represent 90% confidence intervals.

countries to predict the timing of export entry. Due to the limited variation in tariffs, this exercise serves as a check, rather than the core of our analysis. Nevertheless, the combined variation in tariffs over time and across 4-digit sectors is sufficient to yield a strong first stage. We confirm our findings from within-plant trajectories: tariff-induced export entry is associated with marginal costs declining by approximately 25%. In relative terms, this corresponds to approximately one-third of the standard deviation in year-to-year changes in marginal costs across all plant-products in the sample.

We provide evidence that technology upgrading is the most likely explanation for declining marginal costs at export entry. Plant-level investment (especially in machinery) spikes right after export entry. In addition, marginal costs drop particularly steeply for plants that are initially less productive. This is in line with Lileeva and Trefler (2010), who point out that, for the case of investment-exporting complementarity, plants that start off from lower productivity levels will only begin exporting if the associated expected productivity gains are large.

In addition to export entry, we also analyze export expansions of *existing* exporters that are induced by falling export tariffs on Chilean products. Over our sample period, these tariff-induced export expansions lead to a decline in marginal costs by approximately 20% among existing exporters. Since export expansions are accompanied by investment in capital, technology upgrading

is a likely driver of efficiency gains among existing exporters, as well. We also show that in the case of established exporters, pass-through of efficiency gains to customers is more limited than for new export entrants: about three quarters of the decline in marginal costs translate into lower prices, and the remainder, into higher markups. Consequently, TFPR also increases and reflects about one-fourth of the actual efficiency gains. Thus, while the downward bias of TFPR is less severe for established exporters, it still misses a substantial part of efficiency increases.

Why are markups stable around export entry, but increase for established exporters after tariff-induced expansions? This pattern is compatible with a ‘demand accumulation process’ (Foster, Haltiwanger, and Syverson, 2016) – while existing exporters already have a customer base abroad, new entrants may use low prices to attract buyers.⁴ To support this interpretation, we separately analyze the domestic and export price of the same product in a subset of years with particularly detailed pricing information. We find that for export entrants, the export price drops more than its domestic counterpart (19% vs. 8%). There is also some evidence in our data that markups grow as export entrants become more established.⁵

Finally, we examine whether our main findings hold in two additional countries with detailed manufacturing panel data that are suited for our analysis: Colombia (2001-13) and Mexico (1994-2003). Both datasets have been used extensively in studies of international trade, and we show that they are representative of the stylized facts documented in the literature (c.f. Bernard and Jensen, 1999).⁶ We find strong evidence for our main results. As shown in Figure 2 for Colombia and in Figure 3 for Mexico, there is no relationship between TFPR and export entry. On the other hand, marginal costs decline strongly after export entry in both countries. Prices fall hand-in-hand with marginal costs, while markups are relatively stable.⁷ We also show that investment (especially in machinery and equipment) spikes after export entry in both samples. The fact that our main findings hold for exporting plants in three different countries strongly suggests that our main conclusion is broadly applicable: revenue-based productivity measures miss important export-related efficiency gains within manufacturing plants.

Our findings relate to a substantial literature on gains from trade. Trade-induced competition

⁴Foster et al. (2016) provide evidence that supports this mechanism in the domestic market. They show that by selling more today, firms expand buyer-supplier relationships and therefore shift out their future demand.

⁵There is a longer delay between export entry and changes in markups in our data as compared to De Loecker and Warzynski (2012), who document increasing markups right after export entry for Slovenian firms. However, our data confirm De Loecker and Warzynski’s cross-sectional finding that exporters charge higher markups.

⁶One limitation is that – unlike the Chilean data – the Colombian and Mexican data do not provide product-specific variable costs. We therefore cannot exploit this information to derive product-specific markups and marginal costs in multi-product plants. Consequently, we restrict our analysis to the subset of single-product plants, where all inputs are clearly related to the (single) produced output.

⁷We discuss the (quantitatively small) increase of markups after export entry in Colombia in Section 6.

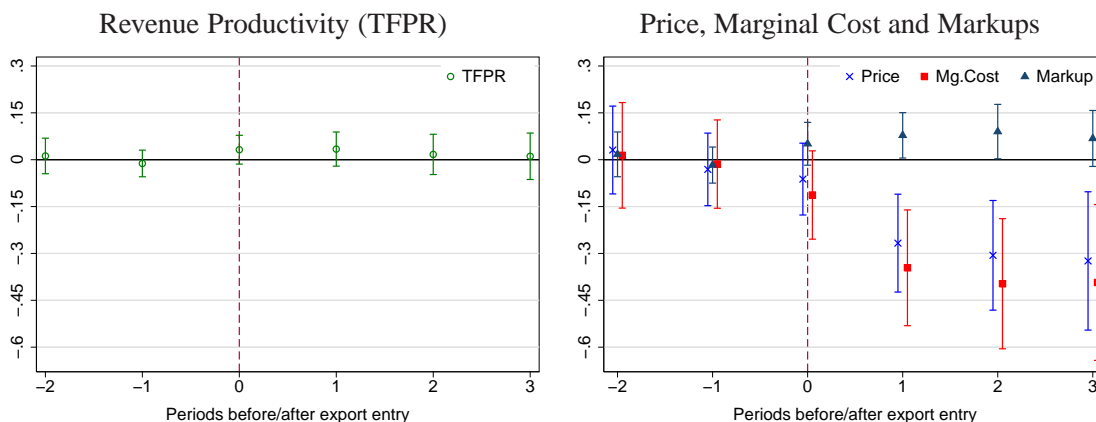


Figure 2: Trajectories for Export Entrants in Colombia

Notes: Data are from the Colombian Annual Manufacturing Survey for the period 2001-13 (described in Appendix B.4). The figure shows the trajectories for our main outcome variables before and after export entry; period $t = 0$ corresponds to the export entry year. The left panel shows the trajectory for revenue productivity (TFPR); the right panel, for marginal cost, price, and markup. All results are for single-product plants. The coefficient estimates are reported in Table A.11. The lines and whiskers represent 90% confidence intervals.

can contribute to the reallocation of resources from less to more efficient producers. Bernard, Eaton, Jensen, and Kortum (2003) and Melitz (2003) introduce this reallocation mechanism in trade theory, based on firm-level heterogeneity. The empirical evidence on this mechanism is vast, and summarizing it would go beyond the scope of this paper.⁸ In contrast, the majority of papers studying productivity *within* firms or plants have found no or only weak evidence for export-related gains. Clerides et al. (1998, for Colombia, Mexico, and Morocco) and Bernard and Jensen (1999, using U.S. data) were the first to analyze the impact of exporting on plant efficiency. Both document no (or quantitatively small) empirical support for this effect, but strong evidence for selection of productive firms into exporting. The same is true for numerous papers that followed: Aw, Chung, and Roberts (2000) for Taiwan and Korea, Alvarez and López (2005) for Chile, and Luong (2013) for Chinese automobile producers.⁹ The survey article by ISGEP (2008) compiles micro level panels from 14 countries and finds nearly no evidence for within-plant productivity

⁸Two influential early papers are Bernard and Jensen (1999) and Pavcnik (2002), who analyze U.S. and Chilean plants, respectively. Recent contributions have also drawn attention to the role of imports. Amiti and Konings (2007) show that access to intermediate inputs has stronger effects on productivity than enhanced competition due to lower final good tariffs. Goldberg, Khandelwal, Pavcnik, and Topalova (2010) provide evidence from Indian data that access to new input varieties is an important driver of trade-related productivity gains.

⁹Alvarez and López (2005) use an earlier version of our Chilean plant panel. They conclude that "Permanent exporters are more productive than non-exporters, but this is attributable to initial productivity differences, not to productivity gains associated to exporting." [p.1395] We confirm this finding when using revenue-productivity.

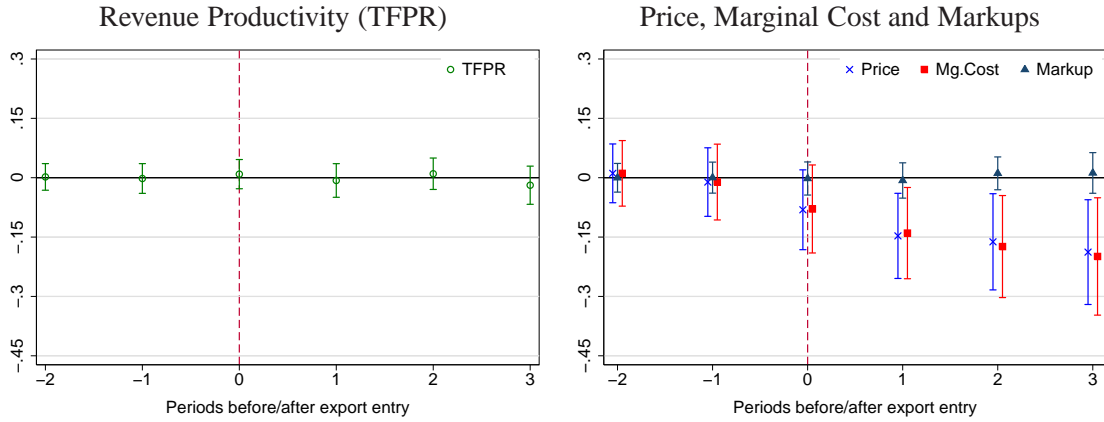


Figure 3: Trajectories for Export Entrants in Mexico

Notes: Data are from the Mexican Annual Industrial Survey for the period 1994-2003 (described in Appendix B.5). The figure shows the trajectories for our main outcome variables before and after export entry; period $t = 0$ corresponds to the export entry year. The left panel shows the trajectory for revenue productivity (TFPR); the right panel, for marginal cost, price, and markup. All results are for single-product plants. The coefficient estimates are reported in Table A.12. The lines and whiskers represent 90% confidence intervals.

increases after entry into the export market.

The few papers that have found within-plant productivity gains typically analyzed periods of rapid trade liberalization, such as De Loecker (2007) for the case of Slovenia and Lileeva and Trefler (2010) for Canada, or demand shocks due to large (and permanent) exchange rate changes such as Park, Yang, Shi, and Jiang (2010).¹⁰ Our results illustrate why it may be more likely to identify within-plant gains in *revenue* productivity during periods of major tariff reductions: especially for established exporters, declining export tariffs have effects akin to a demand shock, which may lead to rising markups in general demand structures such as Melitz and Ottaviano (2008). Then, TFPR will rise because of its positive relationship with markups.¹¹ The downward bias in TFPR can also be tackled by computing quantity productivity (TFPQ). In a paper that follows ours, Lamorgese, Linarello, and Warzynski (2014) document rising TFPQ for Chilean export entrants.¹² Our find-

¹⁰Van Biesebroeck (2005) also documents efficiency gains after export entry – albeit in a less representative setting: among firms in sub-Saharan Africa. These gains are likely due to economies of scale, because exporting lifts credit constraints and thus allows sub-Saharan African firms to grow.

¹¹Potentially, markups could rise even if the actual efficiency is unchanged, causing an upward-bias of TFPR. However, our data suggest that changes in markups generally fall short of actual efficiency gains, so that altogether, TFPR is downward biased.

¹²We discuss below that marginal costs have an advantage over TFPQ in the context of our study: For multi-product plants, *product*-level marginal costs can be computed under relatively unrestrictive assumptions. This allows us to analyze efficiency gains by decomposing prices into markups and marginal costs – all variables that naturally vary at the product level. Disentangling these components also has the advantage that we can analyze pass-through of

ings are compatible with Caliendo, Mion, Opromolla, and Rossi-Hansberg (2015) who show that in response to productivity or demand shocks, firms may reorganize their production by adding a management layer. This causes TFPQ to rise, while TFPR falls because the increase in output quantity leads to lower prices.

Relative to the existing literature, we make several contributions. To the best of our knowledge, this paper is the first to use marginal cost as a measure of efficiency that is not affected by the pricing behavior of exporters, and to document a strong decline in marginal costs after export entry and tariff-induced export expansions.¹³ Second, we discuss in detail the conditions under which declining marginal costs reflect gains in physical efficiency. Third, we show that disentangling the trajectories of prices and physical efficiency is crucial when analyzing export-related efficiency gains: it allows us to quantify the bias of the traditional revenue-based productivity measure. We find that TFPR misses almost all efficiency gains related to export entry, and a substantial share of the gains from tariff-induced export expansions. Consequently, we identify substantial export-related efficiency gains that have thus far passed under the radar. This also applies to the few studies that *have* found export related changes in TFPR within plants: our results suggest that the actual magnitude of efficiency gains is likely larger. Our study thus complements a substantial literature that argues that within-plant efficiency gains should be expected.¹⁴ Fourth, as a corollary contribution, our unique main (Chilean) dataset allows us to verify the methodology for computing marginal costs based on markups (De Loecker et al., 2016): we show that changes in computed plant-product level marginal costs are very similar to those in self-reported average variable costs. Finally, by confirming that our results hold for two additional countries (Colombia and Mexico), we provide strong support for their general validity.

