Productivity Measures for the Colombian Manufacturing Industry

Por: Camila Casas
Alejandra González

Borradores de ECONOMÍA

Núm. 947
2016

tá - Colombia - Bogotá - Colombia - Bogotá - Colombia - Bogotá - Colombia - Bogotá - Colombia - Bogotá - Colombia - Bogotá - Colo

Productivity Measures for the Colombian Manufacturing Industry*

Camila Casas
mcasaslo@banrep.gov.co

Alejandra González
agonzara@banrep.gov.co

Centro de Estudios sobre Economía Industrial e Internacional,
Banco de la República, Cali

Abstract

In this paper we present estimates for the coefficients of a production function, and the corresponding total factor productivity (TFP) for the Colombian manufacturing industry during 2005–2013. We follow several structural microeconometric techniques to estimate the production function parameters. We compare the estimation results across methodologies, as well as their robustness to changes in our estimation sample, variable definitions, and/or weights used to aggregate the estimated firm-level TFP into an industry-level average TFP. Our results show that, in general, all estimation methodologies result in a similar growth pattern during our sample period. Moreover, the general growth trend is not affected greatly by the additional changes mentioned above. However, the *level* of productivity (and the estimated TFP growth) does depend on how the production function is estimated.

JEL Classification: D24; L25; L60

Keywords: total factor productivity; production function; manufacturing industry

*The views expressed in the paper do not represent those of the Banco de la República or its Board of Directors. We are grateful to Oscar M. Valencia and other participants at the Banco de la República internal research seminar for helpful comments and suggestions. All remaining errors are our own.

1 Introduction

Productivity is a key determinant of the dynamics of the manufacturing sector. It is often linked to the creation, growth, and survival of firms, and to several organizational decisions, such as (but not limited to) entry into foreign markets, investment in R&D, and financial decisions. Hence, in order to design effective economic policies, it is of the utmost importance to have an accurate productivity measure. However, to estimate it properly, one needs to estimate the coefficients of a production function and this is not an easy feat. As it was first pointed out by Marschak and Andrews (1944), if the unobserved productivity shocks are correlated with the firm's input choices then standard econometric techniques will yield biased estimates of the production function coefficients, affecting the resulting productivity estimates as well.

In this paper we present estimates for the parameters of a production function for the Colombian manufacturing industry, obtained following several microeconometric methodologies that correct for the simultaneity bias. We then use our estimated parameters to calculate the corresponding total factor productivity (TFP), and we analyze its evolution during 2005–2013. In addition, we study the robustness of our estimations to changes in the estimation methodology, the estimation sample, the way we measure the variables used for the estimation, the way in which we aggregate individual productivity to get the industry total average per year, and the way we group firms to estimate their production function coefficients. For all our exercises, we use the financial data self-reported by firms to the "Superintendencia de Sociedades".

Our findings can be summarized as follows. First, our production function estimates result in an average TFP index that, after an initial increase between 2005 and 2006, declines every year until it reaches its lowest value during the 2009-10 crisis, and has recovered steadily ever since. We observe a similar pattern regardless of the methodology used to estimate the production function parameters. Second, we find that the general growth trend is not affected by changes to our estimation sample, the way we define the labor variable, or the weights used to aggregate firm-level TFP estimates into a manufacturing average. In particular, we find that the production function coefficients are very similar when we estimate sector-specific production functions and when we pool all the manufacturing firms together and estimate a single production function for the industry as a whole. Third, despite the average TFP growth pattern being similar across our exercises, the *level* of the TFP index in any given year and the estimated cumulative growth during our sample period do depend on how we estimate these coefficients.

Methodologically, this paper is related to the literature on the structural estimation of production functions and unobserved TFP. Recently, there has been a renewed interest in the estimation of production functions, and several techniques have been proposed to deal with the endogeneity bias (see Ackerberg, Benkard, Berry, and Pakes (2007) for a review). In this paper, we focus on two methods. In particular, our strategy for estimating firm-level productivity relies alternatively on the *proxy* methods proposed by Olley and Pakes (1996), and later modified Levinsohn and Petrin (2003) and Ackerberg, Caves, and Frazer (2015), and on the share equation method proposed by Gandhi, Navarro, and Rivers (2016).

Our paper is also related to the series of papers that have focused specifically on the Colombian manufacturing sector. Several methodologies have been applied to study the evolution of productivity for the Colombian manufacturing sector, and the relationship between productivity and other firm characteristics. For instance, Clerides, Lach, and Tybout (1998) studies the causal relationship between a firm's labor productivity and its export intensity. Pombo (1999) calculates non-parametric productivity indices to measure the contribution of different factors to growth in a growth accounting setting. Eslava, Haltiwanger, Kugler, and Kugler (2004) studies the relationship between market allocation of resources, and productivity and profitability. For the productivity estimation, these authors take advantage of the availability of rich plant-level price data and estimate the production function coefficients using instrumental variables. López (2006) studies the relationship between productivity and export status, measuring productivity with a linear approximation of the variable costs. And Meléndez and Seim (2006) and Echavarría, Arbeláez, and Rosales (2006) study the relationship between total factor productivity and the trade liberalization, with TFP estimates obtained following the methodology proposed by Levinsohn and Petrin (2003).

We build on this literature in two ways. First, we compare several TFP estimations obtained using the different methodologies mentioned above. This allows us to extract common patterns across estimations, instead of focusing on a single set of results that can be affected by the assumptions of each individual estimation technique. In addition, when we check the robustness of our results, we can analyze the effects of some of these assumptions on the resulting TFP. Second, we use the latest available data to update the existing productivity estimations. The papers cited above cover periods up to the early 2000s. However, as shown in Carranza and Moreno (2013), the manufacturing industry has evolved greatly in recent years. Therefore, it is important to have updated calculations in order to analyze the behavior of productivity in this new setting.

The rest of the paper is organized as follows. Section 2 describes the different structural estimation techniques we follow in this paper. Section 3 describes how the data are collected and presents the basic features of our estimation sample. In Section 4 we present our estimation results, and in Section 5 we discuss the effects of altering our estimation sample, variable definitions, or aggregation method. Section 6 concludes.

2 The Framework to Estimate Productivity

In order to obtain a reliable firm-level TFP measure, we start by estimating the coefficients of a production function. As mentioned above, in order to get consistent estimates for the input coefficients it is necessary to correct for the potential simultaneity between the unobserved productivity and the demand for inputs.[1] In this section, we describe the different methodologies that we follow in order to account for the potential biases in the production function estimation.[2]

The different estimation algorithms presented in this section can be grouped into two categories, according to the way in which they incorporate productivity into the estimation procedure. In the first approach, an observable variable is used to approximate productivity. This idea was originally presented by Olley and Pakes (1996), henceforth OP , and later extended by Levinsohn and Petrin (2003) and Ackerberg, Caves, and Frazer (2015), henceforth LP and ACF. Following the idea that productivity is positively correlated with the demand for inputs, the estimation technique proposed by OP uses a firm's (observed) input demand as a proxy for (unobserved, to the econometrician) productivity shocks. By inverting the input demand function, it is possible to express productivity as a function of only observable variables. This way, the proxy variable is used to control for the endogeneity in the production function.

The second algorithm, proposed by Gandhi, Navarro, and Rivers (2016), henceforth GNR, uses the information implicit in the firm's optimization problem. By transforming

[1]The production function coefficients may suffer from a selection bias as well, if there is a selection of firms and only the most productive ones remain in the market and are observable. However, according to Griliches and Mairesse (1998) and Ackerberg, Benkard, Berry, and Pakes (2007), once one controls for the simultaneity bias, the selection bias is practically negligible if a full and unbalanced panel is used for the estimation.

[2]Fixed effects and instrumental variables are alternative, more traditional methods to control for a simultaneity bias like the one present in the estimation of production functions. These approaches, however, have not yielded satisfactory results in this particular setting (see Ackerberg, Benkard, Berry, and Pakes (2007) for a review).

the first-order condition to express the intermediate input's revenue share as a function of capital, labor, and intermediate inputs (all observable variables), it is possible to estimate the underlying production function parameters while removing the productivity term from the estimation procedure.

Despite the different ways in which these methods use observable data to account for the effect of productivity while eliminating the unobservable term from the estimation itself, all of them can be applied to estimate the same economic model. We present this model next.[3]

In general, we observe an unbalanced panel of firms $j \in \{1, \ldots, J\}$ over periods $t \in \{1, \ldots, T\}$. The relevant variables for our purpose are the firm's output (Y_{jt}), labor (L_{jt}), capital (K_{jt}), and intermediate inputs (M_{jt}), where y_{jt}, l_{jt}, k_{jt} and m_{jt} denote their respective logarithms.

The first component of the model is the production function, that relates inputs to outputs. In its most general form, this function can be expressed as

$$Y_{jt} = F_t(L_{jt}, K_{jt}, M_{jt}) e^{\nu_{jt}} \ , \tag{1}$$

where ν_{jt} is a Hicks-neutral efficiency term, given by the sum of a productivity shock ω_{jt}, observed by the firm at the beginning of every period, plus an idiosyncratic ex-post shock ϵ_{jt}. Under this definition, it is the productivity term ω_{jt} that causes the simultaneity issue.

The second component of the model is the evolution of the two elements of ν_{jt}. The productivity shock ω_{it} is assumed to be a persistent shock, that evolves over time following

[3] For the presentation of the general model, we follow an earlier version (from 2013) of GNR.

an exogenous Markov process.[4] This implies that it can be expressed as

$$\omega_{jt} = h(\omega_{jt-1}) + \eta_{jt} \ , \tag{2}$$

where η_{jt} can be interpreted as the innovation to the firm's persistent productivity in each period, and is orthogonal to ω_{jt-1}. The ex-post shock is assumed to be i.i.d., and without loss of generality can be normalized to $E[\epsilon_{jt}] = 0$.

The third component of the general model is the timing of input decisions. Capital is determined at, or prior to, $t-1$. Depending on the methodology, labor is considered a flexible input determined at t, or it is chosen at $t-1$ or before, prior to the realization of ω_{jt}. Intermediate inputs are assumed to be fully flexible, and are determined at period t once the firm observes the realization of its productivity shock.

In addition to these three elements, a supplementary restriction generally used by the literature is the assumption that intermediate inputs can be expressed as a function of productivity and other state variables, Δ_{jt}.[5] In particular, the demand for intermediate inputs can be written as $M_{jt} = \mathbb{M}_t(K_{jt}, \Delta_{jt}, \omega_{jt})$, where $\mathbb{M}_t$ is strictly monotone in ω_{jt} for any relevant (K_{jt}, Δ_{jt}). It is this assumption that makes it possible to use the observable intermediate inputs as a proxy variable for unobserved productivity. In the case of GNR, this assumption is not necessary but it is still consistent with the optimizing behavior of the firm under fairly weak conditions on the production function.

