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, Ariel Burstein, Leonardo Bursztyn, Donald Davis, Jan De Loecker, Gilles Duranton, Eduardo E. Engel, Edward Leamer, Steve Redding, John Ries, Aaron Tornell, RomainWacziarg, and Pierre-OlivierWeill, as well as seminar audiences at Pontificia Universidad Católica de Chile, RMET 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 2013
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 nonresult is derived from revenue productivity, thus also reflecting prices. Using a census panel of Chilean manufacturing plants, we first confirm the non-result for revenue productivity. We then compute plant-product level marginal cost as an efficiency measure that is not affected by prices. We find within-plant efficiency gains of 15-25%, the same order of magnitude as selection effects across plants. Evidence suggests that technology upgrading in combination with export entry is an important driver behind these gains.

Alvaro Garcia Marin
UCLA Anderson School of Management
110 Westwood Pl
Los Angeles, CA 90095
alvaro.garcia.2013@anderson.ucla.edu

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

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

1 Introduction

While exporting plants are on average significantly more productive than their non-exporting counterparts, empirical studies typically find that new exporters do not increase their productivity over time. This suggests that selection of the most productive plants into exporting, rather than efficiency gains within plants, is responsible for aggregate productivity gains from trade competition. The selection effect across plants has received strong theoretical and empirical support (c.f. Melitz, 2003; Pavcnik, 2002). On the other hand, within-plant productivity gains after export entry are typically found to be small and insignificant (c.f. Clerides, Lach, and Tybout, 1998; Bernard and Jensen, 1999; Wagner, 2007, 2012).[1] This non-result 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.

In this paper, we show that the missing evidence on within-plant efficiency gains after export entry is an artefact of the measure: previous studies have typically used revenue-based productivity, which is affected by changes in prices. If 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] Addressing this caveat by measuring physical productivity is difficult. For example, changes in product quality make physical units of output incomparable – even within products from the same plant. Thus, meaningful results can only be derived for physically homogenous products (Foster et al., 2008) – a small subset of all exported goods. To bypass this issue, we first apply the method pioneered by De Loecker and Warzynski (2012) to derive plant-product level markups in a rich panel of Chilean establishments. Second, 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.[3] This procedure is flexible with respect to the underlying price setting model and the functional form of the production function (e.g., allowing for different degrees of returns to scale). In standard production functions, marginal costs are directly (inversely) related to physical productivity, and are thus a good candidate for analyzing within-plant

---

[1] Two exceptions – albeit in less representative settings – are De Loecker (2007) for Slovenia and Van Biesebroeck (2005) for sub-Saharan Africa.

[2] Recent evidence suggests that this downward bias also affects the link between trade and productivity. Smeets and Warzynski (2013) use a firm level price index to deflate revenue productivity and show that this correction yields significantly 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.

[3] De Loecker, Goldberg, Khandelwal, and Pavcnik (2012) use the same methodology to analyze how trade liberalization in India affected prices, markups, and marginal costs.

efficiency gains after export entry.[4]

We find that gains from exporting are substantial: marginal costs within plant-product categories drop by approximately 15-25% during the first three years after export entry. At the same time, in line with previous findings, *revenue* productivity does not change within exporting plants. This is due to prices falling by a similar magnitude as marginal costs – new exporters pass physical productivity gains on to their customers.[5] Our results are very similar when using 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 (average) cost measures at the plant-product level. This suggests that our findings are not an artefact of the methodology used to calculate marginal costs. Finally, we show that there are no efficiency gains for non-exported products in multi-product plants that enter the international market. Thus, we are confident that our results are not driven by shocks or trends at the plant level that are unrelated to exporting.[6]

To guide the discussion of possible drivers behind our results, we provide a stylized framework that combines the flexible supply-demand structure from Foster et al. (2008) with heterogeneous returns to technology investment as in Lileeva and Trefler (2010). We discuss four channels that may drive export entry: (i) shocks to foreign demand, (ii) productivity shocks, (iii) anticipated learning by exporting, and (iv) investment opportunities in new technologies that become profitable in combination with access to larger markets. Since we find *falling* prices associated with export entry, demand shocks (i) are an unlikely driver.[7] On the other hand, the supply-side mechanisms (ii)-(iv) are all broadly compatible with our empirical observations. However, they imply different causal effects. Productivity shocks as in (ii) mean a selection effect – firms enter the export market as a consequence of higher productivity, as in Melitz (2003), and causality runs from productivity to export entry. The opposite is true in case (iii), where plants enter the export market because they anticipate learning effects. Finally, case (iv) reflects a complementarity between efficiency-enhancing investment and export entry – the fixed cost can only be recovered in a large-enough

---

[4] Marginal costs are not immune to changes in product quality. However, the associated bias is likely to work against finding efficiency gains: exported goods from developing countries are typically of higher quality (c.f. Verhoogen, 2008), which should *raise* marginal costs.

[5] In other words, for an average plant-product, markups do not change significantly during the first years following export entry. Over the medium run, however, we find that export entrants raise their markups, in line with De Loecker and Warzynski (2012) who document increasing markups for Slovenian export entrants. However, our data suggest that this effect is limited, making it unlikely that the efficiency gains observed in marginal costs will be fully reflected in revenue productivity.

[6] This also suggests that within-plant efficiency spillovers to non-exported goods are probably limited.

[7] Increasing returns or falling input prices are also unlikely drivers of the observed drop in marginal cost. Our production function estimates suggest approximately constant returns, and input prices do not change significantly after export entry.

market.[8]

We provide some suggestive evidence that learning-by-exporting and investment complementarity are the most likely drivers of our results. To address selection into exporting after a productivity shock (channel ii), we follow the matching approach by De Loecker (2007), which controls for pre-exporting differences in productivity levels and trends, as well as other characteristics. The fact that our results are robust to this methodology suggests that they are probably not driven by productivity shocks *before* export entry. This leaves (iii) and (iv) as more probable mechanisms. The observation that marginal costs keep falling in the years *after* entry is in line with learning-by-exporting (iii). In addition, several pieces of evidence suggest that investment complementarity (iv) is also important. First, export entry goes hand-in-hand with a decline in marginal costs already in the entry period, which is compatible with a switch to more efficient production modes upon export entry. Second, 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 productivity gains are large. Third, we show that plant-level investment (especially in machinery) spikes immediately before, and during the first years, of export entry.

Our analysis is subject to two important caveats. First, we do not establish causality. While our matching estimation takes a step in this direction, it does not use exogenous variation. Second, we cannot fully disentangle the three supply mechanisms (ii)-(iv). For example, if productivity shocks (ii) occur suddenly (i.e., are not preceded by pre-trends), and if export entry occurs immediately (in the same period as the shock), then we cannot differentiate between (ii) and (iv). Similarly, if learning-by-exporting lowers marginal costs immediately when export entry occurs, then we cannot differentiate between (iii) and (iv). Nevertheless, our main result does not hinge on causality or which exact mechanism is at play: we document substantial within-plant efficiency gains associated with export entry, and these gains are not identified by the commonly used revenue-based productivity measures.

