NBER WORKING PAPER SERIES

EXPORTING AND PLANT-LEVEL EFFICIENCY GAINS:
IT'S IN THE MEASURE

Alvaro Garcia Marin
Nico Voigtländer

Working Paper 19033
http://www.nber.org/papers/w19033

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2013

We would like to thank Roberto Álvarez, Andrew Atkeson, Nick Bloom, Johannes van Biesebroeck, Joaquin Blaum, Ariel Burstein, Leonardo Bursztyn, Lorenzo Caliendo, Donald Davis, Jan De Loecker, Gilles Duranton, Eduardo E. Engel, Juan Carlos Hallak, Amit Khandelwal, Edward Leamer, Steve Redding, Ariell Reshef, Andrés Rodríguez-Clare, Veronica Rappoport, John Ries, Aaron Tornell, Eric Verhoogen, Romain Wacziarg, and Pierre-Olivier Weill, as well as seminar audiences at Columbia, CREI, the FREIT Conference in Virginia, KU Leuven, 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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Exporting and Plant-Level Efficiency Gains: It's in the Measure
Alvaro Garcia Marin and Nico Voigtländer
NBER Working Paper No. 19033
May 2013, Revised October 2014
JEL No. D24,F14,L25,L60

ABSTRACT

While there is strong evidence for productivity-driven selection into exporting, previous research has mostly failed to identify export-related efficiency gains within plants. This non-result is derived from revenue productivity (TFPR), thus also reflecting pricing decisions of exporters. 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 entrants, we find within-plant 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 the timing of export entry. We also find sizeable efficiency gains for tariff-induced export expansions of existing exporters. Only half of these gains are reflected by TFPR, due to a partial rise in markups. Our results thus suggest that 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.

Alvaro Garcia Marin
UCLA Anderson School
110 Westwood Pl
Los Angeles, CA 90095
alvarof.garciam@gmail.com

Nico Voigtländer
UCLA Anderson School of Management
110 Westwood Plaza
C513 Entrepreneurs Hall
Los Angeles, CA 90095
and NBER
nico.v@anderson.ucla.edu

A data appendix is available at:
http://www.nber.org/data-appendix/w19033

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). Technology upgrading, in turn, should lead to observable efficiency increases. Why has the empirical literature struggled to identify such gains?

In this paper, we 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 *physical* 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 physical productivity, while revenue productivity reflects efficiency gains only if markups rise.

We then exploit an unusually rich dataset of Chilean manufacturing plants to analyze the trajectories of marginal cost, markups, and prices around export entry and export expansions. To derive plant-product level markups, we apply the method pioneered by De Loecker and Warzynski

[1]Early 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).

[2]Recent 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.

(2012). Because our dataset comprises physical units as well as revenues for each plant-product pair, we can calculate product prices (unit values). Dividing these by the corresponding markups allows us to identify marginal costs at the plant-product level (De Loecker, Goldberg, Khandelwal, and Pavcnik, 2012). 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).

In order to compare efficiency gains captured by the different measures, we also construct revenue productivity (TFPR) at the plant and at the plant-product level. We examine the relationship between efficiency and exporting, both along the extensive margin (export entry) and the intensive margin (export expansions of established exporters). We first confirm that, in line with most previous findings, the trajectory of TFPR is flat around export entry for average Chilean plants. We then disentangle this pattern and find 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 physical productivity gains, but that these are passed on to their customers. In other words, 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-priory comparable likelihood of entering the export market, but continued to be sold domestically only. In addition, we show that we obtain quantitatively similar results when using reported variable cost measures 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 variable costs. We also discuss that our results are unlikely to be confounded by changes in product quality.[3] We then exploit tariff changes 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: Chile did not undergo major trade liberalization during our sample period. 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

[3]The 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.

costs declining by 20-30%.

We provide evidence that technology upgrading is the most likely explanation for declining marginal costs at export entry. This is supported by several patterns in the data. For example, we show that plant-level investment (especially in machinery) spikes immediately before, and during the first years, of 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.

We then turn to continuing exporters. In the sample overall, we find no correlation between export expansions and efficiency measures or markups within established exporting plants. Because of the relatively stable trade costs over our sample period, the ups and downs of existing exporters are likely due to transitory export demand shocks that are insufficient to trigger investment in new technology. In fact, when we restrict the sample to a sub-period when many industries experienced falling export tariffs, the relationship between export sales and plant efficiency becomes stronger. To exploit tariff-driven variations in exporting more systematically, we then use a 2SLS approach with export tariffs at the detailed industry level as instruments. We find strong evidence that export expansions that are induced by tariff declines lead to lower marginal costs (by approximately 10% over our sample period), and that this link works via investment in capital. This suggests that *permanent* changes in trade costs – in the form of stable tariff declines – induce investment in new technology and thereby increase efficiency.

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 one half of the decline in marginal costs translate into lower prices, and the remaining half, into higher markups. Consequently, TFPR also increases and reflects one half 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.

We discuss the differences between export entry and expansions of existing exporters. First, why is the former, but not the latter, associated with efficiency gains even in the absence of tariff declines? One interpretation is that the decision to enter the export market for the first time reflects (at least in expected terms) a permanent change in production, and thus incentivizes investment in new technology. Temporary increases in sales of existing exporters, on the other hand, are too short-lived to render technology upgrading profitable. Second, why are markups stable around export entry, but increase for established exporters after tariff-induced expansions? This pattern is compatible with 'demand building' (Foster, Haltiwanger, and Syverson, 2012) – while existing exporters already have a customer base abroad, new entrants may use low prices to attract customers.[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 significantly more than its domestic counterpart (22% vs. 8%). There is also some evidence in our data that markups grow as export entrants become more established.[5]

Our findings relate to a substantial literature on gains from trade. Trade-induced competition 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.[6] 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 weak) empirical support for this effect, while reporting 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.[7] The survey article by ISGEP (2008) compiles micro level panels from 14 countries and finds nearly no evidence for within-plant productivity 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).[8] Our results illustrate why it may be more likely

[4]Foster et al. (2012) 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.

[5]There 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.

[6]Two 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.

[7]Alvarez 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.

[8]Van Biesebroeck (2005) also documents productivity 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.

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 one-to-one relationship with markups.[9] 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.[10]

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.[11] Second, we show that disentangling the trajectories of prices and 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 about half the gains from tariff-induced export expansions. Consequently, we identify substantial export-related efficiency gains that have thus far passed under the radar.[12] Our study thus complements a substantial literature that argues that within-plant efficiency gains should be expected.[13] Finally, our unique dataset allows us to verify the methodology for computing marginal costs based on markups (De Loecker et al., 2012): we show that changes in computed plant-product level marginal costs are very similar to those in self-reported average costs.

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 dataset, and Section 4 presents our

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

[10]We discuss below that marginal costs have an important advantage over TFPQ in the context of our study: For multi-product plants (the majority of exporters), *product*-level marginal costs can be computed under relatively unrestrictive assumptions. This allows for our analysis of efficiency by decomposing prices into markups and marginal costs – all variables that naturally vary at the product level.

[11]De Loecker et al. (2012) document a fall in the marginal cost of Indian firms following a decline in *input* tariffs.

[12]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.

[13]Case 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.

empirical results for export entrants, and Section 5, for continuing exporters. Section 6 discusses our results and draws conclusions.

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- and at the 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 are a valid efficiency measure.

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).[14] This measure has an important shortcoming, which can be illustrated by its decomposition into prices, P, and physical productivity (or efficiency), A, assuming that the true A is known: $\ln(\text{TFPR}) = \ln(P)+\ln(A)$. If prices are unrelated to efficiency, using TFPR as a proxy for A merely introduces noise, and TFPR is unbiased. However, when prices respond to 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 and A, so that TFPR will underestimate physical productivity.

Despite these 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.[15] While some corrections to the estimation of production functions have been proposed, only a few studies have derived A directly.[16] In addition, even if quantities are known, they cannot

[14] Some 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 TFP when different plants combine inputs in different proportions (see Syverson, 2011).

