Economic Inquiry

TRADE LIBERALIZATION AND FIRM PRODUCTIVITY: ESTIMATION METHODS MATTER

JOHN KEALEY, PAU S. PUJOLAS and CÉSAR SOSA-PADILLA*

In this paper, we show that the relationship between trade liberalization and firm productivity is sensitive to the method used to estimate the production function. We estimate the productivity of Colombian manufacturing plants using the methods of Levinsohn and Petrin, Ackerberg et al., and Gandhi et al. and at times come to surprisingly different conclusions about firm productivity growth after the liberalization. Results from a growth decomposition exercise and from a quantile regression model reinforce the dissimilarity of results across methods. (*JEL* F13, 14, D24, C14)

I. INTRODUCTION

Contemporary theories of international trade tend to imply that import competition improves the productivity of domestic firms. From this perspective, a trade liberalization either encourages innovation among local producers who do not wish to see their market share erode, or forces the less productive firms to shut down freeing up inputs for more productive firms to operate. As a result, the home country's industrial landscape becomes more productive. Relatedly, recent decades have seen the emergence of multiple methods to estimate production functions. In this paper, we employ three of these methods and find a surprising result: they all imply different relationships between firm productivity and trade liberalization.

The methods we consider are Levinsohn and Petrin (2003, henceforth LP), Ackerberg, Caves, and Frazer (2015, henceforth ACF), and Gandhi, Navarro, and Rivers (2017, henceforth GNR), and we examine whether they yield consistent conclusions vis-à-vis firm-level productivity growth. Using (the well-known) data from the Colombian manufacturing sector (see Roberts and Tybout 1997), we find that different methods substantially change the estimated empirical relationship between trade liberalizations and firm productivity. We arrive to this result from two separate exercises: a productivity-growth decomposition and a quantile-regression analysis.

Melitz and Polanec (2015) propose a decomposition procedure that allows for the empirical isolation of four contributing factors to aggregate productivity changes: reallocation from less productive to more productive firms, growth in productivity of incumbent firms, exit of unproductive firms, and entry of more productive firms. We apply the decomposition to the productivity estimates of our three methods and investigate whether different estimation techniques assign different importance to these three channels. We find that any judgment about the relative

*We thank the editor (Maggie Chen) and three anonymous referees for valuable comments and insights. We are especially grateful to Jeff Racine for many useful discussions in the beginnings of this project, and to Marinho Bertanha and Jeff Thurk for comments and suggestions. We also thank Ana Fernandes, Mark Roberts, Kim Ruhl, and James Tybout for making data available to us, and to David Rivers for sharing programs with us. For comments and suggestions, we thank participants at the 2015 Canadian Economic Association Annual Conference and the 5th Annual Meeting of the Argentine Regional Science Association. John Kealey gratefully acknowledges support from the Social Sciences and Humanities Council of Canada. The usual disclaimer applies. E-mails: john.kealey.econ@gmail.com, pujolasp@mcmaster.ca, and csosapad@nd.edu.

Kealey: Researcher, Department of Economics, McMaster University, Hamilton, ON L8S4M4, Canada. Phone 905 525-9140, E-mail john.kealey.econ@gmail.com

Pujolas: Professor, Department of Economics, McMaster University, Hamilton, ON L8S4M4, Canada. Phone 905 525-9140 ext. 23819, E-mail pujolasp@mcmaster.ca

Sosa-Padilla: Professor, Department of Economics, University of Notre Dame, Notre Dame, IN 46556. Phone 574 401-3998, Fax 574 631-4783, E-mail csosapad@nd.edu

ABBREVIATIONS

ACF: Ackerberg, Caves, and Frazer (2015)
ERP: Effective Rate of Protection
GNR: Gandhi, Navarro, and Rivers (2017)
LP: Levinsohn and Petrin (2003)
OLS: Ordinary Least Squares

Economic Inquiry
(ISSN 0095-2583)
doi:10.1111/ecin.12767

contributions of incumbent firms, exiting firms, and new entrants to industry-level productivity growth ultimately depends on the underlying specification of the production function. The contributions of each channel of productivity growth vary notably across estimation procedures. Moreover, the rank correlations for each source/channel of productivity growth across estimation methods are modest and sometimes even negative.

Most empirical literature finds a positive association between trade liberalizations and firm productivity. For instance, Tybout and Westbrook (1995), Pavcnik (2002), Schor (2004), Fernandes (2007), Topalova and Khandelwal (2011), and Hu and Liu (2014) demonstrate that the liberalization of trade policy has generally coincided with productivity growth at the firm level in Mexico, Chile, Brazil, Colombia, India, and China, respectively. The empirical focus of these studies tends to be the conditional mean of firm productivity, given different levels of trade protection; that is, most authors employ linear regression methods to evaluate whether there exists a rather general relationship between trade policy on one hand, and the conditional expectation of firm productivity on the other. However, these results fail to shed light on whether different types of firms, ranging from the least to the most efficient producers of a particular good, exhibit similar responses to changes in the policy environment. Thus, in this paper, we opt for a quantile regression approach that is better able to reflect trends in the *distribution* of firm productivity, as opposed to just focusing on its conditional mean. We find that the trade liberalization-firm productivity link is not robust. The sign of this link depends on the specific estimation method, and the strength of the link increases as we move to the right tail of the productivity distribution: more productive firms gain more (or lose less) from trade liberalizations than less productive firms. Overall, we find that the linear regression estimates can be misleading when trying to empirically assess this link.

Our paper is related to the work of Van Biesebroeck (2008). He uses five different methods to estimate firm productivity using data from Colombia and Zimbabwe. He finds that the different methods produce surprisingly similar productivity estimates. This may seem, at first, contradictory to our findings. We argue in Section IV.D that this is in fact not the case, and that we view our work as complementary to Van Biesebroeck's (2008) contribution.

A. *Layout of the Paper*

Section II provides a thorough summary of the three different methods that we employ to estimate firm-level productivity in the Colombian manufacturing sector. Section III describes the input, output, and trade policy data that are used in the analysis in Section IV, where we discuss the coefficient estimates that we obtain under several specifications of our quantile regression model, and the results of the Melitz-Polanec decomposition exercise that we perform for a long list of manufacturing industries. Section V concludes.

II. REVIEW OF METHODS

In this section, we provide a thorough overview of three different strategies for the identification and estimation of firm-level productivity. These approaches, which are presented in the chronological order of their appearance in the productivity literature, were originally proposed by LP, ACF, and GNR, and are now in widespread use in a number of different subfields of empirical economics. In what follows, we adopt the convention whereby lower-case (upper-case) letters are used to denote the log (level) values of the variables in the production model.

