Heterogeneous Firms: Skilled-Labor Productivity and the Destination of Exports*

Jorge Balat† — Johns Hopkins University
Irene Brambilla‡ — UNLP
Yuya Sasaki§ — Johns Hopkins University

April 19, 2016

Preliminary and Incomplete. Do not circulate.

Abstract

This paper studies a systematic link between the choice of export destinations and technology differences across firms. Our premise is that firms differ in the relative the efficiency with which they can utilize skilled labor. In a context in which quality provision is skill-intensive and consumers in high income countries are more willing to pay for quality, exporting firms that are more efficient in the use of skilled labor export relatively more to high income destinations.

The contribution of the paper is twofold. First, we propose a new estimation method of production functions that allows for heterogeneity in the production function coefficients across firms and addresses the aggregation problem when firms are multiproduct. The estimation strategy is based on an extension of the structural control variable approach (Olley and Pakes (1996); Levinsohn and Petrin (2003)) to multi-dimensional heterogeneous parameters. Second, we provide an empirical measure of capability of quality production and show that it is a determinant of the choice of exports, export destinations, and quality using firm-level data from Chile.

*We benefited from useful comments by Dan Ackerberg and Elena Krasnokutskaya and participants at Hebrew, Johns Hopkins, Mannheim, Maryland, Penn State, Pittsburgh, Toronto, Toulouse School of Economics, Vanderbilt, Econometric Society Winter Meeting 2014, and the International Panel Data Conference 2014.

†Corresponding author. Department of Economics, Mergenthaler Hall 459, 2400 N. Charles St., Baltimore, MD 21218. Email: jorge.balat@jhu.edu

‡Department of Economics, Universidad Nacional de La Plata, Calle 6 N777, La Plata, 1900. Email: irene.brambilla@econo.unlp.edu.ar

§Department of Economics, Mergenthaler Hall 460, 2400 N. Charles St., Baltimore, MD 21218. Email: sasaki@jhu.edu

1 Introduction [XXX]

Production technology differences across firms have been shown to affect exporting behavior and firm performance. Melitz (2003) describes a selection mechanism in which more productive firms charge lower prices, sell more, earn higher variable profits, and are thus able to cover the fixed costs to enter export markets. Empirically, many papers have shown that there is a relationship between total factor productivity and the decision to export (Bernard and Bradford Jensen (1999), Clerides, Lach, and Tybout (1998), and Tybout (2003)).

More recently, it has been argued that product quality, in addition to firm productivity, is also a source of heterogeneity in exporting behavior among firms. Specifically, that firms that are capable of producing high-quality output tend to export more and to high-income markets of destination. The argument is that high-income countries value quality more and are willing to pay for high-quality goods.[1] With the increasing availability of detailed customs data, there is vast empirical evidence supporting these claims. A positive correlation between income in the destination market and product quality has been reported by Manova and Zhang (2012), Bastos and Silva (2010), Görg, Halpern, and Muraközy (2010) and Martin (2012) for firms in China, Portugal, Hungary and France. Production and export of high-quality output is also associated with employing skilled workers and high-quality intermediate inputs. Verhoogen (2008), Brambilla, Lederman, and Porto (2012), and Brambilla and Porto (2015) establish a causal relationship between high-income exports and wages (and skills) in Mexico, Argentina, and for a panel of 82 countries. Bastos, Silva, and Verhoogen (2014) establish a similar relationship between high-income exports and the quality of intermediate inputs in Portugal.

While the evidence linking high-income exports and quality of output and skilled workers is substantial, little attention has been paid to the firm-level heterogeneity in technology from which these differences in behavior might arise, and to how efficiency in quality production is a determinant of the choice of export destinations. In this paper we intend to fill that gap by studying the link between the choice of export destinations and technology differences across firms. Our idea is based on three premises: firms differ in their skilled-labor productivity, that is, the efficiency with which they can utilize skilled labor (relative to unskilled labor); product quality is higher when more skilled workers and higher quality

[1] See, for example, Verhoogen (2008), Hallak and Sivadasan (2013), Baldwin and Harrigan (2011), Johnson (2012), Brambilla, Lederman, and Porto (2012), and Feenstra and Romalis (2012). An additional channel linking high-income markets and high-quality output is the "Washington apples" effect (Hummels and Skiba (2004)).

intermediate inputs are used; and, finally, high-income countries value quality more. Therefore, firms that are more efficient in the use of skilled labor will export a larger share of their output to high-income destinations.

To formally establish these relationships —and to form the basis of our estimation methodology— we set up a dynamic model of firm behavior. Our model is based on the model of Olley and Pakes (1996) (OP, henceforth), but we extend it in several important ways. Our model and estimation method integrate three features into the existing framework of the productivity and production function estimation literature: it allows for firm heterogeneity that is not Hicks-neutral (heterogeneous labor coefficients), it develops a strategy to deal with output quantity information and multiproduct firms, and it provides a framework to study input and output quality choice and endogenous input and output prices.

Our methodological contribution is twofold. First, we contribute to the production function estimation literature by allowing for non-additive productivity shocks. Our estimation strategy is closely related to the homogeneous coefficient cases of OP, Levinsohn and Petrin (2003) (LP, henceforth), and Ackerberg, Caves, and Frazer (2016) (ACF, henceforth). In OP, investment is used as a proxy to control for unobserved TFP, which varies across firms and is correlated with input use. In our baseline setting, we have four sources of unobserved heterogeneity across firms and thus we use four proxies.[2] The proxies are raw materials, electricity, fuel, and output quality. We estimate the firm-level labor and capital productivity coefficients non-parametrically, joint with the coefficients on intermediate inputs, by GMM in a manner that is analogous to ACF. The identification of the coefficients relies on the structural nature of the model, but requires only one parametric assumption, the production technology.[3]

Second, we contribute by proposing a strategy to deal with the aggregation problem that arises from multiproduct firms in the estimation of a production function in physical units. Previous methods relied on using total revenue at the firm level as the left-hand-side variable. This has been shown to produce biased estimates of the production parameters due to unobserved differences in quality and demand shocks across firms (Klette and Griliches

[2]The possibility of using additional proxies to control for more than one source of unobserved heterogeneity across firms has been discussed by Ackerberg, Benkard, Berry, and Pakes (2005), although not in the context of non-additive productivity shocks.

[3]We assume technology is Cobb-Douglas, as most of the productivity literature. Some exceptions are De Loecker and Warzynski (2012) and De Loecker, Goldberg, Khandelwal, and Pavcnik (2015) who work with translog specifications, and Doraszelski and Jaumandreu (2013) and Doraszelski and Jaumandreu (2015) who estimate CES production functions.

(1996); Foster, Haltiwanger, and Syverson (2008); De Loecker (2011)). We exploit data on output by product (sales and unit prices) from single product firms to estimate (physical) rates of transformation between different products that allow us to add up quantities of different products for multiproduct firms.[4] These transformation rates are technology-based and not contaminated with demand-side information.

It is worth noting that the estimation method we propose is general and can be applied to other settings without relying on the underlying trade model. Relaxing the assumption of a common production technology across firms—even within well specified industries—serves two fundamental objectives.[5] First, the functional form choice restricts the type of questions that can be addressed. In our application, we exploit differences in relative skilled-labor productivity across firms and thus the specification of the production function has to allow for these differences. Second, even in cases where the researcher is only interested in recovering the TFP shocks, a misspecified production function may produce biased and inconsistent TFP estimates. In other words, if the estimation method ignores heterogeneity in, for example, the labor coefficients, the TFP estimates will pick this heterogeneity up.

In terms of our empirical contribution, we use a panel of Chilean firms spanning the period 1996–2006. The dataset includes information on revenue, employment of skilled and unskilled workers, and use of materials, fuels and electricity. Furthermore, the survey collects information on quantities and unit prices of products and materials. We have also been able to match the survey information with customs data containing information on value of exports by country of destination. We combine the two data sources to test the theory that differences in technology lead to differences in exporting behavior across destinations. We show that there is a link between differences in the production technology across firms and their input and output quality choices and their choices of export destinations. In particular, firms that have a technological advantage in the production of quality (i.e., firms that are more efficient in the use of skilled labor) use (relatively) more skilled labor —our measure of labor quality— and higher quality intermediate inputs, they produce higher quality output, and export a larger share of their output to high-income destinations.

The paper is organized as follows. In Section 2 we provide a model of firm behavior that establishes a link between the relative efficiency in the use of skilled labor and the

[4] De Loecker, Goldberg, Khandelwal, and Pavcnik (2015) also use single product firms to deal with the aggregation problem.

[5] This is an idea that dates back to the beginning of the estimation of production function literature. For instance, in the seminal paper by Marschak and Andrews (1944) they begin by saying "[the production functions] will be assumed to involve parameters that vary from firm to firm and from year to year [...]".

destination of exports. In Section 3 we propose an estimation procedure to recover the production function parameters. In Section 4 we explore the empirical relation between differences in technology and export choices using firm-level data from Chile. Section 5 concludes.

2 Model

In this section we set up a model of firm behavior. The objective of the model is twofold. First, we use the theoretical model to highlight the channels underlying our main hypothesis, namely, that differences in technology across firms that are not Hicks-neutral are a determinant of skill use, input and output quality, and the destination of exports. Second, the model of firm behavior provides the basis for the structural estimation of the technology parameters (TFP, output elasticities, and transformation rates across products). The estimates of the technology parameters allow us to test our main hypothesis.

We build upon the dynamic model of Olley and Pakes (1996) of firm investment decisions and total factor productivity and incorporate the following key features: multi-dimensional heterogeneity across firms in the production technology, product differentiation and a demand side, choice of product quality, and exports. The dynamic aspect of the firm decision problem is not essential to the argument that technology differences explain skill use, quality and exporting behavior, however, we keep this feature to be consistent with the productivity estimation literature. We start by describing the consumer and firm domestic-economy problem in Sections 2.1 and 2.2 and add export decisions in Section 2.3.

2.1 Consumers

On the demand side we assume a discrete choice framework with multinomial logit preferences as in Verhoogen (2008) and Brambilla, Lederman, and Porto (2012). Preferences are defined over varieties that are given by firm-product pairs, and which in turn are vertically and horizontally differentiated. The utility that individual i derives from choosing product r by firm j is given by

$$U_{ijrt} = \alpha(x_t)\theta_{jt} - p_{jrt} + \epsilon_{ijrt}, \tag{2.1}$$

where θ_{jt} is the quality of all varieties produced by firm j, p_{jrt} is the price of variety jr, and t denotes time. The coefficient α denotes the valuation for quality, which is assumed to be a strictly increasing weakly concave function of income x_t.[6] The variable ϵ_{ijt} is an individual-level random demand shock that adds horizontal differentiation into the model. We assume that it follows an iid type-I extreme value distribution and that it is the only source of heterogeneity among consumers.

The aggregate demand function q^d for variety jr takes the usual logit form (in logs)

$$q^d_{jrt}(p_{jrt}, \theta_{jt}) = \log W_t + \alpha(x_t)\theta_{jt} - p_{jrt}, \tag{2.2}$$

where W_t is an exogenous demand shifter given by the number of consumers at time t divided by a logit inclusive value term that aggregates all available choices.[7]

2.2 Firms

On the production side, we assume that there is a large fixed number of small heterogeneous firms. Markets are monopolistically competitive. Firms choose the number of products that they offer, output price, output quality, investment, skilled and unskilled labor, and intermediate inputs. We abstract from entry-exit decisions, which means we assume no sunk costs of entry or fixed costs for the first product.[8]

The production function for physical output has a two-tier structure. We follow the convention in the production function literature and denote with lower case letters variables in logs. The lower-tier is given by a Cobb-Douglas production technology that combines skilled labor (l^s_{jt}), unskilled labor (l^u_{jt}), capital (k_{jt}), and a vector of three intermediate inputs ($\mathbf{m}_{jt} = (m^1_{jt}, m^2_{jt}, m^3_{jt})'$) to produce a *composite input* (y_{jt}). The three intermediate inputs are raw materials, electricity and fuels. The production function in logs is given by

$$y_{jt} = \beta^0_{jt} + \beta^{l^s}_{jt} l^s_{jt} + \beta^{l^u}_{jt} l^u_{jt} + \beta^k_{jt} k_{jt} + \mathbf{m}'_{jt}\beta^m_t + \eta_{jt}. \tag{2.3}$$

The terms β^0_{jt} and η_{jt} denote a Hicks-neutral productivity shock and unforeseen random i.i.d. shock. As it is standard in the productivity literature, we assume that Hicks-neutral productivity, β^0_{jt}, differs across firms and time periods. We further allow for heterogeneity

[6]For simplicity we assume that all individuals share the same income level. We later allow income to vary across countries of destination. The utility function can be microfounded by defining primitive preferences over income.

[7]The inclusive value term is $\sum_{j'r' \in V_t} \exp(\alpha(x_t)\theta_{j'r't} - p_{j'r't})$, where V_t is the set of available firm-product varieties at time t.