With this general framework in mind, we now describe each methodology in more detail.

OP starts from a dynamic model of firm behavior that determines both input demand and

[4]Usually this process is assumed to be first-order. GNR's methodology can be generalized to allow for higher-order Markov processes, but proxy methods can only be applied if productivity is assumed to follow a first-order process.

[5]State variables Δ_{jt} may include the firm's age, or the labor force (when L_{jt} is assumed to be chosen before t).

shutdown decisions. In this model, each period the firm observes its productivity shock ω_{jt}, and it decides whether to exit or continue operating. If the firm decides to exit, it receives a liquidation payoff θ. If it continues operating, it has to choose the level of inputs (e.g., labor, raw materials), and the level of investment, i_{jt}, to adjust its capital stock. Both the exit decision and the investment level depend on the firm's perceptions of the distribution of future market structures given the realization of ω_{jt}. When making these two decisions, the state variables are capital and the firm's age, a_{jt}, while labor is assumed to be a variable, static factor. This implies that capital is predetermined and is affected by the distribution of the productivity conditional on past realizations, while the labor choice is affected only by the current productivity value.

In line with these assumptions, OP develops a multi-stage, semiparametric estimation algorithm to recover the parameters of the production function, assumed to be Cobb-Douglas.[6] For the estimation, OP uses investment as the proxy variable. This means that investment is treated like a function of the firm's state variables, strictly increasing in ω for any relevant (k, a), and that for strictly positive values of investment this function can be inverted to express productivity as a function of investment, capital and age—all observable variables.

In the first stage, the unobservable productivity term is replaced by its proxy function, and the production function is rewritten to express output as a function of labor and a nonparametric function of investment, capital, and the firm's age. The particular model to be estimated in this first stage is

$$y_{jt} = \beta_l l_{jt} + \phi_t(i_{jt}, k_{jt}, a_{jt}) + \epsilon_{jt} \ , \tag{3}$$

a semiparametric regression model that yields an estimate for the labor coefficient β_l and an

[6]In the model presented in the paper, the output is measured with value-added. However, the algorithm can be modified to use gross output instead. In this case, the equations that follow would be modified to explicitly include intermediate inputs as an additional flexible, static production factor. In our estimations we follow the model as presented in the paper and measure output with value-added.

approximation of $\phi_t(\cdot)$.

In the second stage, in order to correct for selection, OP computes a probability of survival P_t that depends on the firm's state variables: observed characteristics like the capital stock and the firm's age, and the unobservable productivity. Again, productivity ω is approximated with a function of investment, capital and age, and each firm's survival probability is estimated nonparametrically as a function of these three observable variables.

For the third and final stage of the estimation, the algorithm uses the estimates of β_l, $\phi_t(\cdot)$ and P_t obtained in the earlier stages to approximate the contribution of capital, age, and productivity to the firm's output *net* of labor:

$$y_{jt} - \widehat{\beta}_l l_{jt} = \beta_k k_{jt} + \beta_a a_{jt} + \omega_t + \epsilon_{jt} \quad \forall t. \tag{4}$$

In addition, for this stage the unobservable productivity term ω is expressed as in equation (2), such that $y_{jt+1} - \widehat{\beta}_l l_{t+1}$ can be written as a function of k_{jt+1}, a_{jt+1}, the productivity of the previous period (which in turn will depend on k_{jt}, a_{jt}, the estimated probability of survival $\widehat{P}_t$, and $\widehat{\phi}_t(\cdot)$), and the productivity innovation term, η_{t+1}. Given that current capital and age are orthogonal to this innovation, they can be interacted with η to construct moment conditions, and β_k and β_a are estimated by minimizing the sum of squared residuals.

One weakness of OP's methodology is its use of investment as the proxy variable. A particular drawback of this variable is that, in practice, investment is rarely a continuous variable. Due to the existence of adjustment costs, firms often either invest large amounts, or do not invest at all. In this sense, investment may not be a good proxy: it is possible to observe zero investment if changes in productivity are small enough. A second related drawback is the loss of efficiency and potential selection bias introduced to the estimation if one drops all the observations with zero investment levels. As mentioned above, OP's proxy function is valid only for positive investment levels, and hence any observation with zero investment cannot be used for the estimation of the production function coefficients.

In order to address these issues, LP modifies OP's algorithm to use different variables as proxies for the unobservable productivity shock. In particular, intermediate inputs (raw materials, electricity, or fuels) are used alternatively as the proxy for ω_{jt}. LP claims that intermediate inputs are better proxies than investment for two reasons. First, since these inputs are easily adjustable, they may respond better to productivity changes. And second, most firms report a positive consumption of inputs so there is no efficiency loss from dropping available information.

Besides this modification, the estimation strategy proposed by LP is very similar to the one proposed by OP, described above. First, the intermediate input chosen as the proxy variable (say raw materials, m_{jt}) is expressed as a monotonic function of productivity and capital (the state variables), and this function is inverted to approximate the unobservable productivity with a function of materials and capital. Then, this proxy function is replaced in the production function, once again assumed to be Cobb-Douglas, and a semiparametric model is estimated in order to recover estimates for the coefficients of the flexible inputs (except the one used as the proxy), and for the nonparametric function $\phi_t(m_{jt}, k_{jt})$.[7] In the second stage, the output net of the contribution of labor, electricity, and fuels is expressed as a function of materials, capital, and productivity, which, again, is expressed as a function of its lagged value plus an innovation η. Hence, this net output can be written as a function of current and future values of m and k, the estimate for $\widehat{\phi}$, and the productivity innovation η. Since this innovation is assumed to be orthogonal to the information set before t, it is interacted with the current capital and lagged consumption of raw materials to form the sample analogues to the moment conditions. Parameters β_m and β_k are then estimated by minimizing the distance between these moments and zero.[8]

[7]In this case, the flexible inputs include raw materials, electricity and fuels in addition to labor, and the latter is split between skilled and unskilled, so a total of four coefficients are estimated in the first stage.

[8]Note that in this algorithm there is no intermediate stage to correct for a potential selection bias towards more productive firms. LP (and ACF and GNR) abstract from the exit decision in OP's dynamic model.

A second shortcoming of OP's methodology stems from the assumptions on the timing of input decisions. Since labor is assumed to be fully flexible and chosen every period once the firm observes its realization of the productivity shock ω_{jt} (simultaneously with m_{jt} and/or i_{jt}), it makes sense to assume that the choice of both labor and the proxy variable depend on the same sate variables. However, if this is the case, then there are no independent sources of variation to identify β_l separately from the nonparametric function ϕ in the first stage. ACF modify OP's and LP's methodologies to address what they call the "collinearity issue".

In terms of the general model's components, ACF modifies the timing of labor selection. In particular, it is assumed that period t's labor is chosen sometime between $t-1$ (when capital is chosen), and t (when materials and/or investment are chosen). Besides breaking the collinearity described above, this timing assumption makes sense if it takes time for the firm to train workers, or there are hiring and firing costs, such that labor is not fully flexible. With this change, it is also necessary to modify the estimation algorithm, since labor becomes a state variable at time t. ACF presents a two-stage procedure in which the first stage's purpose is to separate the productivity term ω from the ex-post shock ϵ, while all the production function coefficients are estimated in the second stage. For the estimation ACF assumes that the production function is a Cobb-Douglas function, where output is measured with value-added. Since intermediate inputs are assumed to be fully flexible and chosen after the firm observes ω_{jt}, these are used as the proxy variable.[9]

In the first stage, m_{jt} is expressed as an increasing function of productivity (given labor and capital), and this function is inverted to express productivity as a function of observable variables (m_{jt}, l_{jt}, k_{jt}). This proxy function is then replaced in the production function, such

However, as mentioned in the beginning of this section, the potential selection bias from this omission is not important since we are correcting for the simultaneity bias, and we use a full, unbalanced panel for our estimations.

[9] ACF state that the estimation algorithm can be modified to use investment as the proxy variable.

that the output can be expressed by:

$$y_{jt} = \phi_t(l_{jt}, k_{jt}, m_{jt}) + \epsilon_{jt} \ . \tag{5}$$

A nonparametric estimation of the function $\phi(\cdot)$ allows one to isolate the ex-post shock $\widehat{\epsilon}$, and to get an estimated value of the expected income $\widehat{\phi}$. In the second stage, ACF follows the previously described algorithms, and rewrites the production function as a function of k, l, and productivity, where the latter is in turn rewritten as a function of its lagged value (calculated as $\widehat{\phi} - \beta_l l - \beta_k k$), and an innovation η. Then, the innovation η is interacted with the current capital and lagged labor to construct sample moments, and the coefficients are estimated with GMM using these moments.

Despite the timing modifications introduced by ACF, GNR argues that the proxy methods described above cannot identify the production function coefficients when it contains flexible inputs, since there are no sources of exogenous variation other than the inputs included in the function. Hence, there are no instruments from outside the production function that can be used to identify the flexible inputs coefficient(s).[10] Moreover, GNR states that estimating value-added production function does not solve the identification problem, and leads to overestimating the degree of productivity heterogeneity across firms.[11]

In order to address both the identification problem and the overstating of productivity heterogeneity that results from estimating value-added production functions, GNR introduces an alternative methodology for the estimation of gross-output production functions. This methodology starts from an optimization model in which the choice of a flexible input is assumed to be endogenous. For a general production function $y_{it} = f_t(L_{jt}, K_{jt}, M_{jt}) + \omega_{jt} + \epsilon_{jt}$,

[10]GNR presents a formal proof for why the production function is not identified that formalizes these ideas.

[11]When estimating the coefficients of a value-added specification, one controls for the variation of some inputs (K and L), but part of the observed output heterogeneity across firms is the mechanical result of including the (heterogeneous) intermediate inputs on the left-hand side of the production function.

where capital and labor are determined prior to period t and the current choice of the intermediate input does not have any dynamic implications, GNR derives the first-order condition, and re-writes it in terms of observable variables, such that the unobserved productivity term is not present in the estimating equation (although it is accounted for through the input choice). Then, the information contained in the first-order condition can be used in a completely nonparametric way.