We observe that new exporters pass on most efficiency gains to customers in the form of lower prices, which is accompanied by a strong increase in quantity: within the first three years after entry, the price of the exported product falls by 10-20%, and quantity sold increases by the same margin. In a subset of years with more detailed pricing information, we separately analyze the domestic and export price of the same product. For 'young' export entrants (maximum three years

---

[8] Strictly speaking, mechanism (iv) is not causal, because the investment in new technology would have raised productivity regardless of export status. However, since the investment is not profitable in the domestic market alone, export entry and productivity increases are closely associated.

of exporting), we find that the export price drops significantly more than its domestic counterpart (22% vs. 7%). One explanation for this behavior is ’demand capital’ building, as in the model of Fishman and Rob (2003).[9] Foster, Haltiwanger, and Syverson (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. Applied to export entrants, ’demand building’ would imply lower prices charged to attract foreign buyers, as we observe in the data.

Finally, we gauge the magnitude of the observed within-plant efficiency gains after export entry, comparing them to the well-documented exporter revenue-productivity premium in the cross-section. This premium is about 13% in our sample, which is very similar to the exporter premium reported for other countries (c.f. Bernard and Jensen, 1999). When focusing exclusively on ’young’ exporters, we find a very similar figure. In other words, revenue productivity does not *change* after export entry. This suggests that selection in the spirit of Melitz (2003) is the driving force behind the exporter revenue-productivity premium. On the other hand, the drop in marginal costs that we identify reflects efficiency gains *in addition* to the typically documented selection effect. Our results suggest that these within-plant gains are of a similar magnitude as the between-plant differences.[10]

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.[11] In contrast, another prominently discussed channel has received astonishingly little empirical support: on balance, exporting does not appear to have important effects on productivity *within* firms or plants. Such gains can be expected because exporters face tougher competition, have stronger incentives to innovate since they serve a larger market, and because they have access to expertise from international buyers (Grossman and Helpman, 1991).[12] Clerides et al. (1998, for Colombia, Mexico, and Morocco) and

---

[9] When consumers have different search costs, Fishman and Rob’s (2003) model implies that low-cost firms charge low prices in order to attract more flexible (low search cost) customers who currently buy from high-price firms.

[10] If new exporters raise their markups over time, the marginal-cost based efficiency gains will also be reflected in revenue productivity. We show that this effect is limited and use this observation to gauge overall magnitudes.

[11] 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.

[12] 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

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.[13] 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 exception are the papers by Van Biesebroeck (2005) and De Loecker (2007), which document evidence for learning-by-exporting based on revenue-productivity. Both derive their results in potentially unrepresentative environments: Sub-Saharan Africa and Slovenia during its transition from communism to a market economy.[14] Finally, Smeets and Warzynski (2013) show that controlling for pricing heterogeneity yields larger export-related productivity gains. Using price indexes to deflate revenue-based productivity measures, they correct for the bias that arises when more productive firms charge lower prices. However, in contrast to our marginal cost based approach, this price-based methodology does not disentangle the behavior of markups and efficiency after export entry; it also analyzes the productivity of firms overall, while we focus on the exported product itself and show that spillovers to non-exported products within multi-product plants are probably limited.

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 pricing behavior, and to document a strong decline in marginal costs after export entry. We also show that the usual non-result for revenue productivity holds in our panel of Chilean plants. Thus, a second contribution of this paper is to point to substantial export-related efficiency gains that have thus far passed under the radar. Third, we shed light on possible drivers of the observed efficiency gains. The evidence points towards a complementarity between export entry and investment in technology. Learning by exporting probably also plays a role. Fourth, we show that export entry affects markups with some delay in the context of our Chilean plant data: during the first 2-3 years

---

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. In a more systematic fashion, Bustos (2011) shows that rising export revenues – driven by exogenous changes in tariffs – foster firms' investment in new technology.

[13] 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.

[14] In Van Biesebroeck's findings, exporting lifts credit constraints and thus allows sub-Saharan African firms to grow and profit from scale economies. Syverson (2011) points out that these results may reflect heterogenous treatment effects, with firms that gain most from scale economies sorting into exporting.

of exporting there is nearly complete pass-through of efficiency gains to customers. Thereafter, markups grow as exporters become more established.[15] The observed trajectory of markups is in line with our explanation that initially, export entrants seek to attract customers by charging low prices. This interpretation is also supported by our observation that – for newly exported products that are also sold domestically – the price drop is particularly steep in the international market. As their 'demand stock' grows, exporters begin to raise markups, as argued by Foster et al. (2012) in the context of domestic market entry. 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 empirical results. Section 5 discusses possible mechanisms that may drive the observed efficiency gains. Section 6 concludes.

2 Empirical Framework

In this section, we discuss our efficiency measures and explain how we compute them. Our first measure of efficiency is plant-level *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 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 measure of efficiency.

2.1 Revenue vs. Physical Total Factor Productivity

Revenue-based total factor productivity is the most widely used measure of efficiency. It is calculated as the residual between total revenues and the estimated contribution of production factors (labor, capital, and material inputs).[16] This measure has an important shortcoming, which can be

---

[15] The delayed increase in markups after export entry is more pronounced in our findings than in 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.

[16] Some authors have used labor productivity – i.e., revenues per worker – as a proxy for efficiency (for recent surveys see Wagner, 2007, 2012). This measure is affected by the use of non-labor inputs and is thus inferior to TFP

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 prices and A, so that TFPR will underestimate physical productivity.

Given these shortcomings, why has the literature not used physical productivity to analyze productivity gains from exporting? One practical caveat is the lack of information on physical quantities.[17] While some corrections to the estimation of production functions have been proposed, only a few studies have derived A directly.[18] 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. They find a negative correlation between prices and A. This is consistent with more efficient businesses having lower marginal costs and, in turn, charging lower prices. As a consequence, changes in TFPR understate true efficiency gains.

Even if quantities are known so that A can be calculated, the measure is problematic. Product quantities cannot readily be compared because quality may change. As Foster et al. (2008) recognize, it is essentially impossible to isolate changes in quality from A. This is the reason why these authors restrict their analysis to a set of homogeneous products (e.g., concrete and gasoline) that are arguably not subject to significant changes in quality. An additional problem emerges for multi-product plants, where disaggregate use of inputs is typically not reported for individual products. Thus, A has to be computed at the plant level, which requires the aggregation of quantities. This is problematic because goods produced by multi-product plants often differ in their physical and functional attributes. For example, if a furniture manufacturer produces both tables and chairs, the sum of the two does not provide a meaningful index of quantity. Also, products are usually measured in different units, and the correspondence between these is often non-trivial. To circumvent these issues, we propose marginal cost as a measure of efficiency. Since we compute marginal cost at the product level, we avoid the aggregation issues for multi-product plants. In addition, we

---

when different plants combine inputs in different proportions (see Syverson, 2011).

[17] 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).