The rest of the paper is organized as follows. Section 2 discusses our use of marginal cost as a measure of efficiency and its relationship to revenue productivity; it also illustrates the empirical framework to identify the two measures. Section 3 describes our datasets. Section 4 presents our empirical results for Chilean export entrants and Section 5, for continuing exporters. Section 6 provides evidence for Colombian and Mexican export entrants. Finally, Section 7 discusses our results and draws conclusions.

efficiency gains.

¹³De Loecker et al. (2016) document a fall in the marginal cost of Indian firms following a decline in *input* tariffs.

¹⁴Case studies typically suggest strong export-related efficiency gains within plants. For example, Rhee, Ross-Larson, and Pursell (1984) surveyed 112 Korean exporters, out of which 40% reported to have learned from buyers in the form of personal interactions, knowledge transfer, or product specifications and quality control. The importance of knowledge transfer from foreign buyers to exporters is also highlighted by the World Bank (1993) and Evenson and Westphal (1995). López (2005) summarizes further case study evidence that points to learning-by-exporting via foreign assistance on product design, factory layout, assembly machinery, etc.

2 Empirical Framework

In this section, we discuss our efficiency measures and explain how we compute them. Our first measure of efficiency is *revenue-based* total factor productivity (TFPR) – the standard efficiency measure in the literature that analyzes productivity gains from exporting. We discuss why this measure may fail to detect such gains, and show how we calculate TFPR at the plant-product level. Our second measure of efficiency is the marginal cost of production, which can be derived at the plant-product level under a set of non-restrictive assumptions. We also discuss the relationship between the two measures, and under which conditions marginal costs reflect physical efficiency.

2.1 Revenue vs. Physical Total Factor Productivity

Revenue-based total factor productivity is the most widely used measure of efficiency. It is calculated as the residual between total revenues and the estimated contribution of production factors (labor, capital, and material inputs).¹⁵ TFPR has important shortcomings, which we illustrate by considering a standard Hicks-neutral production function for physical output (Q_{it}) of a given plant i in period t . Output is produced using a vector of inputs \mathbf{X}_{it} . We use a log-linear representation of the production function, with lowercase letters denoting the logarithms of the variables, and we adopt the notation from De Loecker and Goldberg (2014): $q_{it} = \mathbf{x}'_{it}\alpha + \omega_{it}$, where α is a vector of output elasticities and ω_{it} is physical efficiency (TFPQ). In most empirical studies, output quantities are unobserved, so that researchers rely on plants' revenues (R_{it}). In this case, (log) revenues are given by:

$$r_{it} = \mathbf{x}'_{it}\alpha + \underbrace{(\omega_{it} + p_{it})}_{\pi_{it}} \quad (1)$$

where p_{it} is the (log) price of output and π_{it} is revenue productivity. Equation (1) highlights an important shortcoming of revenue productivity. When revenues are used as output variable, the residual term π_{it} reflects both output prices and physical efficiency: $\pi_{it} = p_{it} + \omega_{it}$.¹⁶ Thus,

¹⁵Some authors have used labor productivity – i.e., revenues per worker – as a proxy for efficiency. This measure is affected by the use of non-labor inputs and is thus inferior to TFPR (c.f. Syverson, 2011).

¹⁶For illustration, we assume in subsections 2.1 and 2.2 that the production function coefficients α are known. In practice, the estimation of α is subject to input and output price biases (De Loecker and Goldberg, 2014). We discuss how we address these biases in our estimation below; De Loecker and Goldberg (2014) present conditions under which the two biases cancel each other when output revenues and input revenues are used in estimating the production function. In addition, note that even when α is known, π_{it} is affected by (unobserved) input prices if input *values* are used in (1) to proxy for x_{it} . To see this, let \mathbf{z}_{it} denote input prices. Then $\pi_{it} = \omega_{it} + p_{it} - \alpha\mathbf{z}_{it}$, which corresponds to the profitability residual in the relationship between sales and input expenditures (see De Loecker and Goldberg, 2014, for detail). In the interest of parsimony, we abstract from this issue here. Below, we explain why differences in input prices are unlikely to affect our findings: all our results control for plant-product fixed effects (which absorb differences in input prices across plants); in addition, we find that input prices are also constant within plants around

if output prices respond to a producer’s efficiency, TFPR is biased. For example, when facing downward-sloping demand, firms typically respond to efficiency gains by expanding production and reducing prices. This generates a negative correlation between p_{it} and ω_{it} , so that TFPR will underestimate physical efficiency. Empirical studies attempt to address this bias by deflating revenues with industry price indexes when computing TFPR. However, the downward bias of TFPR persists *within* industries, reflecting the difference between individual plants’ prices and the corresponding industry price index.¹⁷

Despite the shortcomings of TFPR, the majority of studies have used this measure to analyze productivity gains from exporting. One practical reason is the lack of information on physical quantities.¹⁸ While some corrections to the estimation of production functions have been proposed, only a few studies have derived ω_{it} directly.¹⁹ To circumvent some of the issues related to computing ω_{it} , we propose marginal costs as our main measure of efficiency. Next, we discuss under which conditions declining marginal costs reflect efficiency gains.

2.2 Marginal Cost as a Measure of Efficiency, and its Relationship to TFPR

In standard production functions, marginal costs are inversely related to physical efficiency ω_{it} .²⁰ To illustrate this relationship, we use the generic functional form $mc(\omega_{it}, \mathbf{z}_{it})$, where \mathbf{z}_{it} is an input price vector. The derivatives with respect to the two arguments are $mc_1 < 0$ and $mc_2 > 0$. Next, we can use the fact that prices are the product of markups (μ_{it}) and marginal costs to disentangle

the period of export entry.

¹⁷It is important to note that TFPR is not always inferior to TFPQ (or marginal costs); instead, the applicability of the different measures depends on the context. For example, when analyzing misallocation as in Hsieh and Klenow (2009), TFPR is the more appropriate measure. In this framework, with downward-sloping iso-elastic demand and CRS technology, high-TFPQ firms charge lower prices that exactly offset their TFPQ advantage, equalizing TFPR. This provides a useful benchmark: in the absence of distortions, TFPR should be the same across plants in an industry, even if their TFPQ differs. At the same time, the Hsieh-Klenow framework also illustrates the shortcomings of TFPR: in the absence of distortions, plants with higher TFPQ are larger and make higher aggregate profits – these differences are not reflected by TFPR.

¹⁸Data on physical quantities have only recently become available for some countries (c.f. De Loecker et al., 2016; Kugler and Verhoogen, 2012, for India and Colombia, respectively).

¹⁹Melitz (2000) and De Loecker (2011) discuss corrections to the estimation of the production function to account for cross-sectional price heterogeneity in the context of a CES demand function. Gorodnichenko (2012) proposes an alternative procedure for estimating the production function that models the cost and revenue functions simultaneously, accounting for unobserved heterogeneity in productivity and factor prices. Hsieh and Klenow (2009) recover ω_{it} using a model of monopolistic competition for India, China, and the United States. Foster et al. (2008) obtain ω_{it} using product-level information on physical quantities from U.S. census data for a subset of manufacturing plants that produce homogeneous products. Eslava et al. (2013) and Lamorgese et al. (2014) compute TFPQ and use it to analyze gains from trade. Finally, Dhyne, Petrin, Smeets, and Warzynski (2016) derive TFPQ to study the effect of Chinese import competition on plant-product efficiency in Belgium.

²⁰For now, we assume that – in addition to the coefficients α being known – all input and output quantities and prices are observed. Also, we focus on the plant level and ignore the underlying product-level variation within plants. We introduce this dimension below when we turn to the actual estimation.

TFPR (assuming Hicks-neutrality – as is standard in the estimation of productivity):

$$\pi_{it} = p_{it} + \omega_{it} = \mu_{it} + mc(\omega_{it}, \mathbf{z}_{it}) + \omega_{it} \quad (2)$$

Deriving log-changes (denoted by Δ) and re-arranging yields a relationship between efficiency gains and changes in TFPR, markups, and marginal costs:

$$\Delta\omega_{it} = \Delta\pi_{it} - \Delta\mu_{it} - \Delta mc(\omega_{it}, \mathbf{z}_{it}) \quad (3)$$

In order to simplify the interpretation of (3) – but not in the actual estimation of $mc(\cdot)$ – we make two assumptions. First, that the underlying production function exhibits constant returns to scale (CRS). This assumption is supported by our data, where the average sum of input shares is very close to one (see Table A.6 in the appendix). This first assumption implies that we can separate $\Delta mc(\omega_{it}, \mathbf{z}_{it}) = \Delta\phi(\mathbf{z}_{it}) - \Delta\omega_{it}$, where $\phi(\cdot)$ is an increasing function of input prices (see the proof in Appendix A.1). Second, we assume that input prices are unaffected by export entry or expansions, i.e., they are constant conditional on controlling for trends and other correlates around the time of export entry: $\Delta\phi(\mathbf{z}_{it}) = 0$.²¹ Our dataset allows us to calculate input prices, and we show below in Section 4.5 that these do not change with exporting activity.

With constant input prices, we obtain three simple expressions that illustrate the relationship between physical efficiency gains and changes in marginal costs, markups, and TFPR:

1. $\Delta\omega_{it} = -\Delta mc_{it}$, i.e., rising efficiency is fully reflected by declining marginal costs. Note that this is independent of the behavior of markups. Using this equality in (3) also implies:
2. $\Delta\pi_{it} = \Delta\mu_{it}$, i.e., revenue productivity rises if and only if markups increase. For example, even if ω_{it} rises (and mc_{it} falls), TFPR will not grow if markups remain unchanged. And vice-versa, if markups rise while ω_{it} stays the same, TFPR will increase. This underlines the shortcomings of TFPR as a measure of efficiency – it can both fail to identify actual efficiency gains but may also reflect spurious gains due to demand-induced increases in markups.
3. $\Delta\pi_{it} = \Delta\omega_{it}$ if $\Delta\mu_{it} = -\Delta mc_{it}$, i.e., changes in revenue productivity reflect the full efficiency gains if markups rise in the same proportion as marginal costs fall, i.e., if the output price remains constant and pass-through of efficiency gains is zero.

We use these insights when interpreting our empirical results below. For young exporters, the evidence points towards constant markups. Thus, all efficiency gains are passed on to customers,

²¹This also implies that the relationships between TFPR, TFPQ, and MC that we derive in this subsection hold if TFPR is defined using physical inputs or based on input values: the difference between the two approaches is $\alpha\Delta\mathbf{z}_{it} = 0$ (see footnote 16, and recall that for now we take α as given).

so that they are reflected only in marginal costs, but not in TFPR. For more mature exporters there is some evidence for declining marginal costs together with rising markups, meaning that at least a part of the efficiency gains is also reflected in TFPR.

2.3 Estimating Revenue Productivity (TFPR)

To compute TFPR, we specify a Cobb-Douglas production function with labor, capital, and materials as production inputs, and deflate all nominal variables using 4-digit industry specific deflators provided by ENIA.²² We opt for the widely used Cobb-Douglas specification as our baseline because it allows us to use the same production function estimates to derive TFPR and markups (and thus marginal costs). This ensures that differences in the efficiency measures are not driven by different parameter estimates.²³ 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.²⁴ 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.

The estimation of the production function follows Akerberg, Caves, and Frazer (2015). This methodology controls for the simultaneity bias that arises because input demand and unobserved productivity are positively correlated.²⁵ Following De Loecker et al. (2016), we (i) allow exporting to affect current productivity either directly, or through a complementarity with investment in physical capital, and (ii) correct for selection bias that occurs because plant switching from single- to multi-product may be correlated with productivity.²⁶ Appendix A.3 provides technical details

²²To keep our baseline estimation comparable to previous studies, we do not deflate material inputs by *plant-specific* deflators from the Chilean ENIA (which are not available in other datasets). This gives rise to a potential (well-documented) input price bias (see De Loecker et al., 2016). Nevertheless, our baseline estimates with a Cobb-Douglas production function are immune to this bias (see footnote 23). In addition, in Appendix A.4 we show alternative results where we proxy for input prices using output prices and market share as suggested by De Loecker et al. (2016), and in Appendix A.5 we use plant-specific input price deflators to deflate input expenditures. Results in both cases are very similar to our baseline results.

²³As discussed below, TFPR needs to be estimated based on output measured in terms of revenues, while deriving markups based on revenues (rather than quantities) can lead to biased results. In our baseline Cobb-Douglas case, this bias does not affect our results because production function coefficients are constant and are therefore absorbed by plant-product fixed effects. Consequently, the Cobb-Douglas specification allows us to use the *same* production function coefficients to estimate both TFPR and markups (and thus marginal costs). In Appendix C.1 we show that the more flexible translog specification (where fixed effects do not absorb the bias) yields very similar results.

²⁴The 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.

²⁵We follow LP in using material inputs to control for the correlation between input levels and unobserved productivity.

²⁶We 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.

on the production function estimation.