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

[16] Melitz (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 A using a

readily be compared – a problem that is particularly severe for multi-product plants. To circumvent these issues, we propose marginal cost as a 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 efficiency (physical productivity) A. To illustrate this relationship, we use the generic functional form $MC(A_{it}, \mathbf{w}_{it})$, where $\mathbf{w}_{it}$ is an input price index, and the subscripts i and t denote plants and years, respectively. 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 TFPR (assuming Hicks-neutrality – as is standard in the estimation of productivity):

$$\text{TFPR}_{it} = p_{it}A_{it} = \mu_{it} \cdot MC(A_{it}, \mathbf{w}_{it}) \cdot A_{it} \tag{1}$$

Deriving percentage changes (denoted by $\triangle$) and re-arranging yields a relationship between efficiency gains and changes in TFPR, markups, and marginal costs:[17]

$$\triangle A_{it} = \triangle\text{TFPR}_{it} - \triangle\mu_{it} - \triangle MC(A_{it}, \mathbf{w}_{it}) \tag{2}$$

In order to simplify the interpretation of (2) – but not in the actual estimation of $MC(\cdot)$ – we make two assumptions. First, that the underlying production function exhibits constant returns to scale. This assumption is supported by our data, where the average sum of input shares is very close to one (see Table A.1 in the appendix). This first assumption implies that we can separate $\triangle MC(A_{it}, \mathbf{w}_{it}) = \triangle\phi(\mathbf{w}_{it}) - \triangle A_{it}$, where $\phi(\cdot)$ is an increasing function of input prices (see the proof in Appendix A). 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: $\triangle\phi(\mathbf{w}_{it}) = 0$. This assumption is stronger than the previous one and requires some discussion. Our dataset allows us to calculate input prices, and we show below in Section 4.5 that these do not change significantly with exporting activity – if anything, they show a slight increase, biasing our results against finding declining marginal costs. This is compatible with previous findings that more successful exporters typically produce high-quality goods that

model of monopolistic competition for India, China and the United States. Foster et al. (2008) obtain A using product-level information on physical quantities from U.S. census data for a subset of manufacturing plants that produce homogeneous products. Finally, Eslava et al. (2013) and Lamorgese et al. (2014) compute TFPQ and use it to analyze gains from trade.

[17] We slightly abuse notation, employing $\triangle$ to represent changes in the logarithm of variables, for example, $\triangle A_{it} = d\ln(A_{it}) = dA_{it}/A_{it}$.

require more expensive inputs (Manova and Zhang, 2012). Therefore, $\mathbf{w}_{it}$ would tend to increase for more successful export entrants, and efficiency gains $\triangle A_{it}$, inferred from any given $\triangle MC$, would be larger if we allowed also for rising costs of inputs (since $\triangle A_{it} = \triangle\phi(\mathbf{w}_{it}) - \triangle MC$).

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

1. $\triangle A_{it} = -\triangle MC$, 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 (2) also implies:
2. $\triangle \text{TFPR}_{it} = \triangle \mu_{it}$, i.e., revenue productivity rises if and only if markups increase. For example, even if A_{it} rises (and MC falls), TFPR will not grow if markups remain unchanged. And vice-versa, if markups rise while A_{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 markup.
3. $\triangle \text{TFPR}_{it} = \triangle A_{it}$ if $\triangle \mu_{it} = -\triangle MC$, i.e., changes in revenue productivity reflect the full efficiency gains if markups rise in the same proportion as marginal costs fall. Because $p_{it} = \mu_{it} \cdot MC$, this will be the case if prices are constant while marginal costs fall.

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, 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 first have to estimate the revenue production function. We specify a Cobb-Douglas production function with labor (l), capital (k), and materials (m) as production inputs. 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/marginal costs. This ensures that differences in the efficiency measures are not driven by different parameter estimates.[18] Following De Loecker et al. (2012), we estimate a separate production function for each 2-digit

[18] As discussed below, TFPR needs to be estimated based on output in revenues, while deriving markups based on revenues (rather than quantities) can lead to biased results. In the Cobb-Douglas case, this bias does not affect our results because it is absorbed by plant-product fixed effects. Consequently, the Cobb-Douglas specification allows us to use estimates for both TFPR and markups (and thus marginal costs) based on the *same* production function coefficients. In Appendix D we show that the more flexible translog specification (where fixed effects do not absorb the bias) confirms our baseline results, which also implies that the bias is limited.

manufacturing sector (s), using the subsample of single product plants.[19] The reason for using single-product plants is that we do not observe how inputs are allocated to individual outputs within multi-product plants. For the set of single product plants, no assumption on the allocation of inputs to outputs is needed, and we can estimate the following production function with standard plant-level information:

$$q_{it} \quad = \quad \beta_l^s l_{it} + \beta_k^s k_{it} + \beta_m^s m_{it} + \omega_{it} + \varepsilon_{it} \tag{3}$$

where all lowercase variables are in logs; q_{it} are revenues of single-product plant i in year t, ω_{it} is TFPR, k_{it} denotes the capital stock, m_{it} are material inputs, and ε_{it} represents measurement error as well as unanticipated shocks to output. Estimating (3) yields the sector-specific vector of coefficients $\boldsymbol{\beta}^s = \{\beta_l^s, \beta_k^s, \beta_m^s\}$.

When estimating (3) we follow the methodology by Ackerberg, Caves, and Frazer (2006, henceforth ACF), who extend the framework of Olley and Pakes (1996, henceforth OP) and Levinsohn and Petrin (2003, henceforth LP). This methodology controls for the simultaneity bias that arises because input demand and unobserved productivity are positively correlated.[20] The key insight of ACF lies in their identification of the labor elasticity, which they show is in most cases unidentified by the two-step procedure of OP and LP.[21] We modify the canonical ACF procedure by specifying an endogenous productivity process that can be affected by export status and plant investment. In addition, we include interactions between export status and investment in the productivity process. Thus, the procedure allows exporting to affect current productivity either directly, or through a complementarity with investment in physical capital. This reflects the corrections suggested by De Loecker (2013); if productivity gains from exporting also lead to more investment (and thus a higher capital stock), the standard method would overestimate the capital coefficient in the production function, and thus underestimate productivity (i.e., the residual). Finally, using the set of single-product plants may introduce selection bias because plant switching from single- to multi-product may be correlated with productivity. Following De Loecker et al. (2012), we correct for this source of bias by including the predicted probability of remaining single-product, $\hat{s}_{it}$, in

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

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

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

the productivity process as a proxy for the productivity switching threshold.[22] Accordingly, the law of motion for productivity is:

$$\omega_{it} \quad = \quad g(\omega_{it-1}, d^x_{it-1}, d^i_{it-1}, \hat{s}_{it-1}) + \xi_{it} \tag{4}$$

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

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

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

where $\phi_t(\cdot) = \beta^s_l l_{it} + \beta^s_k k_{it} + \beta^s_m m_{it} + h_t(m_{it}, l_{it}, k_{it}, \boldsymbol{x}_{it})$, with $h_t(\cdot)$ denoting the inverse material demand that we use to proxy for the unobserved productivity term.[23] The vector $\boldsymbol{x}_{it}$ contains all other variables that affect material demand (time and product dummies, reflecting aggregate shocks and specific demand components). Using the estimate of expected output, productivity can be computed for any candidate coefficient vector $\tilde{\boldsymbol{\beta}}^s$ as $\omega_{it}(\tilde{\boldsymbol{\beta}}^s) = \hat{\phi}_t - (\beta^s_l l_{it} + \beta^s_k k_{it} + \beta^s_m m_{it})$. Estimating $\omega_{it}(\tilde{\boldsymbol{\beta}}^s)$ non-parametrically as a function of its own lag $\omega_{it-1}(\tilde{\boldsymbol{\beta}}^s)$ and prior exporting and investment status (d^x_{it-1}, d^i_{it-1}), the productivity innovation can be recovered for each candidate $\tilde{\boldsymbol{\beta}}^s$.[24]

In the second stage, all coefficients of the production function are identified through GMM using the moment conditions

$$\mathbb{E}\left(\xi_{it}(\boldsymbol{\beta}^s)\mathbf{Z}_{it}\right) = 0 \tag{5}$$

where ξ_{it} is the productivity innovation term from (4), and $\mathbf{Z}_{it}$ is a vector of variables that comprises lags of all the variables in the production function, and the current capital stock. These variables are valid instruments – including capital, which is chosen before the productivity innovation is observed. Equation (5) thus says that for the optimal $\boldsymbol{\beta}^s$, the innovation in productivity is uncorrelated with the instruments $\mathbf{Z}_{it}$.

Given the estimated coefficients for each product category s (the vector $\boldsymbol{\beta}^s$), revenue productivity can be calculated both at the plant level and for individual products within plants. For the

[22] We estimate this probability for 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.