[8]In terms of estimation, disregarding sunk costs implies not correcting for the issues of selection pointed out by Olley and Pakes (1996).

in the production technology. The labor and capital coefficients —β^{lu}_{jt}, β^{ls}_{jt}, and β^{k}_{jt}— are allowed to vary across firms and time as well, to capture differences in efficiency in the utilization of skilled labor, unskilled labor, and capital. The feature that is key to our model is the heterogeneity in efficiency in the use of skilled and unskilled labor. We assume that the Hicks-neutral productivity and the labor and capital output elasticities are random variables that evolve according to independent first-order Markov stochastic processes with stochastically increasing transition function.[9]

In the upper-tier the composite input is transformed into different products. We denote the (log) units of composite input that firm j assigns to the production of product r by y_{jrt} and the (log) units of output of product r by q_{jrt}. We assume that the composite input can be transformed into product r at a constant rate μ^r_t, that is, $\exp(q_{jrt}) = \mu^r_t \exp(y_{jrt})$. These assumptions imply that the rate at which product r can be transformed into product r' is given by $\mu^r_t/\mu^{r'}_t$, or, put it in other words, that the *adjusted* output quantities $\exp(q_{jrt})/\mu^r_t$ and $\exp(q_{jr't})/\mu^{r'}_t$ are expressed in units of equivalence in terms of input use.

Once quantities of all products are measured in the same units we can then aggregate output at the firm level. Aggregating across the set of products that firm j produces, denoted by R_{jt}, we write the production function as

$$\log\left(\sum_{r\in R_{jt}} \frac{\exp(q_{jrt})}{\mu^r_t}\right) = \beta^0_{jt} + \beta^{lu}_{jt} l^u_{jt} + \beta^{ls}_{jt} l^s_{jt} + \beta^k_{jt} k_{jt} + \mathbf{m}'_{jt}\beta^m_t + \eta_{jt}. \tag{2.4}$$

Notice that this specification does not rule out economies of scope since the (heterogeneous) productivity parameters may be positively correlated with the number of products. The composite input y_{jt} is also interpreted as total firm output expressed in units of equivalence. In equation (2.4), the TFP term (β^0_{jt}), the output elasticities (β^{lu}_{jt}, β^{ls}_{jt}, β^{k}_{jt}, β^m_t), and the rates of transformation (μ^r_t) are parameters to be estimated.

Let the set of all existing products be denoted by R_t, where the total number of products is then given by the cardinality of the set, $|R_t|$. For simplicity, we assume that the firm chooses the *number* of products to produce but not *which* products. Specifically, we assume that each firm faces a random sorting of the products in R_t. The sorted products are an

[9] While the first-order Markov assumption is key, these processes need not be exogenous, nor independent of each other. In Section XXX, we discuss how this assumption can be relaxed. In particular, it is straightforward to allow all four processes to depend on past realizations of all four productivity shocks. Moreover, investment, capital, expenditure in R&D, and exports may also affect the evolution of productivity, as in Aw, Roberts, and Xu (2011), Doraszelski and Jaumandreu (2013) and De Loecker (2013). These variables can be easily added into the model and the estimation method.

$|R_t| \times 1$ vector that we denote by $\overline{R}_{jt}$. The difference in $\overline{R}_{jt}$ across firms is the order of the products. We define the first product in $\overline{R}_{jt}$ as the firm's core product, which can be produced at no fixed cost. In addition to the core product, firms can also choose to produce additional products in the order given by $\overline{R}_{jt}$ by paying a fixed cost. The fixed cost is increasing in the number of products that each firm produces (i.e., the fixed cost of producing n products is F_n, with $F_n > F_{n-1}$ and $F_1 = 0$, where, again, product 1 is the core product.) Let R_{jt} denote the set of products actually produced by firm j. Notice that from a production point of view, products are determined by their transformation rates. Thus, we define $\overline{\mu}_{jt}$ and μ_{jt} as the vectors of transformation rates that correspond to the sets of products $\overline{R}_{jt}$ and R_{jt}; that is, the sorting of transformation rates faced by firm j and the transformation rates of the products that firm j actually chooses to produce.

While the exogeneity assumption on $\overline{R}_{jt}$ might seem restrictive at first, notice that it does not rule out firms' comparative advantages in the production of different products. In particular, we can allow for serial correlation in the sorting of products, resulting in persistence in comparative advantage over time. The exogeneity in the sorting of products implies that firm actions do not affect their comparative advantage.

Our model allows for vertical differentiation in output. Output quality depends on the quality of the inputs used in the production of physical output. There is no production of quality per se, in the sense that there are no inputs affected to the production of quality. We assume that firms choose a single quality (θ_{jt}) for all their products. Quality is a deterministic increasing function of the ratio of skilled to unskilled workers ($l^s_{jt} - l^u_{jt}$) —our measure of the quality of labor—, the quality of raw materials utilized in production (θ^{m1}_{jt}), and TFP (β^0_{jt}). In the production of shirts, for example, output quality depends on the quality of the fabric and the skill level of the workers but not on yards of fabric or hours of work.[10] We write the quality production function as

$$\theta_{jt} = v_t \left(l^s_{jt} - l^u_{jt}, \theta^{m1}_{jt}, \beta^0_{jt} \right), \tag{2.5}$$

where v is increasing in all of its arguments ($v_i > 0 \; \forall i$) and the cross derivatives are non-negative ($v_{ii'} \geq 0 \; \forall (i, i') \;\; i' \neq i$). Graphically, in the ($l^s_{jt} \times l^u_{jt}$) plane, we can think of quantity as fixed along an isoquant and quality as fixed along a ray through the origin. In Figure 1, for a given level of capital, intermediate inputs and quality of raw materials, output

[10] Kugler and Verhoogen (2012) provide empirical support for complementarity in input and output quality using a panel of firms from Colombia. Verhoogen (2008) and De Loecker, Goldberg, Khandelwal, and Pavcnik (2015) also assume complementarity in inputs and output quality.

Figure 1: Level Curves for Quantity and Quality

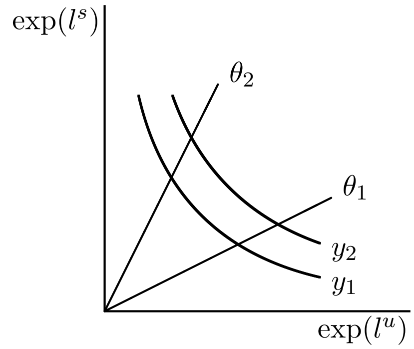

Note: Figure depicts level curves from the production functions for output quantity and quality conditional on fixed values for capital, intermediate inputs, and quality of raw materials. Output quantity is increasing in skilled and unskilled labor ($y_2 > y_1$); whereas output quality is increasing in the ratio of skilled to unskilled labor ($\theta_2 > \theta_1$).

quantity is increasing in both types of labor ($y_2 > y_1$) and output quality is increasing in the ratio of skilled to unskilled labor ($\theta_2 > \theta_1$).

We assume that the adjustment of the capital stock is costly, that there are firing and hiring costs, and that capital, labor, and product lines are subject to "time to build", "time to train", and "time to develop", so that investment, hiring and firing, and product decisions become effective in the next period. These assumptions make the problem dynamic.[11] The adjustment cost of capital and labor are given by the functions $G_t^k(k_{jt}, k_{jt+1})$, $G_t^s(l_{jt}^s, l_{jt+1}^s)$, and $G_t^u(l_{jt}^u, l_{jt+1}^u)$. The unit prices of intermediate inputs are given by $p_t^{m1}(m_{jt}^1, \theta_{jt}^{m1})$, $p_t^{m2}(m_{jt}^2)$ and $p_t^{m3}(m_{jt}^3)$; thus nesting cases of linear and non-linear prices while assuming that pricing schedules are the same across firms. Notice that the price of raw materials (m_{jt}^1) further depends on its quality (θ_{jt}^{m1}).

The objective of the firm is to maximize the present value of the stream of current and future profits. We can split the firm decision problem into a static component and a dynamic component. In the static problem, firms maximize current profits taking as given decisions from $t-1$ which include current capital, skilled and unskilled labor, and the set of products that the firm will produce, as well as the technology parameters (input productivities and TFP). They choose the quantity of intermediate inputs and output quality. Because all firm-

[11]These assumptions imply that in each period there is the sunk cost of predetermined labor, which is assumed to be unavoidable because of binding contracts. Firms could choose not to hire any workers for the next period if they anticipate bad productivity shocks, which implies zero output. This scenario is, however, not akin to exiting the market in the sense that there are no costs of resuming production later on.

product pairs enter the output demand symmetrically, once we condition on the number of products, the only channel through which the sorting of products affects profits is through their different rates of transformation. We can thus write the static choices of a firm that produces n products as a function of the vector μ_t, of size $n \times 1$. The policy function for the static choices is given by

$$(m_{jt}^1, m_{jt}^2, m_{jt}^3, \theta_{jt})' = g_{nt}(k_{jt}, l_{jt}^s, l_{jt}^u, \beta_{jt}^0, \beta_{jt}^{l^s}, \beta_{jt}^{l^u}, \beta_{jt}^k, \mu_{jt}). \tag{2.6}$$

The policy function is indexed by t to reflect changing conditions in input and output markets as well as in the output elasticity of intermediate inputs; these changes are common to all firms. It is also indexed by the number of products n, to reflect the different vector lengths of μ_{jt} across firms.

The intermediate inputs and output quality decisions implicitly define output prices, quality of raw materials, and the price of intermediate inputs ($\mathbf{p}_{jt}$, θ_{jt}^{m1}, p_{jt}^{m1}, p_{jt}^{m2}, p_{jt}^{m3}), which depend on the same state variables as (2.6), and where $\mathbf{p}_{jt}$ is the vector of prices for each output that the firm produces (thus, it has the same dimension as μ_{jt} and R_{jt}.) Notice that under the model assumptions output and input prices differ endogenously across firms.

In the dynamic problem firms choose investment (i_{jt}), skilled and unskilled labor, and the number of products, all of which will become operative at time $t+1$.[12] The firm-level state variables that enter the dynamic problem are predetermined capital and labor, next period's sorting of products, and the technology parameters (TFP, and the input coefficients). Notice that whereas the choice of products depends on the full vector of product sorting $\overline{R}_{jt+1}$, the input choices can be written as conditional on the number of products and the set of products actually produced. The policy variables for the dynamic problem are thus

$$R_{jt+1} = h_t^R(k_{jt}, l_{jt}^s, l_{jt}^u, \beta_{jt}^0, \beta_{jt}^{l^s}, \beta_{jt}^{l^u}, \beta_{jt}^k, \overline{\mu}_{jt+1}). \tag{2.7}$$

$$(i_{jt}, l_{jt+1}^s, l_{jt+1}^u)' = h_{nt}(k_{jt}, l_{jt}^s, l_{jt}^u, \beta_{jt}^0, \beta_{jt}^{l^s}, \beta_{jt}^{l^u}, \beta_{jt}^k, \mu_{jt+1}). \tag{2.8}$$

The function h is indexed by the number of products n, and both functions h and h^R are indexed by time. As will become clear later on, splitting the dynamic choices in this manner is useful for the estimation strategy.

[12] Investment in the current period determines capital according to the usual law of motion (variables in logs): $\exp(k_{jt+1}) = (1-\delta)\exp(k_{jt}) + \exp(i_{jt})$, where δ is the depreciation rate.

Recall that the objective of the model is to establish a link between technology differences, quality choice, and the destination of exports. Before introducing exports we start by arguing that firms that are more efficient in the use of skilled labor produce higher quality output.

Proposition 1 (Skilled-intensive technology implies higher quality)**.** *Given k_{jt}, β^0_{jt} and $\overline{\mu}_{jt}$, output quality θ_{jt} is, on average, increasing in the ratio of skilled to unskilled labor coefficients $\beta^{l^s}_{jt}/\beta^{u^s}_{jt}$.*

The intuition is straightforward: output quality θ_{jt} is increasing in both the skilled ratio ($l^s_{jt} - l^u_{jt}$) and quality of raw materials (θ^{m1}_{jt}), while in turn a firm that is relatively more efficient in the use of skilled labor chooses both a higher skilled ratio and higher quality of raw materials. The choice of higher quality of raw materials follows from the complementarity between the skilled ratio and the quality of intermediate inputs in output quality production. See the Online Appendix for a formal proof.

2.3 Exports

We now introduce exports into the model. Firms sell across different countries of destination $c \in C$, including the domestic market. The willingness to pay for quality varies across destinations. As in Verhoogen (2008) and Brambilla, Lederman, and Porto (2012), quality valuation is increasing in destination income ($\alpha'(x_{ct}) > 0$). The utility that consumer i in country c derives from product r produced by firm j is given by

$$U_{ijrct} = \alpha(x_{ct})\theta_{jt} - p_{jrct} + \epsilon_{ijrct}. \tag{2.9}$$

To simplify we assume that firm quality is the same across products and destinations whereas firm prices are allowed to vary. The dynamic and static problems are analogous to Section 2.2. At time t firms choose investment, skilled and unskilled labor, and number of products (all of which become operative at $t + 1$), intermediate inputs, product quality, and prices at each destination. Choosing prices is equivalent to choosing the level of exports.