Given the estimated output elasticities for each sector s (a vector β^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 input expenditure m_{it} , where lower case letters indicate the natural logarithm of a variable. We then compute plant-level TFPR, $\hat{\pi}_{it}$:

$$\hat{\pi}_{it} = r_{it} - (\beta_l^s l_{it} + \beta_k^s k_{it} + \beta_m^s m_{it}) \quad (4)$$

where r_{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 $\hat{\pi}_{it}$ is not affected by increasing or decreasing returns. When computing *plant*-level TFPR in multi-product plants, we use the vector of coefficients β^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:

$$M_{ijt} = s_{ijt}^{TVC} \cdot M_{it} \quad \text{where} \quad s_{ijt}^{TVC} = \frac{TVC_{ijt}}{\sum_j TVC_{ijt}} \quad (5)$$

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 β^s that corresponds to product j :

$$\hat{\pi}_{ijt} = r_{ijt} - (\beta_l^s l_{ijt} + \beta_k^s k_{ijt} + \beta_m^s m_{ijt}) \quad (6)$$

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

2.4 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 :²⁷

$$\underbrace{\mu_{ijt}}_{\text{Markup}} \equiv \frac{P_{ijt}}{MC_{ijt}} = \underbrace{\left(\frac{\partial Q_{ijt}(\cdot)}{\partial V_{ijt}} \frac{V_{ijt}}{Q_{ijt}} \right)}_{\text{Output Elasticity}} / \underbrace{\left(\frac{P_{ijt}^V \cdot V_{ijt}}{P_{ijt} \cdot Q_{ijt}} \right)}_{\text{Expenditure Share}}, \quad (7)$$

where P (P^V) denotes the price of output Q (input V), and MC is marginal cost. According to equation (7), 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 (7) we use materials (M) as the flexible input to compute the output elasticity.²⁸ Note that in our baseline estimation (due to its use of a Cobb-Douglas production function), the output elasticity with respect to material inputs is given by the constant term β_m^s . Ideally, β_m^s should be estimated using physical quantities for inputs and output, as in De Loecker et al. (2016). However, as discussed above, this would render our results for TFPR and marginal cost less comparable, since differences could emerge due to the different parameter estimates. The Cobb-Douglas case allows us to compute markups based on revenue-based estimates of β_m^s , without introducing bias in our within-plant/product analysis (see Section 2.5 for detail). Thus, our baseline results use the *same* elasticity estimates to compute both TFPR and markups.

The second component needed in (7) – 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 (5) to obtain the value of material inputs $P_{ijt}^V \cdot V_{ijt} = M_{ijt}$. 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 (7).²⁹ This procedure yields plant-product-year spe-

²⁷Note that the derivation of equation (7) 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.

²⁸In 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 firing cost in Chile can be found in Montenegro and Pagés (2004).

²⁹By 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

cific 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. Table A.13 in the appendix shows the average and median markup by sector.

2.5 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 β^s have been estimated correctly. Then, in order to back out TFPQ by using (4), 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.³⁰ 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 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 (7) 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.³¹ 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 (7), still assuming that β^s has been correctly estimated. The output elasticity with respect to material inputs is given by β_m^s , 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 (5)

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.

³⁰Exceptions, where input quantities are available, include Ornaghi (2006), Davis, Grim, and Haltiwanger (2008), and Lamorgese et al. (2014).

³¹A shortcoming of this more aggregate approach is that plant-level output price indexes do not account for differences in product scope (Hottman, Redding, and Weinstein, 2016).

to back out plant-product specific input expenditure shares. Thus, plant-product specific markups can be immediately calculated in our Chilean data.³²

We now turn to the estimation of β^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_m^s from the vector β^s . Note that measurement error of β_m^s will affect the estimated *level* of markups, but not our within-plant results: because we analyze *log-changes* at the plant-product level, $\ln(\beta_m^s)$ cancels out. In other words, the estimated *log-changes* in markups in (7) are only driven by the observed material expenditure shares, but not by the estimated output elasticity β_m^s .³³ Contrast this with the computation of TFPQ, which uses all coefficients in β^s , multiplying each by the corresponding physical input (or deflated input expenditures) in (4). In this case, analyzing *log-changes* in TFPQ will not eliminate errors and biases in the level of β^s .

We discuss further issues related to marginal cost and TFPQ in the appendix. Appendix A.2 discusses the implications of deviations from CRS. We show that in the presence of increasing returns, marginal costs will tend to overestimate actual efficiency gains. In this case, TFPQ is the preferable efficiency measure (subject to the concerns discussed above), since its estimation allows for flexible returns to scale. Throughout the empirical sections, we thus present results based on TFPQ as a robustness check. Appendix A.4 and A.5 discuss the estimation of quantity-based production functions, and Appendix A.6 shows that marginal costs and TFPQ are equally affected by investment in new technology (even if only TFPQ directly takes the capital stock into account).

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). In addition, we confirm our main results using plant-level panel data from Colombia (for the period 2001-2013) and from Mexico (for 1994-2003). A key advantage of the Chilean data is that multi-product plants are required to report product-specific total variable costs. These are crucial for the calculation of plant-product level markups and marginal costs in multi-product plants, as described in Section 2.4. In the Colombian and Mexican samples, this information is not available. In order to keep the methodology consistent, we thus restrict attention to single-product plants in these countries, where all

³²Note 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.

³³This is also the reason why we can use estimates of β^s from the *revenue* production function, i.e., the same coefficients used to compute TFPQ. Note that for the more flexible translog specification, β_m^s itself depends on the use of inputs by each plant and may thus vary over time. We show in Appendix C.1 that our results are nevertheless robust to this specification.

inputs are clearly related to the single output. Correspondingly, the Chilean ENIA is our main dataset, and we describe it in detail below. The Colombian and Mexican datasets are described in Appendix B.4 and B.5, and we compare the three datasets in Appendix B.7. Overall, the sectoral composition of the three datasets is similar, but export orientation is markedly stronger for Mexican manufacturing firms, where almost 40% of all plants are exporters, as compared to 20% and 25% in the Chilean and Colombian samples, respectively.

3.1 Detail on ENIA Data

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.³⁴ 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. Appendix B.1 describes the procedures we use to clean the Chilean ENIA data and to generate a consistent plant-product dataset. After these adjustments, our sample consists of 118,178 plant-product-year observations over the period 1996-2007. In Appendix B.2, we confirm that our data replicate some well-documented systematic differences between exporters and non-exporters.³⁵

³⁴For 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."

³⁵Following Bernard and Jensen (1999), we show that within their respective sectors, exporting plants are larger both in terms of employment and sales, are more productive (measured by revenue productivity), and pay higher wages. This is in line with the exporter characteristics documented by Bernard and Jensen (1999) for the United States, Bernard and Wagner (1997) for Germany, and De Loecker (2007) for Slovenia, among others. Using product-level data, we also find that markups are higher among exporters, confirming the findings in De Loecker and Warzynski (2012). Our Colombian and Mexican data show very similar patterns (see Appendix B.4 and B.5).

3.2 Definition of Export Entry

The time of entry into export markets is crucial for our analysis. We impose four conditions for product j , produced by plant i , to be classified as an export entrant in year t : (i) product j is exported for the first time at t in our sample, which avoids that dynamic efficiency gains from previous export experience drive our results, (ii) product j is sold domestically for at least one period before entry into the export market, i.e., we exclude new products that are exported right away, (iii) product j continues to be reported in ENIA for at least two years after export entry, which ensures that we can compute meaningful trajectories, and (iv) product j is the first product exported by plant i . The last requirement is only needed for multi-product plants. It rules out that spillovers from other, previously exported products affect our estimates. Under this definition we find 861 export entries in our ENIA sample (plant-products at the 7-digit level), and approximately 7% of active exporters are new entrants. For our auxiliary Colombian and Mexican data, the construction of export entry is described in detail in Appendix B.6.

4 Efficiency Gains of Export Entrants in Chilean Manufacturing

In this section we present our empirical results for new export entrants in Chile. We show the trajectories of revenue productivity, marginal costs, and markups within plant-products around the time of export entry. We verify that our results hold when we use propensity score matching to construct a reference group for export entrants, and when we use tariff changes to predict export entries. We also provide suggestive evidence that the observed efficiency gains are driven by a complementarity between exporting and investment.

4.1 New Export Entrants: Plant-Product Trajectories

To analyze trajectories of various plant-product characteristics, we estimate the following regression for each plant i producing product j in period t :

$$y_{ijt} = \alpha_{st} + \alpha_{ij} + \underbrace{\sum_{k=-2}^{-1} T_{ijt}^k}_{\text{Pre-Trend}} + \underbrace{\sum_{l=0}^L E_{ijt}^l}_{\text{Post-Entry Trend}} + \delta_{ijt}^{exit} + \varepsilon_{ijt}, \quad (8)$$

where y_{ijt} refers to TFP, marginal cost, markup, or price; α_{st} are sector-year effects that capture trends at the 4-digit level, and α_{ij} are plant-product fixed effects (at the 7-digit level). We include two sets of plant-product-year specific dummy variables to capture the trajectory of each variable y_{ijt} before and after entry into export markets. First, T_{ijt}^k reflects pre-entry trends in the two periods before exporting. Second, the post-entry trajectory of the dependent variable is reflected by E_{ijt}^l ,

which takes value one if product j is exported l periods after export entry.³⁶ Finally, the dummy δ_{ijt}^{exit} allows for changes in trajectories when plant-products exit the export market.

Table 1 (Panel A) reports the coefficients of estimating (8) for the sub-sample of export entrants (and Figure 1 above visualizes the results). TFPR is virtually unrelated to export entry, with tight confidence intervals around zero. This result is in line with the previous literature: there are no apparent efficiency gains of export entry based on TFPR. The trajectory of marginal costs shows a radically different pattern. After entry into the export market, marginal costs decline markedly. According to the point estimates, marginal costs are about 12% lower at the moment of entry, as compared to pre-exporting periods. This difference widens over time: one period after entry it is 20%, and after 3 years, 26%. These differences are not only economically but also statistically highly significant. In relative terms, the observed decline in marginal costs after export entry corresponds to approximately one-third of the standard deviation in year-to-year changes in marginal costs across all plant-products in the sample. The trajectory for prices is very similar to marginal costs. This results because markups remain essentially unchanged after export entry. The pattern in markups coincides with the one in TFPR, in line with our theoretical results in Section 2. Finally, physical quantities sold of the newly exported product increase by approximately 20%.

Reported Average Variable Costs and TFPQ

One potential concern with respect to our marginal cost results is that they rely on the correct estimation of markups. If we underestimate the true changes in markups after export entry, then the computed marginal cost would follow prices too closely. We can address this concern by using the unique feature that plants covered by ENIA report the variable production cost *per product*, as well as the number of units produced. The questionnaire defines total variable cost per product as the product-specific sum of raw material costs and direct labor involved in production. It explicitly asks to exclude transportation and distribution costs, as well as potential fixed costs. Consequently, dividing the reported total variable cost by the units produced of a given product yields a reasonable proxy for its average variable cost. Figure 4 plots our computed marginal costs against the reported average variable costs (both in logs), controlling for plant-product fixed effects, as well as 4-digit sector-year fixed effects (that is, the figure plots the within plant-product variation that we exploit empirically). The two measures are very strongly correlated. This lends strong support to the markup-based methodology for backing out marginal costs by De Loecker et al. (2016).

Panel B of Table 1 shows that reported average variable costs (AVC) decrease after export entry,

³⁶Due to our relatively short sample, we only report the results for $l = 0, \dots, 3$ periods after export entry. However, all regressions include dummies E_{ijt}^l for all post-entry periods. Also, in order to make trajectories directly comparable across the different outcomes, we normalize all coefficients so that the average across the two pre-entry periods (-1 and -2) equals zero.

Table 1: Within Plant-Product Trajectories for Export Entrants in Chile

Periods After Entry	-2	-1	0	1	2	3	Obs/ R^2
<i>Panel A: Main Outcomes</i>							
TFPR	-.0029 (.0193)	.0029 (.0159)	-.0061 (.017)	.0017 (.0212)	.0264 (.0263)	.0159 (.0269)	3,330 .535
Marginal Cost	.0406 (.0651)	-.0406 (.0498)	-.1207** (.0614)	-.1997*** (.0676)	-.2093*** (.0787)	-.2583*** (.0927)	3,330 .792
Markup	-.012 (.0219)	.012 (.0174)	-.0042 (.0189)	.011 (.0233)	.0359 (.0288)	.0189 (.0311)	3,330 .492
Price	.0286 (.0634)	-.0286 (.0491)	-.1248** (.0582)	-.1887*** (.0665)	-.1735** (.0738)	-.2394*** (.0897)	3,330 .804
Physical Quantities	-.0437 (.0913)	.0437 (.0667)	.1899*** (.0719)	.2672*** (.0905)	.1923* (.1045)	.2098* (.1198)	3,330 .822
<i>Panel B: Additional Efficiency Measures</i>							
Reported AVC	.0297 (.0642)	-.0297 (.0511)	-.1286** (.0600)	-.1838*** (.0672)	-.1904** (.075)	-.2535*** (.0918)	3,330 .795
TFPQ	-.0389 (.0732)	.0389 (.0536)	.118** (.0600)	.1646** (.0683)	.1768** (.0803)	.1937** (.0945)	3,330 .798

Notes: The table reports the coefficient estimates from equation (8). All regressions are run at the plant-product level (with products defined at the 7-digit level); they control for plant-product fixed effects and 4-digit sector-year fixed effects. A plant-product is defined as an export entrant if it is the *first* product exported by a plant and is sold domestically for at least one period before entry into the export market. Section 4.1 provides further detail. For comparability, we normalize all coefficients so that the average across the two pre-entry periods (-1 and -2) equals zero. Standard errors (clustered at the plant-product level) in parentheses. Key: *** significant at 1%; ** 5%; * 10%. TFPR = Revenue productivity; TFPQ = Quantity Productivity; AVC = Average variable cost (self-reported).

closely following the trajectory that we identified for marginal cost. Export entry is followed by a decline in reported AVC by 13% in the period of entry, growing to 18% after one year, and to 25% three periods after entry. These results confirm that the documented efficiency gains after export entry are not an artefact of the estimation procedure for marginal costs.