[18] 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. Katayama, Lu, and Tybout (2009) show that revenue-based output can lead to productivity mismeasurement and incorrect interpretations of how heterogeneous producers respond to shocks. Hsieh and Klenow (2009) recover A using a model of monopolistic competition for India, China and the United States.

analyze trends in marginal costs rather than levels, which allows comparisons across producers of different products. In the following 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 reflect 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):

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

Deriving percentage changes (denoted by $\triangle$) and re-arranging yields a relationship between efficiency gains and changes in TFPR, markups, and marginal costs:[19]

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

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, 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 5 that these do not change significantly after export entry. In addition, there is reason to believe that, if anything, assuming constant input prices implies a downward bias in efficiency gains inferred from $\triangle MC(\cdot)$. More successful exporters typically produce high-quality goods that 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

---

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

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 production function. We follow 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, specifying an endogenous productivity process, where past export-status is allowed to impact current productivity. This reflects the correction suggested by De Loecker (2013);

---

[20] We follow LP in using material inputs to control for the correlation between input levels and unobserved productivity. Our approach for estimating the production function is explained in detail in Appendix B.

[21] The main technical difference is the timing of 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.

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).

We estimate a translog production function with labor (l), capital (k), and materials (m) as production inputs. While the translog specification nests the typically used Cobb-Douglas production function, it is more flexible, allowing for varying degrees of economies of scale, as well as complementarities between the inputs. We follow De Loecker et al. (2012) in using the subset of single-product plants to estimate the following function within 2-digit product categories:[22]

$$\begin{aligned} q_{it} \quad = \quad & \beta_l l_{it} + \beta_k k_{it} + \beta_m m_{it} + \beta_{ll} l_{it}^2 + \beta_{kk} k_{it}^2 + \beta_{mm} m_{it}^2 + \beta_{kl} l_{it} k_{it} \\ & + \beta_{mk} m_{it} k_{it} + \beta_{lm} l_{it} m_{it} + \beta_{lmk} l_{it} m_{it} k_{it} + d_{it}^x + \omega_{it} + \varepsilon_{it} \end{aligned} \tag{3}$$

where all lowercase variables are in logs; q_{it} are revenues of plant-product i in year t, d_{it}^x is an export dummy, ω_{it} is plant-level productivity, and ε_{it} represents measurement error as well as unanticipated shocks to output.[23] Given the estimated coefficients for each product category s, revenue productivity is computed as

$$\hat{\omega}_{it} = \hat{q}_{it} - \hat{f}_s(k_{it}, m_{it}, l_{it}) \tag{4}$$

where $\hat{f}_s(\cdot)$ represent the estimated contribution of the production factors to total output in s. Note that the estimated production function allows for returns to scale, so that the residual $\hat{\omega}_{it}$ is not affected by increasing or decreasing returns at the sector level. When computing TFPR in multi-product plants, we use the f_s that corresponds to the product category s of the predominant product produced by plant i.

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.

---

[22] The reason for using single-product plants is that we do not observe how inputs are allocated across outputs within a plant, which makes the estimation of (3) at the product level unfeasible for multiple-product plants. For the set of single product plants, no assumption on the allocation of inputs to outputs is needed, and the estimation of (3) can be performed with standard plant level information. 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.

[23] We include an export dummy in the production function following De Loecker and Warzynski (2012). This allows exporters to produce under a different technology.

The methodology for deriving marginal costs follows the production approach proposed by Hall (1986), recently revisited by De Loecker and Warzynski (2012). This approach computes markups without relying on detailed market-level demand information; it only requires standard plant-level data on input use and production output. 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:[24]

$$\underbrace{\mu_{ijt}}_{Markup} \equiv \frac{P_{ijt}}{MC_{ijt}} = \underbrace{\left(\frac{\partial Q_{ijt}(\cdot)}{\partial V_{ijt}}\frac{V_{ijt}}{Q_{ijt}}\right)}_{\text{Output Elasticity}} \Big/ \underbrace{\left(\frac{P^V_{ijt}\cdot V_{ijt}}{P_{ijt}\cdot Q_{ijt}}\right)}_{\text{Revenue Share}}, \tag{5}$$

where P (P^V) denotes the price of output Q (input V), and MC is marginal cost. According to equation (5), the markup can be computed by dividing the output elasticity of product j (with respect to the flexible input) by the share of the flexible input cost in the sales of product j. We use materials as the flexible input to compute the first component in (5) – the output elasticity – based on our estimates of (3).[25] The second component needed in (5) – the expenditure share for material inputs – is observed in our data.[26] Because markups are computed at the plant-product level, and prices (unit values) are observed at the same level, we derive marginal costs at the plant-product level in each year.[27] Appendix C provides further detail on the estimation of marginal costs.

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, using to the International Standard Industrial Classification (ISIC), revision 2. It contains detailed infor-

---

[24] More precisely, the first order condition with respect to V is $\frac{\partial \mathcal{L}}{\partial V} = P^{\nu} - \lambda \frac{\partial Q(\cdot)}{\partial V} = 0$, where the Lagrange multiplier λ equals the marginal cost of production. Manipulating this expression yields (5).

[25] 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).

[26] To derive the expenditure share of material inputs in multi-product plants, we follow Foster et al. (2008) in assuming that plants allocate their inputs proportionately to the share of each product in total revenues.

[27] In contrast, TFPR is computed at the plant level. The reason is that TFPR is derived using data on inputs, which are only available at the plant level, while marginal costs follow from prices and markups, which can be computed or observed at the product level within each plant.

mation 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 cost of production, the number of units produced and sold, and the fraction of production that is exported. 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.[28] 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, investment, demand for raw materials, sales, or product quantities. Second, whenever our analysis involves quantities of production, we have to carefully account for possible changes in the unit of measurement. For example, wine producers change in some instances from "bottles" to "liters." Total revenue is generally unaffected by these changes, but the derived unit values (prices) have to be corrected. This procedure is needed for about 1% of all plant-product observations; it is explained in Appendix D. 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 D for detail). After these adjustments, our sample consists of 109,210 plant-product-year observations.

---

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

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.[29]

3.3 Validity of the Sample

Before turning to our empirical results, we check whether our data replicate previously established stylized facts – main differences between exporters and non-exporters. Following Bernard and Jensen (1999), we run the regression

$$\ln(y_{ist}) \quad = \quad \alpha_{st} + \delta \; d_{ist}^{exp} + \gamma \; \ln(L_{ist}) + \varepsilon_{ist} \; , \tag{6}$$

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.[30] 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: within their respective sectors, exporting plants are more productive (measured by revenue productivity), larger both in terms of employment and sales, pay higher wages, and are more capital intensive. 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 6, we also find that markups are higher among exporters, confirming the findings in De Loecker and Warzynski (2012).

---

[29] Note that there are no export entries in the first year of our sample, 1996, because we do not observe domestic sales or export volumes prior to that date.

[30] Throughout the empirical analysis, 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 effects.

[Insert Table 1 here]

4 Empirical Results

In this section we present our empirical results. We first show the dynamics of revenue productivity and marginal costs within plants. Our main finding is that TFPR does not change after export entry, while marginal costs drop substantially. We then use propensity score matching to address selection and pre-exporting trends. As a consistency check, we also use reported expenditure data to calculate average costs at the plant-product level. Finally, we check the robustness of our results.