We further assume that there are no fixed or sunk cost of exporting. This implies that all firms sell domestically and to all destinations. Sunk costs of entry and fixed costs of production and exports are found to be empirically relevant in the IO and trade literature and our ignoring them here serves the purpose of simplifying the algebra as well as pointing out that the mechanisms described in this model do not depend on selection derived from participation decisions based on fixed or sunk cost. Differences in technology lead firms to differ in the share of production that they ship to different destinations. This is the variation across firms that we exploit both in the model and in the empirical results.

We now turn to the link between technology, quality, and income across destinations. Because quality valuation is increasing in income, firms that choose a higher output quality sell a higher share of their exports to high-income (or high-quality-valuation) markets. This is because, for a given price, country income operates as a demand shifter which is in turn increasing in output quality.

Proposition 2 (Higher quality implies larger high-income share)**.** *The share in sales of firm j of countries with income above the sales-weighted average income across destinations (high-income countries) is increasing in output quality.*

Proof. Let $\varphi_{jct} = \frac{\exp(y_{jct})}{\sum_{c' \in C} \exp(y_{jc't})}$ denote the share of destination c in total sales of firm j. The derivative of the share φ_{jct} with respect to quality θ_{jt} is

$$\frac{\partial \varphi_{jct}}{\partial \theta_{jt}} = \varphi_{jct} \left(\alpha(x_{ct}) - \sum_{c'} \varphi_{c'} \alpha(x_{c't}) \right). \tag{2.10}$$

If c is a high-income country, then $x_{ct} - \sum_{c'} \varphi_{jc't} x_{c't} > 0$. Because α is a strictly increasing function, $\alpha(x_{ct}) - \alpha\left(\sum_{c'} \varphi_{c'} x_{c't}\right) > 0$. By Jensen inequality, because α is a weakly concave function, $\alpha(x_{ct}) - \sum_{c'} \varphi_{c'} \alpha(x_{c't}) > 0$, which implies that the derivative (2.10) is strictly positive and proves that the share of high income countries in sales is increasing in quality. □

The mechanisms described in Propositions 1 and 2 create a link between skilled-intensive technology and exports to high income countries whereby firms that are relatively more efficient in the use of skilled labor tend to export more to high-income destinations. Our objective is to test empirically that technology differences across firms are a determinant of output quality and the destination of exports.

3 Estimation of the Technology Parameters

In this section we discuss the estimation of the production technology parameters: the TFP shock (β^0_{jt}), the output elasticities (β^{lu}_{jt}, β^{ls}_{jt}, β^k_{jt}, β^m_t), and the rates of transformation (μ^r_t). With estimates of these parameters at hand, we later test the model predictions on technology and firm decisions on input use, quality and exports in Section 4 using firm level data from Chile. Our estimation strategy has two novel components: the estimation of heterogeneous input coefficients for labor and capital, and the estimation of transformation rates that allow us to use information on physical units of output instead of revenue.

Relaxing the assumption of homogeneous labor and capital coefficients serves the fundamental objective of allowing a broader type of questions. In our application, we exploit differences in relative labor productivity across firms and thus the specification of the production function has to allow for these differences. Doraszelski and Jaumandreu (2015) also allow for differences in labor and capital productivity across firms to estimate labor biased technical change. Additionally, even in the case where the researcher is only interested in recovering the TFP shocks, a misspecified production function may produce biased and inconsistent TFP estimates, as TFP estimates pick up other heterogeneity.

In the homogeneous coefficient methods of Olley and Pakes (1996) (OP) and Levinsohn and Petrin (2003) (LP), the estimation is based on using a proxy (derived from inverting a structural policy function) to control for unobserved TFP, and exploiting the Markov process for TFP. Our estimation method extends that logic to control for higher-dimensional heterogeneity in the production function. We thus write the unobserved firm-level technology parameters (β^0_{jt}, $\beta^{l^s}_{jt}$, $\beta^{l^u}_{jt}$, β^k_{jt}) as a function of four proxies. Suitable proxies are all of the choice variables in the static and dynamic policy functions (2.6) and (2.8). We describe a baseline case in which the proxies are the three intermediate inputs (raw materials, electricity, and fuel) and output quality. Alternatively, forward skilled labor, forward unskilled labor, and investment can also be used as proxies with minor modifications to the estimation algorithm to account for the different vector of transformation rates that enter (2.6) and (2.8). The heterogeneous input coefficients are estimated as non-parametric functions of state variables and proxies. We estimate all input coefficients jointly, as in ACF, instead of sequentially as in OP and LP.

Our strategy is also closely related to De Loecker, Goldberg, Khandelwal, and Pavcnik (2015) (DLGKP) in that we use information on output quantity by product. Estimating production functions typically involves using deflated revenue instead of physical output. Because price deflators are common to all firms in an industry, estimates of the productivity parameters confound the effects of physical productivity with demand shocks, quality and mark-ups (DLGKP; Klette and Griliches (1996); Foster, Haltiwanger, and Syverson (2008); De Loecker (2011)). With the increasing availability of firm-product level data on quantities and unit values, it is possible to use more detailed firm information instead of industry-deflated revenue.

When dealing with multiproduct firms there are three issues we need to address. First, there is a problem when using quantity information in that units of different products are not comparable even within well specified industries (i.e., the proverbial comparison of apples and oranges). In the case of single-product firms producing exactly the same product, a production function with quantity on the left-hand side could be estimated using the regular methods. When products differ, or in the case of multiproduct firms, some assumptions are needed to make units of physical output comparable. When revenue is used on the left-hand side, firm-level prices are used implicitly as transformation rates, which makes physical units comparable at the cost of contaminating quantities with demand-side information. Instead, we estimate transformation rates that are technology-based and make the assumption that heterogeneous products can be written in terms of constant units of equivalence (equation 2.4). Second, from typical firm-level input-output data it is not possible to determine the quantities of inputs used in the production of each product.[13] The assumption of the two-tier production technology introduced in Section 2.2 together with the assumption that the (constant) rates of transformation are common across firms, allow us to write and estimate the production functions at the firm-level rather than at the firm-product level. Finally, there is the issue of endogeneity of product choice. We assume that transformation rates are constant across firms, and that the core and subsequent sorting of products is given exogenously to each firm. The constant transformation rates across firms rule out Roy-type selection in product choice. However, comparative advantage in the production of different products is not ruled out, and is instead represented by the exogenous sorting of products for each firm. Serial correlation in the sorting of products results in persistence in comparative advantage over time. The exogeneity in the sorting of products implies that firm actions do not affect their comparative advantage.

We now turn a discussion of the estimating assumptions we make. The first assumption has to do with the structure of the demand system. We assume that the sorting of products affects instantaneous firm profits (and therefore input and quality decisions) only through the transformation rates. In other words, the policy functions can be written in terms of the number of products and the transformation rates of the products actually produced by the firm, instead of as a function of the full sorting of products $\overline{R}_{jt}$. In our model this

[13]Typically, even when output data is available at the product level, input data is aggregated at the firm level.

condition holds because of the particular demand structure that we assume. More generally, the condition holds for any demand structure in which products enter the utility function symmetrically except for an iid demand shock and endogenous choices such as quality. Additional demand systems that satisfy the latter structure are, for example, a demand model with homogeneous products or a symmetric CES utility function with horizontal product differentiation. Formally, we write the assumption of symmetry of the demand system as

$$
\begin{aligned}
(SYM) \qquad g_{nt}(k_{jt}, l^s_{jt}, l^u_{jt}, \beta^0_{jt}, \beta^{l^s}_{jt}, \beta^{l^u}_{jt}, \beta^k_{jt}, \mu_{jt}) &= g_t(k_{jt}, l^s_{jt}, l^u_{jt}, \beta^0_{jt}, \beta^{l^s}_{jt}, \beta^{l^u}_{jt}, \beta^k_{jt}, \overline{R}_{jt}) \\
h_{nt}(k_{jt}, l^s_{jt}, l^u_{jt}, \beta^0_{jt}, \beta^{l^s}_{jt}, \beta^{l^u}_{jt}, \beta^k_{jt}, \mu_{jt}) &= h_t(k_{jt}, l^s_{jt}, l^u_{jt}, \beta^0_{jt}, \beta^{l^s}_{jt}, \beta^{l^u}_{jt}, \beta^k_{jt}, \overline{R}_{jt}).
\end{aligned}
$$

The relevance of (SYM) is that for the inversion strategy to work, it is necessary to be able to write the four unobserved technology parameters as a function of observed variables or estimable parameters. The full sorting of products $\overline{R}_{jt}$ is an unobserved variable, whereas the number of products is observed and the rates of transformation are estimable parameters.

The remaining estimating assumptions are similar to the ones in the production function estimation literature. The policy function for intermediate inputs and quality (system (2.6)), is assumed to be invertible (conditional on the state variables k_{jt}, l^s_{jt}, l^u_{jt}, and μ_{jt}), so that we can write the firm-level technology parameters as

$$
(INV) \qquad (\beta^0_{jt}, \beta^{l^s}_{jt}, \beta^{l^u}_{jt}, \beta^k_{jt})' = g^{-1}_{nt}(k_{jt}, l^s_{jt}, l^u_{jt}, \mu_{jt}, \mathbf{m}_{jt}, \theta_{jt}).
$$

This assumption is an extension of the invertibility condition of the homogeneous coefficients literature (OP, LP, ACF, DLGKP) to a multi-dimensional case. The extension is not straightforward as it requires functional independence of the intermediate input demands in (2.6). Intuitively, to work as separate proxies the input demands need to provide independent information on the technology parameters. Functional independence of the input demands may come from two sources. First, raw materials are vertically differentiated whereas electricity and fuels are not. Because output quality is increasing in the ratio of the labor coefficients β^{l^s}/β^{l^u} (Proposition 1), the quantity and quality of raw materials react differently to values of the labor coefficients compared to electricity and fuels. Second, the price of electricity usually takes the form of a two-part tariff. This again leads to different intermediate input use upon different realizations of TFP and labor productivity. More generally, a condition under which intermediate inputs are not functionally dependent is when two of the price schedules $p_t^{m1}(m_{jt}^1, \theta_{jt}^{m1})$, $p_t^{m2}(m_{jt}^2)$ and $p_t^{m3}(m_{jt}^3)$ are non-linear.[14]

Following the literature, the unforeseen i.i.d. shock, η_{jt}, is assumed to be orthogonal to the firm's input and quality choices. That is,

$$(IND) \qquad \mathrm{E}[\eta_{jt}|k_{j\tau}, l_{j\tau}^s, l_{j\tau}^u, \beta_{j\tau}^0, \beta_{j\tau}^{l^s}, \beta_{j\tau}^{l^u}, \beta_{j\tau}^k, \mu_{j\tau}] = 0.$$

Total factor productivity, β_{jt}^0, is assumed to follow a first-order Markov process. Specifically, letting $\mathcal{I}_{jt} = \sigma(k_{jt}, l_{jt}^s, l_{jt}^u, \beta_{jt}^0, \beta_{jt}^{l^s}, \beta_{jt}^{l^u}, \beta_{jt}^k, \mu_{jt}, k_{jt-1}, l_{jt-1}^s, l_{jt-1}^u, \beta_{jt-1}^0, \beta_{jt-1}^{l^u}, \beta_{jt-1}^{l^s}, \beta_{jt-1}^k, \mu_{jt-1}, \cdots)$ denote the information available to firm j at the end of time period t, we assume that the stochastic process for β_{jt}^0 satisfies

$$(MAR) \qquad \mathrm{E}[\beta_{jt}^0 \mid \mathcal{I}_{jt-1}] = \mathrm{E}[\beta_{jt}^0 \mid \beta_{jt-1}^0].$$

The assumption that the evolution of β_{jt}^0 depends solely on β_{jt-1}^0 is made for expositional simplicity. Variables in $t-1$ such as the skilled and unskilled labor coefficients, investment, capital, labor, output, expenditure in R&D, and exports may also affect firm-level productivity, as in Aw, Roberts, and Xu (2011), Doraszelski and Jaumandreu (2013) and De Loecker (2013). The addition of these variables into the estimation method is straightforward as long as the stochastic process remains first-order Markov. In the empirical implementation we add lagged input coefficients and lagged endogenous variables to the conditioning set in (MAR). The statistical significance of the conditioning variables can be tested empirically as part of the estimation.

Intuitively, our estimation method extends the framework of homogeneous input coefficients to estimate labor and capital coefficients that vary across firms, and transformation rates that allow us to add up differentiated products. We exploit the structural nature of the model to write the firm-level coefficients as non-parametric functions of observed variables, which in turn requires four proxies and invertibility of a system of four equations (INV). The symmetry of the demand system assumption (SYM) is needed so that an estimable vector of rates of transformation can be used as a control for the sorting of products, entering as an argument in (INV). The independence (IND) and Markov (MAR) assumptions are identical to the homogeneous coefficients literature (OP, LP, ACF, DLGKP).