Another concern is that the decline in marginal (and average) costs may be driven by increasing returns to scale in combination with expanded production after export entry. Our production function estimates suggest that this is unlikely; we find approximately constant returns to scale in most sectors – the mean sum of all input shares is 1.023 (and weighted by plants in each sector, the average is 1.009).³⁷ Nevertheless, we also compute TFPQ as an alternative efficiency measure

³⁷Table A.6 in the appendix reports further details, showing output elasticities and returns to scale for each 2-digit sector in our ENIA sample. Table A.6 also shows that returns to scale are very similar when we instead estimate a more flexible translog specification. The translog case allows for interactions between inputs, so that output elasticities




Figure 4: Estimated Marginal Cost and Reported Average Variable Cost

Notes: The figure plots plant-product level marginal costs computed using the methodology described in Section 2 against plant-product level average costs reported in the Chilean ENIA panel (see Section 3). The underlying data include both exported and domestically sold products, altogether 109,612 observations. The figure shows the relationship between the two cost measures after controlling for plant-product fixed effects (with products defined at the 7-digit level) and 4-digit sector-year fixed effects. The strong correlation thus indicates that *changes* in computed marginal cost at the plant-product level are a good proxy for changes in actual variable costs.

that allows for flexible returns to scale (but is subject to the caveats discussed in Section 2.5).³⁸ The last row of Table 1 shows that the trajectory for TFPQ is very similar to marginal costs. This suggests that our results are not confounded by deviations from CRS.

4.2 Matching Results

Our within-plant trajectories in Table 1 showed a slight (statistically insignificant) decline in prices and marginal costs of new exported products before entry occurs (in $t = -1$). This raises the concern of pre-entry trends, which would affect the interpretation of our results. For example, price and marginal cost could have declined even in the absence of exporting, or export entry could be the result of selection based on pre-existing productivity trajectories. In the following we address this issue by comparing newly exported products with those that had a-priori a similar likelihood of being exported, but that continued to be sold domestically only (De Loecker, 2007).

depend on the use of inputs. Consequently, if input use changes after export entry, this could affect elasticities and thus returns to scale. To address this possibility, we compute the average elasticities for 2-digit sectors using i) all plants, and ii) using only export entrants in the first three periods after entry. Both imply very similar – approximately constant – returns to scale, as shown in columns 5 and 6 in Table A.6. In addition, Table A.15 splits our Chilean sample into sectors with above- and below-median returns to scale and shows that the decline in marginal costs after export entry are actually somewhat stronger in the subset with below-median returns to scale. Thus, it is unlikely that our main results are driven by increasing returns to scale.

³⁸The estimation procedure for TFPQ is described in Appendix A.4.

This empirical approach uses propensity score matching (PSM) in the spirit of Rosenbaum and Rubin (1983), and further developed by Heckman, Ichimura, and Todd (1997). Once a control group has been identified, the average effect of treatment on the treated plant-products (ATT) can be obtained by computing the average differences in outcomes between the two groups.

All our results are derived using the nearest neighbor matching technique. Accordingly, treatment is defined as export entry of a plant-product (at the 7-digit level), and the control group consists of the plant-products with the closest propensity score to each treated observation. We obtain the control group from the pool of plants that produce similar products as new exporters (within 4-digit categories), but for the domestic market only. To estimate the propensity score, we use a flexible specification that is a function of plant and product characteristics, including the level and trends in product-specific costs before export entry, lagged product-level TFPR, the lagged capital stock of the plant, and a vector of other controls in the pre-entry period, including product sales, number of employees (plant level), and import status of the plant.³⁹ Appendix A.8 provides further detail on the methodology. Once we have determined the control group, we use the difference-in-difference (DID) methodology to examine the impact of export entry on product-level TFPR, marginal cost, and markups. As Blundell and Dias (2009) suggest, using DID can improve the quality of matching results because initial differences between treated and control units are removed.

Table 2 shows the matching estimation results. Since all variables are expressed in logarithms, the DID estimator reflects the difference in the *growth* of outcomes between newly exported products and their matched controls, relative to the pre-entry period ($t = -1$).⁴⁰ When compared to the previously reported within-plant-product trajectories, the PSM results show a slightly smaller decline in marginal costs at export entry (6.5% vs. 12.1%) – which is to be expected if the PSM procedure corrects for pre-trends. However, for later periods, decreases in marginal costs are the same as documented above: the difference in marginal cost relative to the control group grows to 11% in the year after entry, to 20% after two years, and to 27% three periods after entry. Our alternative efficiency measures – reported average variable costs and TFPQ – confirm this pattern. Changes in TFPR after export entry are initially small and statistically insignificant. However, after three periods, TFPR increases by about 9% more for export entrant products than for the matched

³⁹Following Abadie, Drukker, Herr, and Imbens (2004), we use the 5 nearest neighbors in our baseline specification. The difference in means of treated vs. controls are statistically insignificant for all matching variables in $t = -1$. We include import status to account for the possibility that input trade liberalization drives export entry as in Bas (2012). As a further check, we also replicated our within-plant trajectories in Table 1, controlling for log imports at the plant level. Results are virtually unchanged (available upon request).

⁴⁰For example, a value of 0.1 in period $t = 2$ means that two years after export entry, the variable in question has grown by 10% more for export entrants, as compared to the non-exporting control group.

control products. This suggests that, eventually, efficiency gains are partially reflected in TFPR – we discuss this pattern in more detail below in Section 4.6.

Table 2: Matching Results: Exported Entry and Efficiency Gains in Chilean Manufacturing

Periods After Entry	0	1	2	3
<i>Panel A. Main Outcomes</i>				
TFPR	-.0164 (.0183)	-.0352 (.0236)	.0152 (.0298)	.0887** (.0396)
Marginal Cost	-.0647* (.0347)	-.110** (.0439)	-.199*** (.0657)	-.269*** (.0882)
Markup	.00379 (.0216)	-.0193 (.0246)	.0415 (.0300)	.0506 (.0401)
Price	-.0609** (.0305)	-.129*** (.0420)	-.158** (.0609)	-.218*** (.0719)
<i>Panel B. Additional Efficiency Measures</i>				
Reported AVC	-.0834** (.0345)	-.157*** (.0437)	-.153** (.0689)	-.263*** (.0777)
TFPQ	.0470 (.0320)	.0956** (.0429)	.151** (.0667)	.339*** (.0946)
Treated Observations	261	179	128	75
Control Observations	1,103	752	534	299

Notes: Period $t = 0$ corresponds to the export entry year. Coefficients reflect the differential growth of each variable with respect to the pre-entry year ($t = -1$) between export entrants and controls, all at the plant-product level. The control group is formed by plant-products that had a-priori a similar likelihood (propensity score) of becoming export entrants, but that continued to be sold domestically only. We use the 5 nearest neighbors. Controls are selected from the pool of plant-products in the same 4-digit category (and same year) as the export entrant product. The specification of the propensity score is explained in Section 4.2 and in Appendix A.8. Robust standard errors in parentheses. Key: *** significant at 1%; ** 5%; * 10%.

TFPR = Revenue productivity; TFPQ = Quantity Productivity; AVC = Average variable cost (self-reported).

4.3 Robustness and Additional Results

In this subsection we check the robustness of our results to alternative specifications and sample selection. Due to space constraints, we present and discuss most tables with robustness checks in Appendix C, and we summarize the main takeaways here.

Balanced Sample of Entrants

To what extent does unsuccessful export entry drive our results? To answer this question, we construct a balanced sample of export entrants, including only plant-products that are consistently

exported for four subsequent years. Table 3 shows the propensity score matching results for this balanced sample. The main pattern is unchanged. TFPR results are quantitatively small and insignificant in the first two years of exporting, but now there is stronger evidence for increases in TFPR in later periods (which coincide with increasing markups). Marginal costs drop markedly after export entry – by approximately 20-30%. The main difference with Table 2 is that marginal costs are now substantially lower already at the time of export entry ($t = 0$). This makes sense, given that we only focus on ex-post successful export entrants, who will tend to experience larger efficiency gains. In addition, in our baseline matching results (Table 2), efficiency continued to increase over time. This may have been driven by less productive products exiting the export market, so that the remaining ones showed larger average differences relative to the control group. In line with this interpretation, the drop in marginal costs is more stable over time in the balanced sample. Our alternative efficiency measures TFPQ and reported AVC show the same pattern (Panel B of Table 3). In sum, the results from the balanced sample confirm our full sample estimates and suggest relatively stable efficiency gains over time.

Table 3: Matching Results for Chile: Balanced Sample

Periods After Entry	0	1	2	3
Panel A. Main Outcomes – Balanced Sample				
TFPR	.0335 (.0299)	.0421 (.0348)	.112*** (.0355)	.109*** (.0380)
Marginal Cost	-.190** (.0839)	-.234** (.0887)	-.308*** (.0933)	-.225** (.0877)
Markup	.0266 (.0369)	.00565 (.0401)	.110*** (.0382)	.0594 (.0414)
Price	-.151* (.0782)	-.210** (.0795)	-.189** (.0870)	-.152** (.0724)
Panel B. Additional Efficiency Measures – Balanced Sample				
Reported AVC	-.227** (.0919)	-.268*** (.0843)	-.242** (.0977)	-.220*** (.0813)
TFPQ	.183** (.0831)	.269*** (.0850)	.348*** (.100)	.318*** (.0911)
Treated Observations	70	71	70	70
Control Observations	275	277	276	278

Notes: The results replicate Table 2 for the sample of plant-products that are observed in each period $t = -2, \dots, 3$ (balanced panel). See the notes to Table 2 for further detail. Robust standard errors in parentheses. Key: *** significant at 1%; ** 5%; * 10%.

Single-Product Plants

In order to estimate product-level TFPR, marginal costs, and markups, we had to assign inputs to individual products in multi-product plants. This is not needed in single-product plants, where all inputs enter in the production of one final good. Table A.14 uses only the subset of single-product plants to estimate the trajectories following equation (8).⁴¹ Despite the fact that the sample is smaller, results for single-product plants remain statistically highly significant and quantitatively even larger than for the full sample. Marginal costs fall by 24-40% after export entry, and this magnitude is confirmed by TFPQ and reported average costs. There is also evidence for increases in TFPR and markups in later periods, but these are quantitatively much smaller than the changes in marginal costs.

Further Robustness Checks

In our baseline matching estimation, we used the 5 nearest neighbors. Table A.17 shows that using either 3 or 10 neighbors instead does not change our results. Next, we investigate to what extent our results change if we deviate from the Cobb-Douglas specification in our baseline productivity estimation. In Table A.18, we present plant-product level estimates based on the more flexible translog production function, which allows for a rich set of interactions between the different inputs. Again, there is no significant change in TFPR after export entry. In Panel B and C of Table A.18 we use the production function coefficients based on the translog specification to compute markups and marginal costs. This has to be interpreted with caution: because the translog production function is estimated based on revenues *and* allows for varying input shares over time, it gives rise to a potential bias in the coefficient estimates (see Appendix A.7 for further discussion). In contrast to the Cobb-Douglas specification, this bias is not constant over time and thus not absorbed by fixed effects in within-plant/product analysis. Nevertheless, the bias is probably of minor importance: we obtain very similar results for markups and marginal costs as in the baseline specification. In the same table, we also demonstrate that our results are the same as in the baseline when we estimate a quantity production function for the Cobb-Douglas case. Finally, Appendix C.4 shows that results are also relatively similar when analyzed at the plant level. Appendix C discusses the additional robustness checks in greater detail.

4.4 Export Entry Predicted by Tariff Changes

In the following, we attempt to isolate the variation in export entry that is driven by trade liberalization. This strategy helps to address endogeneity concerns – in particular, that unobservables may drive both export entry and improvements in efficiency. We follow a rich literature in international

⁴¹For single-product plants, the product index j in y_{ijt} is irrelevant in (8). In line with our methodology for plant-level analyses, we include sector-year fixed effects at the 2-digit level (see footnote 25).

trade, using tariff changes to predict export entry. Before presenting the results, we discuss the limitations of this analysis in the context of our Chilean data.