4.1 Within Plant Trajectories

We begin by analyzing the trajectories for price, marginal cost, markups and revenue productivity for the sub-sample of new export entrants. For each plant i producing good j in period t, we estimate the following regression:

$$y_{ijt} = \alpha_{st} + \alpha_{ij} + \underbrace{\sum_{k=-2}^{-1} T_{ijt}^{k}}_{Pre-Trend} + \underbrace{\sum_{l=0}^{L} E_{ijt}^{l}}_{Entry-Effect} + \varepsilon_{ijt} \ , \tag{7}$$

where y_{ijt} refers to the characteristic of product j – either price, marginal cost, markup, or (plant-level) 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).[31] 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 entry.[32]

Figure 1 visualizes the results of estimating (7) for the sub-sample of export entrants. The figure shows the point estimates for each outcome variable, together with whiskers representing the 90% confidence interval. Time on the horizontal axis is normalized to zero for the entry period. The left panel of the figure shows that TFPR within plants is virtually unaffected by exporting, with tight confidence intervals around zero (see Table 2 for the corresponding estimates). This result

---

[31] Sector fixed effects at the 4-digit level correspond to approximately 200 aggregate product categories. For TFPR, the product index j in y_{ijt} is irrelevant since revenue productivity is computed at the plant level. We include sector-year fixed effects at the 2-digit level for TFPR regressions (see footnote 30).

[32] 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.

is in line with the previous literature: there are no apparent efficiency gains when TFPR is used as a measure of efficiency. 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 (reported in Table 2) 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 25%. These differences are not only economically but also statistically significant.

[Insert Figure 1 and Table 2 here]

Table 2 reports the corresponding coefficients. The trajectory for prices is very similar to marginal costs, while changes in markups are close to zero with small standard errors. Physical quantities sold increase by approximately 20% after export entry. Finally, there is 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 next section we address this issue.

4.2 Matching Results

We apply propensity score matching (PSM) in the spirit of Rosenbaum and Rubin (1983), and further developed by Heckman, Ichimura, and Todd (1997). We attempt to isolate the causal effect of exporting 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). Once such a control group is 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 lagged marginal cost before export entry (MC_{t-1}) of the product, lagged plant-level TFPR, the capital stock of the plant (k_{t-1}), as well as a pre-trend in marginal costs (ΔMC_{t-1}), and a vector of other controls (Z_{t-1}).[33] Appendix E provides further detail. Once we have determined the control

---

[33] Other controls include product sales, number of employees, and import status of the plant. Following Abadie,

group, we use the difference-in-difference (DID) methodology to evaluate the impact of exporting on TFPR, prices, marginal cost, and markups. As Blundell and Dias (2009) suggest, using DID can improve the quality of matching results because initial differences between treated and control units are removed.

Table 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 counterfactuals, relative to the pre-entry period ($t = -1$).[34] These results confirm the within-plant pattern documented above: changes in TFPR after export entry are quantitatively small and statistically insignificant for most periods; the same is true for markups.[35] Price and marginal cost, on the other hand, both 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 price (marginal cost) relative to the control group grows from 4% (3%) in the period of export entry, to 20% (17%) after two years, and to more than 28% (17%) three periods after entry.

[Insert Table 3 here]

4.3 Robustness and Additional Results

In this subsection we address potential concerns about the validity of our results. We test if they are driven by the estimation of marginal costs, or by products exiting export markets. We also perform a check using domestically sold products of export entrants.

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.[36] We can address this concern by using data reported in ENIA, which allow us 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

---

Drukker, Herr, and Imbens (2004), we use the three nearest-neighbors. Our results are very similar when using 1 or 5 nearest neighbors instead. The difference in means of treated vs. controls are statistically insignificant for all matching variables in $t = -1$.

[34] 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 exporters, as compared to the non-exporting control group.

[35] The increase in TFPR after three periods goes hand-in-hand with higher markups. This suggests that, eventually, efficiency gains are partially reflected in TFPR (although marginal costs continue to fall). We discuss this in detail in Section 5.3.

[36] 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.

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 should thus be a reasonable proxy for marginal 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. This validates reported average costs as an alternative measure of efficiency. In addition, it also 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.

[Insert Figure 2 here]

Row 5 of Table 3 shows that average costs decrease after export entry, closely following the trajectory that we identified for marginal cost.[37] Export entry is followed by a decline in average costs of 7% in the period of entry, growing to 14% after one year, and to 19% three periods after entry. Table A.2 in the appendix shows that *within* plant-product estimates for reported average costs are also similar. These results confirm that the efficiency gains we documented based on marginal costs are not an artefact of the computation of this measure.

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 20% on average). 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 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. Compatible with this interpretation, the drop in costs and prices is more stable over time in the balanced sample. In sum, the results from the balanced sample confirm our full sample estimates and suggest more stable efficiency gains over time.

[Insert Table 4 here]

---

[37] The only exception is 2 periods after export entry, when the average cost coefficient is not statistically significant.

Results for Domestic Products

Next, we perform a placebo test of our results, analyzing marginal costs of domestically sold products that are produced by multi-product export entrants.[38] Specifically, we compare the non-exported products of new export-entry plants to similar goods produced by other non-exporting plants.[39] Table 5 shows that the differences in marginal costs are small and statistically insignificant. This suggests that our results are specific to the exported products, and are probably not driven by more general efficiency trends at the plant level.[40]

[Insert Table 5 here]

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 three nearest neighbors. Table A.3 shows that using either one or five neighbors instead does not change our qualitative results. Another potential concern is that we compute TFPR at the plant level, while marginal costs are calculated for plant-products. To show that this distinction does not drive our results, Table A.4 uses only the subset of single-product plants. The trajectory of marginal costs is very similar to its counterpart in the full sample, albeit measured with larger standard errors due to the substantially reduced sample size. Finally, we estimated the production function in Section 2.3 using revenues as a measure of output. If output prices are correlated with unobserved demand shocks, this may introduce a bias in the estimated output elasticities, which in turn would affect the computation of markups and marginal costs (see Appendix B.1 for a detailed discussion). One remedy is to use a (more restrictive) Cobb-Douglas production function, where the bias affects only the level, but not the trend of the output elasticities (De Loecker and Warzynski, 2012). Thus, in the Cobb-Douglas case the bias is less relevant for our main results, because we exploit trends rather than levels. Table A.5 shows that our results for markups and marginal costs are almost identical for the baseline translog and the Cobb-Douglas case, suggesting that the omitted price bias is probably not important. Appendix F discusses the additional robustness checks in greater detail.

---

[38] Recall that, by our definition, only the first exported product of each plant qualifies as export entry.

[39] The control group is selected by propensity score matching within the same 4-digit categories.

[40] This result also suggests that spillovers of export-related efficiency gains across products are probably limited. In fact, if our ’placebo test’ had delivered strong declines in marginal costs of domestically sold products, this could have been due to both spillovers or spurious trends at the plant level.