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

[24] Following Levinsohn and Petrin (2003), we approximate the law of motion for productivity (the function $g(\cdot)$ stated in (4)) with a polynomial.

former, we use the plant-level aggregate labor l_{it}, capital k_{it}, and material inputs m_{it}. We then compute plant-level TFPR, $\hat{\omega}_{it}$:

$$\hat{\omega}_{it} = q_{it} - (\beta_l^s l_{it} + \beta_k^s k_{it} + \beta_m^s m_{it}) \tag{6}$$

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

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

$$M_{ijt} = s_{ijt}^{TVC} \cdot M_{it} \quad \text{where} \quad s_{ijt}^{TVC} = \frac{TVC_{ijt}}{\sum_j TVC_{ijt}} \tag{7}$$

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

$$\hat{\omega}_{ijt} = q_{ijt} - (\beta_l^s l_{ijt} + \beta_k^s k_{ijt} + \beta_m^s m_{ijt}) \tag{8}$$

where q_{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. The first order condition of a plant'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:[25]

$$\underbrace{\mu_{ijt}}_{Markup} \equiv \frac{P_{ijt}}{MC_{ijt}} = \underbrace{\left(\frac{\partial Q_{ijt}(\cdot)}{\partial V_{ijt}} \frac{V_{ijt}}{Q_{ijt}}\right)}_{\text{Output Elasticity}} \Big/ \underbrace{\left(\frac{P^V_{ijt} \cdot V_{ijt}}{P_{ijt} \cdot Q_{ijt}}\right)}_{\text{Expenditure Share}}, \tag{9}$$

where P (P^V) denotes the price of output Q (input V), and MC is marginal cost. According to equation (9), the markup can be computed by dividing the output elasticity of product j (with respect to the flexible input) by the cost of the flexible input, relative to the sales of product j.

In our computation of (9) we use materials (M) as the flexible input to compute the output elasticity – based on our estimates of (3).[26] 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^s_m. Ideally, β^s_m should be estimated using physical quantities for inputs and output in (3). 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^s_m, 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 (9) – 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 approximation described in equation (7) to obtain the value of material inputs $P^V_{ijt} \cdot V_{ijt} = M_{ijt}$. Since total product-specific revenues $P_{ijt} \cdot Q_{ijt}$ are reported in our data, we can then compute the plant-product specific expenditure shares needed in (9).[27]

Because markups are computed at the plant-product level, and prices (unit values) are observed at the same level, we can derive marginal costs at the plant-product level in each year. To avoid that extreme values drive our results, we only use observations within the percentiles 1 and 99 of the markup distribution. The remaining markup observations vary between (approximately) 0.5

[25] More precisely, the first order condition with respect to V is $\frac{\partial \mathcal{L}}{\partial V} = P^V - \lambda \frac{\partial Q(\cdot)}{\partial V} = 0$, where the Lagrange multiplier λ equals the marginal cost of production. Manipulating this expression yields (9).

[26] In principle, labor could be used as an alternative. However, in the case of Chile, labor being a flexible input would be a strong assumption due to its regulated labor market. A discussion of the evolution of job security and firing cost in Chile can be found in Montenegro and Pagés (2004).

[27] By using each product's reported variable cost shares to proxy for product-specific material costs, we avoid shortcomings of a prominent earlier approach: since product-specific cost shares were not available in their dataset, Foster et al. (2008) had to assume that plants allocate their inputs proportionately to the share of each product in total *revenues*. This is problematic because differential changes in markups across different products will affect revenue shares even if cost shares are unchanged. De Loecker et al. (2012) avoid this issue by using an elaborate estimation technique to identify product-specific material costs; this is not necessary in our setting due to the reported variable cost shares.

and 4. In Table A.2 we show 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 $\boldsymbol{\beta}^s$ have been estimated correctly.[28] Then, in order to back out TFPQ by using (6), both output and inputs need to be transformed into physical quantities, using price indexes. 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 variable costs (which is sufficient to derive the product-specific material expenditure share needed in (9) to compute markups), it does not contain information on how to assign input *quantities* to products. Thus, assigning m_{it}, l_{it}, and k_{it} to individual products is prone to errors – or, in the case of capital, conceptually questionable. In light of these limitations, most studies compute TFPQ at the plant or firm level, thus not allowing for a product-specific analysis of export entry. An additional shortcoming of this more aggregate approach is that plant-level price indexes do not account for differences in product scope (Hottman, Redding, and Weinstein, 2014).

Contrast this with the computation of markups in (9), still assuming that $\boldsymbol{\beta}^s$ has been correctly estimated. The output elasticity 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 (7). Thus, no elaborate procedures with price indexes is needed. Also, we only need to proportionately assign the expenditure share of *material* inputs to individual products, but not of capital and labor: while it is reasonable to assume that product-specific reported variable costs reflect materials (and also labor), this is more of a stretch for capital.

We now turn to the estimation of $\boldsymbol{\beta}^s$, which is challenging and may introduce further error. When using a Cobb-Douglas production function, this issue is less severe for markups than for TFPQ. The computation of markups uses only β_m^s from the vector $\boldsymbol{\beta}^s$. 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

[28] To compute TFPQ, the elasticities in the production function (3) must be estimated in quantities. Estimating this vector is challenging in itself: When estimating the production function (3), product-specific output and inputs have to be adjusted by proper price indexes. In addition, if input quantities are not available and input expenditure is used instead, the estimation of the production function coefficients is biased (see De Loecker et al., 2012). Although this bias may be corrected using proxies for input price variation, such proxies are typically unavailable and researchers need to rely on output price variation the as main driver of input prices. We discuss this in more detail in Appendix D.

log-*changes* in markups in (9) are only driven by the observed material expenditure shares, but not by the estimated output elasticity β^s_m.[29] Contrast this with the computation of TFPQ, which uses all coefficients in $\boldsymbol{\beta}^s$, multiplying each by the corresponding physical input in (6). In this case, analyzing log-changes in TFPQ will not eliminate errors and biases in the level of $\boldsymbol{\beta}^s$.

Finally, since we study efficiency gains in the context of investment-exporting complementarity, it is also worthwhile to discuss how investment in new technology affects TFPQ and marginal cost. In particular, one may worry that while TFPQ explicitly accounts for the effect of fixed-cost investment in capital equipment, our estimation may (wrongly) identify declining marginal cost even if the technology itself does not change. We show that this is not the case if plants minimize costs, and under our assumptions from Section 2.2 that i) input prices do not change with export activity and ii) constant returns to scale. In the following discussion, we assume that the input elasticities $\boldsymbol{\beta}^s$ have been correctly estimated for the quantity production function, so that changes in physical output can be readily computed using the quantity-equivalent of (3). Suppose that a plant raises its capital stock by $\triangle k$ (in log changes), adding the same type of machines, so that true efficiency is unchanged ($\triangle\omega = 0$).

Because it is minimizing costs, the plant will maintain its expenditure shares for all other inputs (material m and labor l) proportional to the respective input elasticities. Under constant input prices, this implies that $\triangle k = \triangle m = \triangle l$. Thus, due to constant returns, total output increases by $\triangle q = \triangle k$, and (6) correctly implies that TFPQ is unchanged. Next, we turn to marginal costs. Recall that we use materials as the variable input V. Also, for the moment, hold output prices P fixed. The first term of (9) – the material input elasticity – is unchanged. In the second term, the quantity of the flexible input V has increased by $\triangle m$ log points, and physical quantity Q has increased by $\triangle q$ log points. Because $\triangle q = \triangle k = \triangle m$, markups are unchanged for *given* output prices P. However, the latter may have changed during the plant's investment-driven expansion. Suppose that output prices fell by $\triangle p$ log points (e.g., because the plant had to charge lower prices in order to sell its increased output volume). Then (9) implies that the total effect of the investment-driven expansion is a decline in markup by $\triangle\mu = \triangle p$ log points. Log-changes in marginal cost can then be computed as $\triangle mc = \triangle p - \triangle\mu = 0$. Consequently, marginal costs correctly reflect that efficiency has not changed. Finally, the same calculation can be made for investment-driven expansions that raise efficiency by $\triangle\omega > 0$ (e.g., by adding new, more efficient machines). Provided that the new technology uses all inputs in the same proportions as before

[29] This is also the reason why we can use estimates of $\boldsymbol{\beta}^s$ from the *revenue* production function, i.e., the same coefficients used to compute TFPR. Note that for the more flexible translog specification, β^s_m itself depends on the use of inputs and may thus vary over time. We show in Appendix D that our results are nevertheless robust to this specification.

(Hicks-neutrality – a standard assumption in the productivity literature), both TFPQ and marginal costs will drop by $\triangle\omega$.

3 Data

Our data are from a Chilean plant panel for the period 1996–2005, the *Encuesta Nacional Industrial Anual* (Annual National Industrial Survey – ENIA). Data for ENIA are collected annually by the Chilean National Institute of Statistics (INE), with direct participation of Chilean manufacturing plants. 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 4,900 manufacturing plants per year with positive sales and employment information. Out of these, about 20% are exporters, and 70% of exporters are multi-product plants. Within the latter (i.e., conditional on at least one product being exported), exported goods account for 79.6% 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.[30] The CUP categories identify 2,169 different products in the sample. These products – in combination with each plant producing them – form our main unit of analysis. In the following, we briefly discuss how we deal with inconsistent product categories, units of output, and other issues of sample selection.