Limitations of the 2SLS approach

Declines in export tariffs during our sample period (1996-2007) are limited because Chile had already undergone extensive trade liberalization starting in the mid-1970s. Nevertheless, there is some meaningful variation that we can exploit: during the second half of the 1990s, Chile ratified a number of trade agreements with neighboring countries, and between 2003 and 2005, with the United States and the European Union. On average across all destinations, export tariffs for manufacturing products fell from 10.1% in 1996 to 4.5% in 2007 (using total sectoral output in 1996 as constant weights). The European Union and the U.S. were the most important destinations, accounting for 24% and 16% of all exports, respectively, on average over the period 1996-2007. The export tariff decline was staggered over time and thus less dramatic than other countries' rapid trade liberalization (e.g., Slovenian manufacturing export tariffs to the EU fell by 5.7% over a single year in 1996-97). However, we can exploit differential tariff changes across Chilean sectors. These are illustrated in Figure 5 for 2-digit industries. For example, 'clothes and footwear' saw a decline by approximately 10 percentage points, while export tariffs for 'metallic products' fell by as little as 2 p.p. In addition, there is variation in the *timing* of tariff declines across sectors, and the plotted average tariff changes at the 2-digit level in Figure 5 hide underlying variation for more detailed industries. We exploit this variation in the following, using 4-digit ISIC tariff data (the most detailed level that can be matched to our panel dataset).⁴²

This leads to the second limitation of our analysis: as in [Bustos \(2011\)](#), we use industry level tariffs, so that the identifying variation is due to changing export behavior *on average* for plant-products within the corresponding 4-digit tariff categories. The third limitation follows from the staggered pattern of (relatively small) tariff declines over time – as opposed to a short period of rapid trade liberalization. In order to obtain sufficiently strong first stage results, we have to exploit the full variation in tariffs over time. In particular, in most specifications, including year effects – or 2-digit sector-year effects – leaves us with a weak first stage. Consequently, we do not include such fixed effects, so that the full variation in tariffs – across sectors and over time – is exploited. This leads to the possibility that other factors that change over time may drive our results. To

⁴²Chilean tariffs are available at the HS-6 level, but a correspondence to the 7-digit ENIA product code does not exist. The most detailed correspondence that is available matches tariff data to 4-digit ISIC – an industry code that is provided for each ENIA plant. When aggregating export tariffs to the 4-digit level, we use total Chilean exports within each detailed category as weights. For multi-product plants, ENIA assigns the 4-digit ISIC code that corresponds to the plant's principal product. This does not impose an important constraint on our analysis: for the vast majority (85%) of export-entrant multi-product plants in our sample, the principal product (highest revenue) is in the same 4-digit product category as the one that is exported.

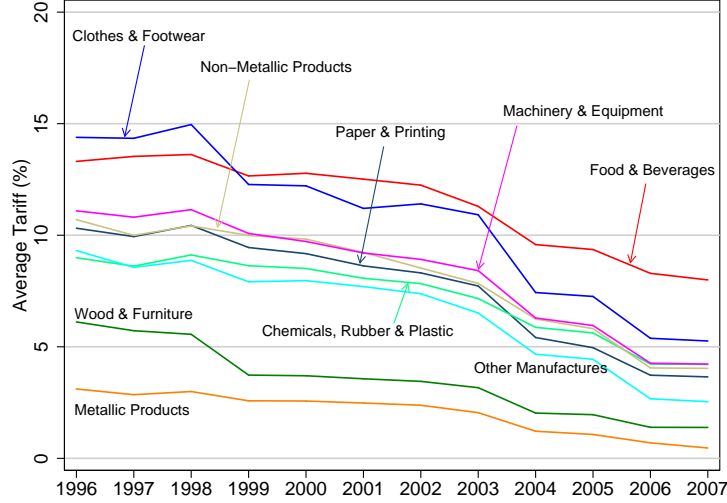


Figure 5: Average Chilean Export Tariffs (2-digit industries)

Notes: The figure plots the average export tariff for all 2-digit ISIC industries. We first compute average tariffs at the 6-digit HS product level across all destinations of Chilean exports, using destination-specific aggregate export shares as weights. We then derive average tariffs at the more aggregate 2-digit ISIC level.

alleviate this concern, we control for total sales of each plant. Thus, our results are unlikely to be driven by sales expansions over time that happen to coincide with trends in tariffs. We perform a number of checks to underline this argument. Nevertheless, in light of the limitations imposed by the data, our 2SLS results should be interpreted as an exploratory analysis.

Empirical setup

We continue to exploit within-plant-product variation, using plant-product fixed effects. In the first stage, we predict export entry based on export tariffs:

$$E_{ijt} = \alpha_{ij} + \beta_1 \tau_{st} + \gamma_1 \ln(\text{sales}_{ijt}) + \varepsilon_{ijt} , \quad (9)$$

where E_{ijt} is a dummy that takes on value one if plant i exports product j in year t , sales_{ijt} are total (domestic and exported) sales, and τ_{st} are export tariffs in sector s (to which product j belongs) in year t , as described in footnote 42. Correspondingly, all standard errors are clustered at the 4-digit sector level s . Because we use plant-product fixed effects α_{ij} , neither established (continuing) exporters nor plant-products that are never exported affect our results. We thus restrict the sample to export entrants as defined in Section 3.2. Note that our analysis is run in levels rather than changes. This allows for tariff declines in different years to affect export behavior – as we

discussed above, Chile’s trade liberalization over our sample period was a staggered process, so that we cannot explore before-after variation over a short time window as in [Bustos \(2011\)](#). In addition, running the analysis in levels with fixed effects (rather than, say, annual changes) allows for flexibility in the timing with which tariff declines affect exporting. For example, if the reaction to lower tariffs gains momentum over time (as in the Canadian case documented by [Lileeva and Trefler, 2010](#)), annual changes would not properly exploit this variation. Finally, we use OLS to estimate (9); probit estimates would be inconsistent due to the presence of fixed effects.

Column 1 in Table 4 presents our first-stage results for export entrant products, showing that declining export tariffs are strongly associated with export entry. The first stage F-statistic is well above the critical value of 16.4 for 10% maximal IV bias. As discussed above, we only exploit the extent to which tariffs predict the *timing* of export entry, by including plant-product fixed effects and restricting the sample to those plant-products that become export entrants at some point over the period 1996-2007. The highly significant coefficient on export tariffs thus implies that export entry is particularly likely in 4-digit sectors (and years) where export tariffs decline more steeply. In other words, plant-products that eventually become exporters are particularly likely to do so when they face lower export tariffs. The magnitude of the first-stage coefficient (-8.403) implies that an extra one-percentage-point decrease in export tariffs (both over time and across 4-digit sectors) is associated with an increase in the probability of exporting by 8.4% among those plant-products that become exporters at some point. Our methodology tackles the endogeneity of export entry in two ways: First, we address the possibility that plant-products that ‘react’ to lower tariffs by export entry differ systematically from those that never start exporting – by restricting the sample to the former. Second, by exploiting only the variation in exporting that is predicted by tariffs, we address the possibility that the timing of export entry may be driven by unobserved productivity trends.

Next, we proceed with the second stage, where we regress several characteristics y_{it} that include marginal costs, markups, and TFPR on predicted export entry \hat{E}_{ijt} :

$$\ln(y_{ijt}) = \alpha_{ij} + \beta_2 \hat{E}_{ijt} + \gamma_2 \ln(\text{sales}_{ijt}) + \vartheta_{ijt} . \quad (10)$$

Columns 2-5 in Table 4 report the second-stage results for our main outcome variables. Marginal costs drop by 27.7% after tariff-induced export entry, and this effect is statistically significant with a p-value of 0.03 (we report weak-IV robust Anderson-Rubin p-values in square brackets, based on [Andrews and Stock, 2005](#)). This estimate is remarkably similar to those presented above in Tables 1-3. On the other hand, neither markups nor TFPR change upon (predicted) export entry, while output prices drop similar to marginal costs. This also confirms our results for within-plant

Table 4: Tariff-Induced Export Entry in Chile. Plant-Product Level Analysis

	First Stage	Second Stage					
	(1)	(2)	Main Outcomes			Additional Outcomes	
Dependent Variable	Export Dummy	TFPR	MC	Markup	Price	Reported AVC	TFPQ
Export Tariff	-8.403*** (1.151)	—	—	—	—	—	—
First Stage F-Statistic	53.09						
Export Dummy	—	.0291 [.608]	-.277** [.0338]	.0268 [.702]	-.255* [.0541]	-.312** [.0228]	.259* [.0525]
Plant-Product FE	✓	✓	✓	✓	✓	✓	✓
log Sales	✓	✓	✓	✓	✓	✓	✓
Observations	2,081	2,081	2,081	2,081	2,081	2,081	2,081

Notes: This table examines the effect of tariff-induced export entry on our main outcome variables, as well as on reported average variable costs (AVC) and TFPQ. We report plant-product results, including only plant-products that become new export entrants (see definition in Section 3.2) at some point over the sample period. Export tariffs (at the 4-digit ISIC level) are used to instrument for the timing of export entry. The first stage results of the 2SLS regressions are reported in col 1, together with the (cluster-robust) Kleibergen-Paap rK Wald F-statistic. The corresponding Stock-Yogo value for 10% maximal IV bias is 16.4. Second stage results (cols 2-7) report weak-IV robust Anderson-Rubin p-values in square brackets (see Andrews and Stock, 2005, for a detailed review). All regressions control for the logarithm of plant sales and include plant-product fixed effects. Standard errors are clustered at the 4-digit ISIC level, corresponding to variation in tariffs. Key: *** significant at 1%; ** 5%; * 10%. TFPR = Revenue productivity; TFPQ = Quantity Productivity; AVC = Average variable cost (self-reported).

trajectories. Our alternative efficiency measures in columns 6 and 7 – reported AVC and TFPQ – also show changes that are quantitatively very similar to those based on marginal costs.

In the appendix, we present a number of additional checks. Table A.19 shows that the reduced-form results of regressing export entry directly on tariffs show the same pattern as the 2SLS estimates. We also show that there is no relationship between export tariffs and *domestic* sales at the plant level (Table A.20). This makes it unlikely that our results are driven mechanically by falling tariffs that coincide with expanding sales over time. In sum, despite the limited variation in tariffs, there is compelling evidence for within-plant efficiency gains after tariff-induced export entry, and for our argument that these gains are not captured by TFPR.

4.5 Interpretation of Export Entry Results and Possible Channels

In the following, we discuss possible channels that may drive the observed trajectories of prices and marginal costs for export entrants. We differentiate between demand- and supply-side explanations. Among the latter, export entry can be driven by selection on pre-exporting efficiency (as

in Melitz, 2003), or by a complementarity between exporting and investment in new technology (c.f. Constantini and Melitz, 2007; Atkeson and Burstein, 2010; Lileeva and Trefler, 2010; Bustos, 2011). In addition, anticipated learning-by-exporting also provides incentives for export entry. We discuss the extent to which each of these explanations is compatible with the patterns in the data.

Demand-driven export entry

If demand shocks – rather than changes in production – were responsible for our results, we should see no change in the product-specific marginal costs, while sales would increase and markups would tend to rise. This is not in line with our empirical observation of falling marginal costs and constant markups. Thus, demand shocks are an unlikely driver of the observed pattern.

Selection on pre-exporting productivity

Firms that are already more productive to start with may enter international markets because of their competitive edge. Consequently, causality could run from initial productivity to export entry, reflecting self-selection. In this case, the data should show efficiency advantages already before export entry occurs. Since we analyze within-plant-product trajectories, such pre-exporting efficiency advantages should either be captured by plant-product fixed effects, or they would show up as declining marginal costs *before* export entry. There is only a quantitatively small decline in marginal costs in our within-plant/product trajectories, and a much stronger drop in the year of export entry (see Figure 1). In addition, our matching estimation is designed to absorb pre-entry productivity differences, and our 2SLS results for tariff-induced export entry are unlikely to be affected by selection. In sum, while we cannot fully exclude the possibility of selection into exporting, it is unlikely to be a major driver of our results.

Learning-by-exporting

Learning-by-exporting (LBE) refers to exporters gaining expertise due to their activity in international markets. LBE is typically characterized as an ongoing process, rather than a one-time event after export entry. Empirically, this would result in continuing efficiency growth after export entry. There is some limited evidence for this effect in our data: Tables 1 and 2 show a downward trend in marginal costs during the first three years after export entry. However, this may be driven by the differential survival of more successful exporters. In fact, the trend in marginal costs is less pronounced in the balanced sample in Table 3. Thus, learning-by-exporting can at best explain parts of our results.

Complementarity between Technology and Exporting

Finally, we analyze the case where exporting goes hand-in-hand with investment in new technology. As pointed out by Lileeva and Trefler (2010), expanded production due to export entry may

render investments in new technology profitable. In this case, a plant will enter the foreign market if the additional profits (due to both a larger market and lower cost of production) outweigh the combined costs of export entry and investment in new technology. This setup implies an asymmetry in efficiency gains across initially more vs. less productive plants (or plant-products in our setting). Intuitively, productive plants are already close to the efficiency threshold required to compete in international markets, while unproductive plants need to see major efficiency increases to render exporting profitable. Thus, we should expect "negative selection" based on initial productivity – plant-products that are initially less productive should experience larger changes in efficiency. This prediction can be tested in the data.