The estimation algorithm involves two steps. We first estimate the rates of transfor-

[14] In cases in which assuming functional independence of the intermediate inputs is unreasonable, or when no disaggregate data on intermediate input use is available, other proxies may be used. See Section 3.3.

mation (μ_t^r). With estimates of the transformation rates we can then add up quantities of different products and estimate the input coefficients ($\beta_{jt}^{l^s}$,$\beta_{jt}^{l^u}$,β_{jt}^{k},β_t^m) and TFP (β_{jt}^0).

3.1 Estimation of the Transformation Rates

The production function regression is given by (2.3). The challenge is that the composite input y_{jt}, which is empirically interpreted as firm-level physical output aggregated across heterogeneous products in units of equivalence, is not observed in the data. We follow a strategy that has some similarities with DLGKP in the sense that we exploit information from single product firms. Using single product firms works because, by definition, there is no need to add up output quantities across products. DLGKP estimate the input coefficients using the set of single product firms and later apply these same parameters to multiproduct firms. This method cannot be applied to our setting because input coefficients are assumed to be heterogeneous across firms. What we do instead is to use single product firms to estimate rates of transformation between products and then estimate the production function parameters using single and multiproduct firms together.

From equation (2.4) we can write the production function for single-product firms as

$$q_{jrt} = \sum_{r \in R_t} \log \mu_t^r D_{jrt} + \beta_{jt}^0 + \beta_{jt}^{l^s} l_{jt}^s + \beta_{jt}^{l^u} l_{jt}^u + \beta_{jt}^k k_{jt} + \mathbf{m}_{jt}' \beta_t^m + \eta_{jt}. \tag{3.1}$$

where D_{jrt} is a dummy variable that takes the value of one if single-product firm j produces product r.

The objective is to estimate the transformation rates μ_t^r using a non-parametric function to control for everything else in the regression, including unobserved firm-level technology. Based on the symmetry and invertibility assumptions (SYM) and (INV), we can write the unobserved technology parameters for the sample of single-product firms as a function of the observed inputs and quality as $\beta_{1t}^0(k_{jt}, l_{jt}^s, l_{jt}^u, \mu_{jt}, \mathbf{m}_{jt}, \theta_{jt})$, $\beta_{1t}^{l^s}(k_{jt}, l_{jt}^s, l_{jt}^u, \mu_{jt}, \mathbf{m}_{jt}, \theta_{jt})$, $\beta_{1t}^{l^u}(k_{jt}, l_{jt}^s, l_{jt}^u, \mu_{jt}, \mathbf{m}_{jt}, \theta_{jt})$, and $\beta_{1t}^k(k_{jt}, l_{jt}^s, l_{jt}^u, \mu_{jt}, \mathbf{m}_{jt}, \theta_{jt})$. These functions depend on the transformation rates μ_{jt}, which are unobserved. In order to write the technology parameters as a function of observed variables only, we can exploit the information of which core product each firm produces. For single-product firms producing the same (observed) core product r, we can write the unobserved technology parameters as $\beta_{r1t}^0(k_{jt}, l_{jt}^s, l_{jt}^u, \mathbf{m}_{jt}, \theta_{jt})$, $\beta_{r1t}^{l^s}(k_{jt}, l_{jt}^s, l_{jt}^u, \mathbf{m}_{jt}, \theta_{jt})$, $\beta_{r1t}^{l^u}(k_{jt}, l_{jt}^s, l_{jt}^u, \mathbf{m}_{jt}, \theta_{jt})$, and $\beta_{r1t}^k(k_{jt}, l_{jt}^s, l_{jt}^u, \mathbf{m}_{jt}, \theta_{jt})$. Plugging these functions into (3.1) we can write the production function for single-product firms as

$$q_{jrt} = \sum_{r \in R_t} \log \mu_t^r D_{jrt} + \beta_{r1t}^0(k_{jt}, l_{jt}^s, l_{jt}^u, \mathbf{m}_{jt}, \theta_{jt}) + \beta_{r1t}^{l^s}(k_{jt}, l_{jt}^s, l_{jt}^u, \mathbf{m}_{jt}, \theta_{jt}) l_{jt}^s + \beta_{r1t}^{l^u}(k_{jt}, l_{jt}^s, l_{jt}^u, \mathbf{m}_{jt}, \theta_{jt}) l_{jt}^u + \beta_{r1t}^k(k_{jt}, l_{jt}^s, l_{jt}^u, \mathbf{m}_{jt}, \theta_{jt}) k_{jt} + \mathbf{m}_{jt}' \beta_t^m + \eta_{jt}, \tag{3.2}$$

or, more compactly, as

$$q_{jrt} = \sum_{r \in R_t} \log \mu_t^r D_{jrt} + \phi_{rt}^1(k_{jt}, l_{jt}^u, l_{jt}^s, \mathbf{m}_{jt}, \theta_{jt}) + \eta_{jt} \tag{3.3}$$

where the unknown function ϕ^1 is the expectation of output conditional on the inputs and proxies. The function ϕ^1 varies across time and across products to reflect the different choices made by single-product firms that face different transformation rates for their core (and only) product.

We estimate regression (3.3) for the set of single-product firms using data on units of output, inputs, proxies, and the indicator variables D. Since the functional form of ϕ^1 is unknown, the regression can be estimated by partially linear least squares (Robinson, 1988) or with a high order polynomial. Estimates of the rates of transformation $\widehat{\mu}^r$ are obtained from the coefficients on the indicator variables.[15]

The estimates of $\widehat{\mu}^r$ are used to add up across heterogeneous products produced by multiproduct firms by previously transforming them into units of equivalence according to

$$y_{jt} = \log \left(\sum_{r \in R_j} \frac{\exp(y_{jrt})}{\mu^r} \right). \tag{3.4}$$

This is a measure of physical output which, unlike revenue, is not contaminated by prices.

3.2 Estimation of the Input Coefficients and TFP

We now turn to the estimation of the input coefficients and TFP. The estimation of the input coefficients and TFP is based on equation (2.4) where the left-hand side variable is firm-level physical output y_{jt} computed using the transformation rates as described in the previous section. The estimation algorithm is again based on writing unobserved technology parameters as a function of firm choices. It can be split into two stages most closely

[15] Transformation rates only need to be identified up to a numéraire product. We assume that the constant in ϕ_{rt}^1 is the same across products and without loss of generality we normalize it to zero.

resembling ACF and DLGKP but allowing for heterogeneous input coefficients.

The first stage is based on the production function regression in equation (2.3) and the assumptions (SYM), (INV) and (IND). We again proceed to write the unobserved parameters as a function of the proxies as $\beta^0_{nt}(k_{jt}, l^s_{jt}, l^u_{jt}, \mu_{jt}, \mathbf{m}_{jt}, \theta_{jt})$, $\beta^{l^s}_{nt}(k_{jt}, l^s_{jt}, l^u_{jt}, \mu_{jt}, \mathbf{m}_{jt}, \theta_{jt})$, $\beta^{l^u}_{nt}(k_{jt}, l^s_{jt}, l^u_{jt}, \mu_{jt}, \mathbf{m}_{jt}, \theta_{jt})$, $\beta^{k}_{nt}(k_{jt}, l^s_{jt}, l^u_{jt}, \mu_{jt}, \mathbf{m}_{jt}, \theta_{jt})$, for $n = 1, \ldots, N_t$, where N_t is the largest number of products produced by one firm. We can control for μ_{jt} using the estimates obtained in the previous section. Plugging these functions into (2.3) we can write the production function as

$$\begin{aligned} y_{jt} \quad = \quad & \beta^0_{nt}(k_{jt}, l^s_{jt}, l^u_{jt}, \mu_{jt}, \mathbf{m}_{jt}, \theta_{jt}) + \beta^{l^s}_{nt}(k_{jt}, l^s_{jt}, l^u_{jt}, \mu_{jt}, \mathbf{m}_{jt}, \theta_{jt}) l^s_{jt} \\ & + \beta^{l^u}_{nt}(k_{jt}, l^s_{jt}, l^u_{jt}, \mu_{jt}, \mathbf{m}_{jt}, \theta_{jt}) l^u_{jt} + \beta^{k}_{nt}(k_{jt}, l^s_{jt}, l^u_{jt}, \mu_{jt}, \mathbf{m}_{jt}, \theta_{jt}) k_{jt} \\ & + \mathbf{m}'_{jt}\beta^m_t + \eta_{jt}, \end{aligned} \tag{3.5}$$

or, more compactly, as

$$y_{jt} = \phi_{nt}(k_{jt}, l^s_{jt}, l^u_{jt}, \mu_{jt}, \mathbf{m}_{jt}, \theta_{jt}) + \eta_{jt} \tag{3.6}$$

where the unknown function ϕ_{nt} is the expectation of output conditional on the inputs and proxies. This function varies across time and across firms that produce a different number of products. Equation (3.6) can be estimated non-parametrically. We denote the estimated non-parametric functions with $\widehat{\phi}_{nt}(\cdot)$.

In the second stage we estimate the input coefficients. This stage is based on assumptions (SYM), (INV) and (MAR). For notational convenience, we define the conditional expectation function Λ, so that $\Lambda(\beta^0_{jt-1}) = \mathrm{E}(\beta^0_{jt} \mid \beta^0_{jt-1})$, and the random variable $\xi_{jt} = \beta^0_{jt} - \Lambda(\beta^0_{jt-1})$, which is interpreted as an unanticipated component of the evolution of TFP. Using symmetry and invertibility to write the unobserved technology parameters as a function of the inputs and proxies we can write the unanticipated component of TFP as

$$
\begin{aligned}
\xi_{jt} \quad = \quad & \phi_{jt} - \beta^{l^s}_{nt}(k_{jt}, l^s_{jt}, l^u_{jt}, \mu_{jt}, \mathbf{m}_{jt}, \theta_{jt}) l^s_{jt} - \beta^{l^u}_{nt}(k_{jt}, l^s_{jt}, l^u_{jt}, \mu_{jt}, \mathbf{m}_{jt}, \theta_{jt}) l^u_{jt} \qquad (3.7) \\
& - \beta^k_{nt}(k_{jt}, l^s_{jt}, l^u_{jt}, \mu_{jt}, \mathbf{m}_{jt}, \theta_{jt}) k_{jt} - \mathbf{m}'_{jt}\beta^m_t \\
& - \Lambda \Big[\phi_{jt-1} - \beta^{l^s}_{nt-1}(k_{jt-1}, l^s_{jt-1}, l^u_{jt-1}, \mu_{jt-1}, \mathbf{m}_{jt-1}, \theta_{jt-1}) l^s_{jt-1} \\
& - \beta^{l^u}_{nt-1}(k_{jt-1}, l^s_{jt-1}, l^u_{jt-1}, \mu_{jt-1}, \mathbf{m}_{jt-1}, \theta_{jt-1}) l^u_{jt-1} \\
& - \beta^k_{nt-1}(k_{jt-1}, l^s_{jt-1}, l^u_{jt-1}, \mu_{jt-1}, \mathbf{m}_{jt-1}, \theta_{jt-1}) k_{jt-1} - \mathbf{m}'_{jt-1}\beta^m_{t-1} \Big]
\end{aligned}
$$

where $\phi_{jt} = \phi_{nt}(k_{jt}, l^s_{jt}, l^u_{jt}, \mu_{jt}, \mathbf{m}_{jt}, \theta_{jt})$.

The input coefficients are estimated by GMM by defining orthogonality conditions $E\left(\xi_{jt}(\beta^{l^s}_{nt}(\cdot), \beta^{l^u}_{nt}(\cdot), \beta^k_{nt}(\cdot), \beta^m_t)\mathbf{z}_{jt}\right) = 0$ with ξ_{jt} defined as in equation (3.7), and plugging in the estimates $\widehat{\phi}_{nt}(\cdot)$ from the first stage into (3.7).[16] The vector $\mathbf{z}_{jt}$ is a vector of instruments. Suitable instruments are k_{jt}, l^u_{jt}, l^s_{jt}, $\mathbf{m}_{jt-1}$ and higher order lags of these same variables. The non-linear search is performed over the parameters β^m_t and the non-parametric functions $\beta^{l^s}_{nt}(\cdot)$, $\beta^{l^u}_{nt}(\cdot)$, $\beta^k_{nt}(\cdot)$, while the non-parametric function $\Lambda(\cdot)$ is estimated in each iteration. Estimates of the heterogenous input coefficients are obtained from the estimates of the non-parametric functions as $\widehat{\beta}^{l^s}_{jt} = \widehat{\beta}^{l^s}_{nt}(k_{jt}, l^s_{jt}, l^u_{jt}, \mu_{jt}, \mathbf{m}_{jt}, \theta_{jt})$, $\widehat{\beta}^{l^u}_{jt} = \widehat{\beta}^{l^u}_{nt}(k_{jt}, l^s_{jt}, l^u_{jt}, \mu_{jt}, \mathbf{m}_{jt}, \theta_{jt})$, and $\widehat{\beta}^k_{jt} = \widehat{\beta}^k_{nt}(k_{jt}, l^s_{jt}, l^u_{jt}, \mu_{jt}, \mathbf{m}_{jt}, \theta_{jt})$.