The algorithm proposed by GNR estimates the production function coefficients in two stages. In the first step, GNR uses the nonparametric first-order condition to identify both the flexible input's elasticity from the observed revenue share of that input, and the ex-post shock to output. In particular, these authors use a standard sieve series estimator, and propose a finite dimensional truncated linear series given by a complete polynomial of second degree. This procedure is analogue to identifying the intermediate input coefficient directly from the revenue share in a Cobb-Douglas production function, where the ex-post shocks to production are the residuals, but is generalized to allow for the estimation of the parameters of more general production functions.[12]

In the second step, GNR's algorithm exploits an additional piece of information provided by the first-order condition—that the flexible input elasticity defines a partial differential equation on the production function, which imposes nonparametric cross-equation restrictions. Invoking the fundamental theorem of calculus, they can integrate this differential equation to obtain information on the production function. Given the polynomial sieve estimator used for the first stage, this integral has a closed-form solution, so it is possible to recover an approximated value for the unobserved productivity (plus the constant of integration). To separate these two, GNR then runs a nonparametric regression of $\omega_{jt}(\beta_k, \beta_l)$ on $\omega_{jt-1}(\beta_k, \beta_l)$ to recover the innovation component $\eta_{jt}(\beta_k, \beta_l)$, and given the orthogonality

[12] GNR illustrates the estimation method with a translog production function, that nests the more restrictive Cobb-Douglas functional form. In this paper, however, we apply GNR's method to estimate a Cobb-Douglas function, since our data does not allow us to consistently estimate all the coefficients of a translog function.

between this innovation component and capital and labor, these terms are then interacted to form the moment conditions. The remaining coefficients, β_k and β_l, are estimated with GMM.

GNR has several advantages over the more traditional proxy methods. First, since the estimation algorithm is developed for a general production function, it can be used to estimate a wide arrange of functional forms. Second, since it is not necessary to invert any input demand function to approximate a firm's productivity, the unobserved term ω_{jt} does not need to be restricted to a scalar, acknowledging that productivity may be multidimensional. And third, although this methodology is consistent with the monotonicity assumption (i.e. the demand for intermediate inputs being strictly increasing in productivity), this assumption is not needed for the estimation. Despite these advantages, in this paper we follow both the proxy methods and the share equation method to estimate the production function coefficients and recover the unobserved productivity of manufacturing firms. We then compare the estimates that we obtain when we follow all the described methodologies, such that we can analyze how the different estimation procedures affect the estimated TFP.

3 Data

In our analysis, we use a firm-level dataset which contains detailed balance sheet and operational information. Our data on firms' production and input consumption come from "Superintendencia de Sociedades," the agency in charge of supervising corporations. Specifically, the data come from the "Sistema de Información y Riesgo Empresarial" (SIREM) database.[13] The data are at an annual frequency and are self-reported by the firms. We

[13]The SIREM includes information for relatively large firms, and for firms in financial trouble. In particular, firms must report their financial data if their assets and/or income (adjusted by inflation) are grater than 30,000 times the current legal monthly minimum wage, if their external liability is grater than the total assets, if the financial expenditures are at least 50% of their income, if their cash flow is negative, or if their losses reduce the net equity below 70% of the social capital.

have access to public information such as balance sheets, as well as to confidential data included in the annexes filed by the firms.[14] These variables include the income obtained from the sales of each product, the use of raw materials, investments, and the capital stock. Additionally, we observe the number of employees and the payroll, broken down by type (executive, administrative, and production workers) and tenure (permanent or temporary).[15]

3.1 Data Description

The data from SIREM include information on firms from several industries. In general, we focus only on manufacturing firms, excluding manufacturers of coke, refined petroleum products, nuclear fuel, and basic metals (which include metals such as gold, silver, platinum, and nickel). We exclude the firms classified in these two manufacturing sectors because they are commodity producers, and therefore their dynamics are probably different from those of the other manufacturing firms. Our data cover the period 2005–2013.

Given our focus on the manufacturing sector, the first step prior to estimation was to define precisely which firms would be considered manufacturers. This step was relevant for multi-product firms that are not limited to manufacturing. In applying this definition, we took advantage of the rich data on income, reported by firms at the product level.[16] For our estimations, we consider as manufacturers only the firms that report having positive income from manufacturing products in all the years they appear in the sample.

In the presence of multi-product manufacturing firms, the second step was to decide

[14]We obtained access to the confidential data through the Banco de la República.

[15]The variables listed above are the relevant ones for our empirical work. The dataset also includes several other variables, such as detailed financial information. The information included in the SIREM dataset is very similar to the one included in the commonly used "Encuesta Anual Manufacturera" (EAM), conducted by the Colombian National Administrative Department of Statistics (DANE). We use the SIREM data for our estimations due to its availability, and because it can be linked to detailed trade data that we plan to use in future work.

[16]In the operational income annex, products are defined according to the International Standard Industrial Classification (ISIC, Revision 3.1), at the 4-digit level.

how to allocate each firm to a specific manufacturing sector.[17] Once again, we used the information on income by product, and we assigned each firm to the sector that includes the product that generated the most income throughout the sample period. Specifically, we added up the (deflated) income per product for 2005–2013, and assigned the firm to the manufacturing sector with the highest share.

With the subset of manufacturing firms clearly defined, the final step was to clean the data, given that the raw data from SIREM contains a large number of missing values and inconsistencies. The cleaning process included removing observations with exorbitant annual growth rates (perhaps confusing thousands with millions of Colombian pesos, or number of employees with payroll), as well as occasional value interpolation when a particular variable was missing for a single year.[18] Once we exclude those observations for which there were missing values for any variable, the resulting dataset contains 26,131 firm-year observations, corresponding to over 4,000 firms. This is the baseline sample we used in our estimations.[19]

Table 2 presents some basic statistics of our SIREM sample. In the first column we observe that, on average, we have around 2,900 manufacturing firms per year. In the remaining columns we report, for the average firm in our sample, the income, capital stock, value of raw materials used, number of workers employed, and the share of these that were production workers. Thus, the average firm had an average annual income of 29.5 billion Colombian pesos of 2005, an average capital stock of 16 billion, used raw materials worth 12.7 billion, and employed 160 workers, of whom 55 percent were production workers.[20,21]

[17] By sector or industry we mean, specifically, a 2-digit industry based on the ISIC (Revision 3.1) classification. See Table 1 for the description of all manufacturing sectors considered.

[18] See Appendix A for details on the data cleaning process.

[19] Our baseline sample includes 88 observations for which the value of raw materials used for production is greater than the operational income, such that our calculated value-added (given by the difference between income and materials) is negative. We drop these observations from our sample whenever value-added is used for the estimation.

[20] The values for income, capital, and raw materials are expressed in billions (thousand millions) of Colombian pesos of 2005. Each variable was deflated using a variable-specific deflator.

[21] These values correspond to the overall manufacturing sector, excluding petrochemicals and basic metals. If we include these two sectors in our overall statistics, the aveage number of firms per year increases to 3,000, and average income, capital and materials increase to 30.6, 17.9, and 13.9 billions, respectively. The

In Table 3 we present the analogous statistics, broken down by industry and averaged over time.[22] From the table it is clear that there is great heterogeneity across sectors. For instance, manufactures of food products and beverage (ISIC 15) and motor vehicles (ISIC 34) have a similar average income, however the average number of employs is 20% larger in the former industry. Moreover, the distribution between production and other workers is quite dissimilar across the two sectors. In a similar fashion, sectors ISIC 15 (foods and beverage) and ISIC 18 (apparel) have work forces similar sizes, despite the latter uses significantly less amounts of their inputs (capital and materials) than the former.

3.2 Representativeness of the Data

We now compare our SIREM data against data from two alternative sources in order to evaluate its representativeness. This is particularly important given the novelty of our dataset and the fact that the data from SIREM are neither census-based nor from a random survey. Still, as we show next, we are able to capture a large share of the universe of Colombian manufacturing firms. First, we compare it with national accounts data, containing the official aggregate estimates for the manufacturing sector. Second, we compare it with the annual survey of manufactures, EAM.

In Table 4 we benchmark our data against data from these two other sources.[23] Given that each database contains a different set of variables, we can only compare the levels of (real) income and of permanent workers.[24] Still, these two variables are probably the most

labor statistics are very similar across the two samples.

[22]There are three industries for which we only have very few observations: ISIC 16 (tobacco), ISIC 30 (office and computing machinery), and ISIC 32 (radio, television and communication equipment). In order to avoid disclosing confidential information, we do not report individual statistics for these industries.

[23]In our comparisons, we use information on all manufacturing sectors (including ISIC 23 and 27, coke, refined petroleum products, and nuclear fuel, and basic metals) because there is not a one-to-one correspondence between the ISIC codes and the sector codes used in the national accounts data. Therefore, we cannot exclude only these two sectors from the industry totals. When we keep firms from these sectors, our sample increases to 26,887 observations (4,990 firms).

[24]In the case of income, we are specifically comparing the estimated value of output of the complete

important ones for our purposes. As can be seen, with our sample we cover more than half of all manufacturing income according to the national accounts, and almost two-thirds of the production from EAM. In terms of employment, for which we can only compare our SIREM data with the EAM data, our sample covers on average over 90 percent of the permanent workers in the EAM data. Thus, based on the information contained in Table 4 we conclude that our data provide a fairly comprehensive picture of the overall Colombian manufacturing sector.

4 Productivity Estimations

In this section, we present our estimates for the production function coefficients and the corresponding TFP for the period between 2005–2013. We follow OP, LP, ACF, and GNR to obtain alternative estimates for the firm-level productivity shock. For most of our empirical exercises, we use the estimations obtained under the ACF and GNR methods. We chose these methods as our baseline, since these represent the latest developments in the literature dedicated to the structural estimations of production functions, with ACF representing the proxy methods used by OP and LP as well. However, as we show below, the production function coefficients are always statistically significant and have the expected signs regardless of the method followed to estimate them, and the TFP general growth pattern we find is robust to the different productivity estimates.

Most of the coefficients we present below were obtained by pooling all manufacturing firms together and estimating a single production function. When firms are grouped together to estimate common production function coefficients, the implicit assumption is that all firms have the same input requirements. Hence, ideally one wants to group together only firms that

manufacturing sector (national accounts), the value of output (EAM), and operational income (SIREM). In all cases, the variables are expressed in billions of pesos of 2005.

have similar production technologies. In practice, this often means estimating sector-specific production functions, with sectors being defined as narrowly as possible. In our case, we can only estimate sector-specific coefficients when we follow ACF's methodology. The other three algorithms yield coefficients that are not statistically significant.[25] For this reason, we pool all firms together for most of our estimations. However, as we show in section 5.5, when we do estimate sector-specific coefficients, both the production function coefficients themselves and the resulting TFP estimates are very similar to the ones we obtain when we group all manufacturing sectors together.