Table 5 provides evidence for this effect, reporting the change in marginal costs for plant-products with low and high pre-exporting productivity.⁴³ We find a steeper decline in marginal costs for plant-products with low pre-exporting productivity, and the difference is particularly pronounced for ‘young’ exporters in the first two years after export entry. This result is in line with a complementarity channel where exporting and investment in technology go hand-in-hand, and where initially less productive plants will only make this joint decision if the efficiency gains are substantial (Lileeva and Trefler, 2010).

Table 5: Marginal Cost by Initial Productivity of Export Entrants in Chile. Matching Results.

Periods After Entry	0	1	2	3
Low Initial Productivity	-.167*** (.0520)	-.193*** (.0649)	-.148* (.0817)	-.276** (.113)
High Initial Productivity	.0335 (.0449)	-.0331 (.0587)	-.247** (.102)	-.262* (.134)
<i>p-value for difference</i>	[.004]	[.07]	[.45]	[.94]
Treated Observations	261	179	128	75
Control Observations	1,103	752	534	299

Notes: The table analyzes heterogeneous effects of export entry on marginal costs at the plant-product level, depending on the product-specific initial productivity. Coefficients are estimated using propensity score matching; see the notes to Table 2 for further detail. We use pre-exporting TFPR to create an indicator for plant-products with above- vs. below-median productivity and then estimate the average treatment of the treated (ATT) effect separately for the two subsets. Period $t = 0$ corresponds to the export entry year. Robust standard errors in parentheses. Key: *** significant at 1%; ** 5%; * 10%. The p-value refers to the null hypothesis of equal coefficients for low and high initial productivity.

⁴³Because marginal costs cannot be compared *across* plant-products, we use pre-exporting TFPR to split them into above- and below median productivity. Also, pre-exporting TFPR can only be computed when the export entry date is known with certainty. Thus, we cannot apply our 2SLS methodology where tariff changes predict the *probability* of export entry. Consequently, we use propensity score matching, applied to the subsamples of plant-products with high and low pre-exporting TFPR.

The complementarity channel is also supported by detailed data on plant investment. ENIA reports annual plant-level investment in several categories, allowing us to analyze the corresponding trends for export entrants. Because investment is lumpy, we examine the trend in the following intervals: the last two years before export entry ("pre-entry"), the entry year and the first two years thereafter ("young exporters"), and three or more years after entry ("old" exporters). In Panel A of Table 6 we present the results. Coefficients are to be interpreted as within-plant changes relative to the industry level (since we control for plant fixed effects and 2-digit sector-year effects). Overall, investment shows a marked upward trend right after export entry. Disentangling this aggregate trend reveals that it is mainly driven by investment in machinery and – to some degree – by investment in vehicles. Investment in structures, on the other hand, is unrelated to export entry. We also confirm this pattern in our auxiliary Colombian and Mexican data, where investment spikes after export entry exclusively for machinery, but not for vehicles or structures (see Table A.31 and Table A.32 in the appendix). The observed time trend in investment is in line with the findings in Bustos (2011).⁴⁴ Overall, our investment data suggest that the observed efficiency gains are driven by a complementarity between investment in new productive technology and export entry.

Alternative Interpretations: Input Prices, and Product Quality

Could marginal costs fall after export entry simply because exporters purchase inputs at discounted prices? Panel B in Table 6 examines this possibility, reporting trends in the average price of all inputs, as well as for a stable basket of inputs (those that are continuously used for at least two periods before and after export entry). The table shows that input prices remain relatively stable after export entry, making it unlikely that this channel confounds our results. It is also unlikely that quality upgrading of exporters is responsible for our results, since higher product quality is associated with *higher* output prices and production costs (c.f. Kugler and Verhoogen, 2012; Manova and Zhang, 2012; Atkin et al., 2014; Fan, Li, and Yeaple, 2015). This is not compatible with the observed decline in output prices, marginal costs, and the relatively stable input prices in our data. In addition, the results from a structural model by Hottman et al. (2016) suggest that quality differences are predominantly associated with TFPR differences, rather than differential costs.

On balance, our findings point to exporting-technology complementarity as an important driver of efficiency gains among export entrants. Importantly, the main contribution of our findings is independent of which exact channels drive the results: we show that there are substantial efficiency gains associated with entering the export market, and that the standard TFPR measure does not

⁴⁴It is possible that the installation of new equipment began before export entry, but was reported only after its completion. For example, the ENIA investment category allows for "assets measured in terms of their (historical) accounting cost of acquisition."

Table 6: Investment and Input Price Trends Before and After Export Entry

Period:	Pre-entry	‘Young’ Exp.	‘Old’ Exp.	Obs./R ²
<i>Panel A. Investment</i>				
Overall	0.169 (0.269)	0.635** (0.271)	0.337 (0.290)	2,761 0.519
Machinery	0.258 (0.264)	0.737*** -0.277	0.447 (0.294)	2,761 0.521
Vehicles	0.469** (0.232)	0.607** (0.253)	0.267 (0.236)	2,761 0.324
Structures	0.240 (0.249)	-0.147 (0.274)	0.0758 (0.269)	2,761 0.486
<i>Panel B. Input Prices</i>				
All inputs	-0.0361 (0.155)	-0.0563 (0.163)	-0.0460 (0.195)	7,120 0.368
Stable inputs	-0.0888 (0.152)	0.0284 (0.142)	-0.0946 (0.252)	2,375 0.339

Notes: This table analyzes investment and input prices before and after export entry. All dependent variables are in logs, and all regressions include fixed effects; thus, coefficients reflect the percentage change in investment (panel A) or input prices (panel B) in each respective year relative to the average across all years. ‘Old Exp.’ groups all periods beyond 2 years after export entry; ‘Young Exp.’ comprises export periods within 2 years or less after export entry; and ‘Pre-Entry’ groups the two periods before entry. Regressions in panel A are run at the plant level and control for plant sales, plant fixed effects, and sector-year effects (at the 2-digit level). Regressions in Panel B are run at the 7-digit input-plant level and control for plant-input fixed effects and 4-digit input sector-year effects. In the first row of Panel B (‘All inputs’), we use all inputs observed in the export entry year; in the second row (‘Stable inputs’), we restrict the sample to the set of inputs that are also used at least two periods before and after export entry. The criteria for defining a plant as entrant are described in the notes to Table 1. Robust standard errors in parentheses. Key: *** significant at 1%; ** 5%; * 10%.

capture these gains because of relatively stable markups during the first years after entry.

4.6 Stable Markups after Export Entry – A Result of ‘Foreign Demand Accumulation’?

We observe that, on average, prices of plant-products fall hand-in-hand with marginal costs after export entry. Understanding why prices fall is important for the interpretation of our results; if they did not change, TFPR would reflect all efficiency gains, eliminating the need for alternative measures. We observed that export entrants charge relatively constant markups (at least in the periods immediately following export entry), so that efficiency gains are passed through to customers. One explanation is that new exporters engage in ‘demand accumulation,’ as described by Foster et al. (2016) – charging lower prices abroad in an attempt to attract customers where ‘demand capital’

is still low. If this is the case, we should expect a stronger decline in export prices as compared to their domestic counterparts, because export entrants are already established domestically, but still unknown to international customers. In the following, we provide supportive evidence for this assertion.

We can disentangle domestic and foreign prices of the same product in a subsample for 1996–2000. For this period, the ENIA questionnaire asked about separate quantities and revenues for domestic and international sales of each product. Thus, prices (unit values) can be computed separately for exports and domestic sales of a given product. Within this subsample, we define ‘young’ export entrants as plant-products within 2 years after export entry and compare their average domestic and foreign prices. We find that within plant-products of ‘young’ exporters, the price of exported goods is about 22% lower than pre-export entry, while the price of the same good sold domestically falls by 8%.⁴⁵ Assuming that the marginal cost of production is the same for both markets, the results provide some evidence that efficiency gains are passed on to both domestic and foreign customers – but significantly more so to the latter. While we cannot pin down the exact mechanism that explains the observed price setting, our observations are in line with ‘demand accumulation’ in foreign markets.

5 Export Expansions of Existing Exporters

We have shown that marginal costs drop substantially after export *entry*, while markups and TFPR remain roughly unchanged. We have interpreted this as evidence for quantitatively important efficiency gains within plants that are not captured by standard productivity measures. Does the same pattern hold for existing exporters – that is, do increases in export *volume* have the same effect as export entry itself? In the following, we examine this question, exploiting export tariff changes.

5.1 Empirical Setup with Existing Exporters

When analyzing existing exporters, we have to switch from the plant-product to the plant level. The reason is that export *sales* – a crucial variable in this analysis – are reported only at the plant level by ENIA (while export status is reported for each product as a dichotomous variable). Before proceeding, we first check whether our previous findings also hold at the plant level. These results are presented in Appendix C.4.⁴⁶ Table A.23 presents within-plant trends after export entry,

⁴⁵To obtain these estimates, we separately regress logged domestic and export prices (at the 7-digit plant-product level) on an exporter dummy, controlling for plant-product fixed effects and 4-digit sector-year effects. Table A.21 in the appendix shows the results. In addition, Table A.22, estimates the effect of export entry on domestic and foreign profit margins after export entry (which is discussed in detail in Appendix C.3).

⁴⁶For multi-product plants, TFPR at the plant level can be calculated with the procedure described in Section 2.3, but aggregating markups and marginal costs to the plant level is less straightforward. We employ the following method, which is explained in more detail in Appendix B.3. First, because our analysis includes plant fixed effects, we can

showing that TFPR increases only slightly, while marginal costs decline substantially. The fact that plant-level results are similar to those at the plant-product level is not surprising, given that the exported product typically accounts for the majority of output in exporting multi-product plants. We run the following regression at the plant (i) level:

$$\ln(y_{it}) = \beta \ln(\widehat{exports}_{it}) + \gamma \ln(domsales_{it}) + \delta_i + \varepsilon_{it} , \quad (11)$$

where y_{it} denotes our standard outcome variables: marginal costs, markups, and TFPR. We use export tariffs to predict plant-level export sales $\ln(\widehat{exports}_{it})$; more precisely, since we include plant fixed effects δ_i , we implicitly use *changes* in tariffs to predict *changes* in exports. As discussed in Section 4.4, we exploit the variation in tariffs over time and across 4-digit sectors – the same limitations as discussed above apply here, too. Next, $domsales_{it}$ denotes total domestic sales. Controlling for $domsales_{it}$ ensures that our results are not driven by plant size and are instead attributable to expansions of exports *relative* to domestic sales.

Throughout our analysis of existing exporters, we report results for different subsamples of plants, according to their overall export share. We begin with the full sample that includes all exporters (i.e., all those with export shares above zero) and then move to plants with at least 10%, 20%, ..., 50% export share. This reflects the following tradeoff: On the one hand, plants that export a larger fraction of their output will react more elastically to changes in trade costs than plants that export little. Thus, estimated effects will tend to increase as we raise the export share cutoff. On the other hand, for plants that already have a high export share there is a smaller margin to increase exports relative to total sales. This will attenuate the effect of falling tariffs. In combination, the two opposing forces should lead first to stronger and then to weaker effects as we increase the export share cutoff. Indeed, we find that results are typically strongest for plants with 20-40% export shares.

5.2 Tariff Changes and Within-Plant Efficiency Gains: 2SLS Results

We obtain a strong first stage when estimating (11) – the first stage F-statistics typically exceed the critical value for a maximal 10% IV bias (detailed first stage results are shown in Appendix Table A.25). In terms of magnitude, tariff declines over our sample period predict increases in export sales by approximately 20-30% among existing exporters (on average across the different specifications). Table 7 presents the second stage of our 2SLS results. These show that tariff-induced ex-

normalize plant-level marginal costs and markups to unity in the last year of our sample, 2007 (or the last year in which the plant is observed). We then compute the annual percentage change in marginal cost at the plant-*product* level. Finally, we compute the average *plant*-level change, using product revenue shares as weights, and extrapolate the normalized plant-level marginal costs. For markups, we use the same product revenue shares to compute a weighted average plant-level markup.

port expansions led to statistically significant efficiency increases, as measured by falling marginal costs (panel A) and rising TFPQ (panel B). To interpret the magnitude of effects, we compute the change in each outcome due to the overall tariff reduction over the sample period (denoted by $\hat{\Delta}$). For example, in col 3, panel A, the effect size of -0.218 is obtained by multiplying the coefficient estimate (-0.845) with the corresponding predicted increase $\hat{\Delta}$ in exports for 1996-2007 from the first-stage regressions in Appendix Table A.25 (0.258). We find that export tariff declines are associated with marginal costs falling by approximately 25% over the sample period; the TFPQ results confirm this magnitude. This is similar to the observed efficiency gains after export entry (15-25% as reported in Table 4). If taken at face value, our results thus suggest that export entry has (on average) a similar effect on productivity as a tariff-induced increase in export volume by 20-30% among existing exporters.