3.1 Sample Selection and Data Consistency

In order to ensure consistent plant-product categories in our panel, we follow three steps. First, we drop plant-product-year observations whenever there are signs of unreliable reporting. In particular, we exclude plant-product-year observations that have zero values for total employment, demand for raw materials, sales, or product quantities. Second, whenever our analysis involves quantities of production, we have to carefully account for possible changes in the unit of mea-

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

surement. For example, wine producers change in some instances from "bottles" to "liters." Total revenue is generally unaffected by these changes, but the derived unit values (prices) have to be corrected. This procedure is needed for about 1% of all plant-product observations; it is explained in Appendix B. Third, a similar correction is needed because the product identifier in our sample changes in the year 2001. We use a correspondence provided by the Chilean Statistical Institute to match the new product categories to the old ones (see Appendix B for detail). After these adjustments, our sample consists of 109,210 plant-product-year observations.

3.2 Definition of Export Entry

The time of entry into export markets is crucial for our analysis. We observe the exporting history of each plant-product pair from 1996 to 2005. We impose three requirements for product j, produced by plant i, to classify 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, and (iii) 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 772 export entries (plant-products at the 7-digit level), and approximately 7% of active exporters are new entrants.

3.3 Validity of the Sample

Before turning to our empirical results, we check whether our data replicate some well-documented systematic differences between exporters and non-exporters. Following Bernard and Jensen (1999), we run the regression

$$\ln(y_{ist}) \quad = \quad \alpha_{st} + \delta \, d_{ist}^{exp} + \gamma \, \ln(L_{ist}) + \varepsilon_{ist} \; , \tag{10}$$

where y_{ist} denotes several characteristics of plant i in sector s and period t, d_{ist}^{exp} is an exporter dummy, L_{ist} is total plant-level employment, and α_{st} denotes sector-year fixed effects.[31] The coefficient δ reports the exporter premium – the percentage-point difference of the dependent variable between exporters and non-exporters. Panel A in Table 1 reports unconditional exporter premia, while Panel B controls for plant-level employment. The results are similar for both specifications:

[31] Whenever we use plant-level regressions, we control for sector-year effects at the 2-digit level. When using the more detailed plant-product data, we include a more restrictive set of 4-digit sector-year dummies. These correspond to approximately 200 product categories.

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 in column 5, we also find that markups are higher among exporters, confirming the findings in De Loecker and Warzynski (2012).

4 Efficiency Gains of Export Entrants

In this section we present our empirical results for new export entrants. We show the trajectories of revenue productivity, marginal costs, and markups within plant-products around the time of export entry. Our main finding is that TFPR does not change after export entry, while marginal costs drop substantially. Markups are also constant, indicating that efficiency gains are passed on to customers. We show that the same results hold when we focus on export entries that are predicted by tariff declines, and we provide suggestive evidence that the observed efficiency gains are driven by a complementarity between export entry and investment.

4.1 New Export Entrants: Within Plant Trajectories

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

$$y_{ijt} = \alpha_{st} + \alpha_{ij} + \underbrace{\sum_{k=-2}^{-1} T_{ijt}^{k}}_{Pre-Trend} + \underbrace{\sum_{l=0}^{L} E_{ijt}^{l}}_{Entry-Effect} + \varepsilon_{ijt} \; , \tag{11}$$

where y_{ijt} refers to price, marginal cost, markup, or TFPR; 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).[32] 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.[33]

Figure 1 visualizes the results of estimating (11) for the sub-sample of export entrants. The

[32] For *plant-level* TFPR, the product index j in y_{ijt} is irrelevant in (11). Since plants are classified at the 4-digit SIC level, we include sector-year fixed effects at the 2-digit level (see footnote 31).

[33] Due to our relatively short sample, we only report the results for $l = 0, ..., 3$ periods after export entry. However, all regressions include dummies E_{ijt}^{l} for all post-entry periods.

figure shows the point estimates for each outcome variable, together with the 90% confidence intervals. Time on the horizontal axis is normalized such that zero represents the entry period. The left panel of the figure shows the trajectories of TFPR at the plant- and plant-product level. Both are virtually unaffected by export entry, with tight confidence intervals around zero during the first two periods after entry ($t = 0$ and $t = 1$).[34] This result is in line with the previous literature: there are no apparent efficiency gains when TFPR is used as a measure of efficiency. In $t = 2$ and $t = 3$, there is some weak evidence for increasing TFPR, which we discuss in more detail below.

The right panel of Figure 1 shows a radically different pattern. After entry into the export market, marginal costs decline markedly. According to the point estimates, marginal costs are about 11% lower at the moment of entry, as compared to pre-exporting periods. This difference widens over time: one period after entry it is 15%, and after 3 years, about 28%. These differences are not only economically but also statistically significant. Table 2 reports the corresponding coefficients. The trajectory for prices is very similar to marginal costs. This results because markups remain essentially unchanged after export entry – only two years after entry, there is a slight increase in markups by about 5%. The pattern in markups coincides with the one in TFPR, in line with our theoretical results in Section 2. This confirms that revenue productivity reflects efficiency gains only if markups rise, i.e., if not all gains are passed on to customers. Physical quantities sold increase by approximately 11-18% after export entry.

Reported Average Costs

One potential concern for 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.[35] We can address this concern by using a unique feature in our dataset to compute an alternative cost measure. Plants covered by ENIA report the total production cost *per product*, as well as the number of units produced. The questionnaire defines total 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, and is thus a reasonable proxy for average variable costs. Figure 2 plots our computed marginal costs against the reported average costs (both in logs), controlling for plant-product fixed effects, as well as 4-digit sector-year fixed effects (i.e., reflecting the within plant-product variation that we exploit empirically). The two measures are very strongly correlated.

[34] The fact that TFPR for the exported product shows a very similar trend as its plant-level counterpart is not surprising, given that the exported product typically accounts for the majority of output in exporting multi-product plants.

[35] For example, suppose that prices actually fall because markups shrink upon export entry, but that noisy data cloud these changes when applying the methodology in section 2. Then we would wrongly attribute the observed decline in prices after export entry to a decline in marginal cost.

This lends strong support to the markup-based methodology for backing out marginal costs by De Loecker et al. (2012). Next, we use average cost as a measure of efficiency and repeat the above estimations.

The last row of Table 2 shows that average costs decrease after export entry, closely following the trajectory that we identified for marginal cost. Export entry is followed by a decline in average costs of 11% in the period of entry, growing to 14% after one year, and to 26% 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.

4.2 Matching Results

Our within-plant trajectories in Table 2 showed a slight (statistically insignificant) decline in price and marginal cost 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). 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.[36] Appendix C provides further detail. 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

[36]Following 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 $t = -1$.

results because initial differences between treated and control units are removed.

Table 3 shows the matching estimation results. Since all variables are expressed in logarithms, the DID estimator reflects the difference in *growth* between newly exported products and their matched controls, relative to the pre-entry period ($t = -1$).[37] These results confirm the within-plant pattern documented above: changes in TFPR after export entry are initially small and statistically insignificant; the same is true for markups. After three periods, TFPR increases slightly, and this goes hand-in-hand with higher markups – both increase by about 8 percentage points more than their counterparts for the matched control products. This suggests that, eventually, efficiency gains are partially reflected in TFPR. Marginal costs, on the other hand, decrease after entry into export markets. When compared to the previously reported within-plant trajectories, the PSM results show somewhat smaller initial differences that grow over time: the difference in marginal cost relative to the control group grows from 0.4% in the period of export entry to 16% in the year after entry, and to 24% three periods after entry.

4.3 Robustness and Additional Results

In this subsection we check the robustness of our results to alternative specifications and sample selection.

Balanced Sample of Entrants

To what extent does unsuccessful export entry drive our results? To answer this question, we construct a balanced sample of exporters, including only plant-products that are exported in each of the first 3 years after export entry. Table 4 shows the results for propensity score matching. The main pattern is unchanged. TFPR results are quantitatively small and mostly insignificant, while marginal costs drop markedly after export entry – by approximately 18-30 percentage points more than for comparable plant-products that did not enter the export market. The main difference with Table 3 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, effects tended 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. In sum, the results from the balanced sample confirm our full sample estimates and suggest relatively stable efficiency gains over time.

[37] For 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 percentage points more for export entrants, as compared to the non-exporting control group.