Finally, the Hicks-neutral parameter β^0_{jt} is estimated using all previous estimates, by noting that from equations (2.3) and (3.6)

$$\beta^0_{jt} = \phi_{jt} - \beta^{l^s}_{jt} l^s - \beta^{l^u}_{jt} - \beta^k_t k_{jt} - \mathbf{m}'_{jt}\beta^m_t. \qquad (3.8)$$

As discussed by Gandhi, Navarro, and Rivers (2013), the coefficients of flexible inputs (the intermediate inputs in our case) are non-parametrically not identified. The identification of the β_m's in our setting relies on the Cobb-Douglas functional form assumption and on the non-linear price schedules for intermediate inputs.[17] The Cobb-Douglas technology is thus a structural assumption in our model and not merely an approximation to the true shape of the production function. The intuition behind why we need non-linear intermediate input prices is simple. With a Cobb-Douglas production function and constant

[16] Alternatively, orthogonality conditions can be defined in terms of the composite error term $\xi + \eta$.

[17] In contrast, ACF assume a fixed coefficient technology and avoid estimating the coefficient on intermediate inputs. In a case of predetermined capital and labor this implies that the demand for the intermediate input does not depend on its price. Also, for the fixed coefficient assumption to work, TFP needs to apply only to labor and capital. See Gandhi, Navarro, and Rivers (2013) for a discussion of output versus value added production functions.

input prices, if inputs are chosen optimally and are perfectly flexible, then all inputs are perfectly collinear with the productivity shocks observed by the firms. It is clear then that we need some differential "friction" across the flexible inputs so that they do not respond in a collinear way to changes in $\beta_{jt}^{l^s}$, $\beta_{jt}^{l^u}$, β_{jt}^{k}, and β_{jt}^{0}. Different non-linear price schedules provide such a friction. This idea is similar in spirit to the one proposed by Bond and Söderbom (2005). Using a Cobb-Douglas production function and constant input prices, they propose the introduction of input adjustment costs to attain identification. Closely related, Doraszelski and Jaumandreu (2015) also attain parametric identification of the coefficients on flexible inputs with constant prices but assuming a CES production function (the "friction" in their case comes from the elasticity of substitution being different from one). While both the method in Doraszelski and Jaumandreu (2015) and our method require a parametric assumption on the production function, the former cannot accommodate non-linear input prices[18] and requires the researcher to observe input prices.

A different concern regarding the identification of the coefficients on flexible inputs has to do with $\mathbf{m}_{jt-1}$ potentially being weak instruments. Differently from the traditional setup in which firms only differ in their TFP shock, in our setting, $\beta_{jt}^{l^s}$, $\beta_{jt}^{l^u}$, and β_{jt}^{k} are serially correlated. Since intermediate inputs depend on the productivity shocks (equation (2.6)), lagged and current intermediate input choices are going to be correlated through channels that are not controlled for in (3.7).

3.3 Discussion

Our model and estimation method integrate three features into the existing framework of the productivity and production function estimation literature: it allows for firm-heterogeneity that is not Hicks-neutral (heterogeneous labor and capital coefficients), it develops a strategy to deal with output quantity information and multiproduct firms, and it provides a framework to study quality choice and endogenous output and input prices. We argue that these three features can be introduced into the model and estimation strategy separately, since they do not rely on one another. Here we discuss the assumptions and implications of each of the three contributions in more detail.

First of all, the estimation of heterogeneous labor coefficients does not rely on having quantity information or a setting of vertical differentiation. The strategy to estimate heterogeneous labor coefficients can be applied to datasets with information on revenue and

[18]This is problematic, since electricity pricing usually takes the form of a two part tariff, and quantity discounts for materials are prevalent in practice.

input use. In this case revenue is used as the left-hand side variable y in regression equation (2.3).

Second, regarding our strategy to deal with output quantity data, it does not rely on labor coefficients being heterogeneous or in a particular demand system, as long as it satisfies (SYM). It can be applied to homogenous coefficient methods such as ACF. The transformation rates are estimated in an initial stage to express output quantity in units of equivalence. The control functions in all stages need to be adapted to having a sole dimension of firm heterogeneity (and one proxy).

Finally, we integrate a demand system and quality choice into the dynamic firm model. We study quality choice and exports in our application and thus we need to be clear about how we treat these variables empirically. Our assumption is that all firms face the same price schedule for intermediate inputs. The implication is that differences in prices of intermediate inputs across firms are endogenous and that thus input prices are not a state variable in equation (2.6), the policy function that is inverted to express unobserved technology as a function of the proxies. This assumption could be relaxed as in LP and DLGKP to allow for exogenous variation in input prices at the regional level. We can actually go further and relax the assumption that the input price schedules could differ only based on observables. In fact, we can allow each input price schedule to depend on a (scalar) iid supply shock that varies across firmsas long as the schedule is strictly monotone in the supply shock. The right-hand-side of (2.6) would now include the unobservable input supply shocks, but note that given the intermediate input choices, the observed input prices are sufficient statistics for the intermediate supply shocks.

The estimation of heterogeneous coefficients can be adapted to other model specifications and assumptions. The production function regression (2.3) could be written as a function of different input combinations, for example, merging skilled and unskilled labor as a sole labor input, and merging materials, electricity and fuel as a sole intermediate input. Likewise the coefficients on skilled labor, unskilled labor and capital need not all be heterogeneous across firms. We adopt the specification that provides the largest flexibility that is compatible with our dataset, and that allows us to test the hypothesis that differences in the productivity of skilled and unskilled workers explains quality and export destinations.

There is also flexibility regarding which variables to use as proxies. We discuss an estimation algorithm that is based on using static choice variables as proxies (three intermediate inputs and quality).[19] It is also possible to use some or all of the dynamic choice variables as proxies (equations 2.6 and 2.8). For example, when data on quality and separate use of electricity and fuel is not available, the four proxies could be materials, electricity+fuel, forward skilled labor, and forward unskilled labor. In the latter case, the algorithm needs to be modified to include the forward transformation rates (μ_{jt+1}) as an additional state variable. Investment can be used as proxy as well, subject to the problem of frequent zeros pointed out by LP. The number of products is also a dynamic choice variable but it is not an adequate proxy as it depends on the unobserved full vector of sorted products, including those that the firm chooses not to produce (equation 2.7). As a general rule, if the number of predetermined inputs is $(n-1)$, it is possible to have n dimensions of heterogeneity across firms (including TFP, β^0_{jt}); while the estimation of n dimensions of heterogeneity requires n current decision variables as proxies.

Other model assumptions are also susceptible to change. In particular, some specifications of our estimation method are amenable to adding certain types of demand shocks. Demand shocks would enter the static policy function (2.6) and because they are unobserved they invalidate using static choice variables as proxies. Dynamic choice variables can be used as proxies instead, as long as the demand shocks are i.i.d.[20] Finally, in our model we include labor adjustment costs to highlight that labor can be treated symmetrically to capital. This assumption is not necessary either for the model or for the estimation method. Labor does need to be predetermined, however, in order to be able to estimate heterogeneous labor coefficients.

4 Data and Results [PRELIMINARY AND INCOMPLETE]

In this section we first describe the data we use, we then present the estimates for the production function parameters and, finally, we discuss the results regarding quality, technology, and export destinations. As a preview of the results, our findings are the following. First, both output and input quality (including both labor and intermediate inputs) are determined by differences in the production technology. In particular, firms that have an advantage in the use of skilled labor (i.e., firms that have an advantage in the production of quality) use (relatively) more skilled labor and higher quality intermediate inputs, and

[19] In the empirical implementation we use output unit values and material unit values as two alternative measures of quality.

[20] In the case of a multinomial logit demand system, the utility function can be written as $U_{ijrct} = \alpha(x_{ct})\theta_{jt} - p_{jrct} + \bar{\epsilon}_{jrct} + \epsilon_{ijrct}$., where $\bar{\epsilon}_{jrct}$ is an i.i.d. variety-destination shock.

produce higher quality output. Second, we show that the main hypothesis that we propose in this paper —namely, that firms that have an advantage in the production of quality tend to export to destinations with higher income— holds in our data.

4.1 Data

Our data is a panel of Chilean manufacturing firms spanning the years 1996–2006. The data comes from the ENIA (Annual National Industrial Survey) and is collected by Chile's National Institute of Statistics (INE). It surveys all manufacturing firms with 10 or more employees. The main module of the survey has detailed information on industry affiliation, revenue, skilled and unskilled labor, investment, and intermediate inputs. This dataset, spanning different years, has been used by Pavcnik (2002), LP, and several other studies. In addition to the (firm-level) aggregated data, survey module 3 collects detailed information on firm's output by product, which is what allows us to estimate the production parameters using units of physical quantity instead of revenue. The module contains product-level information on units by product. Products are defined at the 7-digit-level of disaggregation following INE's own classification system.

We have augmented the input/output data from ENIA with customs data by matching firms using their tax identification numbers. For each firm in the panel, we have information on the value of exports by country of destination. These data are key for our empirical application in which we establish an empirical link between technology and the destination of exports.

Construction of variables.

The construction of the inputs and output variables follows from Lui (1991), LP, and Greenstreet (2005). In particular, our measures of revenue, capital, materials, electricity, and fuels are deflated with their own annual price deflator (constructed by the Banco Central de Chile) and deflated to real 1995 Chilean pesos. Our measure of labor is the number of man-years used in production, and it is broken down into skilled (white-collar) and unskilled (blue-collar) workers. In the case of labor and intermediate inputs, we also purge our measures from differences in quality across firms. For the labor variables, we deflate them by the firm-specific wages relative to the industry. Similarly, for the intermediate inputs, we construct and use a firm-specific price index (see details below).

As mentioned earlier, we complement the input/output data with customs data. We construct a dummy variable that identifies the firms that are engaged in exporting activity in

Table 1: Summary Statistics

	Mean	Std. Dev.	10-th perc.	50-th perc.	90-th perc.	N
Labor force	71	138	13	29	156	65,180
Skilled	26	71	2	11	50	65,180
Unskilled	45	96	6	17	101	65,180
Share skilled	0.38	0.34	0.04	0.26	1.00	64,838
Materials	2,121,849	19,520,236	20,396	141,471	2,532,223	65,180
Electricity	80,737	921,310	445	3,068	52,503	65,180
Fuel	48,779	514,963	0	2,240	37,739	65,180
Capital	338,800	2,575,033	2,239	24,284	454,803	63,927
Investment	158,312	3,401,951	0	423	134,971	65,180
Investment>0	0.54	.	.	.	.	65,180
Revenue	3,367,215	25,558,149	36,718	279,152	4,597,998	65,180
Exporter	0.17	.	.	.	.	63,050
Exporter to high income dest.	0.81	.	.	.	.	11,009
Main dest. is high income	0.52	.	.	.	.	11,009
GDP dest (mean)	12,533	10,927	2,069	7,408	28,782	11,009
GDP dest (main)	13,342	13,510	879	7,408	30,888	11,009

a given year. Since we also have rich data at the firm level on destinations, we also construct several measures related to export destinations: a dummy variable that identifies if the firm exports to high-income countries,[21] a dummy variable that equals 1 if the firm's main destination (measured by sales) is a high-income country, average GDP across destinations (sales weighted), and GDP of the main destination. We report summary statistics of the main variables we use in our empirical analysis in Table 1.

Proxies for quality of output and intermediate inputs.

Our theoretical model endogenizes firm's choice of output and intermediate inputs qualities (θ_{jt} and θ^m_{jt}, respectively). Unfortunately, these variables are not directly observed in the data. But, to the extent that prices reflect quality, we can use them to construct proxies for quality that can be later used in the quality choice regressions.

We construct these proxies as follows. We work with two alternatives that exploit the fact that, in our data, we observe multiple products and intermediate inputs for a given firm-year. In the first alternative, the quality of output (intermediate inputs) is estimated as a firm-year fixed effect from an OLS regression of firm-product (firm-intermediate inputs)

[21] We experimented with different definitions: high-income as classified by the World Bank, GDP above the mean, GDP above the 75th percentile. We obtain similar results regardless of which definition we use.

prices (in logs) that also includes product-year (input-year) dummies. That is, we run the following regressions

$$p_{jrt} = \psi_{jt,\mathrm{FE}} + \widetilde{\psi}_{rt} + \epsilon_{jrt} \tag{4.1}$$

$$p^m_{jmt} = \psi^m_{jt,\mathrm{FE}} + \widetilde{\psi}^m_{mt} + \epsilon^m_{jmt} \tag{4.2}$$

where p_{jrt} is the price (in logs) of product r produced by firm j at time t, p^m_{jmt} is the price (in logs) of intermediate input m used by j at t, $\psi_{jt,\mathrm{FE}}$ and $\psi^m_{jt,\mathrm{FE}}$ are firm-year fixed effects, and $\widetilde{\psi}_{rt}$ and $\widetilde{\psi}^m_{mt}$ are product-year and intermediate input-year fixed effects, respectively. The coefficients ψ_{FE} and ψ^m_{FE} are thus interpreted as the average deviation in prices of output and intermediate inputs at the firm-level relative to the market average.