Next, we turn to the results for markups and TFPR (panel C and D in Table 7, respectively). Both variables increase statistically significantly with tariff-induced export expansions among firms that export more than 10% of their output (cols 2-6). Nevertheless, TFPR captures only about one quarter of the efficiency gains reflected by marginal costs and TFPQ: tariff declines over our sample period raised TFPR by approximately 5%. The increase in markups is very similar, in line with our result in Section 2. Our results for tariff-induced export expansions thus also imply that about three-quarters of the efficiency gains reflected by lower marginal costs are passed on to customers in the form of lower prices.

In Appendix C.2 we present a number of consistency checks. Table A.26 shows the reduced-form results corresponding to Table 7. We confirm the 2SLS results: lower tariffs lead to significant declines in marginal costs, and to significant (but relatively smaller) increases in markups and TFPR. Next, Table A.27 shows that falling export tariffs are *not* associated with changes in domestic sales. This suggests that we identify a pattern that is specific to trade, and not driven by a general expansion of production. In Table A.28 we show that input prices are largely unchanged following tariff-induced export expansions. Finally, Table A.29 shows that tariff-induced export expansions are also associated with increases in capital stock. This is compatible with our interpretation that investment in new technology is responsible for the observed efficiency increases.

The fact that for existing exporters some of the increased efficiency is captured by TFPR marks an important difference to the results on export entry, where markups and TFPR remained largely unchanged. The core of the difference is related to pricing behavior: while new export entrants pass efficiency gains on to their international customers, established exporters raise markups. Related to our discussion in Section 4.6, existing exporters may face relatively less elastic demand because they already have an established customer base. This may explain why efficiency increases translate – at least partially – into higher markups for established exporters. This interpretation is also

Table 7: Tariff-Induced Export Expansions of Exporting Plants in Chile – 2SLS

	(1)	(2)	(3)	(4)	(5)	(6)
Export Share	>0%	>10%	>20%	>30%	>40%	>50%
<i>Panel A. log Marginal Cost Index</i>						
log Exports (predicted)	-.692**	-.55**	-.845***	-.919***	-.879***	-.822***
<i>weak-IV robust p-value:</i>	[.0215]	[.0183]	[.001]	[.0011]	[.0017]	[.0078]
$\hat{\Delta}$ MC [‡]	-.119	-.130	-.218	-.242	-.245	-.244
First Stage F-Statistic	8.92	24.27	21.59	20.56	19.46	11.91
Observations	6,996	4,089	3,257	2,815	2,443	2,137
<i>Panel B. log TFPQ</i>						
log Exports (predicted)	.734**	.52**	.759***	.728***	.677**	.627**
<i>weak-IV robust p-value:</i>	[.0126]	[.0382]	[.0057]	[.0089]	[.0102]	[.0301]
$\hat{\Delta}$ TFPQ [‡]	.124	.122	.196	.192	.189	.186
First Stage F-Statistic	8.746	24.12	21.58	20.55	19.43	11.91
Observations	6,988	4,083	3,256	2,814	2,442	2,137
<i>Panel C. log Average Markup</i>						
log Exports (predicted)	.0235	.22***	.227***	.262***	.223***	.145***
<i>weak-IV robust p-value:</i>	[.78]	[.0081]	[.0004]	[.0001]	[.0001]	[.0004]
$\hat{\Delta}$ Markup [‡]	.003	.042	.047	.057	.052	.036
First Stage F-Statistic	10.44	25.19	24.55	22.34	20.31	12.87
Observations	9,855	5,744	4,570	3,974	3,454	3,015
<i>Panel D. log TFPR</i>						
log Exports (predicted)	.0461	.182**	.172**	.195***	.163**	.11
<i>weak-IV robust p-value</i>	[.469]	[.0114]	[.0134]	[.0053]	[.0115]	[.195]
$\hat{\Delta}$ TFPR [‡]	.009	.043	.044	.053	.047	.034
First Stage F-Statistic	10.44	25.19	24.55	22.34	20.31	12.87
Observations	9,855	5,744	4,570	3,974	3,454	3,015
For all regressions:						
Plant FE	✓	✓	✓	✓	✓	✓
log Domestic Sales	✓	✓	✓	✓	✓	✓

Notes: This table examines the effect of within-plant export expansions due to falling export tariffs on plant-level marginal costs (panel A), TFPQ (panel B), markups (panel C), and TFPR (panel D). The regressions in columns 1-6 are run for different samples, according to the plants' export shares: col 1 includes all plants with positive exports, col 2 those whose exports account for more than 10% of total sales, col 3, 20%, and so on. The first stage regresses plant-level log exports on sector-specific export tariffs. Export tariffs vary at the 4-digit ISIC level. The first stage regression results are reported in Table A.25 in the appendix. Each panel above reports the second-stage coefficients for the respective outcome variable, together with the weak-IV robust Anderson-Rubin p-values in square brackets (see Andrews and Stock, 2005, for a detailed review). We also report the (cluster-robust) Kleibergen-Paap rK Wald F-statistic for the first stage. The corresponding Stock-Yogo value for 10% maximal IV bias is 16.4. For multi-product plants, the dependent variables in panels A, B, and C reflect the product-sales-weighted average, as described in Appendix B.3. All regressions control for the logarithm of plant-level domestic sales and include plant fixed effects. Standard errors are clustered at the 4-digit ISIC level, corresponding to the level at which tariffs are observed. Key: *** significant at 1%; ** 5%; * 10%.

[‡] In each panel of the table, $\hat{\Delta}$ denotes the predicted change in the corresponding dependent variable due to export tariff reductions over the sample period (tariffs declined by 5.6 p.p. on average (sales-weighted) in 1996-2007).

in line with models such as Melitz and Ottaviano (2008), where lower tariffs have an effect akin to a demand shock for existing exporters, inducing them to raise markups.⁴⁷

6 Evidence from other Countries: Colombia and Mexico

In this section, we repeat our main empirical analysis for two additional countries: Colombia (2001-13) and Mexico (1994-2003). Both provide datasets with similarly detailed coverage as the Chilean ENIA, and these datasets have been used extensively in studies of international trade.⁴⁸ Appendices B.4 and B.5, respectively, describe the Colombian and Mexican data in detail and show that the standard stylized facts documented above for Chile hold in these samples, as well. Appendix B.6 discusses export entry in the two samples, and Appendix B.7 compares them to the Chilean ENIA, showing that the sectoral composition in all three samples is similar. In terms of export orientation, Chile and Colombia are also comparable, with about 20-25% of all plants being exporters. Mexican manufacturing plants, on the other hand, export more of their output – about 39% (which may in part be due to larger plants being overrepresented in the Mexican sample).

One important limitation is that – unlike the Chilean ENIA – the Colombian and Mexican data do not provide product-specific variable costs. We therefore cannot use equation (5) to compute product-specific material shares in multi-product plants – the basis to derive product-specific markups and marginal costs. We thus restrict our analysis for Colombia and Mexico to the subset of single-product plants, where all inputs are clearly related to the (single) produced output. Fortunately, both datasets include a large number of single-product plants, with almost 20,000 plant-year observations each (as compared to 25,000 for Chile). This allows us to compare the single-product results for Chile (shown in Table A.14) to those obtained for Colombia and Mexico, using exactly the same methodology.⁴⁹

We begin by describing the within-plant trajectories for Colombia in Figure 2.⁵⁰ TFPR remains

⁴⁷An alternative explanation for lower pass-through among existing exporters may be related to plant size: established Chilean exporters produce on average 30% more output and have 21% higher employment than new export entrants. Amiti, Itskhoki, and Konings (2016) show that for *domestic* sales, larger firms in Belgium show stronger strategic complementarities in pricing and therefore lower pass-through than smaller firms. However, strategic complementarities are less likely in the context of our findings, which are based on *export* sales to different markets across the globe – in contrast to the relatively small domestic market in Belgium. In addition, when replicating the results from panel C in Table 7 for above- vs. below-median employment, we find that increases in markups are quantitatively similar and – if anything – somewhat larger for smaller plants (see Table A.30 in Appendix C.5).

⁴⁸For example, Kugler and Verhoogen (2012) and Eslava et al. (2013) use the Colombian firm-level data from the Annual Manufacturing Survey (*Encuesta Annual Manufacturera*); Iacovone and Javorcik (2010) and Eckel, Iacovone, Javorcik, and Neary (2015) use data from the Mexican Monthly Industrial Survey (*Encuesta Industrial Mensual*) and from the Annual Industrial Survey (*Encuesta Industrial Anual*).

⁴⁹In all three cases, we estimate (8) for single-product plants, including plant fixed effects. We also include sector-year fixed effects at the 2-digit level, in line with our methodology for plant-level analyses (see footnote 25).

⁵⁰The corresponding within-plant coefficients for Colombia and Mexico are displayed in Tables A.11 and A.12,

essentially unchanged after export entry. Marginal costs, on the other hand, show a steep and highly significant decline by up to 40% after export entry. Markups increase mildly, by less than 10%.⁵¹ TFPQ confirms the magnitude of the marginal cost trajectory.

Figure 3 present the within-plant trajectories for Mexican export entrants. There is no change in TFPR or markups. Marginal costs, on the other hand, decline by 15-20% in the three years after export entry. This is quantitatively smaller than in the case of Colombia, but the results remain statistically significant at the 5% level. The results for TFPQ confirm the efficiency gains reflected by marginal costs. One potential reason for the relatively smaller efficiency gains after export entry is that larger plants are overrepresented in the Mexican data (see Appendix B.5). Larger plants are on average more productive (Syverson, 2011), and we know from the Lileeva and Trefler type test in Section 4.5 that more productive plants tend to see smaller efficiency gains after export entry. In fact, when splitting the Chilean sample into plants with above- and below-median employment, we also find smaller productivity gains for larger plants after export entry (see Table A.16).

Altogether, the results for Colombia and Mexico strongly confirm our findings for Chile: after export entry, plants experience significant efficiency increases, and these are almost entirely passed through to consumers in the form of lower prices. Thus, TFPR remains almost unchanged, which confirms its inferiority to alternative measures such as marginal costs or TFPQ. In Tables A.31 and A.32 in the appendix we show that investment of Colombian and Mexican export entrants spikes after export entry for "young exporters," and that this is almost entirely driven by increasing investment in machinery (as opposed to structures or vehicles). This confirms our findings for Chile, and suggests that an export-investment complementarity is a likely candidate for explaining the observed efficiency gains in Colombia and Mexico, as well.

7 Discussion and Conclusion

Over the last two decades, a substantial literature has argued that exporting induces within-plant efficiency gains. This argument has been made by theoretical contributions in the spirit of Grossman and Helpman (1991) and is supported by a plethora of case studies in the management literature. The finding that exporting induces investment in new technology also suggests that within-plant efficiency gains must exist (Bustos, 2011). A large number of papers has sought to pin down these effects empirically, using firm- and plant-level data from various countries in the developed and developing world. With less than a handful of exceptions, the overwhelming number of studies has

respectively (see Appendix B.8).

⁵¹The fact that markups grow somewhat more than TFPR is discussed in Appendix A.2: Colombian manufacturing shows on average (slightly) increasing returns to scale. In this case, fast expansions of volume (which are also observed for Colombia – see Panel B of Table A.11) can lead to MC overestimating efficiency gains, and to markup changes exceeding TFPR changes.

failed to identify such gains. We pointed out a reason for this discrepancy, and applied a recently developed empirical methodology to resolve it.

Previous studies have typically used revenue-based productivity measures, which are downward biased if higher efficiency is associated with lower output prices. In order to avoid this bias, we estimated marginal costs as a productivity measure at the plant-product level, following the approach by De Loecker et al. (2016). We have documented that marginal costs drop significantly after export entry, while markups remain relatively stable. Thus, productivity gains after export entry are largely passed on to customers in the form of lower output prices. We also showed that the typically used revenue-productivity remains largely unchanged after export entry. These results hold in three different countries that provide sufficiently detailed manufacturing data for our analysis: Chile, Colombia, and Mexico. Thus, our results likely reflect a general pattern, implying that a large number of previous studies has underestimated export-related efficiency gains by focusing on revenue-based productivity.

To support our argument that the observed efficiency gains are indeed trade-related, we used tariff variations in the particularly rich Chilean manufacturing panel. In this context, we distinguished between tariff-induced export entry and expanding foreign sales by established exporters. We found that both are associated with declining marginal costs (and – as a robustness check – with increasing TFPQ). We also compared these results to those based on the typically used TFPR. For tariff-induced export entry, TFPR fails to identify any gains; for tariff-induced export expansions, TFPR gains are statistically significant, but they reflect only one quarter of the productivity gains captured by marginal costs. These differences arise from the behavior of markups: on average, export entrants pass on almost all efficiency gains to customers – markups are unchanged, and therefore TFPR is unchanged. Established exporters, on the other hand, translate part of the efficiency gains into higher markups. These observations are compatible with ‘demand accumulation’ (Foster et al., 2016): new exporters may charge low prices initially in order to attract customers, while established exporters can rely on their existing customer network, so that lowering prices is less vital.