A. *Levinsohn and Petrin's Control Function Method*

Consider a logarithmically transformed Cobb–Douglas production function:

$$(1) \qquad y_{it} = \alpha_k k_{it} + \alpha_l l_{it} + \alpha_m m_{it} + \omega_{it} + \epsilon_{it},$$

where y_{it} is the log of firm i's gross output in period t, k_{it} is the capital stock, l_{it} is the quantity of labor employed by the firm, and m_{it} is an intermediate input variable comprising raw materials and energy consumption. Firm-level productivity is denoted by ω_{it} and ϵ_{it} is a random error term. LP propose a "control function" approach whereby the firm's intermediate input demand is a function of its capital stock and its level of productivity:

$$(2) \qquad m_{it} = m\left(k_{it}, \omega_{it}\right).$$

Assuming that the function $m(\bullet)$ is strictly increasing in ω_{it} holding k_{it} fixed, one can invert (2) to obtain an expression for firm-level productivity:

$$(3) \qquad \omega_{it} = m^{-1}\left(k_{it}, m_{it}\right).$$

Inserting (3) into (1) yields:

$$y_{it} = \alpha_k k_{it} + \alpha_l l_{it} + \alpha_m m_{it} + m^{-1}\left(k_{it}, m_{it}\right) + \epsilon_{it}$$

$$(4)\quad = \alpha_l l_{it} + \theta\left(k_{it}, m_{it}\right) + \epsilon_{it},$$

where $\theta(k_{it}, m_{it}) = \alpha_k k_{it} + \alpha_m m_{it} + m^{-1}(k_{it}, m_{it})$. One can specify $\theta(k_{it}, m_{it})$ as a third-order polynomial in k_{it} and m_{it} and estimate (4) by means of an ordinary least squares (OLS) regression. This yields an estimate of the elasticity of output with respect to labor, $\widehat{\alpha}_l$.

Next, LP's framework assumes that firm-level productivity evolves according to a first-order Markov process:

$$(5)\qquad \omega_{it} = g\left(\omega_{it-1}\right) + \eta_{it},$$

where η_{it} can be interpreted as an unanticipated productivity shock. Using the fitted values $\widehat{\theta}\left(k_{it}, m_{it}\right)$ from the regression in (4), one can obtain the following expression for ω_{it}:

$$(6)\quad \omega_{it}\left(\alpha_k, \alpha_m\right) = \widehat{\theta}\left(k_{it}, m_{it}\right) - \alpha_k k_{it} - \alpha_m m_{it}.$$

Lagged productivity, $\omega_{it-1}(\alpha_k, \alpha_m)$, is analogously defined. We specify (5) as a third-order polynomial $\omega_{it} = \rho_0 + \rho_1\omega_{it-1} + \rho_2\omega_{it-1}^2 + \rho_3\omega_{it-1}^3 + \eta_{it}$ and estimate ρ_0, ρ_1, ρ_2, and ρ_3 for given values of α_k and α_m, which allows us to write the unanticipated productivity shock as a function of the unknown elasticity parameters $\eta_{it}(\alpha_k, \alpha_m)$. LP use the following moment condition to identify the elasticity of output with respect to capital and intermediate inputs:

$$(7)\qquad \mathbb{E}\left[\eta_{it}\left(\alpha_k, \alpha_m\right) | k_{it}, m_{it-1}\right] = 0.$$

Finally, $\widehat{\alpha}_k$ and $\widehat{\alpha}_m$ can be plugged into (6) to obtain firm i's period-t productivity, $\widehat{\omega}_{it}$.

B. Ackerberg, Caves, and Frazer's Value-Added Model

ACF point out that LP's approach suffers from a multicollinearity issue stemming from the likelihood that a firm's labor and intermediate input decisions are both influenced by its level of productivity. They show how this can complicate estimation of α_l in the partially linear model that is depicted in (4), and as an alternative, they propose the following value-added Cobb–Douglas production model:

$$(8)\qquad va_{it} = \alpha_k k_{it} + \alpha_l l_{it} + \omega_{it} + \epsilon_{it},$$

where now, va_{it} denotes firm i's value-added output in period-t. The right-hand side of (8) is the same as in (1), with the exception that the intermediate input variable m_{it} has been omitted. ACF use the same control function as LP that appears in (3), and rewrite (8) as:

$$va_{it} = \alpha_k k_{it} + \alpha_l l_{it} + m^{-1}\left(k_{it}, m_{it}\right) + \epsilon_{it}$$

$$(9)\qquad = \phi\left(k_{it}, l_{it}, m_{it}\right) + \epsilon_{it}.$$

Note that the central difference between the current approach and the one described in section 2.1 lies in the specification of $\phi(k_{it}, l_{it}, m_{it})$ in (9) as opposed to that of $\theta(k_{it}, m_{it})$ in (4). Once again, $\phi(k_{it}, l_{it}, m_{it})$ can be specified as a third-order polynomial in k_{it}, l_{it}, and m_{it} and estimated via OLS. Productivity can then be written as $\omega_{it}\left(\alpha_k, \alpha_l\right) = \widehat{\phi}\left(k_{it}, l_{it}, m_{it}\right) - \alpha_k k_{it} - \alpha_l l_{it}$ and the productivity shock η_{it} in (5) can be expressed in terms of the unknown elasticity parameters $\eta_{it}(\alpha_k, \alpha_l)$ by following the same procedure that was described in the previous subsection. Finally, ACF use the following moment condition to identify α_k and α_l:

$$(10)\qquad \mathbb{E}\left[\eta_{it}\left(\alpha_k, \alpha_l\right) | k_{it}, l_{it-1}\right] = 0.$$

Firm-level productivity is then given by $\widehat{\omega}_{it} = \widehat{\phi}\left(k_{it}, l_{it}, m_{it}\right) - \widehat{\alpha}_k k_{it} - \widehat{\alpha}_l l_{it}$.