Further Robustness Checks

We perform several additional robustness checks in the appendix, and briefly summarize these here. In our baseline matching estimation, we used the 5 nearest neighbors. Table A.3 shows that using either 3 or 10 neighbors instead does not change our results. Next, in order to estimate product-level TFPR, marginal costs, and markups, we had to assign inputs to individual products in multi-product plants, using reported variable cost shares as in equation (7). This is not needed in single-product plants, where all inputs enter in the production of one final good. Table A.4 uses only the subset of single-product plants. This robustness check comes at a cost: export entries by single-product plants represent only about one-fourth of the total number of entries in our sample. Correspondingly, the results are noisier than before. Nevertheless, the magnitude of coefficients confirms our main finding: while changes in TFPR are minuscule, marginal costs fall substantially after export entry. Finally, we investigate whether the non-result for TFPR could be an artefact of us using a Cobb-Douglas specification in the productivity estimation in Section 2.3. In Table A.5 we estimate the more flexible translog production function, which allows for a rich set of interactions between the different inputs. We confirm our main results: there is no significant change in TFPR after export entry. We also use 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 price bias in the coefficient estimates (see Appendix D for detail). 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 analyses. Nevertheless, the bias is probably of minor importance: as shown in Table A.6, we obtain very similar results for markups and marginal costs as in the baseline specification. In the same table, we also show that our results are very similar when estimating a quantity production function for the Cobb-Douglas case. Appendix D 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, for example, that unobservables may drive both export entry and improvements in efficiency. We follow a rich literature in international 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

First, export tariff declines during our sample period are limited because Chile did not undergo major trade liberalization. On average across all destinations, export tariffs for manufacturing products fell from 10.2% in 1996 to 5.6% in 2005 (weighted by volume, with the European Union and the U.S. being the most important destinations, accounting for 28% and 19% of all exports, respectively). The average decline in tariffs is relatively small when compared to periods of trade liberalization in other countries. For example, average export tariffs for Slovenian manufacturing to the EU fell by 5.7% over a single year in 1996-97. Nevertheless, there is some meaningful variation across sectors in Chilean manufacturing that we can exploit, as illustrated in Figure 3 for 2-digit industries. For example, 'clothes and footwear' saw a decline by approximately 7 percentage points, while export tariffs for 'metallic products' fell by as little as 1 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 3 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).[38]

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 (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, our main specifications 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 alleviate this concern, we control for total plant or plant-product sales in all regressions. 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.

[38]Chilean 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 (96%) of export-entrant multi-product plants in our sample, the principal product (highest revenue) is also the one that is exported. Consequently, our main analysis continues to examine export entry at the plant-product level. To check for robustness, we also provide results at the plant level.

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:

$$E_{ijt} = \alpha_{ij} + \beta_1 \tau_{st} + \gamma_1 \ln(sales_{ijt}) + \varepsilon_{ijt} \; , \tag{12}$$

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. 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 did not undergo a major trade liberalization over our sample period, 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 (12); probit estimates would be inconsistent due to the presence of fixed effects.

Column 1 in Table 5 presents our first-stage results for export entrants – in Panel A at the product level, and in Panel B, at the plant level. For the latter, we can drop the subscript i in (12), and the export dummy E_{it} takes on value one in periods t when plant i has entered the export market. Our results imply that declining export tariffs have a strong effect on export entry within 4-digit sectors, and the corresponding first stage F-statistics are well above the critical threshold of 16.4. 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 $\hat{E}_{ijt}$:

$$\ln(y_{ijt}) = \alpha_{ij} + \beta_2 \, \hat{E}_{ijt} + \gamma_2 \ln(sales_{ijt}) + \vartheta_{ijt} \; . \tag{13}$$

Columns 2-4 in panel A of Table 5 report the second-stage results at the plant-product level. Marginal costs drop by approximately 20% upon tariff-induced export entry, and this effect is statistically highly significant (we report weak-IV robust Anderson-Rubin p-values in square brackets, based on Andrews and Stock, 2005). On the other hand, neither markups nor marginal costs change upon (predicted) export entry, confirming our results for within-plant trajectories. The plant-level results in panel B show a very similar pattern. In the appendix, we present a number of additional checks. Table A.7 shows that the reduced-form results of regressing export entry directly on tariffs show the same pattern as the 2-SLS estimates. We also show that there is no relationship between export tariffs and *domestic* sales (Table A.8). 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 revenue productivity.

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 will also raise the odds of 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. This is not the case in our within- plant/product data (see Figure 1), where marginal costs only drop in the year of export entry (see Table 2). In addition, our matching estimation is designed to absorb pre-entry productivity differences,[39] and our 2SLS results for tariff-induced export entry are unlikely to be affected by selection. Of course,

[39]Note that the drop in marginal costs in the period of export entry ($t = 0$) from Table 2 becomes small and insignificant in the matching results in Table 3.

the *sample* of export entrants itself could be selected – with more productive plants being more likely to eventually become exporters as in Melitz (2003). However, this does not affect our results, which are based exclusively on *within*-plant/product variation. In sum, selection based on *pre-exporting* productivity differences is unlikely to drive the observed within-plant-product trajectories.

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: Table 2 shows a downward trend in marginal costs during the first three years after export entry. However, the trend is less pronounced in the matching results in Tables 3 and 4. 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 fixed costs of export entry and new technology. This setup implies that initially less productive plants will require larger efficiency gains. 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 – plants that are initially less productive should experience larger changes in efficiency. This prediction can be tested in the data (Lileeva and Trefler, 2010).

Table 6 provides evidence for this effect, reporting the change in marginal costs for plant-products with low and high pre-exporting productivity.[40] We find a substantially steeper decline for initially less productive plant-products. 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.

[40] Because marginal costs cannot be compared *across* plant-products, we use pre-exporting TFPR to split plants 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 *expected* timing 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. We analyze the corresponding trends for export entrants in Panel A of Table 7. Overall investment shows an upward trend right before and shortly after export entry. Disentangling this aggregate trend reveals that it is driven by investment in machinery, but not in vehicles or structures. The evidence is thus in line with a complementarity between investment in new productive technology and export entry. The fact that investment spikes already before export entry does not conflict with this interpretation – it typically takes some time until newly purchased machinery and equipment is installed and fully integrated into the production process. In addition, the time lag suggests that (on average) export entry is planned and prepared ahead of time, while the cost trajectories documented above imply that efficiency gains coincide with export entry. Overall, our findings suggests a pattern where plant managers first decide to export and perform the necessary investments, and then begin to sell to foreign markets when technology has been updated.

Alternative Interpretations: Returns to Scale, Input Prices, and Product Quality

Economies of scale could potentially also explain declining marginal costs after export entry: if exporting goes hand-in-hand with a general expansion of production, this could raise efficiency even without targeted investment in better technology, or it could lower input prices due to volume discounts. However, our production function estimates suggest approximately constant returns to scale in most sectors – the mean sum of all input shares is 1.018 with a standard deviation of 0.047. Table A.1 in the appendix reports further details, showing output elasticities and returns to scale for each 2-digit sector in our sample. Table A.1 also shows that returns to scale are very similar when we instead estimate a more flexible translog specification.[41]

Could marginal costs fall after export entry simply because exporters purchase inputs at discounted prices? Panel B in Table 7 examines this possibility, reporting trends of the average price of all inputs, as well as for a stable basket of inputs (i.e., those that are continuously used for at least two periods before and after export entry). The table shows that input prices do not decrease after export entry; if anything, inputs become somewhat more expensive, although this trend is statistically weak.

Finally, it is 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

[41] The translog case allows for interactions between inputs, so that output elasticities 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 4 periods after entry. Both imply very similar – approximately constant – returns to scale, as shown in Table A.1.

and Verhoogen, 2012; Manova and Zhang, 2012; Atkin et al., 2014). 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. (2014) 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. In addition, there is some suggestive evidence for learning-by-exporting in the years after entry. Importantly, the 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 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 Building'?

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, so that efficiency gains are passed through to customers. One explanation is that new exporters engage in 'demand building', as described by Foster et al. (2012) – 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 identify 'young' export entrants as plant-products that have been exported for a maximum of 3 years and compare their average domestic and foreign prices before and after export entry. 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%.[42] 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

[42]To 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.9 in the appendix shows the results.

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 building' 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 suggestive evidence for substantial 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. We differentiate between 'stationary' periods with relatively constant export tariffs, and periods of trade liberalization when export tariffs fell. Most of our sample period is characterized by the former – as discussed above, Chile did not undergo major trade liberalization between 1996 and 2005. Nevertheless, there is variation in tariff changes across 4-digit sectors, which we exploit. Thus, the same limitations as described in Section 4.4 apply to the following analysis.

5.1 Export Volume and Efficiency within Plants: OLS Results

We begin by examining the relationship between export volume and plant performance in simple OLS regressions. A caveat is that export sales – a crucial variable in this analysis – are reported at the plant- but not at the product-level. Thus, we run the following analysis at the plant (i) level:

$$\ln(y_{it}) = \beta\ \ln(export_{it}) + \gamma\ \ln(sales_{it}) + \delta_i + \varepsilon_{it}\ , \tag{14}$$

where y_{it} denotes our standard outcome variables: marginal costs, markups, and TFPR.[43] The variable $export_{it}$ reflects total plant-level export revenue, and $sales_{it}$ denotes total plant-level sales; δ_i are plant fixed effects, so that all our results reflect within-plant variation. Controlling for $sales_{it}$ ensures that our results are not driven by plant size and are instead attributable to expansions of exports *relative* to overall sales.