To gauge the quantitative importance of our findings, we compare the observed within-plant efficiency gains after export entry for the different productivity measures. We begin with TFPR. For export entrants, we found no increase in TFPR; and for tariff-induced export expansions of established exporters, the gains over the full sample period are approximately 5% (Table 7). Thus, if we had used the common revenue-based productivity measure, we would have confirmed the predominant finding in the previous literature – little evidence for within-plant efficiency gains. Based on marginal costs, on the other hand, new export entry is accompanied by efficiency increases of 15-25%. In addition, tariff-induced export expansions led to approximately 20% higher efficiency

over our sample period – roughly four times the magnitude reflected by TFPR. Compare this to Lileeva and Trefler (2010), who found that *labor* productivity rose by 15% for Canadian exporters during a major trade liberalization with the US in 1984-96. Since labor productivity is subject to the same (output) price bias as TFPR, the actual efficiency gains may well have been larger – if Canadian exporters, similar to their Chilean counterparts, passed on some of the efficiency gains to their customers in the form of lower prices.

Note that TFPR underestimating export-related efficiency gains is not a foregone conclusion: In principle, TFPR could also *overestimate* actual efficiency gains – if markups rise more than productivity. An extreme example would be exporters that raise their markups when tariffs fall, but do not invest in better technology. While our results suggest that such a strong response of markups is unlikely, we do observe markup increases among existing exporters when tariffs fall. This implies that the output price bias of TFPR is weaker during trade liberalization. One interpretation is that export tariff declines have an effect akin to demand shocks, which creates incentives to raise markups in models with endogenous markups such as Bernard et al. (2003) or Melitz and Ottaviano (2008). Consequently, it is more likely to find TFPR (i.e., markup) increases during periods of falling export tariffs. This may explain why the few studies that have identified export-related within-plant efficiency gains exploited periods of rapid trade liberalization (such as De Loecker, 2007 or Lileeva and Trefler, 2010).

Our results have two important implications for gains from trade: First, they rectify the balance of within-plant efficiency gains versus reallocation across plants. So far, the main effects have been attributed to the latter. For example, Pavcnik (2002) estimates that reallocation is responsible for approximately 20% productivity gains in export-oriented sectors during the Chilean trade liberalization over the period 1979-86. Using marginal cost as a productivity measure that is more reliable than its revenue-based counterparts, we show that export-related within-plant efficiency gains probably have a similar order of magnitude. Second, our results underline the necessity for future empirical studies to use productivity measures that are not affected by changes in output prices – and to re-examine previous findings that used revenue productivity. In particular, future studies should make further progress where our analysis was mostly exploratory due to the limited variation in Chilean export tariffs. Ideally, more detailed tariff changes at the plant- or disaggregated industry-level should be combined with marginal costs as a more reliable proxy for efficiency gains. Finally, our results imply that relatively stable markups are the reason why efficiency gains are not fully translated into higher revenue productivity. Thus, future research should examine the relationship between exporting and markups in more detail.

References

- Abadie, A., D. Drukker, J. L. Herr, and G. W. Imbens (2004). Implementing Matching Estimators for Average Treatment Effects in Stata. *Stata Journal* 4, 290–311.
- Akerberg, D. A., K. Caves, and G. Frazer (2015). Identification Properties of Recent Production Function Estimators. *Econometrica* 83(6), 2411–2451.
- Alvarez, R. and R. López (2005). Exporting and Performance: Evidence from Chilean Plants. *Canadian Journal of Economics* 38(4), 1384–1400.
- Amiti, M., O. Itskhoki, and J. Konings (2016). International Shocks and Domestic Prices: How Large Are Strategic Complementarities? NBER Working Paper 22119.
- Amiti, M. and J. Konings (2007). Trade Liberalization, Intermediate Inputs, and Productivity: Evidence from Indonesia. *American Economic Review* 97(5), 1611–1638.
- Andrews, D. W. K. and J. H. Stock (2005). Inference with Weak Instruments. NBER Working Paper 313.
- Atkeson, A. and A. T. Burstein (2010, 06). Innovation, firm dynamics, and international trade. *Journal of Political Economy* 118(3), 433–484.
- Atkin, D., A. K. Khandelwal, and A. Osman (2014). Exporting and Firm Performance: Evidence from a Randomized Trial. Working Paper.
- Aw, B. Y., S. Chung, and M. J. Roberts (2000). Productivity and Turnover in the Export Market: Micro-level Evidence from the Republic of Korea and Taiwan (China). *World Bank Economic Review* 14(1), 65–90.
- Aw, B. Y., M. J. Roberts, and D. Y. Xu (2011). R&D Investment, Exporting, and Productivity Dynamics. *American Economic Review* 101(4), 1312–44.
- Bas, M. (2012). Input-trade Liberalization and Firm Export Decisions: Evidence from Argentina. *Journal of Development Economics* 97(2), 481–493.
- Bernard, A. and J. Wagner (1997). Exports and Success in German Manufacturing. *Review of World Economics (Weltwirtschaftliches Archiv)* 133(1), 134–157.
- 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. and J. B. Jensen (1999). Exceptional Exporter Performance: Cause, Effect, or Both? *Journal of International Economics* 47(1), 1–25.
- Bernard, A. B., J. B. Jensen, S. J. Redding, and P. K. Schott (2012). The Empirics of Firm Heterogeneity and International Trade. *Annual Review of Economics* 4, 283–313.
- Blundell, R. and M. C. Dias (2009). Alternative Approaches to Evaluation in Empirical Microeconomics. *Journal of Human Resources* 44(3), 565–640.
- 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.
- Caliendo, L., G. Mion, L. Opromolla, and E. Rossi-Hansberg (2015). Productivity and Organization in Portuguese Firms.
- Clerides, S. K., S. Lach, and J. R. Tybout (1998). Is Learning By Exporting Important? Micro-Dynamic Evidence From Colombia, Mexico, and Morocco. *Quarterly Journal of Economics* 113(3), 903–947.
- Constantini, J. A. and M. J. Melitz (2007). The Dynamics of Firm-Level Adjustment to Trade Liberaliza-

- tion. In E. Helpman, D. Marin, and T. Verdier (Eds.), *The Organization of Firms in a Global Economy*. Cambridge: Harvard University Press.
- Davis, S., C. Grim, and J. Haltiwanger (2008). Productivity Dispersion and Input Prices: The Case of Electricity. Working Papers 08-33, Center for Economic Studies, U.S. Census Bureau.
- De Loecker, J. (2007). Do exports generate higher productivity? Evidence from Slovenia. *Journal of International Economics* 73(1), 69–98.
- De Loecker, J. (2011). Product Differentiation, Multiproduct Firms, and Estimating the Impact of Trade Liberalization on Productivity. *Econometrica* 79(5), 1407–1451.
- De Loecker, J. (2013). A Note on Detecting Learning by Exporting. *American Economic Journal: Macroeconomics*. forthcoming.
- De Loecker, J. and P. K. Goldberg (2014). Firm Performance in a Global Market. *Annual Review of Economics* 6(1), 201–227.
- 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. Neary (2015). Multi-Product Firms at Home and Away: Cost-versus Quality-Based Competence. *Journal of International Economics* 95(2), 216–232.
- Eslava, M., J. Haltiwanger, A. Kugler, and M. Kugler (2013). Trade and market selection: Evidence from manufacturing plants in colombia. *Review of Economic Dynamics* 16(1), 135–158.
- Evenson, R. E. and L. E. Westphal (1995). Technological Change and Technology Strategy. In H. Chenery and T. Srinivasan (Eds.), *Handbook of Development Economics*, Volume 3 of *Handbook of Development Economics*, Chapter 37, pp. 2209–2299. Elsevier.
- Fan, H., Y. A. Li, and S. Yeaple (2015). Trade Liberalization, Quality and Export Prices. *Review of Economics and Statistics* 97(5), 1033–1051.
- 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.
- Foster, L., J. Haltiwanger, and C. Syverson (2016). The Slow Growth of New Plants: Learning about Demand? *Economica* 83(329), 91–129.
- Goldberg, P. K., A. K. Khandelwal, N. Pavcnik, and P. Topalova (2010). Imported Intermediate Inputs and Domestic Product Growth: Evidence from India. *Quarterly Journal of Economics* 125(4), 1727–1767.
- Gorodnichenko, Y. (2012). Using Firm Optimization to Evaluate and Estimate Productivity and Returns to Scale. Working paper.
- Grossman, G. M. and E. Helpman (1991). *Innovation and Growth in the Global Economy*. Cambridge, MA: MIT Press.
- Hall, R. E. (1986). Market Structure and Macroeconomic Fluctuations. *Brookings Papers on Economic Activity* 17(2), 285–338.

- Heckman, J. J., H. Ichimura, and P. E. Todd (1997). Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme. *Review of Economic Studies* 64(4), 605–54.
- Hottman, C., S. J. Redding, and D. E. Weinstein (2016). Quantifying the Sources of Firm Heterogeneity. *Quarterly Journal of Economics*. forthcoming.
- Hsieh, C.-T. and P. J. Klenow (2009). Misallocation and Manufacturing TFP in China and India. *Quarterly Journal of Economics* 124(4), 1403–1448.
- Iacovone, L. and B. Javorcik (2012). Getting Ready: Preparation for Exporting. CEPR Discussion Paper 8926.
- Iacovone, L. and B. S. Javorcik (2010). Multi-Product Exporters: Product Churning, Uncertainty and Export Discoveries. *Economic Journal* 120(5), 481–499.
- ISGEP (2008). Understanding Cross-Country Differences in Exporter Premia: Comparable Evidence for 14 Countries. *Review of World Economics (Weltwirtschaftliches Archiv)* 144(4), 596–635. [ISGEP: International Study Group on Exports and Productivity].
- Klette, T. J. and Z. Griliches (1996). The Inconsistency of Common Scale Estimators when Output Prices are Unobserved and Endogenous. *Journal of Applied Econometrics* 11(4), 343–61.
- 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.
- Lileeva, A. and D. Trefler (2010). Improved Access to Foreign Markets Raises Plant-Level Productivity...for Some Plants. *Quarterly Journal of Economics* 125(3), 1051–1099.
- López, R. A. (2005). Trade and growth: Reconciling the macroeconomic and microeconomic evidence. *Journal of Economic Surveys* 19(4), 623–648.
- Luong, T. A. (2013). Does Learning by Exporting Happen? Evidence from the China Automobile Industry. *Review of Development Economics*. forthcoming.
- Manova, K. and Z. Zhang (2012). Export Prices across Firms and Destinations. *Quarterly Journal of Economics* 127(1), 379–436.
- Mas-Colell, A., M. D. Whinston, and J. R. Green (1995). *Microeconomic Theory*. Oxford University Press, USA.
- Melitz, M. (2000). Firm Productivity Estimation in Differentiated Product Industries. Working paper.
- 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.
- Ornaghi, C. (2006). Assessing the effects of measurement errors on the estimation of production functions. *Journal of Applied Econometrics* 21(6), 879–891.
- Park, A., D. Yang, X. Shi, and Y. Jiang (2010). Exporting and Firm Performance: Chinese Exporters and the Asian Financial Crisis. *Review of Economics and Statistics* 92(4), 822–842.
- Pavcnik, N. (2002). Trade Liberalization, Exit and Productivity Improvements: Evidence from Chilean Plants. *Review of Economic Studies* 69, 245–76.
- Rhee, W., B. Ross-Larson, and G. Pursell (1984). *Korea's Competitive Edge: Managing Entry into World Markets*. Johns Hopkins University Press, Baltimore.
- Rosenbaum, P. R. and D. B. Rubin (1983). The central Role of the Propensity Score in Observational Studies for Causal Effects. *Biometrika* 70(1), 41–55.
- Smeets, V. and F. Warzynski (2013). Estimating productivity with Multi-Product Firms, Pricing Heterogeneity and the Role of International Trade. *Journal of International Economics* 90(2), 237–244.
- Syverson, C. (2011). What Determines Productivity? *Journal of Economic Literature* 49(2), 326–365.
- Van Biesebroeck, J. (2005). Exporting Raises Productivity in Sub-Saharan African Manufacturing Firms. *Journal of International Economics* 67, 373–391.
- Verhoogen, E. (2008). Trade, Quality Upgrading and Wage Inequality in the Mexican Manufacturing Sector. *Quarterly Journal of Economics* 123(2), 489–530.
- Wagner, J. (2007). Exports and Productivity: A Survey of the Evidence from Firm-level Data. *The World Economy* 30(1), 60–82.
- Wagner, J. (2012). International trade and firm performance: a survey of empirical studies since 2006. *Review of World Economics (Weltwirtschaftliches Archiv)* 148(2), 235–267.
- Wooldridge, J. M. (2002). *Econometric Analysis of Cross Section and Panel Data*. Cambridge, MA: MIT Press.
- World Bank (1993). *The East Asian Miracle: Economic Growth and Public Policy*. Oxford University Press, New York.