In Table 5 we present the estimated production function coefficients for our baseline, using the largest possible sample in each case.[26] Standard errors, included in parentheses, were estimated with a bootstrap. For comparison purposes, Table 5 also includes the coefficients obtained when we estimate the same production function specifications using OLS instead of the methodologies proposed by ACF or GNR. All coefficients have the expected sign and are statistically significant. Moreover, it is clear from the table that, as the theory predicts, the OLS estimates for labor and raw materials are upward biased, while the capital coefficient is underestimated.[27] Overall, these results highlight the importance of estimating productivity with a method that solves the simultaneity problem.

In terms of the specific functional forms used for our estimations, there are two differences worth mentioning. First, as pointed out in Section 2, ACF's methodology is designed to estimate a value-added function, while GNR's methodology allows us to estimate a gross-

[25]We believe that we cannot estimate sector coefficients precisely because our final sample is relatively small. It is possible that the reason why we can identify the coefficients with ACF, is beacuse this is the methodology that estimates the smallest number of parameters.

[26]We present the full-sample estimations for our baseline methodologies only. The coefficients estimated following OP and/or LP are available upon request.

[27] In general, the sign of the predicted bias depends on the firm's ability to adjust inputs once it observes its productivity shock ω. In the case of relatively flexible inputs, if higher productivity leads to higher input consumption, when productivity is not accounted for in the estimation the coefficients are going to absorb both the input requirements and the productivity effect, and will have a positive bias. If an input cannot be easily adjusted, it is possible that firms with high levels of fixed input(s) withstand worse productivity shocks without exiting the market, and the corresponding coefficients will be biased downwards.

output function. Hence, when we follow ACF our dependent variable is the difference between operational income and the value of used materials, while we use operational income as our dependent variable when we follow GNR.[28] Second, GNR's methodology does not allow us to separate labor into different groups, so we have to aggregate different types of workers into a single labor variable. To account for possible differences between production and other workers, and to make both types more comparable, when aggregating the two groups we weight the latter by the ratio of the average administrative and executive (other) worker wage to the average production worker wage.

Given the estimated production function coefficients, we can calculate the productivity shock for every firm-year observation. To calculate annual averages for the manufacturing industry as a whole, we aggregate the individual TFP weighting each observation with the corresponding firm-level income.[29] Then, to make our different productivity measures comparable, we normalize the resulting estimates so that the value of the TFP index for the overall manufacturing sector equals 100 in 2005.

In Figure 1, we plot the evolution of the estimated TFP for the aggregated manufacturing industry during 2005–2013. We accompany the aggregated index with an estimated 95% confidence interval to help us gouge the precision of our TFP estimations.

To construct the confidence intervals, we begin by estimating the variance of the productivity of each firm. For this calculations, we take advantage of the bootstrap used to estimate standard errors. In particular, since we estimate the production function coefficients 200 times (each time, with a different random sample), we can use the coefficients estimated with each sample and calculate the productivity at the firm level, and we can calculate the variance of the productivity across samples for each firm. Then, we proceed to calculate the variance

[28] This is the reason why the estimation sample is smaller in the first case: when following ACF, we need to discard all the observations with a negative value-added.

[29] Whenever we aggregate firm-level TFP to get annual averages, we do so by weighting individual estimates by income unless otherwise specified.

for each aggregated index. Since the TFP index is a weighted average of firm-level productivities, $TFP_t = (w_{1t}TFP_{1t} + w_{2t}TFP_{2t} + \cdots + w_{Jt}TFP_{Jt})$, then the variance of the index can be calculated as $var(TFP_t) = w_{1t}^2 var(TFP_{1t}) + w_{2t}^2 var(TFP_{2t}) + \cdots + w_{Jt}^2 var(TFP_{Jt})$, where the weights w_{jt} are given by each firm's share of total output. Once we have a measure for the variance of each index, we calculate the upper and lower bounds of the confidence intervals as the aggregated TFP plus or minus two times the standard deviation.

From Figure 1 we can see that when we estimate the production function coefficients following ACF, the resulting TFP index is more volatile than the index calculated using our estimates à la GNR. This is consistent with the idea that value-added production functions overestimate the degree of productivity heterogeneity.

It is worth noting that despite the greater volatility, the general average TFP patterns are similar across estimation methodologies. In Figure 2 we present both indices for comparison purposes. We observe that, in both cases, the average TFP increases between 2005 and 2006, then decreases until it reaches its lowest value during the 2009-10 world crisis, and this drop is then followed by a somewhat steady recovery until 2013, the end of our sample. However, the *level* of the productivity index and the estimated cumulative growth rate during our sample period do depend on which methodology we use to estimate the production function coefficients[30] In particular, our results show that while the average TFP has fully recovered and the index is above 100 when estimated following ACF, it is still below its initial value when estimated following GNR.

Before we analyze the robustness of our estimations, we want to make an important remark regarding our productivity estimates. Since we do not observe physical units of outputs or inputs, our productivity measure is actually what is often referred to as "revenue productivity." Although it cannot be directly interpreted as the physical productivity that

[30] Using the confidence interval we construct for the TFP index à la GNR, we can reject the hypothesis of both indices being equal at standard statistical levels.

often comes to mind (that is, how many shirts a firm can produce with a given amount of cloth, hours of labor, and machinery), it is still a measure of a firm's performance. Moreover, since we are using revenue and expenditure on inputs as the arguments of our production function, we have an omitted price bias and differences in the prices charged or payed by firms may be captured in different productivity measures. We do not have the data necessary to correct these two issues. However, it is worth noting that *all* our estimates are affected by them. Therefore, if they affect our alternative estimation algorithms in similar ways, we believe that our comparisons and the exercises presented below are still valid.

5 Robustness Checks

In this section we test the robustness of the results presented in Section 4. In particular, we analyze the effect on our estimated average TFP of changes to the estimation procedure, changes to the sample, changes to the way we measure the labor variables used for the estimation, changes to the way we weight individual estimations to get the industry totals, and changes to the way we group firms to estimate the production function.

5.1 Changes to the Estimation Methodology

To compare among different methodologies, we begin by defining a fixed sample that can be used to estimate the production function coefficients under OP, LP, ACF, and GNR. This way, we are able to isolate the changes that result from the different methods of estimation from those that stem from different data requirements. Given that the procedure proposed by OP uses investment as the proxy variable, for our fixed sample we exclude the observations for which investment is missing or is equal to zero (1,040 observations). Similarly, we need to exclude the observations for which income is smaller that the expenditure in raw materials

(88 observations), such that we have positive value-added values for the estimations à la ACF and LP.[31] When we omit these observations we are left with 25,015 observations (12 observations have missing investment and negative value-added values).

The coefficients estimated with this fixed sample following the four methodologies are reported in Table 6. All relevant variables have the expected signs and are statistically significant, and the biases of the coefficients with respect to the OLS estimations are in the direction predicted by the literature. Again, we aggregate the calculated firm-level TFP weighting by firms' income to obtain the average productivity, and we normalize it so that it equals 100 in 2005. Figure 3 depicts the average TFP, estimated with the coefficients reported in Table 6. We find that the level of the productivity index can change greatly depending on how the production function coefficients are estimated. In particular, we find the GNR-based TFP to be consistently below other estimations, and that the estimations that use value-added as a measure for output are more volatile than those that use gross output instead. Despite these differences, however, we find that the trend is very similar for the five specifications: we observe a positive growth rate between 2005 and 2006, followed by a decrease in productivity until 2009-10, and a recovery from 2010 on.

5.2 Changes to the Estimation Sample

To complement the previous subsection, we keep the estimation method constant and we analyze how changes in the sample affects our results. In order to isolate the sample effect, we focus on the TFP estimations obtained following ACF's and GNR's methodologies, and we compare this new set of results to our baseline results presented above. The objective of these exercises is to assess whether we are inducing a selection bias when we do not use all the available information.

[31] We estimate two sets of coefficients following LP, using income and value-added alternatively as our output measure.

First, we study the effect of omitting observations with zero or missing values for value-added and/or investment.[32] Second, we analyze the effect of limiting the year-observations for the firms included in our sample. Third, we explore the effect of including firms classified in the sectors of manufacture of coke, refined petroleum products, and nuclear fuel (ISIC 23), and basic metals (ISIC 27), omitted from our baseline scenario.

In figure 4 we compare the average TFP estimated using only the observations with positive investment and/or value-added to our baseline results. If we omit observations without information on value-added, our results are very similar to our baseline. One possible explanation for this is that we only eliminate 88 observations. If we omit observations without information on investment, we observe larger differences with respect to our baseline. In general, if we do not include these observations we underestimate the average TFP. One possible explanation for this behavior is that we may be omitting firms that are already highly productive as a result of past investments, such that they are not investing additional resources. This result highlights the importance of recovering the production function parameters using a method that does not require a positive investment value for the estimation.

In Figure 5 we compare the TFP estimations for each of three samples that differ in the observations that we keep for each firm. The full sample, that is the one we use for our baseline, includes every firm-year observation for which we have information for every variable needed for our estimations. In this full sample, however, we observe some gaps over time for many firms. In order to address the noise that the entry and exit (from our sample) may generate, we also estimate the firms' TFP with two sub-samples. For our intermediate sample, we keep the observations from the largest period without gaps for which the firm reports information. For our small sample, we completely omit the firms that have any gaps in their information and keep only those that report their information continuously.[33]

[32]The elimination of observations with missing or negative value-added is only relevant for the estimations following GNR. Since ACF uses a value-added production function, all the estimations based on this methodology eliminate any observation with negative o missing value-added by default.

[33]Our full sample includes 26,131 firm-year observations. If we keep only the longest block of contin-

In general, the estimations with the intermediate sample are very similar to the estimations with the full sample. On the contrary, when we estimate the production function parameters with the small sample, the resulting average TFP index is consistently higher than our baseline index, although growth rates move in the same direction. One possible explanation for the differences between the TFP estimated using the small and the largest samples is that, if we include only those firms for which we have continuous information without any gaps at any year of our sample period, we may end up with a sample that excludes firms that are close to the size threshold (in terms of assets and/or sales), such that they enter our dataset when they have a good year, but that be relatively unprofitable on average.

For our third exercise of this subsection, we study how the average TFP index is affected when we include the manufacturers of coke, refined petroleum products, nuclear fuel, and basic metals in our sample. Results are depicted in Figure 6. When we include firms from ISIC sectors 23 and 27 the estimated average TFP follows a very similar growth pattern to our baseline results. However, including these firms results in higher average TFP indices after 2007. Nonetheless, the gap between our baseline and alternative TFP measures is reduced over time, and when we estimate productivity following GNR it completely disappears by the end of our sample period.