5 Interpretation of Results: Possible Channels and Magnitude

In this section, we discuss possible channels that may drive the observed time-pattern of prices and marginal cost after export entry. We also separately analyze the trajectories of foreign and domestic prices charged for the same good after export entry. Finally, we shed light on the magnitude of within-plant efficiency gains, relative to the well-known cross-sectional productivity advantage of exporters.

5.1 A Stylized Framework of Export Entry and Within-Plant Efficiency Gains

To guide the discussion, we sketch a stylized theoretical framework that distinguishes between various possible drivers of export entry. We marry two models from the trade and productivity literature: first, to differentiate between idiosyncratic technology and demand effects, we build on the framework by Foster et al. (2008). Second, in order to further differentiate between alternative supply-side channels, we combine this setup with the model by Lileeva and Trefler (2010). In particular, export entry can be affected by initial productivity differences (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; Bustos, 2011). In addition, anticipated learning-by-exporting will also raise the odds of export entry. Finally, in the stylized framework, higher efficiency reduces marginal costs, but has an ambiguous impact on TFPR due to falling prices. This reflects our discussion above that marginal costs are a more appropriate measure for efficiency gains. Appendix G provides a detailed exposition of the theoretical framework. Here, we focus on its central elements and explain the intuition.

For ease of exposition, we assume that each plant produces one product. Profits of plant i depend on market size, ideosyncratic product demand, and marginal cost, as in Foster et al. (2008). Export entry requires the payment of a fixed cost F_E (in annualized terms). Firms enter the foreign market if the additional profits due to exporting exceed the fixed cost. In order to analyze the (bi-directional) relationship between export entry and productivity, we define the export entry wedge ε_i. This variable indicates how far (in percentage terms) plant i's marginal cost is from the export threshold – where annual profits equal F_E. When $\varepsilon_i > 0$, plant i sells only domestically; and plants with $\varepsilon_i \leq 0$ enter the export market. Given this setup, we analyze the conditions under which export entry occurs.

Demand-driven export entry

We begin by analyzing demand-side effects. If foreign demand for plant i's product rises, ε_i falls because exporting is becoming more profitable. If ε_i falls below zero, plant i begins to export in response to the demand shock. Empirically, if demand shocks are responsible for export entry, we should see no change in the product-specific marginal costs, while sales would increase and markups would tend to rise.[41] 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. Similarly, it is unlikely that quality upgrading of exporters is responsible for our results, since higher product quality is associated with higher prices and production costs (c.f. Kugler and Verhoogen, 2012; Manova and Zhang, 2012).

Next, we analyze the supply side, where actual or anticipated changes in marginal costs can induce export entry. We use $\varphi_i > 0$ to denote the percentage drop in marginal cost due to one of three effects: a productivity shock (PS), learning-by-exporting (LBE), and technology-exporting complementarity (TEC). The marginal cost after each effect is given by $MC_i^{post} = MC_i^0/(1+\varphi_i)$, where MC_i^0 is the initial value. We discuss the three cases in the following, always starting with a plant i that is currently not exporting, facing an export entry wedge $\varepsilon_i > 0$.

Productivity shock

Suppose that plant i is affected by a productivity shock that is unrelated to exporting and reduces marginal cost by φ_i^{PS} percent. Export entry will occur if $\varphi_i^{PS} \geq \varepsilon_i$, i.e., if the productivity shock is sufficiently strong to push marginal costs below the entry threshold. In this case, causality runs from productivity to export entry, reflecting self-selection of more productive firms into exporting. If there is a time lag between shock and entry, the data will show efficiency gains *before* entry occurs. This is not the case in our within- plant/product data (see Figure 1). In addition, our matching estimation would absorb pre-entry productivity differences. However, we observe a drop in marginal cost and prices in the year of export entry ($t = 0$).[42] Productivity shocks are compatible with this observation if shock and export entry occur in the same period. Thus, we cannot completely exclude the possibility of selection into exporting, although it would only hold under an extremely quick response of export entry to productivity shocks. Nevertheless, note that even if some of the observed efficiency gains were due to selection, they are not reflected in the standard revenue productivity measure in the literature.

---

[41] Intuitively, markups rise because a demand shock increases the wedge between the demand curve (which is linear in our framework) and a given marginal cost, at the optimal point of production. Note that in the presence of increasing returns, marginal costs could fall due to expanding production as a result of demand shocks. However, our data suggest that this is unlikely, as discussed below.

[42] This effect is more pronounced in the within- plant/product results in Table 2. It becomes small and insignificant in the matching results in Table 3.

Learning-by-Exporting

Next, suppose that after export entry, marginal costs fall by φ_i^{LBE} percent because firm i gains new expertise ("learning-by-exporting").[43] Plant i will begin to export if the anticipated learning effects are sufficiently large: $\varphi_i^{LBE} \geq f(\varepsilon_i, F^E)$, where the function f is increasing in both arguments (see Appendix G).[44] Thus, the larger plant i's export entry wedge, and the higher the fixed cost of exporting, the larger is the required LBE to motivate export entry. 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 evidence for this effect in our data: Table 2 shows a downward trend in marginal costs during the first three years after export entry. This is also confirmed by the matching results in Table 3, but not in the balanced exporter panel in Table 4. Thus, learning-by-exporting can 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. To gain access to the new technology, a plant has to pay the fixed cost F^I. As pointed out by Lileeva and Trefler (2010), expanded production due to export entry may render investments in new technology profitable. In this case, plant i 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 exporting and new technology. This can be expressed as a condition on the drop in marginal cost under the new technology: $\varphi_i^{TEC} \geq g(\varepsilon_i, F^E, F^I)$, where the function g is increasing in all arguments (see Appendix G). This setup implies that initially less productive plants (with a larger export entry wedge ε_i) will require a larger efficiency gain φ_i^{TEC} in order to start exporting. Thus, we should expect "negative selection" based on initial productivity – a prediction that can be tested in the data (Lileeva and Trefler, 2010).

Table 6 provides evidence for this effect, reporting the change in marginal costs (Panel A) and reported average costs (Panel B) for plants with low and high pre-exporting productivity.[45] Both cost measures show a substantially steeper decline for initially less productive plants, as compared to exporters with high initial productivity. The difference in coefficients is statistically significant in most cases (as indicated by the p-values in the table). This result is in line with a complementarity channel where exporting and investment in technology go hand-in-hand, and where initially less

---

[43] In a dynamic setting, learning-by-exporting will lead to falling marginal costs over some time after export entry. We interpret φ_i^{LBE} as the (appropriately discounted) average annual decline in marginal cost after export entry.

[44] Note that in contrast to the case of productivity shocks, F_E now also determines the export entry threshold. The reason is that efficiency gains under LBE occur only in conjunction with export entry, while the productivity shocks discussed above are independent of export status.

[45] Following Lileeva and Trefler (2010), we use pre-exporting TFPR to split plants into above- and below median productivity.

productive plants will only make this joint decision if the efficiency gains are substantial.