In the second alternative, quality is computed as a Stone Index of firm-level price deviations relative to the product ($\overline{p}_{rt}$) or input means ($\overline{p}_{mt}$), using the share in revenue and in cost as weights (λ):

$$\psi_{jt,\mathrm{stone}} = \sum_{r\in R_j} \lambda_{jrt} \left(p_{jrt} - \overline{p}_{rt}\right) \tag{4.3}$$

$$\psi^m_{jt,\mathrm{stone}} = \sum_{m} \lambda_{jmt} \left(p_{jmt} - \overline{p}_{mt}\right). \tag{4.4}$$

4.2 Estimates of the Technology Parameters

For the estimation of the technology parameters we split firms in groups of 2-digit industries of the ISIC Revision 3 classification and consider three time periods. Thus, we allow the function $\phi_{nt}(\cdot)$ in equation (3.6) to vary at the industry-period level to reflect differences in the homogeneous coefficients, demand, and input prices across industries and periods. We estimate the homogeneous parameters (β^m_t) at the industry-period level and the heterogeneous parameters ($\beta^{l^s}_{jt}$, $\beta^{l^u}_{jt}$, β^k_{jt} and β^0_{jt}) at the firm-year level.
[IN WHAT FOLLOWS, WE ASSUME THAT β^k DOES NOT VARY ACROSS FIRMS. HENCE, WE ONLY NEED TO USE 3 PROXIES. WE USE THE 3 INTERMEDIATE INPUTS AS PROXIES]

Using data on output quantity from single product firms and input use from the main module we estimate the rates of transformation across products in the same industry (μ^r_t). With the estimated transformation rates we transform physical output into units of equivalence relative to a numéraire product and add up products in units of equivalence to compute total firm physical output y_{jt}. Just as an illustration, in Table 2 we consider two firms, each producing three products. The first firm belongs to the food industry and produces dried

Table 2: Transformation Rates

Product Code	Description	Unit	$\hat{\mu}$
Example 1: firm 10001, year 1996			
3114908	dried fish	kg	1
3114101	frozen fish	kg	1.06
3114121	frozen seafood	kg	.43
Example 2: firm 10002, year 2006			
3220153	slacks	u	1
3220105	pants	u	.37
3220160	reflective vests	u	.66

Table 3: Production Function Estimates (Revenue)

	Homog. Coeff.		Heterog. Coeff.			
	OLS	ACF	mean	10th perc.	50th perc.	90th perc.
skilled labor	0.154	0.144	0.146	0.132	0.147	0.159
unskilled labor	0.080	0.140	0.085	0.035	0.090	0.124
capital	0.132	0.094	0.090			
materials	0.523	0.578	0.578			
electricity	0.073	0.038	0.168			
fuels	0.033	0.067	0.037			
RTS	0.994	1.062	1.105	1.064	1.107	1.138
nobs	39833	30141	28248			
$\beta^s < 0$ (%)			0.01			
$\beta^u < 0$ (%)			5.82			

and frozen fish and frozen seafood. From our estimates of the transformation rates we see that the firm can transform the inputs used to produce 1 Kg. of dried fish into 943 grams of frozen fish or 2.3 Kg. of frozen seafood. The second firm is in the apparel industry and produces slacks, pants, and reflective safety vests. The firm can transform the inputs used in the production of one pair of slacks into 2.7 pairs of pants, or 1.5 reflective vests.

For comparison, we use both the composite physical output y and revenue as left-hand side variables. Table 3 shows the estimates of the production function parameters using revenue as our measure of output, and Table 4 using the physical output. In each case, we present estimates from OLS and a control function method that assumes homogeneous labor coefficients across firms, and from our method. The tables also present the returns-to-

Table 4: Production Function Estimates (Quantity)

(a) All firms

	Homog. Coeff.		Heterog. Coeff.			
	OLS	ACF	mean	10th perc.	50th perc.	90th perc.
skilled labor	0.041	0.018	0.038	0.021	0.030	0.069
unskilled labor	0.056	0.028	0.057	0.021	0.057	0.093
capital	0.041	0.028	0.031			
materials	0.658	0.693	0.695			
electricity	0.146	0.190	0.151			
fuels	0.050	0.042	0.058			
RTS	0.992	0.998	1.030	1.009	1.027	1.052
nobs	21749	15153	15135			
$\beta^s < 0$ (%)			0.17			
$\beta^u < 0$ (%)			0.13			

(b) Single product firms

	Homog. Coeff.		Heterog. Coeff.			
	OLS	ACF	mean	10th perc.	50th perc.	90th perc.
skilled labor	0.034	0.024	0.026	0.008	0.018	0.054
unskilled labor	0.059	0.045	0.072	0.047	0.076	0.092
capital	0.042	0.023	0.015			
materials	0.664	0.681	0.701			
electricity	0.132	0.159	0.139			
fuels	0.065	0.071	0.066			
RTS	0.995	1.004	1.019	1.003	1.017	1.036
nobs	13924	9341	8736			
$\beta^s < 0$ (%)			0.73			
$\beta^u < 0$ (%)			5.23			

Figure 2: Distributions of Skilled and Unskilled Labor Productivity

(a) Marginal density

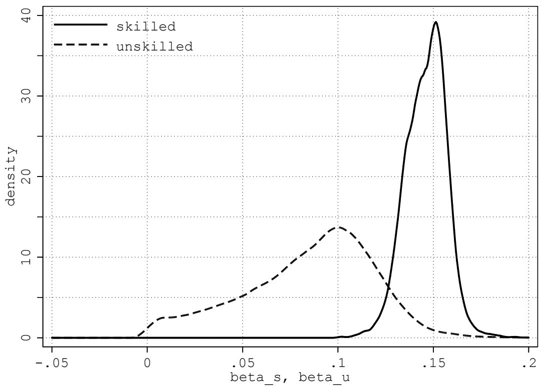

(b) Joint density

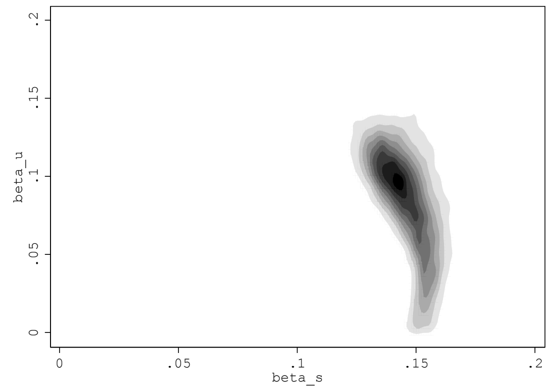

scale (RTS) and, for our method, some moments of the distribution of the labor coefficients and RTS. Note that, due to the endogeneity and nonlinearity of the input choices, there is no reason why the methods that assume that the labor coefficients are homogeneous will even deliver an estimate of the mean of the distribution of the heterogeneous coefficients.[22] Therefore, the usual exercise of comparing the estimates from a method that takes care of the endogeneity problem with those from OLS to see if the OLS bias goes in the right direction is not straightforwardly applicable to our method.

While it has been well established in the literature that firms are heterogeneous in terms of their Hicks-neutral productivity, we now show that there is also non-trivial heterogeneity in terms of firms' labor productivities: β^{l^s} and β^{l^u}. Figure 2 plots the marginal densities of the skilled and unskilled labor productivity parameters in Panel (a) and a contour level map of the joint density in Panel (b). In both panels we use a Kernel estimator using the firm level coefficients recovered by our method as data. While a firm in the 90th percentile of the skilled labor productivity distribution has an output elasticity 20% higher than one in the 10th percentile —all else equal— a firm in the 90th percentile of the unskilled labor productivity distribution has an output elasticity 3.5 times higher than one in the 10th percentile. Consistent with the previous literature we also find substantive heterogeneity in terms of TFP: a firm in the 90th percentile of the distribution of TFP is 45% more productive than one in the 10th percentile, all else equal.

[22] It is easy to show that in a linear model with random coefficients and no endogenous regressors, OLS gives an estimate of the mean of the distributions of the coefficients. The problem, in our setting, is that input choices are endogenous.

Figure 3: TFP

(a) Marginal density (b) BBS vs ACF

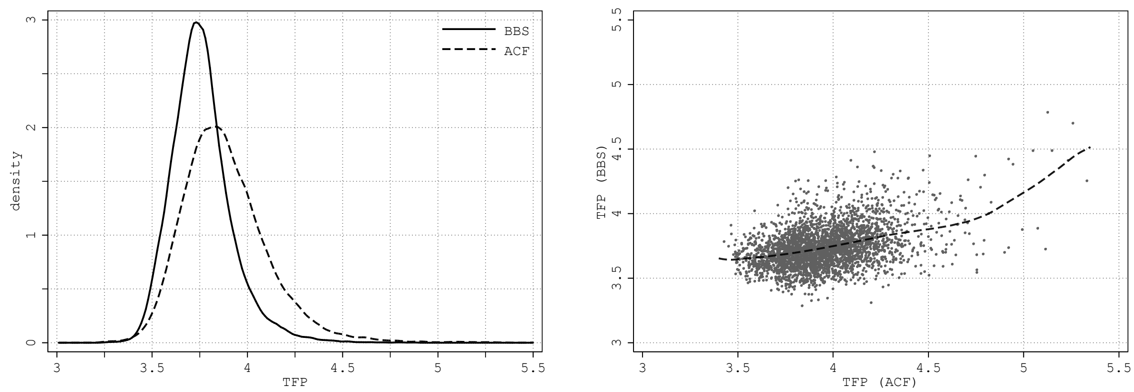

The (Kernel estimate) density of TFP is plotted in Figure 3. There, we also plot the TFP distribution that we obtain from the method that imposes homogeneous labor productivity parameters. Intuitively, the latter distribution has a higher variance since the TFP, which is obtained as a residual, picks up any heterogeneity that has not been accounted for, e.g., in the labor productivities. In Panel (b) we show that there is a positive, though noisy, relationship between the TFP estimates from the two methods. The positive correlation between the two estimates is reassuring, but the relationship also points out that methods that neglect other forms of unobserved heterogeneity might result in biased estimates of TFP. In the Online Appendix, we perform a Monte Carlo exercise to show that if the researcher ignores other sources of heterogeneity, the TFP estimates are biased and inconsistent.

Our estimation method assumes that the three heterogeneous productivity parameters evolve according to a first-order Markvov process, but it poses no assumptions on the contemporaneous correlation structure between them. To further investigate the correlations between the productivity parameters, in Figure 4 we show (pairwise) scatter plots of TFP, skilled, and unskilled productivity. While there appears to be no correlation between TFP and skilled labor productivity (Panel (a)), there is a positive correlation between TFP and unskilled productivity (Panel (b)). Additionally, as we can see from Panel (c) in Figure 4 and Panel (b) in Figure 2, there appears to be a negative correlation between skilled and unskilled productivities.

Figure 4: Correlations between the Productivity Parameters

(a) β^0 and β^s

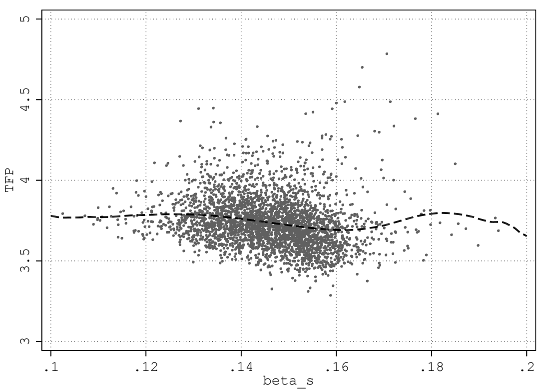

(b) β^0 and β^u

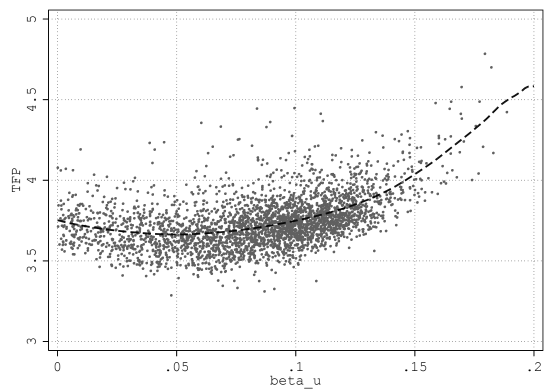

(c) β^u and β^s

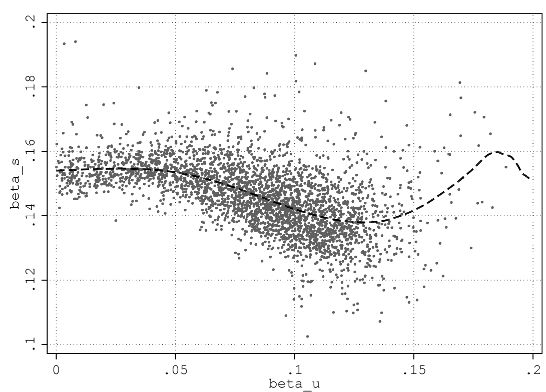

4.3 Technology, Quality, and Export Destinations

Delving into the main hypothesis that we present in this paper, we now look at how differences in the production technology affect exporting behavior and, in particular, export destinations. The logic behind the hypothesis, which was presented in Section 2, builds upon three key pillars: (i) firms that have an advantage in the use of skilled labor will use (relatively) more skilled labor; (ii) firms that use higher quality inputs produce higher quality output; and (iii) countries with higher GDP value quality more and thus are willing to pay for high-quality goods. Putting these three pieces together, our main hypothesis states that firms that have an advantage in the use of skilled labor will end up exporting to countries with higher GDP. We show that the three relationships presented above and the main hypothesis hold in our data from Chile.