5.3 Changes to the Labor Measures

In this subsection we study the effect on aggregate TFP of using alternative labor measures. We tweak our labor definition in two ways. First, we change the way in which we include the number of workers in the production function. In particular, for our estimations à la ACF we re-estimate the production function parameters and calculate the average TFP using

uous information, our intermediate sample is reduced to 24,104 observations. And if we drop firms with information gaps, our small sample is further reduced to 20,217 observations.

only production workers (the ones that are directly related to a firm's output), and using aggregated measures instead of separating production from other workers. These aggregate measures include both a weighted sum, like the one we use when we follow GNR, and a simple sum of production plus other workers. For our estimations à la GNR, we re-calculate the average TFP using production workers only, and replacing the weighted sum with a simple sum.[34]

In our second set of exercises, we analyze the effect of measuring labor with wages instead of using the number of employees. By using wages instead of the number of workers, we account for the actual number of hours of labor, such that we control for part-time workers or workers that were employed only a fraction of each year. For our estimations à la ACF, we compare our baseline results with the average TFP estimated using wages paid to production and other workers separately, to production workers only, and total wages. For our estimations à la GNR, we compare our baseline results with the average TFP estimated using either total wages, or those paid to production workers only.

Again, in order to isolate the effect of the changes we describe below, we keep both the estimation methodology and the sample fixed. In order to keep our sample constant, and since not all firms in our baseline sample report a positive number for each type of workers or include reliable information on wages, for these exercises we only include those firm-year observations for which we have information on all the relevant variables.[35] Therefore, in addition to our baseline TFP index calculated with our full sample, we also present the TFP index calculated with our original labor measure (either production and other workers included separately, or a weighted sum of both types of workers), but using the reduced sample corresponding to each exercise.

[34]Recall that when we following ACF we can separate production workers from other workers, while when following GNR we need to use an aggregate measure of employment.

[35]When we compare different ways to include the number of employees, we are left with 20,175 observations for our estimations à la GNR, and with 20,120 observations for our estimations à la ACF. When we use information on wages, we are left with 24,000 observations and 23,938 observations, respectively.

Figure 7 shows the average TFP estimated with the alternative employee counts described above, and Figure 8 shows the TFP estimated replacing the number of workers with wages. In both figures we include our baseline estimations as a reference, although they are not entirely comparable to our new results since we need to use a smaller sample for these exercises. In general, when we reduce our sample to include only the observations for which we have a positive number of both types of workers and/or information on wages, our new estimates for the average TFP differ from our baseline but are very similar across exercises for any given sample. This suggests that, other than a potential selection bias resulting from the stricter selection of the firms that are included in our estimation sample due to additional data requirements, the effect of changing the way in which we measure labor is relatively small. Although the level of the average TFP series estimated using alternative varies slightly between estimations, growth rates exhibit very similar patterns. Moreover, from Figures 7 and 8 we can conclude that TFP estimations following GNR are less sensitive to changes to the labor measure than those obtained following ACF.

5.4 Changes to the Averaging Weights

The natural way to aggregate firm-level productivity into an industry-level TFP index is to compute an average. Since each firm accounts for a different percentage of the manufacturing industry, we need to construct a weighted average to better approximate industry totals. In general, we weight each firm with its income, since income is our output measure. In this subsection we explore how sensitive the results are to changes in the averaging weights. In particular, we calculate new averages using labor or assets as weights, and we calculate a simple average as well. Once again, we keep our estimation methodology fixed, and we show the TFP that results from estimating the production function following ACF and GNR.

Our results are plotted in Figure 9. For both methodologies, when we weight firm-specific

TFP using assets to construct our yearly average, the resulting TFP index is very similar to our baseline, weighted using income. However, when we construct our index with a weighted average using labor shares, or when we calculate a simple average, the aggregate TFP becomes flatter. In particular, the peaks we observed in 2006 and 2009 are flattened out. If we compare the two estimation methodologies, we find that the TFP index estimated following GNR is less sensible to changes in the way we average firm-specific productivity.

5.5 Estimation by Sectors

As mentioned at the beginning of this section, the estimation of a common production function for different firms implicitly assumes that all firms included in the estimation have the same input requirements. Hence, when estimating production functions, one wants to group firms that have relatively similar technologies. One way of doing this is by estimating a production function specific to each sector, instead of pooling all firms and estimating one for the manufacturing industry as a whole. In this subsection we compare our baseline estimations with the TFP obtained when we estimate industry-specific coefficients for the production function. We focus on estimations obtained following ACF, since this is the only methodology that yields reliable estimators at the 2-digit sector level, and allows us to recover the elasticity of production factors for most sectors.

Our sector-specific estimations include the largest industries (in terms of the number of firms). These represent around 98% of the observations included in our baseline sample. Sectors 16, 32, 33 and 35 are very small, and we have very few observations, so we do not have enough variation to identify the coefficients of the production function for theses sectors. Sector 30 is also very small; however, we group sectors 29, 30 and 31 in one big machinery sector and we compute a joint production function. This allows to identify the common coefficients for the three sectors.

The production function coefficients for each 2-digit sector are reported in Table 7. If we compare each sector-specific coefficient with the corresponding coefficient estimated for the manufacturing industry as a whole, we observe very little variation. To quantify the differences with more precision, we construct a modified version of a coefficient of variation that measures the average deviation of the sector-specific coefficients from the coefficient estimated pooling all manufacturing firms.[36] According to this measure, the average deviation is 0.23, 0.14, and 0.19 for production workers, other workers, and capital coefficients, respectively. These results suggest that our industry-specific coefficients are, on average, very close to our industry-wide, baseline coefficients.

In Figure 10, we compare the average TFP calculated using our baseline estimations for a single production function for the manufacturing industry as a whole, with the TFP we obtain when we use our sector-specific coefficients, reported in Table 7. First, we compute a yearly index by sector, aggregating the firm-specific TFP of firms in each sector with a weighted average using the income shares as weights. Then, we average the TFP index across years, and plot this average for each sector. In general, the average indices are very similar for all sectors. Comparing the two estimations, it is not clear that we are either under or overestimating the TFP index when we calculate it using common, industry-wide coefficients.

In addition, in Figure 11 we plot both TFP indices across years for the four biggest sectors (food products and beverages, wearing apparel, chemicals and chemical products, and rubber, and plastics products). Our results suggest that both industry-wide coefficients and sector-specific coefficients result in indices that exhibit very similar growth patterns across time, although the value of the estimated average TFP can be different in each case.

[36]Given our industry-specific coefficients β^i_j, and the cross-industry (baseline) coefficients β^i, the formula we use to calculate the modified coefficient of variation for each β is $CV = \sigma_{j,i}/\beta^i$, where, $\sigma_{ji} = \sqrt{\frac{\sum_{j=1}^{J}(\beta^i_j - \beta_i)^2}{J}}$ is the standard deviation, $i = \{L^P, L^O, K\}$ denotes each production factor, and $j = \{1, 2, \cdots, J\}$ denotes each one of the 16 sectors for which we have industry-specific coefficients.

The direction of the bias, however, is not clear. Note that the largest differences are in the sectors of apparel and chemicals; in the first case we underestimate the level of the average TFP if we calculate it using the common manufacturing industry coefficients, while in the second case we overestimate it.

6 Conclusion

A robust, reliable estimation of productivity is a fundamental input for the study of the dynamics of the manufacturing sector. In this paper, we apply different structural techniques to estimate the coefficients of the production function, necessary to calculate total factor productivity. We then use our estimates to calculate an average TFP index for the Colombian manufacturing industry, and analyze its evolution during 2005–13. In addition, we study the robustness of our results to changes to the estimation methodology, estimation sample, variable definition, and/or aggregation of firm-level TFP.

We find that, regardless of how the average TFP is estimated, its general evolution across time is very robust. In particular, we find that after an initial increase between 2005 and 2006, average productivity declines every year until it reaches its lowest value during the 2009-10 crisis, and grows steadily during the last years of our sample. However, the level of the TFP index does change across different estimation methodologies. While we observe a positive cumulative growth rate in most cases, when we estimate the production function coefficients following GNR, the resulting average productivity index ends up being smaller in 2013 than it was at the beginning of our sample period.

In terms of the sensitivity of our estimations to different assumptions such as the estimation sample, the functional form of the underlying production function, and/or the way we measure the relevant variables, we find that estimations following GNR are more robust.

Our aggregate TFP measures following their technique are less affected by changes to the sample, the estimation variables, and the way we aggregate individual estimations, than estimations à la ACF.

One drawback of our paper is that our data is not a census, nor are the firms in our sample randomly selected. Although the comparisons between our baseline and the results presented in Sections 5.1 to 5.5 yield interesting insights regarding the effect of several estimation assumptions, it is possible that our findings on the evolution of TFP and its growth over our sample period are not representative of the Colombian manufacturing sector. In future work, we plan to estimate productivity using the more comprehensive plant-level data from the EAM, and compare these results with the ones presented in this paper.

Besides including a larger set of firms, there are a couple of advantages of using this alternative data source. One advantage is that the larger number of observations will probably allow us to estimate sector-specific coefficients following GNR (or other algorithms). Although our results from Section 5.5 suggest that the effect of pooling all firms together and estimating a common production function for the manufacturing sector as a whole instead of estimating industry-specific ones is rather small, it will be interesting to relax this restriction for alternative estimation methodologies. A second advantage of the EAM is the availability of price data that will allow us to control for omitted input and output prices, and to estimate a measure of physical productivity instead of the revenue productivity that we analyze in this paper.

References

Ackerberg, D., L. Benkard, S. Berry, and A. Pakes (2007): "Econometric Tools for Analyzing Market Outcomes," in *The Handbook of Econometrics*, ed. by J. Heckman, and E. Leamer. North-Holland, Amsterdam.

Ackerberg, D., K. Caves, and G. Frazer (2015): "Identification Properties of Recent Production Function Estimators," *Econometrica*, 83(6), 2411–2451.

Carranza, J. E., and S. Moreno (2013): "Tamaño y estructura vertical de la cadena de producción industrial colombiana desde 1990," Borradores de Economía No. 751, Banco de la República.

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

Echavarría, J. J., M. A. Arbeláez, and M. F. Rosales (2006): "La productividad y sus determinantes: el caso de la industria colombiana," Borradores de Economía No. 374 (Bogotá, Banco de la República).

Eslava, M., J. Haltiwanger, A. Kugler, and M. Kugler (2004): "The Effects of Structural Reforms on Productivity and Profitability Enhacing Reallocation: Evidence from Colombia," *Journal of Development Economics*, 75(2), 333–371.