C. Gandhi, Navarro, and Rivers' Nonparametric Identification Strategy

GNR show how one can estimate a production function whose underlying functional form is unknown:

$$(11)\qquad Y_{it} = F\left(K_{it}, L_{it}, M_{it}\right) e^{\omega_{it} + \epsilon_{it}},$$

where the upper-case Y_{it}, K_{it}, L_{it}, and M_{it} denote the output, capital stock, labor, and intermediate input variables in level form. Meanwhile, the productivity and error terms are once again denoted by ω_{it} and ϵ_{it}, respectively. This approach makes use of the firm's first-order condition for its choice of intermediate inputs:

$$(12)\qquad p_M = p_Y F_M\left(K_{it}, L_{it}, M_{it}\right) e^{\omega_{it}} E\left[e^{\epsilon_{it}}\right],$$

where p_M and p_Y are respectively the intermediate input and final output prices and $F_M(K_{it}, L_{it}, M_{it})$ is the partial derivative of the production function with respect to the intermediate input variable. Next, it can be shown that if one subtracts the log of (11) from the log of (12) and subsequently adds the log of M_{it} to both sides of the resulting expression, one obtains:

$$(13)\quad \ln\left(\frac{p_M M_{it}}{p_Y Y_{it}}\right) = \ln\left(\frac{F_M\left(K_{it}, L_{it}, M_{it}\right) M_{it}}{F\left(K_{it}, L_{it}, M_{it}\right)} E\left(e^{\epsilon_{it}}\right)\right) - \epsilon_{it}.$$

The left-hand side of (13) can be computed using firm-level input expenditure and revenue data, while the expression in parentheses on the right-hand side can be approximated by a second-order polynomial in k_{it}, l_{it}, and m_{it} (lower case letters denote the logs of the input variables). The equation can then be estimated by means of a nonlinear least squares regression, and this yields estimates of ϵ_{it}, $E(e^{\epsilon_{it}})$, and $F_M(K_{it}, L_{it}, M_{it})M_{it}/F(K_{it}, L_{ji}, M_{it})$.

As a next step in the process of identifying a firm's production function, GNR make use of the equality

$$\frac{F_M\left(K_{it}, L_{it}, M_{it}\right)}{F\left(K_{it}, L_{it}, M_{it}\right)} = \frac{\partial \ln F\left(K_{it}, L_{it}, M_{it}\right)}{\partial M_{it}}.$$

Integrating both sides of this expression gives us

$$\int \frac{F_M\left(K_{it}, L_{it}, M_{it}\right) M_{it}}{F\left(K_{it}, L_{it}, M_{it}\right)} \frac{\mathrm{d}M_{it}}{M_{it}}$$
$$(14)\qquad = \ln F\left(K_{it}, L_{it}, M_{it}\right) + \mathscr{C}\left(K_{it}, L_{it}\right).$$

Given Equation (13), the expression above makes it possible to identify $\ln F(K_{it}, L_{it}, M_{it})$ up to a constant of integration, which GNR denote by $\mathscr{C}\left(K_{it}, L_{it}\right)$. Combining (14) and the log of (11), the firm-level productivity term ω_{it} satisfies the following equality:

$$\omega_{it} = \ln Y_{it} - \int \frac{F_M\left(K_{it}, L_{it}, M_{it}\right) M_{it}}{F\left(K_{it}, L_{it}, M_{it}\right)} \frac{\mathrm{d}M_{it}}{M_{it}}$$
$$(15)\qquad - \epsilon_{it} + \mathscr{C}\left(K_{it}, L_{it}\right).$$

Lagged productivity, ω_{it-1}, is analogously defined. The constant of integration is modeled as a second-order polynomial in k_{it} and l_{it}. Once again, we can follow the same procedure that was described in Sections II.A and II.B and model the evolution of ω_{it} as a first-order Markov process $\omega_{it} = \rho_1\omega_{it-1} + \rho_2\omega_{it-1}^2 + \rho_3\omega_{it-1}^3 + \eta_{it}$. The moment condition $E(\eta_{it}|K_{it}, L_{it}, Y_{it-1}, K_{it-1}, L_{it-1}) = 0$ identifies the parameters in $\mathscr{C}\left(K_{it}, L_{it}\right)$, yielding an estimate of firm-level productivity ω_{it}.

III. DATA

Our dataset comes from the annual census of Colombian manufacturing plants between the years 1981 and 1991. The data are an unbalanced panel and plants with less than 10 employees are excluded. The census collects data on firm characteristics (like ISIC code, location, and year of start up), employment and labor costs, value of capital and investment, inventories, taxes and subsidies, energy inputs, and other general expenditures.[1] Our final sample contains approximately 61,000 observations from over 11,000 plants in 22 different industries.

A. Production Data

The gross output, value-added, capital stock, and intermediate input variables are all expressed in thousands of Colombian pesos, and are deflated using an industry-by-year price index.[2] Intermediate inputs, which are included in the production functions of LP and GNR but absent from that of ACF, are defined as the total amount of energy and raw materials consumed by a plant in a given year. A plant's value-added production is therefore obtained by subtracting its intermediate input consumption from its gross output. Meanwhile, the labor variable is expressed as the total number of workers that are on a plant's payroll, but with the slight modification that unskilled and skilled laborers are weighted by the ratio of their respective median salaries.[3]

Panel A of Table 1 contains summary statistics of the production data. A considerable amount

1. Most empirical studies in the productivity literature (this paper included) use revenue data to estimate productivity. It is well known that revenue-based measures of productivity may be subject to biases (estimates of the true underlying productivity are confounded by demand shocks and markups). The literature is moving forward by developing methodologies to produce estimates of quantity-based productivity. For example, De Loecker et al. (2016) study the effects of trade liberalization on prices and markups using firm-level data from India. They exploit a rich dataset that allows them to effectively estimate output productivity (i.e., they have data on prices and quantities of firms' products over time) and do not need to rely on revenue measures. Unfortunately, our Colombian manufacturing dataset is not rich enough to allow us to obtain quantity-based productivity estimates.

2. This price index can be directly computed from the information available in the manufacturing census. In particular, both the nominal and the real values of production are recorded for each observation in the panel of manufacturing plants and hence, the ratio of these two variables serves as an industry-level price index.

3. These variables' construction conventions and definitions are commonplace in the literature dealing with plant-level productivity estimation.