[Insert Table 6 here]

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 occur once exporting begins. Overall, this suggests a pattern where plant managers first decide to export, then perform the necessary investment, and finally begin to sell to foreign markets when technology has been updated.

[Insert Table 7 here]

Alternative Interpretations: Returns to Scale and Input Prices

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, several observations in our data contradict this interpretation. First, our production function estimates suggest approximately constant returns to scale in most sectors – the mean (median) sum of all input shares is 1.024 (1.015). Table A.1 in the appendix reports the detailed estimates. Second, the pattern of investment reported in Panel A of Table 7 suggests specific technology upgrading, rather than a general expansion of production capacity. Finally, Panel B in Table 7 examines input prices, 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, which could be due to a shift to higher-quality inputs. In sum, these observations make scale economies an unlikely explanation for the observed decline in marginal cost after export entry.

On balance, our findings point to exporting-technology complementarity as an important driver of within-plant efficiency gains. In addition, there is some suggestive evidence for learning-by-exporting in the years after entry. Importantly, we do not claim that the observed effects are necessarily causal. In fact, the exporting-investment complementarity combines both causal mechanisms, from exporting to technology and vice versa. Our main contribution 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 falling prices.

5.2 'Foreign Demand Building'? Prices of Exported vs. Domestic Goods

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, making our marginal-cost based approach to measure efficiency obsolete. One possible explanation is that firms charge constant markups, so that efficiency gains are passed through completely to both domestic and foreign customers. An alternative explanation predicts differential price changes in the two markets: when applied to the trade context, 'demand building', as described by Foster et al. (2012), implies that export entrants charge lower prices abroad, in an attempt to attract customers where 'demand capital' is still low. Thus, in the context of our analysis, efficiency gains should be passed on to a larger degree in foreign markets, i.e., we should expect a stronger decline in export prices, as compared to their domestic counterpart. In the following, we provide supportive evidence for this mechanism.

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 7%.[46] Assuming that the marginal cost of production is the same for both markets, the results provide some evidence that efficiency gains are passed on to both domestic and foreign customers – but significantly more so to the latter. While we cannot pin down the exact mechanism that explains the observed price

---

[46] 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.6 in the appendix shows the results.

setting, our observations are in line with ’demand building’ in foreign markets.

5.3 Magnitude of Efficiency Gains

In the following, we gauge the magnitude of the observed within-plant efficiency gains after export entry, comparing them to the well-documented revenue productivity gap between exporters and non-exporters in the cross-section. In Table 1 we showed that TFPR of exporters is about 13% higher as compared to all other plants. This is very similar to the figures documented for US plants by Bernard and Jensen (1999). When disentangling ’young’ exporters (maximum three years after entry) from ’old’ ones, the difference in TFPR is small (2.9% higher for ’young’) and statistically insignificant. Therefore, revenue-productivity must be higher for exporters from the beginning; it does not increase substantially during the first years of exporting (see also Figure 1 and Table 2). This suggests that selection on *revenue*-based pre-exporting productivity à la Melitz (2003) is also present in our sample. Thus, the observed exporter TFPR premium of 13% is probably a good proxy for the magnitude of the selection effect.

Before comparing this cross-sectional figure to efficiency trends within plants, we briefly discuss under which circumstances TFPR and marginal costs reflect gains in productivity. As we showed in Section 2.2, constant markups imply that efficiency gains are reflected in marginal costs, but not in TFPR. On the other hand, falling marginal costs together with rising markups mean that at least part of the efficiency gains are also reflected in TFPR. Our data show constant markups during the first years following export entry. This allows for the straightforward interpretation that falling marginal costs reflect *additional* efficiency gains that need to be added to the TFPR-based cross-sectional exporter premium. For more mature exporters, there is also some evidence for rising markups, so that efficiency gains may eventually be partially reflected in TFPR. In this more complex case, the additional efficiency gains identified by our approach are given by the percentage drop in marginal costs, net of the percentage increase in TFPR (both relative to pre-exporting).

We follow two approaches when comparing within-plant marginal cost based efficiency gains to the traditional exporter premium in the cross-section. For simplicity, we refer to the gains not captured by TFPR as "additional gains". First, to obtain a conservative estimate, we assume that additional efficiency gains are only present for ’young’ exporters during their first three years after entry – this is the horizon observed in the data. In the context of our discussion in Section 2.2, this means that for ’old’ exporters, markups rise so much that all within-plant efficiency gains are then reflected by TFPR. That is, after four years of exporting, markups have to rise by the same proportion as marginal costs have dropped. In this (rather extreme) scenario, additional efficiency gains are only due to new export entrants, over the first three years after entry. On average, in our sample the share of ’young’ exporters among all exporting plants is approximately 10%. Multiplying this with the average observed drop in marginal cost from Tables 2 and 3 (20%) implies additional efficiency gains of 2% – about one-sixth of the cross-sectional exporter premium.

Second, we derive an upper bound, assuming instead that the observed additional efficiency gains for new exporters persist over time. For this to be true, the drop in prices observed for ’young’ exporters must not be reversed by rising markups in later years.[47] In this case, only *further* efficiency gains for established exporters would be reflected in TFPR, while the earlier decline in marginal cost for ’young’ exporters (about 20%) has become permanent – it only shows up in the form of lower marginal costs. Whether the actual effect is closer to the lower or the upper bound depends on the extent to which markups rise for ’old’ exporters (and thus translate efficiency gains into growing TFPR). In our data, markups are only slightly larger for ’old’ exporters (with three or more years after entry) – by 2.4%, and this difference is statistically insignificant. Consequently, the actual magnitude is probably closer to the upper bound estimate. Thus, the additional within-plant gains from exporting that are identified by marginal costs are likely of a similar magnitude as those observed for revenue productivity in the cross-section.

Summarizing our findings on selection versus within-plant productivity gains, we document that (i) new exporters have higher TFPR and markups already at the time of entry; (ii) TFPR and markups do not change much during the first years following export entry, but there is some evidence that they rise for more mature exporters; and (iii) marginal costs and prices decline significantly after export entry, and this decline is not reversed for ’old’ exporters. The first fact is in line with more productive firms charging higher markups and sorting into exporting (c.f. Bernard et al., 2003; Melitz and Ottaviano, 2008). The second fact suggests a slower increase of markups over time in our Chilean data, as compared for example to the immediate rise in markups identified by De Loecker and Warzynski (2012) for Slovenian firms. Our result is compatible with ’foreign demand building’ of new exporters, as discussed in Section 5.2. The third fact implies that export entry is also associated with within-plant efficiency gains that are not reflected by TFPR, because they are passed through to customers in the form of lower prices. For ’young’ exporters, pass-through is nearly complete, so that efficiency gains are only reflected in declining marginal costs. For more mature exporters, some of these gains are also reflected by TFPR due to rising markups, but this trend is not strong enough to fully reflect within-firm efficiency gains after export entry. Thus, conventional estimates miss a substantial part of the actual export-related gains from trade.