For each specification, we report results for different subsamples of plants, according to their

[43]For multi-product plants, TFPR can be calculated with the standard procedure, 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.2. First, because our analysis includes plant fixed effects, we can normalize plant-level marginal costs and markups to unity in the last year of our sample, 2005 (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. We also show that results are very similar for single-product plants, where this computation is not needed.

overall export share. We begin with the full sample that includes all continuing exporters (i.e., all those with initial 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 it is more difficult to increase exports relative to sales.[44] 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.

We present our OLS results of estimating (14) in Table 8. There is only weak evidence that increasing export volume is associated with higher efficiency or markups. The OLS coefficients (elasticities) indicate that a doubling of exports within a plant is associated with marginal costs declining by about 4% on average, and markup and TFPR increases of a similar magnitude. The weak OLS results are not surprising. The stationary trade environment of our sample period lacks – on average – systematic incentives for existing exporters to invest in new technology. Thus, OLS results probably reflect mainly short-term responses of exports to *temporary* price or demand shocks, which are not sufficient to incentivize investment in new technology.[45]

Conversely, this argument implies that *permanent* shocks to trade costs (e.g., in the form of lower tariffs) can lead to technology upgrading (see Bustos, 2011, for empirical evidence on this mechanism). One way to illustrate this point in our data is by restricting the OLS regressions to a subsample that examines export expansions (plants with non-declining exports) during years where most sectors saw falling tariffs (2003-05; see Figure 3). Indeed, we find that in this subsample, there is a more pronounced association between export volumes and efficiency. As shown in Table A.11 in the appendix, marginal costs drop by 25-35%, while TFPR and markups increase by up to 10%. In the following, we exploit tariff changes more systematically, by using 2SLS estimation.

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

We now focus on export expansions that are driven by trade liberalization. As in Section 4.4, we use 4-digit industry tariffs to instrument for plant exports, and we include plant fixed effects. Table 9 presents our 2SLS results. The first stage coefficients in panel A show that tariffs are a strong predictor of export expansions – in particular in the subsamples that include plants with at least 20-

[44] In the extreme, for plants that export 100% of their output, export expansions are identical to sales expansions and are thus fully absorbed by including sales as a control.

[45] It is unlikely that the weak relationship between export volume and plant performance is due to the plant- (rather than product-) level data. Table A.10 in the appendix shows that results are similarly weak for single-product plants (while Table A.4 shows strong results for single-product export *entrants*).

40% export share (cols 3-5), where the first stage F-statistic significantly exceeds the critical value for a maximal 10% IV bias (16.4). Panel A also reports the predicted average log-point increase in exports due to export tariff reductions over our entire sample period (we use $\widehat{\triangle}$ to denote predicted changes). According to these estimates, trade liberalization (export tariffs falling by 4.6 p.p. on average) raised exports by up to 10 percent relative to total sales (which are controlled for in the regressions) in 1996-2005. The second-stage results show that tariff-induced export expansions lead to sizeable reductions in marginal costs (panel B). To interpret the magnitude of effects, we compute the change in marginal costs due to the overall tariff reduction over the sample period. For example, in col 3, this effect is obtained by multiplying the predicted increase $\widehat{\triangle}$ in exports from panel A (0.096) with the coefficient estimate from panel B (-1.113). We find that export tariff declines in 1996-2005 reduced marginal costs by approximately 10%. This is smaller than the decline in marginal cost associated with export entry (20-30% as reported in Table 5). If taken at face value, our results thus suggest that export entry has (on average) a stronger effect on efficiency than a moderate (10%) increase in export volume for existing exporters.

Next, we turn to the results for markups and TFPR (panel C and D in Table 9, respectively). Both variables increase statistically significantly for firms that export between 10 and 30% of their output. Nevertheless, TFPR reflects only about one half of the efficiency gains reflected by marginal costs: tariff declines over our sample period raised TFPR by 4-5%. The increase in markups is very similar, in line with our result in Section 2 that changes in markups will reflect changes in TFPR. Our results also imply that about half of the efficiency gains reflected by lower marginal costs are passed on to customers in the form of lower prices.[46]

In the appendix we present a number of consistency checks. Table A.12 shows the reduced-form results corresponding to Table 9. We confirm the 2SLS results: lower tariffs lead to significant declines in marginal costs, and to significant (but relatively smaller) increases in markups and marginal costs. Next, Table A.13 shows that while there is a strong effect of tariff declines on exports (relative to total sales), there is no clear effect on domestic sales – in fact, tariff declines are associated with somewhat lower domestic sales relative to total sales. This suggests that by controlling for total sales in all regressions, we identify a pattern that is specific to trade, and not driven by a general expansion of production. In Table A.14 we show 2SLS results when including 2-digit sector-year fixed effects. As discussed above, these soak up most of the identifying variation in tariffs in the first stage, leading to low F-statistics and thus potentially unreliable 2SLS results. Nevertheless, the second-stage results for exporters with export shares above 20 and 30% are similar in magnitude to our main results, and statistically significant (based on weak-IV robust

[46] This is implicit in Table 9, where markups increase by less than the drop in marginal costs. Statistical tests (not reported) show that the decline in *prices* is also statistically significant for cols 2-4.

p-values). In Table A.15 we show that input prices are largely unchanged following tariff-induced export expansions. Finally, Table A.16 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 some of the increased efficiency is now captured by TFPR marks an important difference to the results on export entry, where markups and TFPR remained unchanged. The core of the difference is related to pricing behavior: while new export entrants pass on most efficiency gains 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 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.

6 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 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 prices. Using a detailed Chilean plant-product level panel over the period 1996-2005, we showed that this bias is likely at work.

In order to avoid the effect of lower prices on the efficiency measure, we used marginal costs, which is directly (negatively) associated with quantity-productivity in standard production functions. We estimated marginal costs at the plant-product level following the approach by De Loecker et al. (2012). When using this measure to analyze export-related efficiency gains, we have distinguished between new export entrants and expanding foreign sales by established exporters. In addition, within the latter category, we have differentiated ‘stationary periods’ (i.e., with relatively stable tariffs) from periods of trade liberalization. For these cases, we have analyzed the relationship between exporting and efficiency – measured by marginal costs. We find that export entry and tariff-related expansions are both associated with increasing efficiency. However, this is not true for increases in export volume during 'stationary periods'. This suggests that efficiency gains occur when plants anticipate *permanent* changes in their production behavior – due to first-time export entry or durable tariff declines. Our data suggest that in these cases, firms invest in technology to improve efficiency. On the other hand, the ups and downs in export volume of established exporters during 'stationary' periods likely reflect transitory demand shocks that do not lead to technology upgrading.

We also compared the results when using marginal costs as efficiency measure to those based on the typically used TFPR. Our results suggest that actual tariff-induced efficiency gains are larger than those reflected by TFPR. For export entrants, TFPR fails to identify any gains, and for tariff-induced export expansions TFPR gains are only half the size as compared to those 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 observation are compatible with 'demand building' (Foster et al., 2012): 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.

Next, we gauge the magnitude of the observed within-plant efficiency gains after export entry, comparing the alternative 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 below 5% (Table 9). 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. However, our results imply that TFPR is an inferior measure for export-related efficiency gains. It fails to identify gains associated with export entry, and it underestimates gains related to export expansions of existing exporters. Based on marginal costs, new export entry is accompanied by efficiency increases of 15-25%. In addition, tariff-induced export expansions led to approximately 10% higher efficiency over our sample period (Table 9) – double 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.

In sum, our main finding is that a large number of previous studies has probably underestimated the effect of trade on efficiency changes within plants. Note that this result is not a foregone conclusion: In principle, TFPR could also *overestimate* actual efficiency gains – if markups rise more than efficiency. 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 significant markup increases among existing exporters when tariffs fall. This implies that the 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 an incentive 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; 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 an efficiency 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 disaggregate 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 markups and export expansions 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.

Ackerberg, D. A., K. Caves, and G. Frazer (2006). Structural Identification of Production Functions. Manuscript, UCLA.

Alvarez, R. and R. López (2005). Exporting and Performance: Evidence from Chilean Plants. *Canadian Journal of Economics 38*(4), 1384–1400.

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.

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.

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 Liberalization. In E. Helpman, D. Marin, and T. Verdier (Eds.), *The Organization of Firms in a Global Economy*. Cambridge: Harvard University Press.