Gandhi, A., S. Navarro, and D. Rivers (2016): "On the Identification of Prodduction Functions: How Heterogeneous is Productivity?," (Unpublished; University of Western Ontario).

Griliches, Z., and J. Mairesse (1998): "Production Functions: The Search for Identification," in *Econometrics and Economic Theory in the Twentieth Century: The Ragnar*

Frisch Centennial Symposium, ed. by Z. Griliches. Cambridge University Press, Cambridge, MA.

Levinsohn, J., and A. Petrin (2003): "Estimating Production Functions Using Inputs to Control for Unobservables," *Review of Economic Studies*, 70(2), 317–342.

López, R. (2006): "Impacto de las Exportaciones en la Productividad del Sector Manufacturero Colombiano," Archivos de Economía No. 299 (Bogotá, Departamento Nacional de Planeación).

Marschak, J., and W. H. Andrews (1944): "Random Simultaneous Equations and the Theory of Production," *Econometrica*, 12(3), 143–205.

Meléndez, M., and K. Seim (2006): "La productividad del sector manufacturero colombiano y el impacto de la política comercial: 1977-2001," *Desarrollo y Sociedad*, 57, 1–41.

Olley, G. S., and A. Pakes (1996): "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 64(6), 1263–1297.

Pombo, C. (1999): "Productividad industrial en Colombia: una aplicación de números índices," *Revista de Economía de la Universidad del Rosario*, 2(1), 107–139.

Appendix

A Data Cleaning

The original SIREM dataset includes over 223,000 firm-year observations for the period 2005–2013, with an average of around 25,000 firms per year. In this appendix, we describe how we cleaned the data to construct our dataset.

In order to select manufacturers, we started by looking at the data on income by product. Out of the 223,623 observations, 203,096 have data on income, and 43,068 report income from a manufactured product. As mentioned in Section 2.2, we defined as manufacturers those firms with a positive income from manufactured products for *every* year they appear in our sample. This selection criterion left us with 36,968 observations, corresponding to 5,760 firms.

Once we defined the subset of manufacturing firms clearly, we proceeded to clean the data in several steps, and we reduced our sample to those firms for which we had complete, consistent information for all the variables we need for our TFP estimation (operational income, capital stock, value of the raw materials used by each firm, and number of workers).

First, we eliminated the firms for which we had no information on capital or raw materials. (Given the way we selected manufacturing firms, we had complete information for income, by construction). We eliminated 2,456 observations corresponding to those firms that did not have information on raw materials throughout the sample, and 38 additional observations that did not have information on capital stock.

Next, we identified firms with exorbitant annual growth rates (500 percent or more) for income, capital, and/or raw materials. Firms are requested to report these variables in thousands of pesos. However, an informal look at the database suggests that, in some cases, firms might have been mixing reporting units: sometimes they appear to have reported these variables in pesos or in millions of pesos, thereby introducing noise to our sample. By looking at firms with growth rates above this threshold, we were able to identify observations that seemed to mix reporting units. In these cases, we either multiplied or divided the reported value by 1,000 (or the appropriate number) to make it comparable to the observations for the same firms in different years. Overall, we changed 58 income observations, 147 raw materials observations, and 66 capital observations (in some cases to zero, when it was not clear how to "fix" a suspicious observation).

Next, if a firm was missing information for a single year for any of these three variables (but not all), we filled the gaps by interpolating the information of the adjacent years. We were able to approximate 382 missing values for raw materials, six for capital, and two for income.[37] Of course, this approach is not valid if we are missing values for the first or the final years of a firm. In these cases, we eliminated 2,349 observations due to missing information on raw materials, and 82 observations due to missing information on the capital stock. In addition, we dropped 323 observations for which the information gaps were longer than two years, such that we were unable to approximate the value of capital or raw materials (four were missing information for capital, and 319 were missing information for raw materials).

The calculation of growth rates for the number of employees allowed us to identify a different kind of mix-up for labor variables: in some cases, the number of employees and the value of wages seemed to be transposed. We identified 171 observations for which this seemed to be a problem, and interchanged the values manually. We did this for every category (male/female, permanent/temporary, and production/administrative/executive workers). For some of these 171 observations, we had to fix more than one labor variable.

[37]The two zeros in the income series were the result of changes made in the previous step.

After cleaning labor variables, we calculated the total number of employees per firm-year and eliminated those observations with zero workers. We did not attempt to fill one-year gaps by interpolating this variable, since we do not have a clear way of distributing employees into the different categories used in our estimations. In this step, we eliminated 4,813 observations.

In addition to these fixes, we identified nine firms with inexplicably high growth rates for some variable, but for which it was not clear that there is a problem with the reporting unit, nor it was clear how to properly interpolate to obtain plausible values for all variables. We dropped the corresponding 19 observations.

This cleaning process left us with 26,888 observations, corresponding to 4,989 firms. If we eliminate all the observations from firms classified as manufacturers of coke, refined petroleum products, and nuclear fuel (ISIC 23), or as manufacturers of basic metals (ISIC 27), we are left with 26,131 observations, corresponding to 4,878 firms. This is the sample we use throughout the paper for our estimations.

B Figures

Figure 1: TFP Index (Baseline)