TABLE 1
Summary Statistics

	Full Sample	Five Main Industries				
		Food Proc.	Apparel	Metal Fabric.	Textiles	Wood Proc.
Panel A						
Number of firms	11,311	2,244	1,885	994	724	378
Employment size						
Mean	64.8	63.6	51.9	52.5	81.1	34.4
Std. dev.	146.8	140.8	120.5	83.8	139.7	74.2
Firm age (years)						
Mean	15.1	17.5	9.9	15.07	13.49	15.26
Std. dev.	12.1	12.7	8.5	10.9	9.8	10.3
Frac. of exporters	0.23	0.22	0.25	0.25	0.20	0.13
Panel B						
Tariffs						
Min	0.0	0.0	0.657	0.252	0.403	0.358
Max	1.331	0.653	1.217	0.797	1.331	0.735
ERP						
Min	0.163	0.791	0.734	0.585	0.826	0.649
Max	2.033	1.470	1.900	0.988	2.033	1.182

Notes: Panel A reports the main firm characteristics for the aggregate and the five main industries. Panel B reports moments of our trade policy variables.

of heterogeneity is found in our final sample regarding age and employment size. On average, the plants in our data are 15.14 years old, and have 64.77 employees. Roughly 23% of the plants are exporters (mostly in the apparel and metal fabrication industries).

B. Trade Policy Data

We measure trade policy (or "trade barriers") in two different ways. First, we use the Colombian government's import tariff schedule that is available for each of the 71 unique 4-digit ISIC codes that are represented in the census. For the 11-year period that runs from 1981 to 1981, tariff data are missing for 1982 and 1989–1991, and so the first specification of the regression model is estimated using a 7-year subsample of the original dataset. Second, in addition to the tariff data, we also use the effective rate of protection (ERP) as a trade policy indicator. This is intended to reflect the dual impact of protectionism, that is, reduced competition from abroad on one hand and increased imported input costs on the other. The ERP is computed as $(va_d - va_w)/va_w$, where va_d and va_w, respectively, denote manufacturers' value-added under distorted domestic (d) and undistorted world (w) prices. The ERP data are available for 22 unique 3-digit ISIC codes for the years 1981, 1984, 1985, 1990, and 1991, and so once again, the regressions that include the ERP as a predictor are only carried out on a 5-year subsample of the data.

Panel B of Table 1 sheds some light on the extent to which Colombia's trade policy regime underwent reform during the period we study. Minimum and maximum tariffs and ERP values are reported for the overall sample and for the main five 3-digit industries that are covered in the sample (as defined by total number of employees). In many instances, there is substantial liberalization, with some industries experiencing a 50 to 60 percentage point decrease in import tariffs between the mid-1980s and the early 1990s. In fact, in the textile industry, the difference between the minimum and maximum ERP is about 120 percentage points, which constitutes quite an aggressive policy reform over a relatively short period of time.

There are two clear advantages of using our dataset. The first one is that this dataset has been used in several previous studies (e.g., Eslava et al. 2004; Eslava, Haltiwanger, and Kugler 2013; Fernandes 2007; Rivers 2013; Roberts and Tybout 1997) and GNR, and it is uniformly recognized to provide a rich description of the Colombian manufacturing sector.

The second advantage is that the time coverage of our dataset includes years of substantial changes in the Colombian trade policy. Following Fernandes (2007), we identify two main phases in Colombia's trade policy during the 1980s: (1) a *protectionist* phase (1981–1984) and (2) a *liberalization* phase (1985–1991). During the *protectionist* phase a number of tariff cuts that had been implemented in the previous decade were reversed. For instance, in between 1981 to 1984, the mean tariff on imported manufactured goods went from 33% to roughly 53%, although some tariffs, such as the ones that applied to textile and apparel imports, climbed as high as 105%. In the *liberalization* phase, the government brought down the average tariff on manufactured goods to approximately 38%, with no tariff category exceeding 70%. By the late 1980s, most tariffs had fallen back to their 1981 levels or even lower. These changes in trade policy are crucial to identify the reaction of plant productivity to liberalization policies.

IV. RESULTS

We build toward our main result in three steps. The first step is to report on the properties of the productivity estimates ($\omega_{i,t}$) under the three different methodologies outlined in Section II: one can interpret this step as a *static* comparative exercise in which we show that in fact the estimates are significantly different. The second step is to take a more *dynamic* view: we decompose productivity growth and show that the sources of this growth are also different across estimation methods. The third and final step is to directly test for the trade liberalization-productivity link: here we find that this link (so frequently studied and scrutinized in the literature) is not robust to the method used to estimate productivity. We close the presentation of our results discussing how they relate to the findings in Van Biesebroeck (2008).

A. *Different Methodologies, Different Estimates*

Our estimates of plant-level productivity are obtained separately for each of the 22 industries in the dataset. The bulk of the analysis will focus on the main five industries (as described in Section III), but results for all the industries are presented in Appendix S1, Supporting information. Whenever possible, we present results for the full sample.[4]

Table 2 presents measures of dispersion for each of the three sets of estimates and measures of association between them. The first thing to notice is that we observe a considerable degree of dispersion in most of the main five industries. Additionally, regarding the ranking of dispersions, we see that in almost all the cases the estimate with the least dispersion is GNR. Related, we observe that consistently across industries the coefficient of variation of ACF exceeds that of GNR.

TABLE 2
Descriptive Statistics for Our Three Productivity Measures

	Coeff. of Variation	Spearman Rank Correlations		
		LP	ACF	GNR
Food proc.				
LP	0.65	1	0.36	−0.40
ACF	1.28		1	−0.79
GNR	0.26			1
Apparel				
LP	0.35	1	0.31	0.25
ACF	0.59		1	0.89
GNR	0.15			1
Metal fabric.				
LP	0.10	1	0.20	−0.01
ACF	0.53		1	0.89
GNR	0.16			1
Textiles				
LP	1.74	1	0.64	0.38
ACF	0.49		1	0.80
GNR	0.20			1
Wood proc.				
LP	0.14	1	−0.42	−0.71
ACF	0.38		1	0.82
GNR	0.15			1
Average				
LP	0.37	1	0.21	−0.03
ACF	0.63		1	0.75
GNR	0.20			1

Notes: All coefficients of variation and Spearman rank correlations are computed within 3-digit industries. The “Average” numbers correspond to simple means across the 22 3-digit industries.

We also present Spearman rank correlations in Table 2. Some patterns emerge. First, the pair of “similar” estimation methods (i.e., LP and ACF, both relying on parametric assumptions about the production function) produce productivity rankings that exhibit in general low to mild positive correlation: the rank correlation between these estimates is positive in four out of five of the main industries and it is in general low. In fact, on average, the rank correlation between these two estimation methods is below 0.25.

Second, the rank correlations between ACF (the *frontier* method using functional form assumptions) and GNR estimates (which rely on firms’ first-order conditions) is on average positive and high. However in some industries (e.g., Food processing, the largest industry by number of employees) there is a negative rank correlation. This suggests a far from perfect association.