---

[47] This does not mean that markups must remain constant; for example, if marginal costs continue to fall, markups can rise while prices remain unchanged. There is some support for this in Table 3: three years after export entry, markups increase while prices are constant (since marginal costs fall at the same time). The part of efficiency gains missed by standard revenue-based measures is then the drop in marginal cost (28.2%), net of the rise in TFPR (13.7%), i.e., approximately 15% efficiency gains.

6 Conclusion

Over the last two decades, case studies and contributions in the management literature have provided strong suggestive evidence for within-firm productivity gains from exporting. 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 point out a reason for this discrepancy, and apply 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 show that this bias is likely at work – new exporters charge significantly lower prices. This is in line with evidence from new entrants in domestic markets. These typically lack connections with customers, whom they seek to attract by charging prices close to marginal costs (c.f. Foster et al., 2012).

In order to avoid the effect of lower prices on the efficiency measure, we use marginal cost, which is directly (negatively) associated with quantity-productivity in standard production functions. We estimate marginal costs at the plant-product level following the approach by De Loecker et al. (2012) – by first calculating markups under an unrestrictive set of assumptions and then deriving marginal costs as the ratio of price over markup. As a first step, we show that with the standard approach used in previous studies (revenue-based productivity), we do not find evidence for productivity gains after export entry in our panel of Chilean plants.

We then show that export entry is followed by a substantial decline in marginal costs within plants – approximately 15-25%. Prices follow a similar trajectory after export entry, suggesting that new exporters pass on most efficiency gains to their customers. This explains why previous revenue-based studies have failed to identify these gains. We also shed some light on the underlying mechanisms, guided by a stylized theoretical framework of export entry. Demand-side forces are unlikely candidates for explaining the observed efficiency gains, because markups are unchanged after export entry, while prices fall. Similarly, selection into exporting based on pre-existing productivity differences is probably not at the core of our results. Using propensity score matching, we construct a control group of plant-products with comparable initial characteristics and productivity trends. When using this comparison group, we find very similar results as within plants.

We identify two likely drivers of within-plant efficiency gains after export entry: first and foremost, a complementarity between export entry and technology investment. For this case, the theoretical framework (following Lileeva and Trefler, 2010) implies negative selection on pre-exporting productivity. We show that this prediction is born out by our data: initially less productive plant-products observe larger efficiency gains after export entry. In addition, disaggregate investment data show a hike in machine and equipment purchases around the period of export entry. Second, the fact that marginal costs continue to fall over 2-3 years after export entry suggests that learning-by-exporting may be another driver of the observed efficiency gains.

In sum, our results suggest a revision of the evidence for productivity gains from trade. So far, the main effects have been attributed to reallocation of resources across plants, in the spirit of Melitz (2003). For example, Pavcnik (2002) estimates that these reallocation effects are 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 within-plant efficiency gains after export entry have been unjustly discarded – they are probably of the same order of magnitude as gains from reallocation.

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., C. L. Benkard, S. Berry, and A. Pakes (2007). Econometric Tools for Analyzing Market Outcomes. In J. Heckman and E. Leamer (Eds.), *Handbook of Econometrics*, Volume 6 of *Handbook of Econometrics*, Chapter 63. Elsevier.

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.

Atkeson, A. and A. T. Burstein (2010, 06). Innovation, firm dynamics, and international trade. *Journal of Political Economy 118*(3), 433–484.

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.

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.

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.

Fishman, A. and R. Rob (2003). Consumer Inertia, Firm Growth and Industry Dynamics. *Journal of Economic Theory 109*(1), 24–38.

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.

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].

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

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.

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.

Ottaviano, G., T. Tabuchi, and J.-F. Thisse (2002). Agglomeration and Trade Revisited. *International Economic Review 43*(2), 409–436.

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

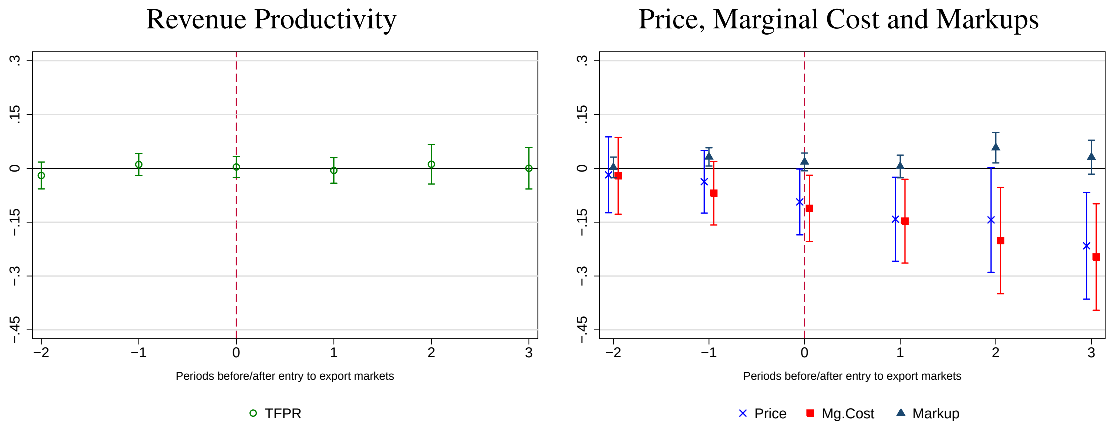
Revenue Productivity

Price, Marginal Cost and Markups

Figure 1: Price, Marginal Cost and TFPR Trajectories for New Exported Products

Notes: The left panel shows the estimated within plant trajectory for revenue productivity, and the right panel, for price, marginal cost and markup before and after export entry. Period $t = 0$ corresponds to the export entry year. For each plant-product, export entry occurs at period $t = 0$. The trajectories correspond to the estimated coefficients of equation (7), as reported in Table 2. 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. Section 4.1 provides further detail.

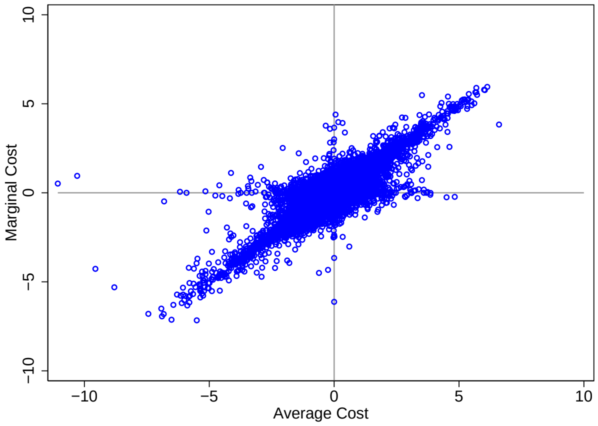

Figure 2: Estimated Marginal Cost and Reported Average Cost

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 actual variable costs.