(a) ACF

(b) GNR

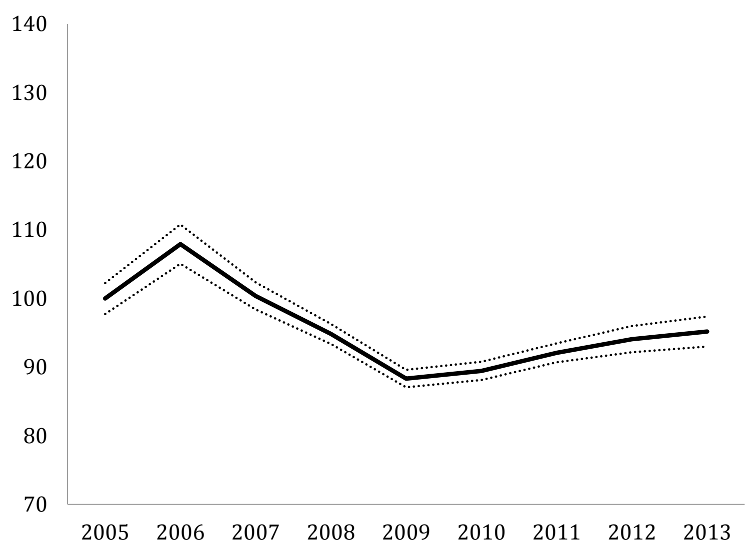

Figure 2: Baseline TFP Index Comparison

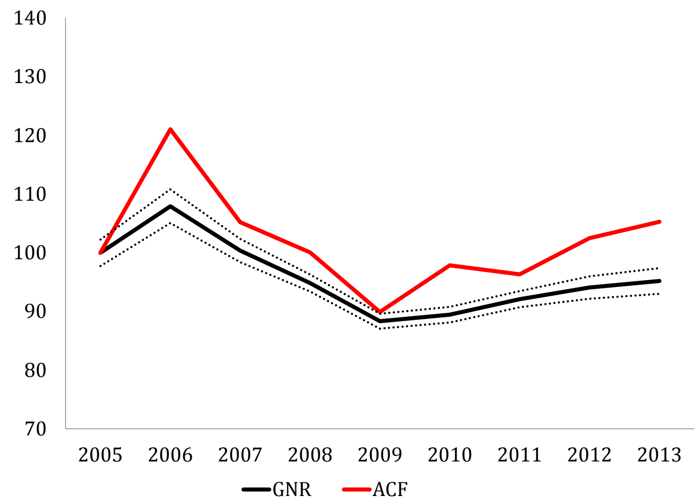

Figure 3: Alternative Estimation Methodologies

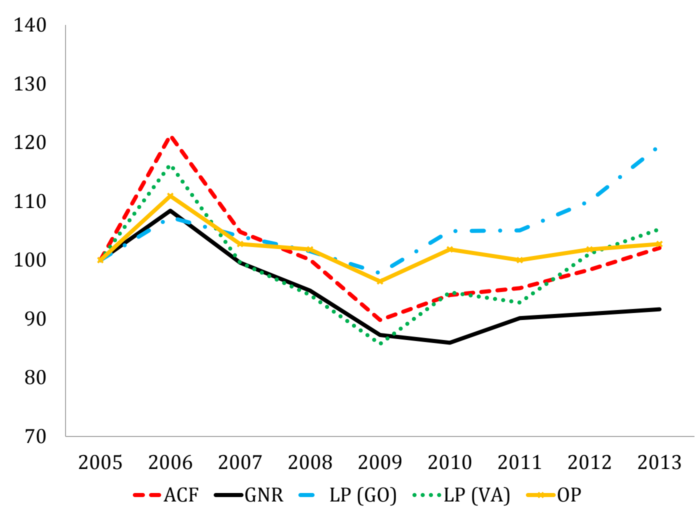

Figure 4: Alternative Samples: Positive Investment and/or Value-Added

(a) ACF

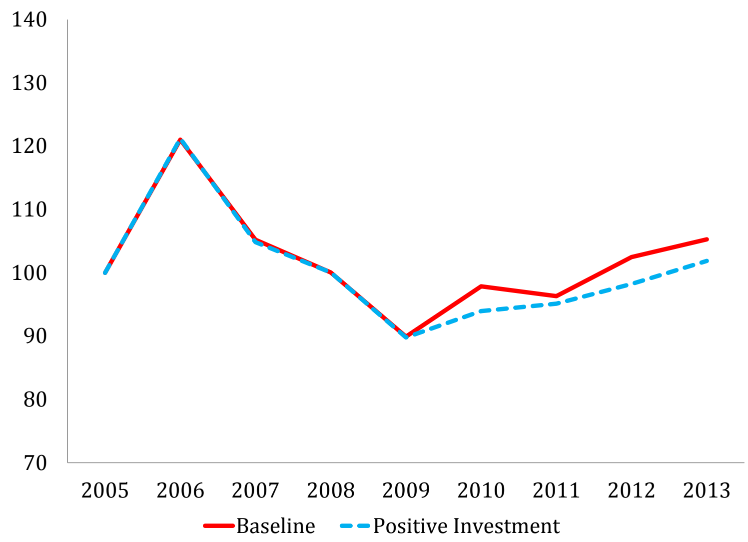

(b) GNR

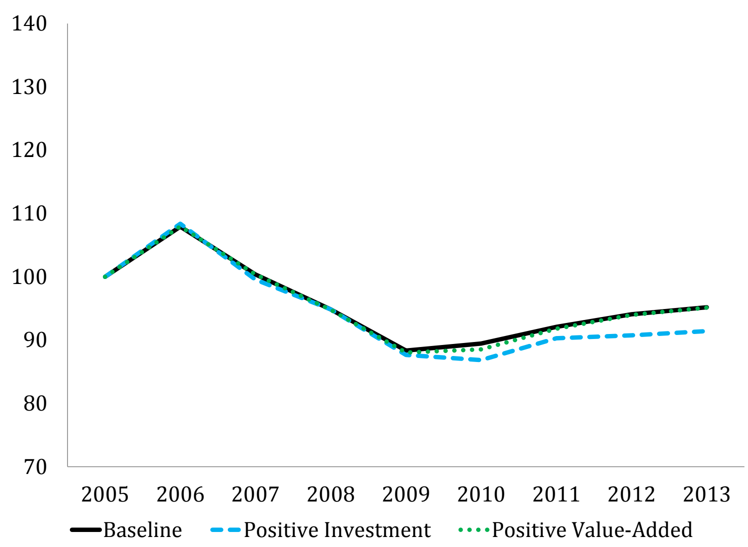

Figure 5: Alternative Samples: Sample Size

(a) ACF

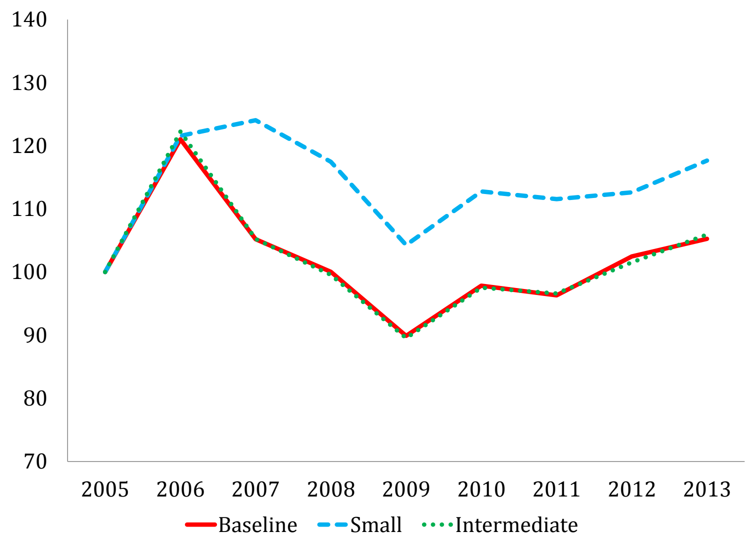

(b) GNR

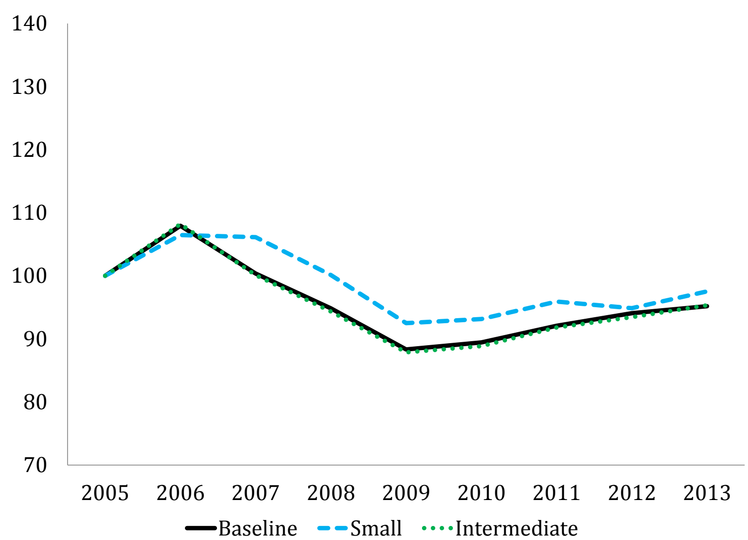

Figure 6: Alternative Samples: Including Petrochemicals and Basic Metals

(a) ACF

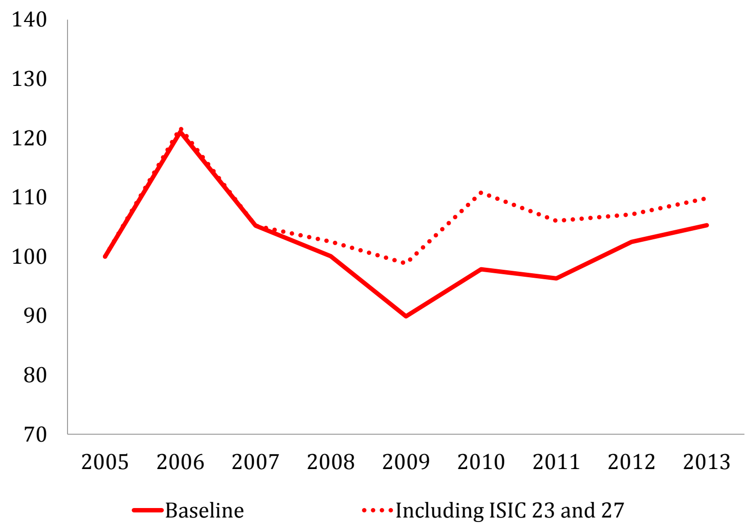

(b) GNR

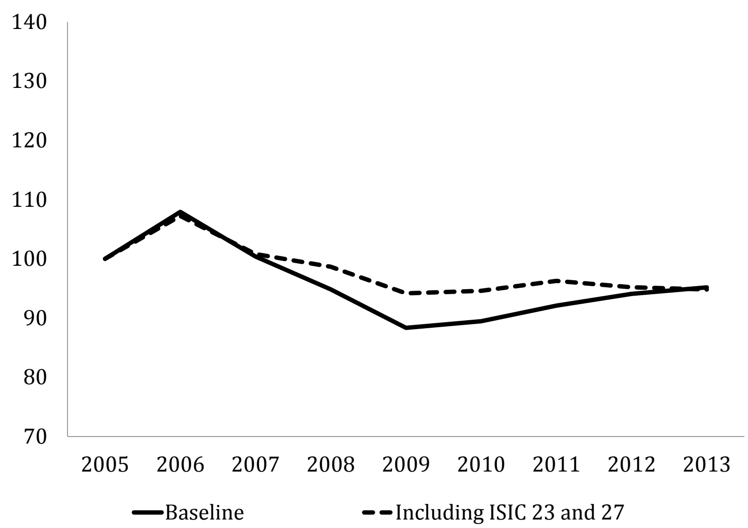

Figure 7: Alternative Labor Measures: Number of employees

(a) ACF

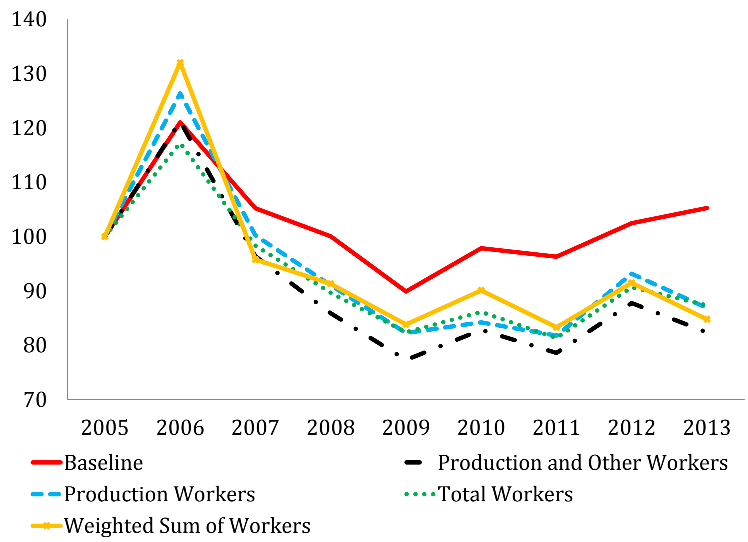

(b) GNR

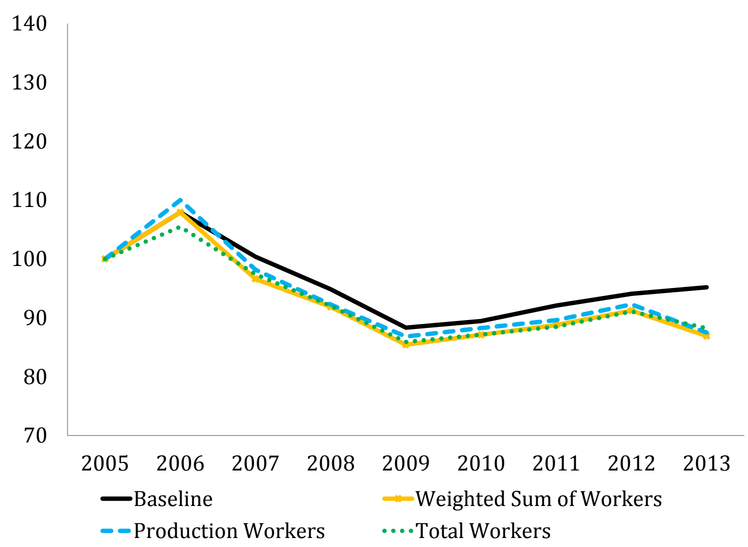

Figure 8: Alternative Labor Measures: Wages

(a) ACF

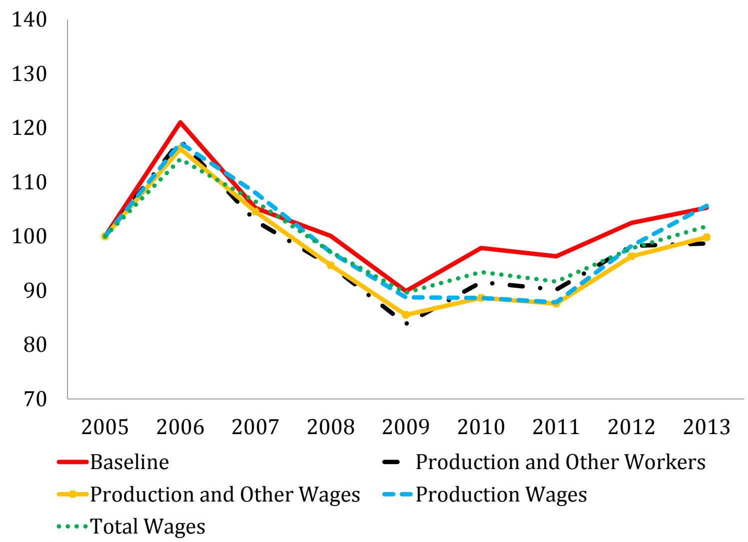

(b) GNR

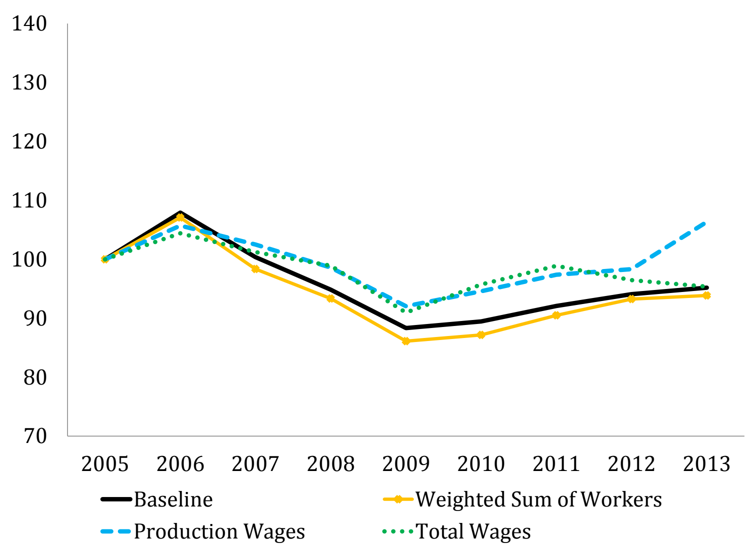

Figure 9: Alternative Aggregation Weights

(a) ACF

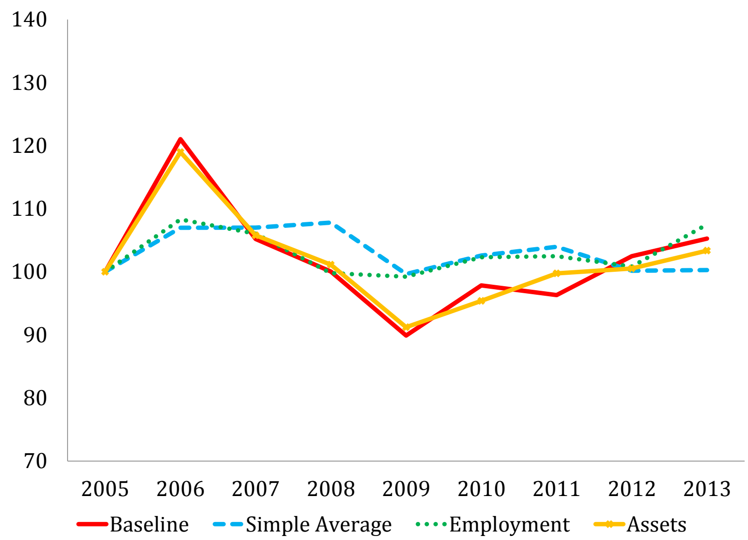

(b) GNR

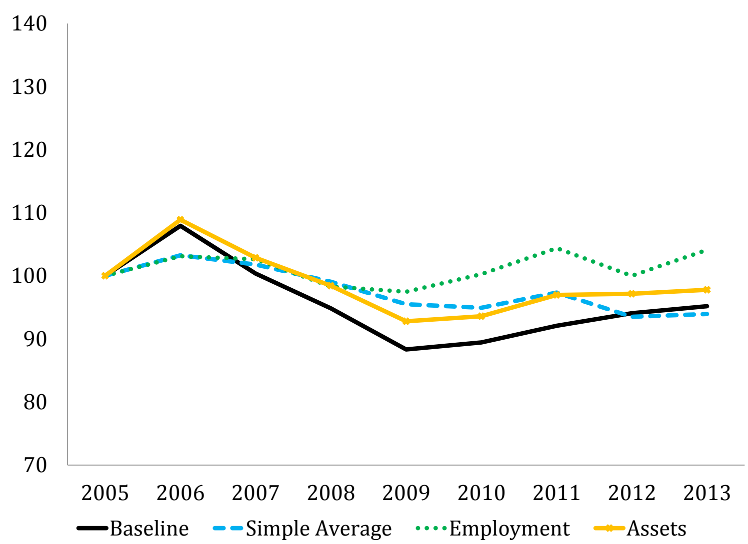

Figure 10: Average TFP Index by Sector

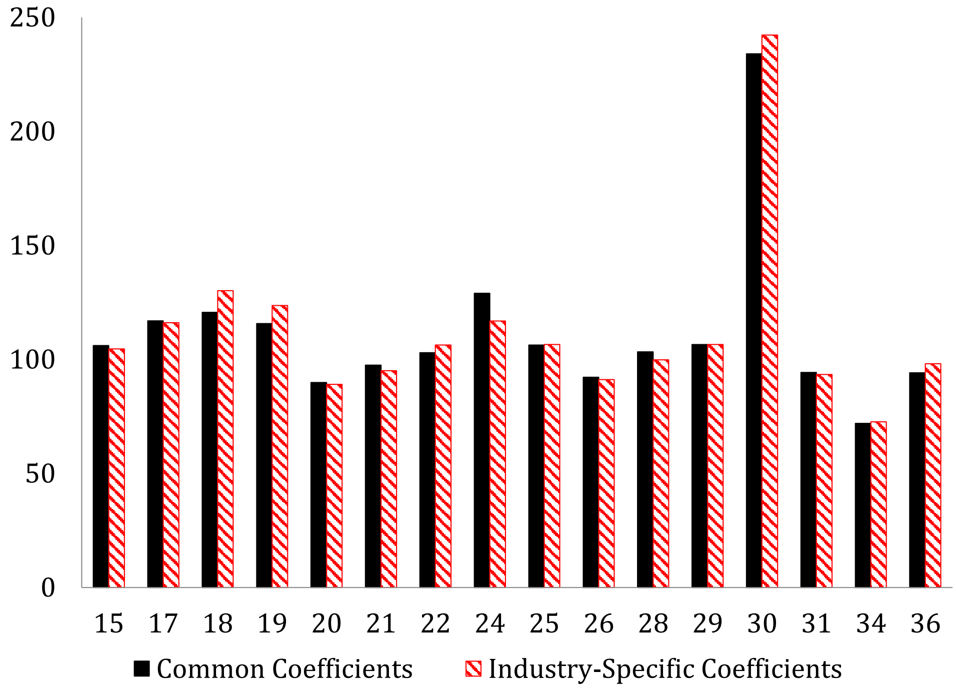

Figure 11: Sector-Specific TFP Index

(a) Food products and beverages

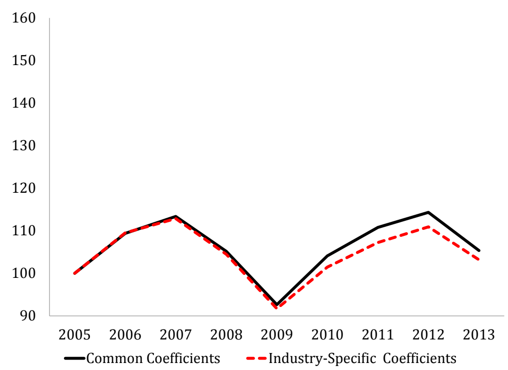

(b) Wearing apparel

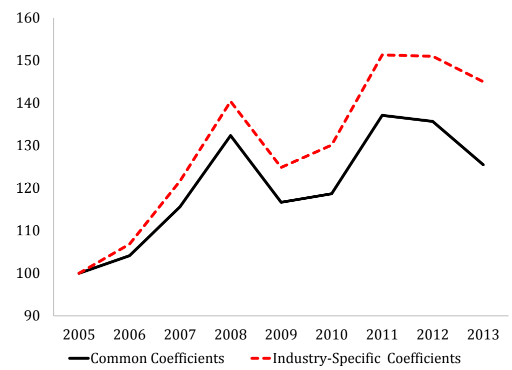