4. In these instances, the numbers pertaining to the “Full Sample” are simple means of the 3-digit industries.

Third, overall it seems that the LP estimates are poorly associated with the ones obtained from first-order condition methods (GNR). The rank correlations between LP and GNR are negative in three of the main industries, positive in the other two, and virtually zero on average. We believe this is an important finding, as most of the literature studying the trade liberalization-productivity link uses either LP or variants of it. We believe this particular finding calls for a reassessment of the strength of this link using different estimation approaches.

B. Decomposition of Aggregate Productivity Changes

Analyzing the effect of trade policy on firm productivity necessarily involves dynamics. When a policy change is implemented, economic actors take time to adjust to the new environment and hence if the policy is to have an effect on productivity (or any other endogenous outcome for that matter) then time necessarily needs to go by. Building from this basic insight, it is then due diligence to investigate whether different estimation methods imply different patterns of productivity growth.

Melitz and Polanec (2015) propose a decomposition of industry-level productivity changes into three categories: surviving firms, new entrants, and exiting firms. In the present context, let $t \in \{H, L\}$ denote a time period that is characterized by either a high (H) or a low (L) tariff regime, and let $j \in \{S, X, E\}$ denote the group to which firm i belongs, namely either survivors (S), exiters (X), or entrants (E). Note that in the Colombian manufacturing data, the high-tariff period generally precedes the low-tariff period, and hence the exiting firms and new entrants only appear in the sample in periods H and L, respectively. Let ω_{ijt} denote firm i's productivity and let s_{ijt} represent its share of industry-level output under tariff regime t, where the subscript j serves to indicate that firm i belongs to group j. Thus, group j's share of aggregate output in period t is given by $s_{jt} = \sum_i s_{ijt}$ and its aggregate productivity is computed as $\Phi_{jt} = \sum_i \frac{s_{ijt}}{s_{jt}} \omega_{ijt}$. Then, aggregate industry-level productivity under the tariff regimes H and L can be written as:

$$\Phi_H = s_{SH}\Phi_{SH} + s_{XH}\Phi_{XH},$$
$$\Phi_L = s_{SL}\Phi_{SL} + s_{EL}\Phi_{EL}.$$

This gives rise to the following decomposition of the change in aggregate productivity $\Delta\Phi$ when an industry's trade policy regime switches from H to L:

$$\begin{aligned}\Delta\Phi &= \left(\Phi_{SL} - \Phi_{SH}\right) + s_{EL}\left(\Phi_{EL} - \Phi_{SL}\right) \\ &\quad + s_{XH}\left(\Phi_{SH} - \Phi_{XH}\right) \\ &= \Delta\overline{\omega}_S + \Delta cov_S + s_{EL}\left(\Phi_{EL} - \Phi_{SL}\right) \\ &\quad + s_{XH}\left(\Phi_{SH} - \Phi_{XH}\right),\end{aligned} \tag{16}$$

where $\Delta\overline{\omega}_S$ denotes the change in the mean productivity of surviving firms, Δcov_S denotes the change in the covariance of surviving firms' productivity and their share of total output, and $s_{EL}(\Phi_{EL} - \Phi_{SL})$ *and* $s_{XH}(\Phi_{SH} - \Phi_{XH})$ respectively capture the effects of entry of more productive firms and exit of less productive firms in the intervening period between the high tariff and low tariff regimes.

Table 3 contains the results from the decomposition exercise for the LP, ACF, and GNR measures of productivity, comparing years 1981 and 1991.[5] All of the reported values have been normalized by setting $\Phi_H = 1$ for each industry.[6]

Panel A of Table 3 shows the four components of the Melitz-Polanec growth decomposition (as well as the total productivity growth) for the main five industries in our dataset. We can see that there is a fair amount of heterogeneity when it comes to finding the main source of productivity growth. Under the LP estimation, we find negative growth in two out of five of the main industries, and in one out of five for the ACF estimation. In regard to the decomposed growth estimates, we find that efficiency gains among surviving firms ($\Delta\overline{\omega}_S$), efficient reallocation of market share among incumbents (Δcov_S), and the exit of inefficient firms ($s_{XH}(\Phi_{SH} - \Phi_{XH})$) tend to play a more important role than the entry of productive firms into the market ($s_{EL}(\Phi_{EL} - \Phi_{SL})$). While the latter is characterized by a positive sign in less than half of the industries in our sample, its magnitude is generally very small, and hence we conclude that it rarely makes any noteworthy contribution to industry-level productivity growth.[7]

5. Two different samples—respectively comprising the years in which import tariffs and the effective rate of protection attain their max and min values—can be used for the analysis. We present results for ERP in the main text and leave results for tariffs for Appendix S1. Using the tariff subsample renders very similar results.

6. The max of both tariffs and the ERP tends to be observed in the mid-1980s, while the min tends to be observed in either the late 1980s (tariffs) or the early 1990s (ERP), due to differences in data availability.

7. Table 3 is showing only the four components of productivity growth for the main five industries, but Appendix S1 has this decomposition for each of our 22 3-digit industries.

TABLE 3
Productivity Growth Decomposition (1981–1991)

	$\Delta\Phi$	$\Delta\overline{\omega}_S$	$\Delta \mathrm{cov}_S$	$s_{XH}(\Phi_{SH} - \Phi_{XH})$	$s_{EL}(\Phi_{EL} - \Phi_{SL})$
Panel A					
Food proc.					
LP	0.078	0.024	0.059	−0.001	−0.003
ACF	−0.358	−0.128	−0.246	−0.019	0.036
GNR	0.018	−0.001	0.000	0.024	−0.005
Apparel					
LP	0.108	0.021	0.068	0.041	−0.022
ACF	0.350	0.067	0.182	0.142	−0.041
GNR	0.019	−0.001	0.013	0.015	−0.008
Metal fabric.					
LP	−0.017	0.005	−0.02	−0.001	0.000
ACF	0.289	−0.028	0.332	0.003	−0.019
GNR	0.013	−0.025	0.029	−0.001	0.01
Textiles					
LP	−0.013	0.018	−0.008	−0.014	−0.008
ACF	0.117	0.071	0.095	−0.020	−0.029
GNR	0.017	0.003	0.020	0.006	−0.012
Wood proc.					
LP	0.077	0.001	0.082	0.023	−0.029
ACF	0.674	0.085	0.640	0.068	−0.118
GNR	0.092	0.023	0.063	0.017	−0.011
Panel B					
$\rho(LP, ACF)$	0.260	−0.112	0.390	0.571	0.819
$\rho(LP, GNR)$	0.133	−0.032	0.065	0.294	0.444
$\rho(ACF, GNR)$	0.221	0.596	0.328	0.609	0.320
Panel C					
LP	0.545	0.727	0.545	0.500	0.455
ACF	0.727	0.455	0.682	0.545	0.409
GNR	0.682	0.500	0.682	0.682	0.455

Notes: Panel A has the Melitz–Polanec decomposition for the five main industries (where the values in columns 2 to 4 add up to the value in column 1). Panel B shows the Spearman rank correlation across the three estimation methods, for the total effect and each of its four components. Panel C shows the proportion of productivity changes (for the total effect and each of its four components) which have the expected positive sign predicted by the theory.