While point (i) in the previous paragraph is new in the empirical trade literature, points (ii) and (iii) have already been shown using data from other countries. These studies look at the complementarities between input and output qualities and at the relationship between export decisions —in particular, the decision to export to high-income countries— and the quality of the output. Notice, however, that these relationships are between endogenous outcomes. Our contribution is to go one step forward and show the following. First, that output and input quality are both determined by differences in the production technology. Thus, firms that have a technological advantage in the production of quality will produce higher quality output. Second, that the firms that have an advantage in the production of quality tend to export to destinations with higher income, namely, our main hypothesis.

Complementarities in Quality and Technology Determinants.

It has already been shown in the literature that the production of high quality output is associated with employing more skilled workers and high quality intermediate inputs. Verhoogen (2008), Brambilla, Lederman, and Porto (2012), and Brambilla and Porto (2015) establish a causal relationship between high quality output and wages (and skills) in Mexico, Argentina, and for a panel of 82 countries. Bastos, Silva, and Verhoogen (2014) establish a similar positive relationship between the quality of the output product and the quality of intermediate inputs in Portugal.

We show that similar relationships can be found in our data from Chile. In Table 5, we see that our measure of output quality, $\psi_{jt,FE}$, (see above for details) is positively correlated with (i) our measure of labor quality (the ratio of skilled to total labor force: $\frac{L^s_{jt}}{L^s_{jt}+L^u_{jt}} \in [0,1]$)

Table 5: Complementarities in Quality and Skills

	Output Quality			Mat. Quality
	(1)	(2)	(3)	(4)
Skilled ratio	0.353***		0.323***	0.107***
	(0.062)		(0.061)	(0.033)
Materials quality		0.243***	0.241***	
		(0.017)	(0.017)	
Time FE	yes	yes	yes	yes
Industry FE	yes	yes	yes	yes
Obs.	28790	27947	27947	29755

and (ii) with our measure of intermediate inputs quality, $\psi^m_{jt,FE}$ (see above for details).[23] All regressions include year and industry fixed effects. Moreover, we also show that labor quality and intermediate inputs quality are positively correlated (column (4)), suggesting complementarities in the quality of inputs.

The previous relationships involve endogenous variables. In the model we present in Section 2, firm's choices of input and output quality (θ^m_{jt} and θ_{jt}, respectively) are endogenous outcomes that ultimately depend on the technology parameters. We show that next. First, we establish that both labor and intermediate inputs quality are positively correlated with the skilled productivity ratio ($\frac{\beta^s_{jt}}{\beta^s_{jt}+\beta^u_{jt}} \in [0,1]$) our measure of technological advantage in the production of quality. In Table 6 we see that, the higher the skilled productivity ratio, the higher the quality of labor (column (1)) as measured by the skilled labor ratio, and the higher the quality of the intermediate inputs (column (5)).[24] These positive relationships are still present once we also control for TFP (columns (2) and (6)). We also see a positive effect of TFP on both the skilled labor ratio and intermediate inputs quality.[25] Finally, consistent with the result in Proposition 1, we see from column (7) that the higher the skilled productivity ratio, the higher the quality of the output product. In column (8) we also control for TFP and the quality of the intermediate inputs. As expected, output quality is positively associated with skilled productivity ratio, TFP, and the quality of the intermediate inputs.

[23] We obtain virtually the same qualitative and quantitative results if we use our alternative measures of quality: $\psi_{jt,stone}$ and $\psi^m_{jt,stone}$. See the Online Appendix for these results.

[24] In the case of materials quality, the relationship is increasing in the skilled productivity ratio except at the high end of the ratio.

[25] TFP is in logs and ranges from 3.45 to 4.2.

Table 6: Technology and Quality

	Skilled Ratio				Mat. Qual.		Output Qual.	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Skilled prod. ratio	0.245***	0.306***	0.200***	0.263***	2.384***	3.917***	0.202*	0.321***
	(0.023)	(0.024)	(0.025)	(0.026)	(0.801)	(0.848)	(0.115)	(0.119)
Skilled prod. ratio sq					−1.719***	−2.713***		
					(0.572)	(0.597)		
TFP		0.765***		0.754***		3.400***		3.398**
		(0.142)		(0.139)		(1.204)		(1.570)
TFP sq		−0.083***		−0.082***		−0.407***		−0.409**
		(0.018)		(0.018)		(0.156)		(0.204)
Exporter			0.029***	0.027***				
			(0.006)	(0.006)				
Materials quality								0.101***
								(0.011)
Time FE	yes	yes	yes	yes	yes	yes	yes	yes
Industry FE	yes	yes	yes	yes	yes	yes	yes	yes
Obs.	36459	36459	35224	35224	29755	29755	29317	28419

Figure 5: Skilled and Skilled Productivity Ratios

(a) Skilled Ratio ($L^s/(L^s + L^u)$) (b) Skilled Productivity ($\beta^s/(\beta^s + \beta^u)$)

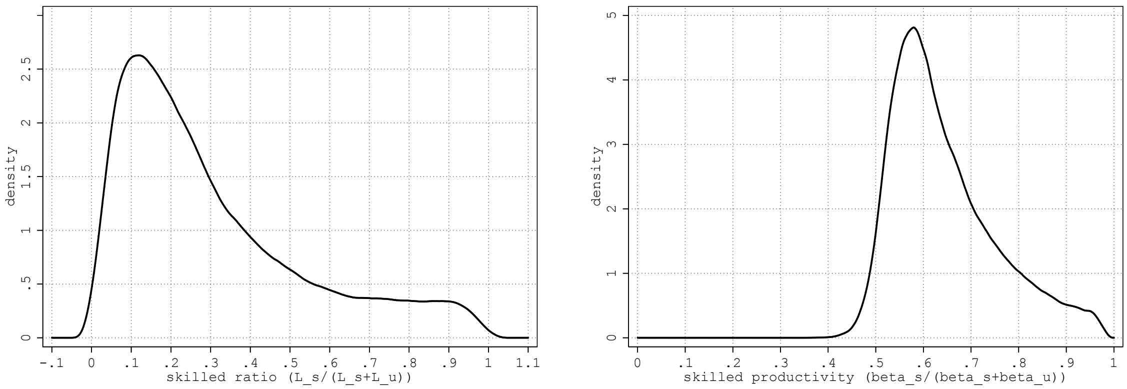

Quality, Exporting Behavior and Technology Determinants.

Technology differences across firms have been shown to affect exporting behavior. For example, Melitz (2003) describes a selection mechanism in which more productive firms (in terms of TFP) engage into exporting activity. Empirically, many papers have shown that there is a relationship between total factor productivity and the decision to export (Bernard and Bradford Jensen (1999), Clerides, Lach, and Tybout (1998), and Tybout (2003)). In recent years, it has been argued that product quality, in addition to firm productivity, is also a source of heterogeneity in exporting behavior among firms. Specifically, that firms that are capable of producing high-quality output tend to export more and to high-income markets of destination. The argument is that high-income countries value quality more and are willing to pay for high-quality goods.[26] With the increasing availability of detailed customs data, there is vast empirical evidence supporting these claims. A positive correlation between income in the destination market and product quality has been reported by Manova and Zhang (2012), Bastos and Silva (2010), Görg, Halpern, and Muraközy (2010) and Martin (2012) for firms in China, Portugal, Hungary and France.

We can also establish similar results in our data. First, in Panel (a) of Figure 5 we perceive a great deal of heterogeneity across firms in terms of our definition of labor quality. What it is more important for us, is that firms engaged in exporting tend to use a higher ratio of skilled labor (see Figure 6, Panel (a)). Similarly, firms that export to high-income

[26] Verhoogen (2008); Hallak and Sivadasan (2013); Baldwin and Harrigan (2011); Johnson (2012); Brambilla, Lederman, and Porto (2012); Feenstra and Romalis (2012). An additional channel linking high-income markets and high-quality output is the "Washington apples" effect (Hummels and Skiba (2004)).

Figure 6: Distribution of Skill Ratio

(a) Exporter (b) High Income Exporter

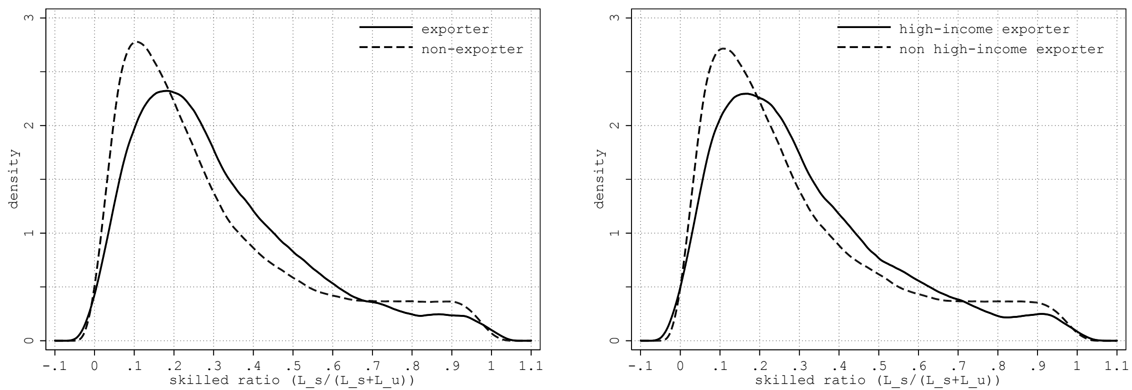

destinations tend to use a higher quality of labor (see Figure 6, Panel (b)). To further investigate the connections between exporting behavior and input and output quality we present regression results in Table 7. The table shows that there is, in fact, a positive correlation between exporter status and the quality of the output (column (1)), and exporter status and the quality of the inputs (both labor —column (2)— and intermediate inputs —column (3)). The correlations remain when we include all the variables at the same time (column (5)). Moreover, we also observe similar correlations if, instead of looking at whether the firms exports or not, we look at whether the firm exports to a high-income destination (columns (6)-(10)).[27]

Once again, the relationships described in the previous paragraph involve endogenous outcomes. Next, we show that the exporting behavior can be explained from technological differences across firms above and beyond TFP. We start by documenting some differences across exporter and non-exporter firms. In Figure 7, Panel (a), we notice that while exporters (i) tend to have higher skilled labor productivity (and are somewhat more homogeneous in terms of it) and (ii) tend to have lower unskilled labor productivity (and are somewhat more heterogeneous in terms of it), they are clearly more likely to have a higher skilled labor productivity ratio, our key technological heterogeneity that we will exploit in the trade regressions below.[28] Consistent with previous studies, in Table 8 we show that higher TFP not only increases the likelihood that a firm becomes an exporter (columns (2)

[27] We obtain qualitatively similar results if we use the ratio L^s/L^u for our graphs and regressions. See the Online Appendix for these results.

[28] We obtain qualitatively similar results if we use the ratio β^s/β^u. See the Online Appendix for these results.