TABLES

Table 1: Plant-Level Stylized Facts

	(1)	(2)	(3)	(4)	(5)	(6)
	Plant Size		Productivity	Capital Intensity	Wages	Markup
Dependent Variable	ln(workers)	ln(sales)	ln(TFPR)	ln(capital/worker)	ln(wage)	ln(markup)
Panel A: Unconditional Premia						
Export Dummy	1.403*** (.0844)	2.227*** (.179)	.122*** (.0307)	.907*** (.148)	.402*** (.0383)	.108*** (.0203)
Sector-Year FE	✓	✓	✓	✓	✓	✓
R^2	.26	.30	.99	.18	.24	.08
Observations	42,264	42,070	42,228	42,264	42,261	95,501
Panel B: Controlling for Employment						
Export Dummy	·	.648*** (.0884)	.136*** (.0228)	.750*** (.127)	.201*** (.0294)	.0620*** (.0183)
Sector-Year FE		✓	✓	✓	✓	✓
R^2		.70	.99	.19	.30	.09
Observations		42,070	42,228	42,264	42,261	95,501

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 6 are computed at the plant-product level; correspondingly, the coefficients reflect the difference in markups between exported products and those that are only sold domestically. Clustered standard errors (at the sector level) in parentheses. 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
Revenue TFP	-.0198	.0109	.00386	-.00574	.0115	.000249	2,752
	(.0228)	(.0188)	(.0180)	(.0217)	(.0336)	(.0352)	.584
Price	-.0180	-.0373	-.0936*	-.142**	-.144	-.216**	2,671
	(.0645)	(.0532)	(.0560)	(.0714)	(.0890)	(.0906)	.857
Marginal Cost	-.0206	-.0692	-.111**	-.147**	-.201**	-.247***	2,671
	(.0653)	(.0540)	(.0563)	(.0712)	(.0904)	(.0904)	.848
Markup	.00260	.0319**	.0179	.00529	.0575**	.0312	2,671
	(.0176)	(.0155)	(.0152)	(.0193)	(.0258)	(.0288)	.575
Physical Quantities	.0248	.145**	.187***	.203**	.133	.203	2,671
	(.0859)	(.0719)	(.0687)	(.0836)	(.108)	(.132)	.857

Notes: Regression output corresponds to the estimation of equation (7), including only new export entrants. The regression for TFPR is run at the plant level; it controls for plant fixed effects and sector-year effects (at the 2-digit level). The 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. Thus, additional products exported by multi-product plants do not enter our analysis. Section 4.1 provides further detail. Standard errors (clustered at the plant-product level) in parentheses. Key: ** significant at 1%; ** 5%; * 10%.

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

Periods After Entry	0	1	2	3
Revenue TFP	-.0103	.0135	.0401	.137***
	(.0209)	(.0275)	(.0411)	(.0497)
Price	-.0272	-.103*	-.166**	-.174
	(.0310)	(.0522)	(.0646)	(.104)
Marginal Cost	-.0429	-.102*	-.203***	-.282**
	(.0373)	(.0598)	(.0716)	(.121)
Markup	.0102	.000293	.0483	.123**
	(.0205)	(.0279)	(.0359)	(.0487)
Reported Average Cost	-.0705**	-.141***	-.0883	-.193*
	(.0342)	(.0500)	(.0732)	(.107)
Treated Observations (Min/Max)	183 / 186	124 / 131	75 / 81	35 / 37
Control Observations (Min/Max)	512 / 524	346 / 366	218 / 230	99 / 107

Notes: Coefficients correspond to the differential growth of each variable with respect to the pre-entry year ($t = -1$) between export entrants and controls. The control group is formed by plant-products that had a-priori a similar likelihood (propensity score) of being exported as export entrants, but that continued to be sold domestically only. Controls are selected from the pool of plants that produce the same product as new exporters. The specification of the propensity score is explained in section 4.2 and in Appendix E. In this table we match each entrant with the 3 nearest neighbors. Period $t = 0$ corresponds to the export entry year. The criteria for defining a plant-product as entrant is explained in the notes to Table 2. 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
Revenue TFP	-.0278	-.00629	.0309	.0899*
	(.0501)	(.0563)	(.0460)	(.0508)
Price	-.118	-.109	-.198**	-.139
	(.0701)	(.106)	(.0965)	(.111)
Marginal Cost	-.213**	-.117	-.276**	-.255**
	(.0875)	(.124)	(.102)	(.120)
Markup	.0503	.0176	.0795*	.118**
	(.0476)	(.0528)	(.0436)	(.0557)
Treated Observations	35 / 35	34 / 35	34 / 35	34 / 35
Control Observations	99 / 103	99 / 101	99 / 102	100 / 101

Notes: The results replicate Table 3 for the sample of plant-products that are observed in each period $t = -2, ..., 3$ (balanced panel). Coefficients correspond to the differential growth of each variable with respect to the pre-entry year ($t = -1$) between entrants and controls. The criteria for selecting controls is explained in the notes to Table 3. Period $t = 0$ corresponds to the entry year. The criteria for defining a plant as entrant can be found in the notes to Table 2. Robust standard errors in parentheses. Key: ** significant at 1%; ** 5%; * 10%.

Table 5: Domestic Goods Sold by Export Entrants

Periods After Entry	0	1	2	3
Marginal Cost	-.00671	-.108	-.0407	-.125
	(.0547)	(.0889)	(.153)	(.161)
Treated Observations	125	68	46	12
Control Observations	356	194	131	36

Notes: The coefficients show the differential growth of marginal cost with respect to the pre-entry year ($t = -1$) between treated and control groups. The treated group contains non-exported products produced by export entrants (thus, treated observations include only multi-product plants). The control group is selected using the criteria explained in the note to Table 3. Period $t = 0$ corresponds to the export entry year. The criteria for defining a plant-product as export entrant are stated in the notes to Table 2. Robust standard errors in parentheses.

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

Periods After Entry	0	1	2	3
Panel A: Marginal Cost				
Low Initial Productivity	-.161***	-.186**	-.335**	-.352*
	(.0534)	(.0827)	(.129)	(.186)
High Initial Productivity	.0656	-.0393	-.123	-.234
	(.0499)	(.0839)	(.0834)	(.162)
p-value	0.002	0.216	0.171	0.635
Treated Observations	184	129	80	37
Control Observations	518	363	229	107
Panel B: Reported Average Cost				
Low Initial Productivity	-.108**	-.184***	-.256**	-.424***
	(.0519)	(.0598)	(.100)	(.131)
High Initial Productivity	-.0371	-.104	.0176	-.0275
	(.0452)	(.0779)	(.0989)	(.151)
p-value	0.094	0.866	0.001	0.005
Treated Observations	183	124	75	36
Control Observations	513	346	218	103

Notes: The table analyzes heterogenous effects of export entry, depending on initial productivity. Coefficients correspond to the average effect of entry for entrants with initially low pre-exporting productivity, relative to high pre-exporting productivity entrants. Outcome variable is the growth of marginal cost (Panel A) and average costs (Panel B), with respect to the pre-entry period ($t = -1$). We use pre-exporting TFPR to split plant-products into above- and below- median productivity. The criteria for selecting controls can be found in the notes to Table 3. Period $t = 0$ corresponds to the export entry year. The criteria for defining a plant as entrant are described in the notes to Table 2. The p-value corresponds to a test of different coefficient sizes for low vs. high initial productivity plants. 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, including only new export entrants. "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%.