Panels B and C of Table 3 shed light on the consistency of the results of the decomposition exercise across the LP, ACF, and GNR productivity measures. Panel B has Spearman rank correlations of the three estimates of each of the growth components. Here, we observe one of this paper's more interesting results, namely that there is far less uniformity than might originally have been anticipated in the dynamics of the LP, ACF, and GNR estimates as Colombia shifted from a protectionist to a more liberalized trade policy regime. The Spearman correlations are quite modest and in some cases, are actually negative. The decomposition procedure shows particularly different outcomes under the LP and GNR approaches. We also find that, while the ACF and GNR estimates correlated very well in the results shown in Table 2, their implied changes in productivity are very different.[8]

8. These dissimilarities are puzzling and we believe more research on these different methods should be done to understand the key differences among them.

Finally, in Panel C, we report the frequency with which the aggregate and decomposed estimates of firm productivity growth exhibit a positive sign, as might be predicted by modern trade theory. The first column shows that aggregate ACF productivity experiences positive change with the greatest frequency; in this instance, the sign of $\Delta\Phi$ is greater than zero in nearly three-quarters of the industries that appear in the sample. On the other hand, aggregate LP productivity growth is positive only half of the time.

Overall, we find that any judgment about the relative contributions of incumbent firms, exiters, and new entrants to industry-level productivity growth ultimately depends on the underlying specification of the production function. If we wish to evaluate the performance of firms and industries subsequent to trade policy reforms, it is therefore imperative that we keep in mind the sensitivity of the Melitz-Polanec framework in (16) to the choice of a Cobb–Douglas functional form versus a more flexible nonparametric alternative.

TABLE 4
Quantile Regressions

	LP		ACF		GNR	
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Panel A: Tariffs						
q_{10}	−0.0067	−0.0014	0.0440	0.1258	0.2352	0.0789
	(0.0016)	(0.0047)	(0.0224)	(0.0637)	(0.0238)	(0.0637)
q_{25}	−0.0150	0.0031	−0.0302	0.0638	0.1934	0.0454
	(0.0031)	(0.0083)	(0.0187)	(0.0522)	(0.0168)	(0.0402)
q_{50}	−0.0477	0.0049	−0.1485	−0.0323	0.1211	−0.0077
	(0.0053)	(0.0135)	(0.0204)	(0.0529)	(0.0108)	(0.0319)
q_{75}	−0.0964	−0.0427	−0.2249	−0.2303	0.0600	−0.0775
	(0.0091)	(0.0210)	(0.0261)	(0.0722)	(0.0084)	(0.0300)
q_{90}	−0.1419	−0.1410	−0.3740	−0.2699	0.0420	−0.0723
	(0.0181)	(0.0414)	(0.0379)	(0.1139)	(0.0093)	(0.0367)
3-digit ISIC effects	Yes	Yes	Yes	Yes	Yes	Yes
Year effects	Yes		Yes		Yes	
Import-competing		Yes		Yes		Yes
Panel B: ERP						
q_{10}	−0.0011	−0.0005	−0.0057	0.0212	0.0399	−0.0610
	(0.0013)	(0.0032)	(0.0152)	(0.0409)	(0.0176)	(0.0376)
q_{25}	−0.0055	0.0085	−0.0122	0.0128	0.0266	0.0017
	(0.0022)	(0.0062)	(0.0143)	(0.0334)	(0.0124)	(0.0309)
q_{50}	−0.0158	0.0020	−0.0506	−0.0210	−0.0008	0.0056
	(0.0038)	(0.0102)	(0.0152)	(0.0358)	(0.0081)	(0.0215)
q_{75}	−0.0335	−0.0138	−0.1110	−0.0973	−0.0006	−0.0375
	(0.0064)	(0.0152)	(0.0193)	(0.0455)	(0.0061)	(0.0197)
q_{90}	−0.0627	0.0104	−0.1317	−0.1612	−0.0075	−0.0174
	(0.0137)	(0.0275)	(0.0315)	(0.0747)	(0.0061)	(0.0225)
3-digit ISIC effects	Yes	Yes	Yes	Yes	Yes	Yes
Year effects	Yes		Yes		Yes	
Import competing		Yes		Yes		Yes

Notes: Dependent variable is log of plant-level productivity in both panels. Panel A has results when the trade policy indicator is the tariff rate, and Panel B when we use ERP. Data available for years 1981, 1983–1988 (Panel A); and 1981, 1984, 1985, 1990, and 1991 (Panel B). Number of Observations: 38,297 (for specifications 1, 3, and 5) and 8,478 (for specifications 2, 4, and 6) for Panel A, and 27,177 (for specifications 1, 3, and 5) and 6,287 (for specifications 2, 4, and 6) for Panel B. Standard errors are in parentheses.

C. Regression Analysis

In Table 4, we report coefficient estimates for a number of different specifications of a quantile regression model in which the dependent variable is the log of firm productivity and the main explanatory variable is a 3-digit industry-level measure of trade policy (either import tariffs, Panel A, or ERP, Panel B). For each of the LP, ACF, and GNR measures of productivity (and each trade policy variable), we estimate two specifications: one that includes both an industry and a time dummy, and another one where consideration is limited to industries that are categorized as “import-competing.” This latter categorization has been applied in previous studies that examine the empirical link between trade policy and productivity, most notably in Pavcnik (2002), who defines an industry as import competing if the ratio of imports to total output exceeds a particular threshold. The author experiments with different cutoff values and finds that her results remain fairly consistent when the ratio lies between 0.10 and 0.25. In her final analysis, she settles on 0.15, which is the value that we use here as well.

We also present the coefficients of the main regression specification graphically, in Figure 1. This will help visualizing the results and also ease the comparison with the OLS results (which due to space constraints we do not include in Table 4). It is convenient to remember here that a *negative* coefficient for our trade policy variables (tariff or ERP) means that trade liberalizations (decreases in either tariffs or ERP) are productivity *enhancing*. Based on our regression results we make the following remarks.