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., P. K. Goldberg, A. K. Khandelwal, and N. Pavcnik (2012, August). Prices, Markups and Trade Reform. Working paper.

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

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.

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 (2012). The Slow Growth of New Plants: Learning about Demand? working paper.

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 (2014). What is 'Firm Heterogeneity' in Trade Models? The Role of Quality, Scope, Markups, and Cost.

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

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.

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.

FIGURES

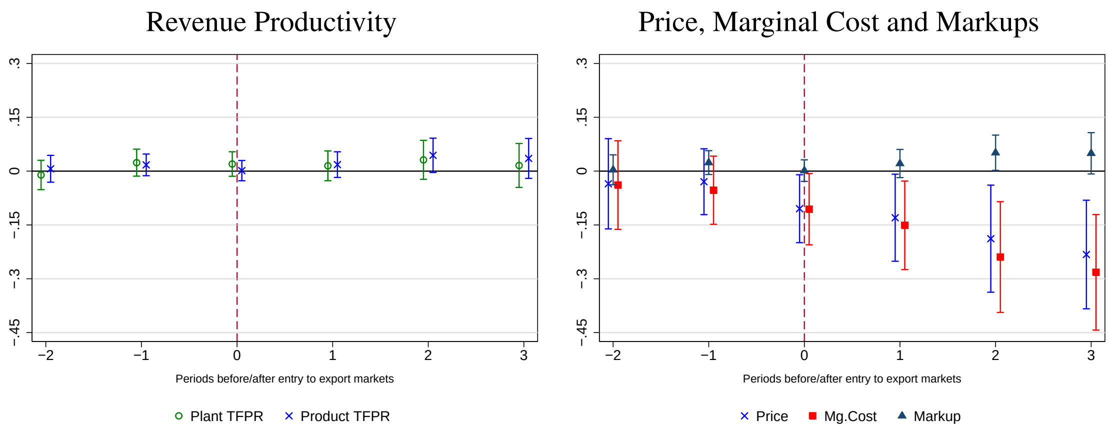

Figure 1: Trajectories for Export Entrants

Notes: The left panel shows the trajectories for revenue productivity at the plant- and at the plant-product level. The right panel shows plant-product trajectories for price, marginal cost, and markup. Period $t = 0$ corresponds to the export entry year. A 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. The trajectories are estimated using equation (11), and coefficients are reported in Table 2. Section 4.1 provides further detail.

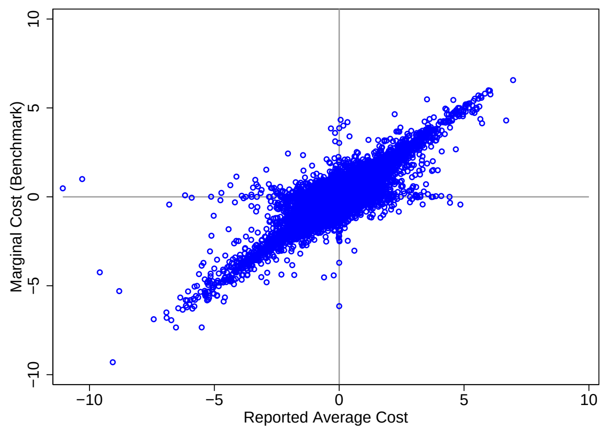

Figure 2: Estimated Marginal Cost and Reported Average Cost

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 for a detailed description). The underlying data include both exported and domestically sold products, altogether 98,688 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.

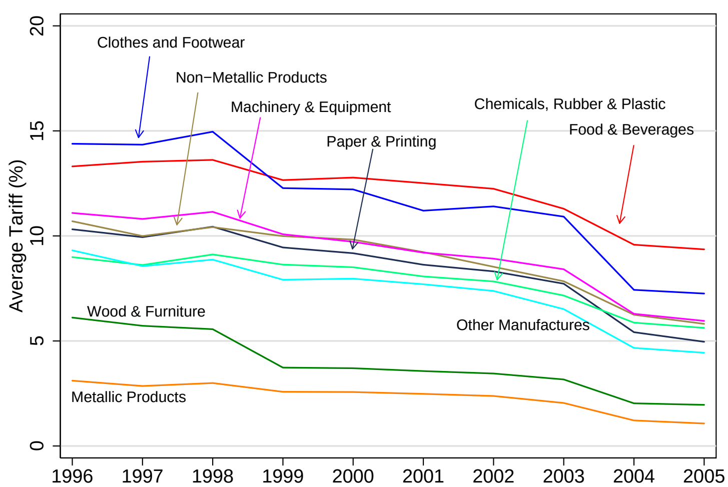

Figure 3: Average Chilean Export Tariffs (2-digit industries)

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 revenue-weighted average tariffs at the more aggregate 2-digit ISIC level.

TABLES

Table 1: Plant-Level Stylized Facts

	(1)	(2)	(3)	(4)	(5)
	Plant Size		Productivity	Wages	Markup
Dependent Variable	ln(workers)	ln(sales)	ln(TFPR)	ln(wage)	ln(markup)
Panel A: Unconditional Premia					
Export dummy	1.403***	2.227***	.133***	.402***	.0352**
	(.083)	(.176)	(.025)	(.039)	(.010)
Sector-Year FE	✓	✓	✓	✓	✓
R^2	.26	.30	.99	.24	.09
Observations	42,264	42,070	42,228	42,261	86,199
Panel B: Controlling for Employment					
Export dummy		.648***	.157***	.201***	.0354***
		(.0885)	(.0237)	(.0303)	(.010)
Sector-Year FE	✓	✓	✓	✓	✓
R^2	.26	.30	.99	.24	.09
Observations	42,264	42,070	42,228	42,261	86,199

Notes: The table reports the percentage-point difference of the dependent variable between exporting plants and non-exporters in a panel of 8,500 (4,900 average per year) Chilean plants over the period 1996-2005. All regressions control for sector-year effects at the 2-digit level; the regressions in Panel B also control for the logarithm of workers. Markups in column 5 are computed at the plant-product level. Standard errors (in parentheses) are clustered at the plant (col 1-4) and plant-product (col 5) level. Key: *** significant at 1%; ** 5%; * 10%.

Table 2: Within Plant-Product Trajectories for New Exported Products

Periods After Entry	-2	-1	0	1	2	3	Obs/R^2
Panel A: Plant Level							
TFPR	-.0108	.0236	.0198	.0149	.0313	.0157	2,029
	(.0249)	(.0230)	(.0209)	(.0253)	(.0331)	(.0374)	.624
Panel B: Product Level							
TFPR	.00645	.0174	.00140	.0181	.0440	.0354	2,309
	(.0228)	(.0184)	(.0172)	(.0218)	(.0292)	(.0339)	.541
Marginal Cost	-.0389	-.0533	-.106*	-.151**	-.240**	-.282***	2,309
	(.0753)	(.0579)	(.0605)	(.0753)	(.0941)	(.0982)	.831
Markup	.00366	.0238	.00141	.0212	.0512*	.0496	2,309
	(.0254)	(.0203)	(.0183)	(.0239)	(.0300)	(.0352)	.820
Price	-.0352	-.0295	-.105*	-.130*	-.188**	-.232**	2,309
	(.0767)	(.0560)	(.0577)	(.0740)	(.0910)	(.0922)	.831
Physical Quantities	.00376	.0249	.111	.157*	.170	.179	2,309
	(.0978)	(.0740)	(.0710)	(.0861)	(.108)	(.119)	.809
Reported Average Cost	-.00238	-.0534	-.114*	-.138*	-.183*	-.258**	2,309
	(.0771)	(.0619)	(.0599)	(.0764)	(.0982)	(.103)	.517

Notes: The table reports the coefficient estimates from equation (11). The regression for plant-level TFPR controls for plant fixed effects and sector-year effects (at the 2-digit level). All remaining 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. Standard errors (clustered at the plant level in panel A, and at the plant-product level in panel B) in parentheses. Key: *** significant at 1%; ** 5%; * 10%.

Table 3: Matching Results: Estimated Trajectories for New Exported Products

Periods After Entry	0	1	2	3
TFPR	-.0327	.0351	.0632	.0856*
	(.0233)	(.0327)	(.0439)	(.0502)
Marginal Cost	-.00354	-.162***	-.144**	-.239**
	(.0358)	(.0603)	(.0628)	(.114)
Markup	-.0355	.0317	.0686	.0789
	(.0264)	(.0384)	(.0416)	(.0604)
Treated Observations (Min/Max)	137/142	79/82	57/59	30/31
Control Observations (Min/Max)	592/612	338/353	241/249	120/125

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 C. The number of treated and control observations differ across dependent variables; the minimum (Min) and maximum (Max) number of observations are reported. Robust standard errors in parentheses. Key: *** significant at 1%; ** 5%; * 10%.