Table 7: Quality and Exports

	Exporter					High-income Exporter				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Output quality	0.010***				0.006**	0.009***				0.006**
	(0.003)				(0.003)	(0.003)				(0.003)
Materials quality		0.010***			0.008***		0.009***			0.007***
		(0.003)			(0.003)		(0.003)			(0.003)
Skilled ratio			0.136***		0.060***			0.100***		0.037***
			(0.015)		(0.014)			(0.013)		(0.012)
Log sales				0.115***	0.113***				0.098***	0.098***
				(0.003)	(0.003)				(0.003)	(0.003)
Time FE	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
Industry FE	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
Obs.	29410	28775	35224	35224	28332	29410	28775	35224	35224	28332

Figure 7: Labor Productivity and Exporting

(a) Exporter (b) High Income Exporter

Skilled Labor Productivity

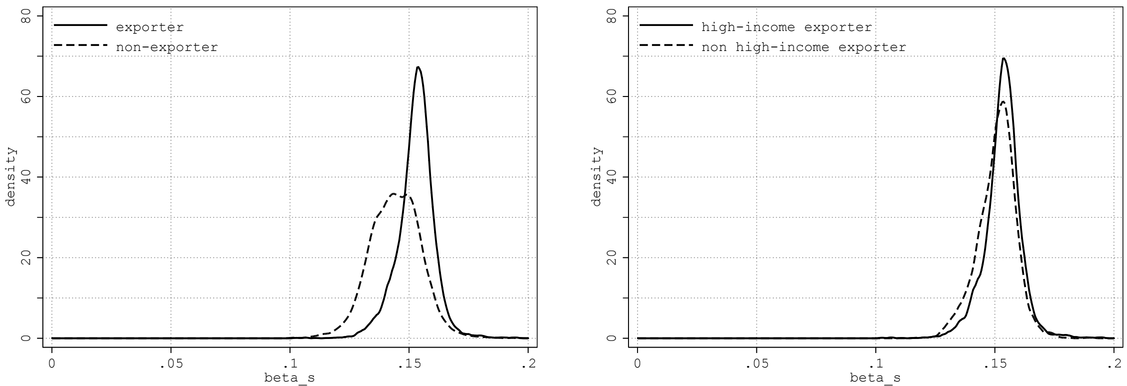

Unskilled Labor Productivity

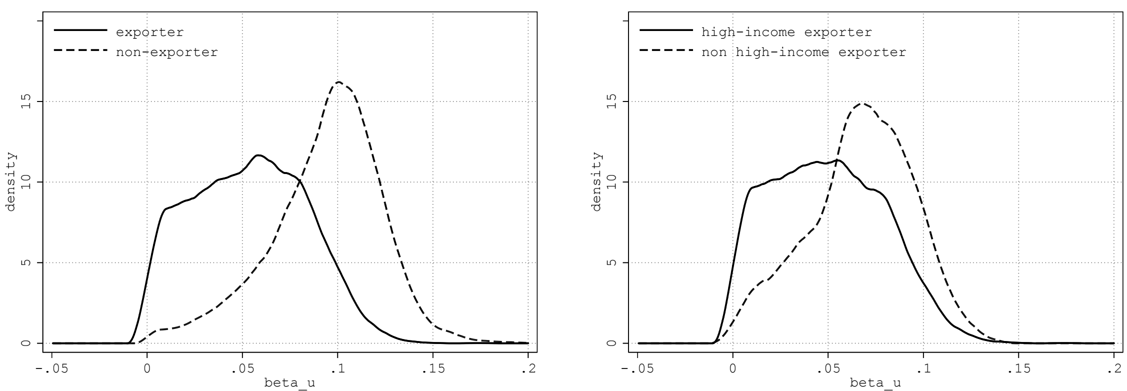

Skilled Productivity ($\beta^s/(\beta^s+\beta^u)$)

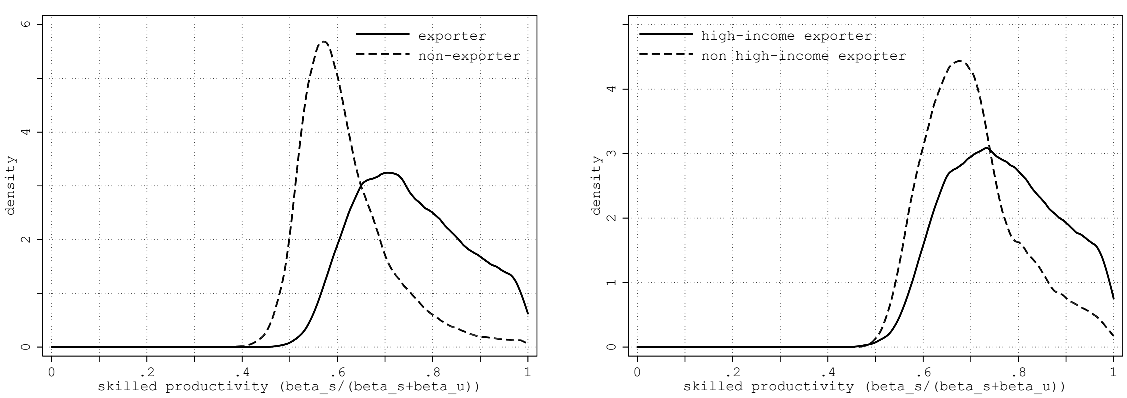

and (3)) but also the size of the firm's exports (columns (5) and (6)). We go beyond differences in TFP to also claim that a higher skilled productivity ratio increases the likelihood that a firm becomes an exporter and, conditional on exporting, it also increases the size of the firm's exports (columns (1) and (4)), even after controlling for TFP and sales (columns (3) and (6)).

To sum up, we have shown so far the three key components that drive our main hypothesis: (i) firms that have an advantage in the use of skilled labor will use more skilled labor (relative to unskilled labor); (ii) firms that use higher quality inputs produce higher quality output; and (iii) countries with higher GDP value quality more and thus will be more likely to buy higher quality products.

Now, moving on to our main hypothesis, we want to show that firms that have an advantage in the use of skilled labor end up exporting to countries with higher GDP. We begin by looking at differences in the distribution of the skilled productivity ratio across exporters to high-income destinations and other destinations. We see from Panel (b) in Figure 7 that exporters to high-income destinations are more likely to have a higher skilled productivity ratio. Next, we show robust evidence to our claim that labor productivity is a key determinant of export destinations. We consider several outcome variables related to export destinations. We find that skilled productivity ratio is positively correlated with: the decision to export to high-income countries, even if we condition on exporting (Table 9); the share of sales to high-income destinations (Table 10); the GDP of the main destination, and the average GDP (weighted by sales) of the export destinations (Table 11).

Table 8: Technology and Exports

	Exporter			Log Exports		
	(1)	(2)	(3)	(4)	(5)	(6)
Skilled prod. ratio	1.625***	1.659***	1.364***	11.565***	11.726***	5.927***
	(0.044)	(0.047)	(0.083)	(0.330)	(0.332)	(0.652)
TFP		0.073***	0.074***		1.066***	0.483**
		(0.023)	(0.023)		(0.217)	(0.201)
Log sales			0.025***			0.567***
			(0.006)			(0.058)
Time FE	yes	yes	yes	yes	yes	yes
Industry FE	yes	yes	yes	yes	yes	yes
Obs.	35224	35224	35224	8134	8134	8134

Table 9: Technology and Export Destinations (I)

	High-income Exporter			High-income Exporter (cond. on exporting)		
	(1)	(2)	(3)	(4)	(5)	(6)
Skilled prod. ratio	1.439***	1.483***	1.414***	0.884***	0.883***	0.793***
	(0.044)	(0.046)	(0.080)	(0.075)	(0.075)	(0.135)
TFP		0.094***	0.094***		−0.011	−0.021
		(0.021)	(0.021)		(0.043)	(0.046)
Log sales			0.006			0.009
			(0.005)			(0.012)
Time FE	yes	yes	yes	yes	yes	yes
Industry FE	yes	yes	yes	yes	yes	yes
Obs.	35224	35224	35224	6899	6899	6899

Table 10: Technology and Export Destinations (II)

	Share of Sales to HI Dest.			Share of Sales to HI Dest. (cond. on exporting)		
	(1)	(2)	(3)	(4)	(5)	(6)
Skilled prod. ratio	0.277***	0.282***	0.327***	0.120**	0.113*	0.352***
	(0.022)	(0.024)	(0.042)	(0.060)	(0.060)	(0.126)
TFP		0.011	0.011		−0.057	−0.029
		(0.013)	(0.013)		(0.046)	(0.045)
Log sales			−0.004			−0.024**
			(0.003)			(0.011)
Time FE	yes	yes	yes	yes	yes	yes
Industry FE	yes	yes	yes	yes	yes	yes
Obs.	35224	35224	35224	6899	6899	6899

Table 11: Technology and Export Destinations (III)

	Mean Dest. GDP			Main Dest. GDP		
	(1)	(2)	(3)	(4)	(5)	(6)
Skilled prod. ratio	0.683***	0.656***	1.424***	0.557**	0.532**	1.495***
	(0.194)	(0.193)	(0.352)	(0.226)	(0.226)	(0.436)
TFP		−0.251**	−0.160		−0.224	−0.111
		(0.114)	(0.117)		(0.140)	(0.145)
Log sales			−0.077**			−0.097**
			(0.032)			(0.038)
Time FE	yes	yes	yes	yes	yes	yes
Industry FE	yes	yes	yes	yes	yes	yes
Obs.	6899	6899	6899	6899	6899	6899

References

ACKERBERG, D., L. BENKARD, S. BERRY, AND A. PAKES (2005): "Econometric Tools for Analyzing Market Outcomes," in *Handbook of Econometrics*, ed. by J. J. Heckman, and E. Leamer, vol. 6. Elsevier.

ACKERBERG, D., K. CAVES, AND G. FRAZER (2016): "Identification Properties of Recent Production Function Estimators," forthcoming in Econometrica.

AW, B. Y., M. J. ROBERTS, AND D. Y. XU (2011): "R&D Investment, Exporting, and Productivity Dynamics," *American Economic Journal: Microeconomics*, 101(4), 1312–1344.

BALDWIN, R., AND J. HARRIGAN (2011): "Zeros, Quality, and Space: Trade Theory and Trade Evidence," *American Economic Journal: Microeconomics*, pp. 60–88.

BASTOS, P., AND J. SILVA (2010): "The quality of a firm's exports: Where you export to matters," *Journal of International Economics*, 82(2), 99–111.

BASTOS, P., J. SILVA, AND E. VERHOOGEN (2014): "Export Destinations and Input Prices," working paper, NBER.

BERNARD, A. B., AND J. BRADFORD JENSEN (1999): "Exceptional Exporter Performance: Cause, Effect, or Both?," *Journal of International Economics*, 47(1), 1–25.

BOND, S., AND M. SÖDERBOM (2005): "Adjustment Costs and the Identification of Cobb Douglas Production Functions," working paper, Nuffield College.

BRAMBILLA, I., D. LEDERMAN, AND G. PORTO (2012): "Exports, Export Destinations, and Skills," *American Economic Review*, 102(7), 3406–38.

BRAMBILLA, I., AND G. PORTO (2015): "High-Income Export Destinations, Quality and Wages," *Journal of International Economics*.

CLERIDES, S. K., S. LACH, AND J. R. TYBOUT (1998): "Is Learning by Exporting Important? Micro-Dynamic Evidence from Colombia, Mexico, and Morocco," *The Quarterly Journal of Economics*, 113(3), 903–947.

DE LOECKER, J. (2011): "Product Differentiation, Multi-Product Firms and Estimating the Impact of Trade Liberalization on Productivity," *Econometrica*, 5, 1407–1451.

——— (2013): "Detecting Learning by Exporting," *American Economic Journal: Microeconomics*, 5(3), 1–21.

De Loecker, J., P. K. Goldberg, A. K. Khandelwal, and N. Pavcnik (2015): "Prices, Markups and Trade Reform," forthcoming in Econometrica.

Doraszelski, U., and J. Jaumandreu (2013): "R&D and Productivity: Estimating Endogenous Productivity," *Review of Economic Studies*, 80, 1338 – 1383.

——— (2015): "Measuring the Bias of Technological Change," Working Paper.

Feenstra, R. C., and J. Romalis (2012): "International Prices and Endogenous Quality," working paper, NBER.

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.

Gandhi, A., S. Navarro, and D. Rivers (2013): "On the Identification of Production Functions: How Heterogeneous is Productivity?," Working paper.

Görg, H., L. Halpern, and B. Muraközy (2010): "Why do within firm-product export prices differ across markets?," working paper.

Greenstreet, D. (2005): "Exploiting Sequential Learning to Estimate Establishment-Level Productivity Dynamics and Decision Rules," Ph.D. Thesis, University of Michigan.

Hallak, J. C., and J. Sivadasan (2013): "Product and Process Productivity: Implications for Quality Choice and Conditional Exporter Premia," *Journal of International Economics*, 91(1), 53–67.

Hummels, D., and A. Skiba (2004): "Shipping the good apples out? An empirical confirmation of the Alchian-Allen conjecture," *Journal of Political Economy*, 112(6).

Johnson, R. C. (2012): "Trade and Prices with Heterogeneous Firms," *Journal of International Economics*, 86(1), 43–56.

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), 343–361.

Kugler, M., and E. Verhoogen (2012): "Prices, Plant Size, and Product Quality," *The 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), 317–342.

Lui, L. (1991): "Entry-Exit and Productivity Changes: An Empirical Analysis of Efficiency Frontiers," Ph.D. Thesis, University of Michigan.

Manova, K., and Z. Zhang (2012): "Export Prices across Firms and Destinations," *The Quarterly Journal of Economics*, 127(1), 379–436.

Marschak, J., and W. Andrews (1944): "Random Simultaneous Equations and the Theory of Production," *Econometrica*, 12, 143–205.

Martin, J. (2012): "Markups, quality, and transport costs," *European Economic Review*, 56, 777–791.

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

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

Tybout, J. (2003): "Plant- and Firm-level Evidence on the 'New' Trade Theories," in *Handbook of International Trade*, ed. by E. K. Choi, and J. Harrigan. Oxford: Basil-Blackwell.

Verhoogen, E. A. (2008): "Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector," *The Quarterly Journal of Economics*, 123(2), 489–530.