REMARK 1. *The sign of the quantile regression coefficient estimates displays a fair amount of sensitivity to the manner in which the production function has been specified. Columns 1–4*

FIGURE 1
Trade Policy Coefficients

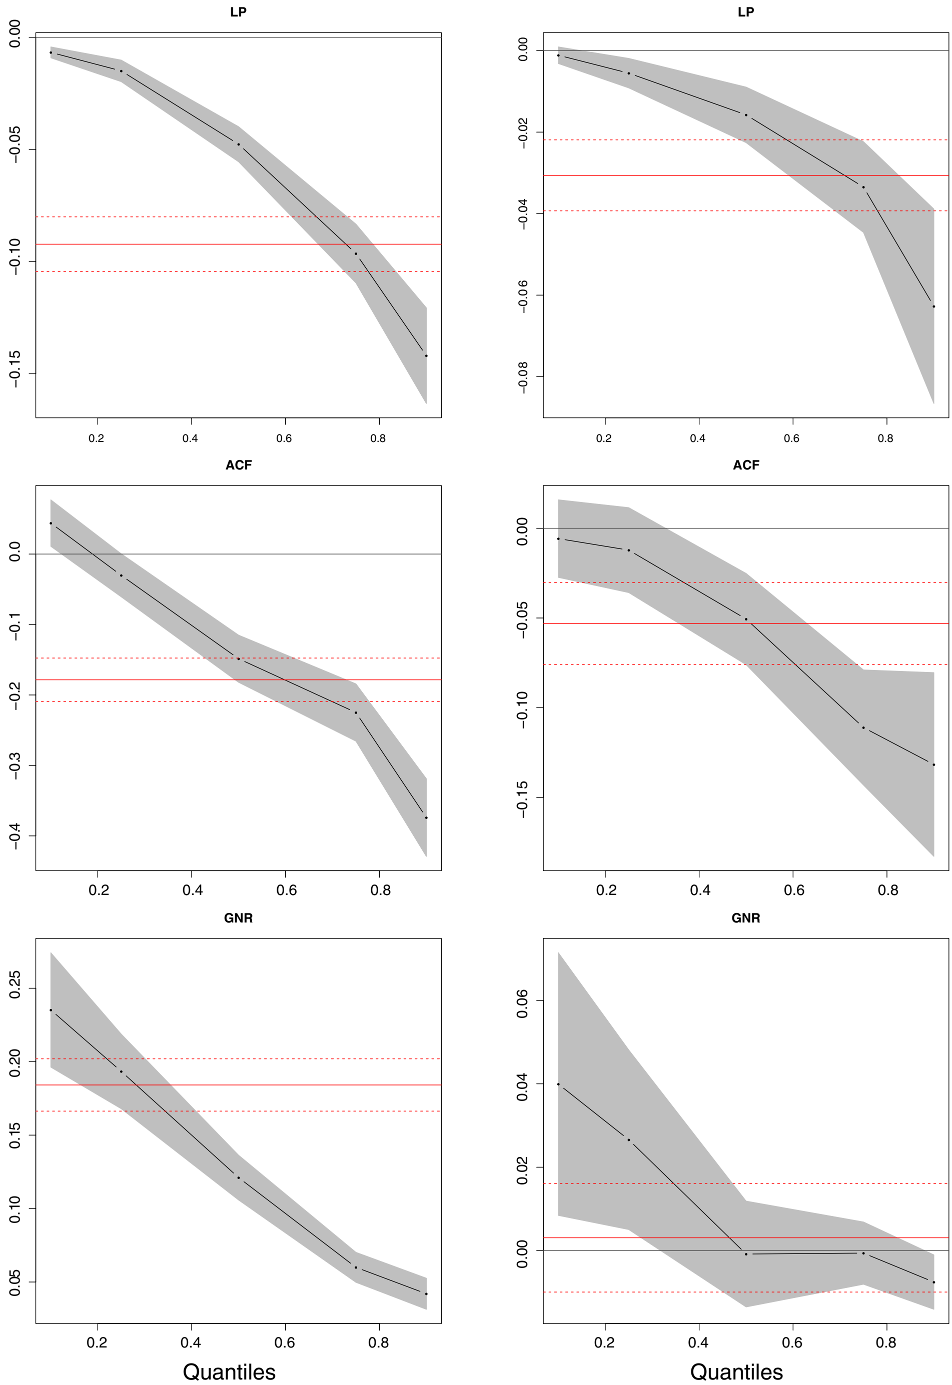

Notes: The solid black lines show the different marginal effects of changes in tariffs (left) and ERP (right) at different productivity quantiles while the gray areas are the confidence bands (i.e., coefficients for the trade policy variable in specifications 1, 3, and 5 in Table 4). The horizontal red lines show the OLS coefficients with their confidence bands.

use LP and ACF productivity as the dependent variables. Most coefficient estimates (but especially the median, upper quartile, and top decile) are negative. This pattern is true both using tariffs or ERP as the trade policy regressor. However, columns 5 and 6, which use GNR, often observe a positive *association between trade barriers and firm productivity. In fact, when these coefficients are negative they are generally insignificant.*

REMARK 2. *Regardless of whether we use LP, ACF, or GNR productivities, regardless of whether we focus on tariffs or ERP as our trade policy variable, there is relatively a more negative association between trade barriers and firm-level productivity in the right tail than in the left tail of the distribution of firms. This phenomenon is evident from the negative slope of the dash-dotted black lines in Figure 1 which show that quantile coefficients are getting more negative (or less positive in the case of GNR) as we move to the right tail of the productivity distribution. Hence, much of the productivity growth takes place among those firms that were already the most productive ones in their respective industries, irrespective of the estimation method.*

REMARK 3. *When we restrict our attention to import-competing industries (specifications 2, 4, and 6), which constitute about one-quarter of the sample, GNR estimates have a negative coefficient at the higher quantiles. We believe this is an important finding: using GNR productivity we obtain that the only group of firms that obtain significant increases in productivity are already highly productive in already import-competing industries.*

REMARK 4. *There is an interesting point of divergence depending on whether we use the tariff rate or the ERP as the explanatory variable. While the former yields positive and statistically significant coefficient estimates for the lower-half of the distribution of GNR productivity, the latter gives rise to coefficient estimates that are either negative or quite small in magnitude relative to their standard errors (or both). This suggests that of the three different measures of productivity that are considered in this paper, the one that relies on the most flexible (i.e., nonparametric) specification of the production function exhibits a more ambiguous statistical relationship with the indicators of trade protection.*