Table 4: Matching Results: Balanced Sample

Periods After Entry	0	1	2	3
TFPR	.0571	.0954	.101	.110**
	(.0401)	(.0610)	(.0640)	(.0509)
Marginal Cost	-.179*	-.318**	-.275*	-.221*
	(.0921)	(.144)	(.138)	(.129)
Markup	.0443	.0766	.0845	.0901
	(.0502)	(.0745)	(.0623)	(.0595)
Treated Observations	31	31	31	31
Control Observations	119	119	119	119

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

Table 5: Tariff-Induced Export Entry

	First Stage	Second Stage		
	(1)	(2)	(3)	(4)
Dependent Variable	Export Dummy	MC	Markup	TFPR
Panel A: Plant-Product Level				
Export Tariff	-8.600***	—	—	—
	(1.322)			
First Stage F-Statistic	42.31			
Export Dummy	—	-.209**	-.0144	.0454
		[.011]	[.812]	[.214]
For all regressions:				
Plant-Product FE	✓	✓	✓	✓
log Sales	✓	✓	✓	✓
Observations	1,761	1,761	1,761	1,761
Panel B: Plant Level				
Export Tariff	-8.084***	—	—	—
	(1.024)			
First Stage F-Statistic	62.38			
Export Dummy	—	-0.338*	-0.0679	-0.0219
		[.094]	[.294]	[.627]
For all regressions:				
Plant FE	✓	✓	✓	✓
log Sales	✓	✓	✓	✓
Observations	1,333	1,333	1,333	1,333

Notes: This table examines the effect of tariff-induced export entry on marginal costs, markups, and TFPR. In panel A we show plant-product results, while in panel B we show plant-level results. The samples include only plant-products (panel A) or plants (panel B) that become new export entrants at some point between 1997 and 2005. 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 2-4) report weak-IV robust Anderson-Rubin p-values in square brackets (see Andrews and Stock, 2005, for a detailed review). For multi-product plants in panel B, the dependent variables in cols 2 and 3 reflect the product-sales-weighted average, as described in Appendix B.2. All regressions control for the logarithm of plant sales and include plant-product (panel A) and plant (panel B) fixed effects. Standard errors are clustered at the 4-digit ISIC level, corresponding to variation in tariffs. Key: *** significant at 1%; ** 5%; * 10%.

Table 6: Differential Effect on Marginal Cost for Initially Low and High Productivity Entrants

Periods After Entry	0	1	2	3
Low Initial Productivity	.0131	-.302***	-.155**	-.296
	(.0528)	(.0926)	(.0722)	(.180)
High Initial Productivity	-.0189	-.0358	-.135	-.197
	(.0489)	(.0743)	(.0979)	(.151)
Treated Observations	142	82	59	31
Control Observations	681	471	294	112

Notes: The table analyzes heterogenous effects of export entry, depending on initial productivity. Coefficients are estimated using propensity score matching; see the notes to Table 3 for further detail. We use pre-exporting TFPR to split plant-products into above- and below- median productivity. Period $t = 0$ corresponds to the export entry year. Robust standard errors in parentheses. Key: *** significant at 1%; ** 5%; * 10%.

Table 7: Investment and Input Price Trends Before and After Entry

Period:	Before	Pre-Entry	'Young' Exp.	'Old' Exp.	Obs/R^2
Panel A: Investment					
Overall	.1131	.4051	.4426	.2916	2,612
	(.431)	(.311)	(.287)	(.425)	.54
Machinery	.2453	.5428*	.5718*	.3181	2,612
	(.432)	(.313)	(.291)	(.436)	.55
Vehicles	.0631	.0501	.0708	.0772	2,612
	(.374)	(.242)	(.230)	(.361)	.37
Structures	-.0123	.1289	-.1395	.5261	2,612
	(.422)	(.303)	(.274)	(.455)	.46
Panel B: Input Prices					
All inputs	-.151	-.0099	.190	.0558	8,078
	(.179)	(.172)	(.148)	(.200)	.44
Stable inputs	-.225	-.146	-.0171	-.00252	2,912
	(.202)	(.230)	(.210)	(.203)	.35

Notes: This table analyzes investment and input prices before and after export entry. 'Old Exp.' groups all periods beyond 3 years after export entry; 'Young Exp.' comprises export periods within 3 years or less after export entry; 'Pre-Entry' groups the two periods before entry, and 'Before' includes all periods prior to that. 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). The coefficients in each column represent the average of the different types of investment (in logs) in each respective period. 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 2. Robust standard errors in parentheses. Key: *** significant at 1%; ** 5%; * 10%.

Table 8: Existing Exporters: Export Volume and Marginal Costs, Markups, and TFPR

	(1)	(2)	(3)	(4)	(5)	(6)
Export Share	>0%	>10%	>20%	>30%	>40%	>50%
Panel A: log Marginal Cost Index						
log Exports	.00939	-.0285	-.0535	-.0858	-.0618	-.0398
	(.0229)	(.0562)	(.0873)	(.0936)	(.122)	(.189)
R^2	.937	.959	.960	.964	.965	.964
Panel B: log Average Markup Index						
log Exports	-.00498	.0312	.0420	.0747	.0591	.0217
	(.00886)	(.0323)	(.0477)	(.0673)	(.0842)	(.111)
R^2	.703	.703	.688	.691	.695	.685
Panel C: log TFPR						
log Exports	-.00712	.0333	.0667*	.114**	.103	.0535
	(.00775)	(.0286)	(.0398)	(.0544)	(.0724)	(.0958)
R^2	.923	.914	.907	.904	.899	.893
For all regressions:						
Plant FE	✓	✓	✓	✓	✓	✓
log Sales	✓	✓	✓	✓	✓	✓
Observations	4,026	2,372	1,901	1,666	1,456	1,267

Notes: This table examines the within-plant correlations between export volume and marginal costs (panel A), markups (panel B), and TFPR (panel C). For multi-product plants, the dependent variable in panels A and B reflect the product-sales-weighted average, as described in Appendix B.2. 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. All regressions control for the logarithm of plant sales and include plant fixed effects. Standard errors (clustered at the plant level) in parentheses. Key: *** significant at 1%; ** 5%; * 10%.

Table 9: Tariff-Induced Export Expansions and Plant-Level Outcomes – 2SLS

	(1)	(2)	(3)	(4)	(5)	(6)
Export Share	>0%	>10%	>20%	>30%	>40%	>50%
Panel A: First Stage: Plant-level log exports						
Export Tariff	-.735	-1.521**	-2.087***	-1.771***	-1.345***	-.917***
	(.915)	(.627)	(.436)	(.273)	(.347)	(.289)
$\widehat{\triangle}$ Exports '96-'05	.0338	.0700	.0960	.0815	.0619	.0422
First Stage F-Statistic	.645	5.876	22.87	42.20	14.99	1.07
Panel B: Second Stage, log Marginal Cost Index						
log Exports (predicted)	-2.153*	-1.297***	-1.113***	-1.170***	-1.141**	-.564
	[.0766]	[.0016]	[.0006]	[.0007]	[.0166]	[.471]
$\widehat{\triangle}$ MC '96-'05	-.0728	-.0907	-.1068	-.0953	-.0706	-.0238
Panel C: Second Stage, log Average Markup						
log Exports (predicted)	.237	.568**	.478**	.576***	.477	-.364
	[.678]	[.0222]	[.0178]	[.0050]	[.152]	[.531]
$\widehat{\triangle}$ Markup '96-'05	.0080	.0398	.0459	.0469	.0295	-.0153
Panel D: Second Stage, log TFPR						
log Exports (predicted)	.678	.613**	.456**	.590***	.571*	.126
	[.139]	[.0108]	[.0371]	[.0102]	[.0583]	[.854]
$\widehat{\triangle}$ TFPR '96-'05	.0229	.0429	.0438	.0481	.0353	.0053
For all regressions:						
Plant FE	✓	✓	✓	✓	✓	✓
log Sales	✓	✓	✓	✓	✓	✓
Observations	4,026	2,372	1,901	1,666	1,456	1,267

Notes: This table examines the effect of within-plant export expansions due to falling export tariffs on marginal costs (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 results of these 2SLS regressions are reported in panel A, 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 report weak-IV robust Anderson-Rubin p-values in square brackets (see Andrews and Stock, 2005, for a detailed review). For multi-product plants, the dependent variables in panels B and C reflect the product-sales-weighted average, as described in Appendix B.2. Export tariffs vary at the 4-digit ISIC level. All regressions control for the logarithm of plant 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, $\widehat{\triangle}$ denotes the predicted change in the corresponding dependent variable due to export tariff reductions over our entire sample period (4.6 p.p. on average over the period 1996-2005).