REMARK 5. *While Figure 1 helps to visualize most of the remarks above, it is especially useful to contrast the quantile regression coefficients with their OLS counterparts. We believe it clearly illustrates that relying only on the conditional mean can be highly misleading. There are some cases (like GNR estimates using ERP) where the OLS estimate is insignificant but the quantile regression shows significant (and opposite) effects at the 10th and 90th quantiles. On the other hand, there are cases where the OLS coefficients are negative and significant, but the quantile regressions show either positive or insignificant effects (like in the 10th and 25th quantiles for ACF estimates). In sum, we find it to be crucial to study the liberalization-productivity link looking beyond the conditional mean.*

Overall, our findings indicate that the empirical link between trade liberalization and productivity is: (a) not robust to the estimation method used to identify productivity, and (b) very heterogeneous along the distribution of firm productivity (with more productive firms benefiting the most from liberalizations). Additionally, we consider it to be crucial to study the liberalization-productivity link looking beyond the conditional mean.

D. Relationship with Van Biesebroeck's (2008) Findings

As argued in the introduction, our exercises are reminiscent of the work in Van Biesebroeck (2008).[9] He uses five different methods to estimate firm productivity (i.e., index numbers, data envelopment analysis, instrumental variables estimation, stochastic frontiers, and semiparametric estimation) using data from Colombia and Zimbabwe. He finds that the different methods produce surprisingly similar productivity estimates.

From all these five methods, the only one that is (almost) comparable to our estimates is the one based on semiparametric estimation (following Olley and Pakes 1996). Our estimates from the LP method are the most comparable to those (since LP directly builds on the work of Olley and Pakes 1996 and improves upon it). From the Colombian dataset, Van Biesebroeck (2008) only focuses on the Apparel industry, which is

9. See also Van Biesebroeck (2007) for an analysis comparing different production function estimation techniques using simulated data.

only one of the five main industries we discuss in our paper. When we look at Apparel only, then our results are not that different from those in Van Biesebroeck (2008). In terms of dispersions (which we measure using the coefficient of variation), Apparel is one of the industries where our methods give the closest metrics of dispersion (so, they are not so different among each other in this particular industry). In terms of productivity growth, again Apparel is among the industries with the most similar estimates in the Melitz–Polanec decomposition. Moreover, our estimate for the total productivity growth in Apparels using LP is 11% which compares very well with the result in Van Biesebroeck (2008) (12%).

More generally, we are not contending any of the debates revisited in Van Biesebroeck (2008) (i.e., "does learning-by-exporting increase productivity?," "what brings about technological change?," and "what drives aggregate productivity?"). We are, in a complementary way, studying a different (older) debate of whether trade liberalizations are productivity enhancing.

V. FINAL REMARKS

In this paper, we compare existing methods to estimate production functions to analyze whether the link between trade liberalization and increased firm productivity is robust. We find it is not. While we believe that further research is needed to elucidate the exact link between the two, we think that it is important to stress that our results are based only on a (widely studied) dataset: 1980s Colombian manufacturing plants. Performing similar exercises to other datasets seems a natural avenue for future research.

Having obtained different answers to the same question when using different approaches also suggests that estimation methods should be context-specific. In particular, our findings suggest that researchers should investigate which method applies to their particular study—a given method's assumptions may be more reasonable in a particular framework, while a different method may be more apt in another framework.

REFERENCES

Ackerberg, D. A., K. Caves, and G. Frazer. "Identification Properties of Recent Production Function Estimators." *Econometrica*, 83(6), 2015, 2411–51.

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

Eslava, M., J. Haltiwanger, A. Kugler, and M. Kugler. "The Effects of Structural Reforms on Productivity and Profitability Enhancing Reallocation: Evidence from Colombia." *Journal of Development Economics*, 75(2), 2004, 333–71.

Eslava, M., J. Haltiwanger, and A. Kugler. "Trade and Market Selection: Evidence from Manufacturing Plants in Colombia." *Review of Economic Dynamics*, 16(1), 2013, 135–58.

Fernandes, A. "Trade Policy, Trade Volumes and Plant-Level Productivity in Colombian Manufacturing Industries." *Journal of International Economics*, 71(1), 2007, 52–71.

Gandhi, A., S. Navarro, and D. A. Rivers. "On the Identification of Production Functions: How Heterogeneous Is Productivity?" Working Paper, University of Western Ontario, 2017.

Hu, A. G., and Z. Liu. "Trade Liberalization and Firm Productivity: Evidence from Chinese Manufacturing Industries." *Review of International Economics*, 22(3), 2014, 488–512.

Levinsohn, J., and A. Petrin. "Estimating Production Functions Using Inputs to Control for Unobservables." *Review of Economic Studies*, 70(2), 2003, 317–41.

Melitz, M. J., and S. Polanec. "Dynamic Olley-Pakes Productivity Decomposition with Entry and Exit." *RAND Journal of Economics*, 46(2), 2015, 362–75.

Olley, S. G., and A. Pakes. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica*, 64(6), 1996, 1263–97.

Pavcnik, N. "Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Plants." *Review of Economic Studies*, 69(1), 2002, 245–76.

Rivers, D. A. "Are Exporters More Productive than Non-exporters?" Technical Report, CIBC Working Paper Series, 2013.

Roberts, M. J., and J. R. Tybout. "The Decision to Export in Colombia: An Empirical Model of Entry with Sunk Costs." *American Economic Review*, 87(4), 1997, 545–64.

Schor, A. "Heterogeneous Productivity Response to Tariff Reduction: Evidence from Brazilian Manufacturing Firms." *Journal of Development Economics*, 75(2), 2004, 373–96.

Topalova, P., and A. Khandelwal. "Trade Liberalization and Firm Productivity: The Case of India." *Review of Economics and Statistics*, 93(3), 2011, 995–1009.

Tybout, J. R., and M. D. Westbrook. "Trade Liberalization and the Dimensions of Efficiency Change in Mexican Manufacturing Industries." *Journal of International Economics*, 39(1–2), 1995, 53–78.

Van Biesebroeck, J. "Robustness of Productivity Estimates." *Journal of Industrial Economics*, 55(3), 2007, 529–69.

——— "The Sensitivity of Productivity Estimates: Revisiting. Three Important Debates." *Journal of Business & Economic Statistics*, 26(3), 2008, 311–28.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Appendix S1. Results