(c) Chemicals and chemical products

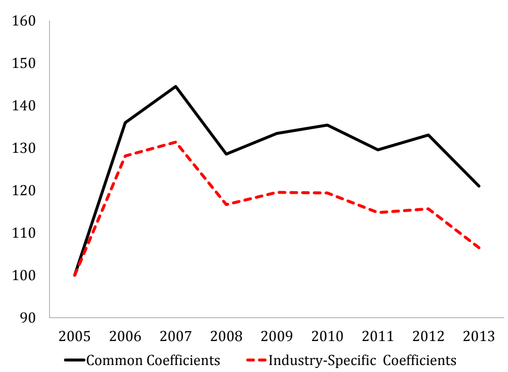

(d) Rubber and plastics products

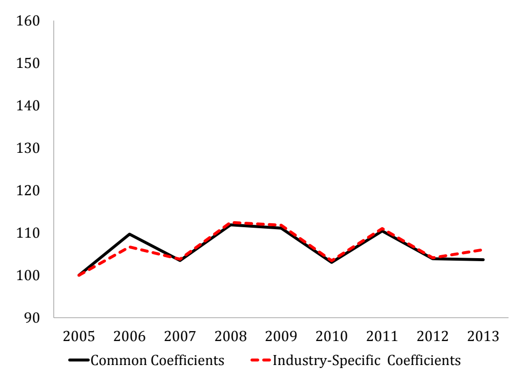

C Tables

Table 1: Industry (ISIC Rev. 3.1) Codes–Section D

Code	Description
15	Manufacture of food products and beverages
16	Manufacture of tobacco products
17	Manufacture of textiles
18	Manufacture of wearing apparel; dressing and dyeing of fur
19	Tanning and dressing of leather; manufacture of luggage, handbags, saddlery, harness and footwear
20	Manufacture of wood and of products of wood and cork, except furniture; manufacture of articles of straw and plaiting materials
21	Manufacture of paper and paper products
22	Publishing, printing and reproduction of recorded media
23	Manufacture of coke, refined petroleum products and nuclear fuel
24	Manufacture of chemicals and chemical products
25	Manufacture of rubber and plastics products
26	Manufacture of other non-metallic mineral products
27	Manufacture of basic metals
28	Manufacture of fabricated metal products, except machinery and equipment
29	Manufacture of machinery and equipment n.e.c.
30	Manufacture of office, accounting and computing machinery
31	Manufacture of electrical machinery and apparatus n.e.c.
32	Manufacture of radio, television and communication equipment and apparatus
33	Manufacture of medical, precision and optical instruments, watches and clocks
34	Manufacture of motor vehicles, trailers and semi-trailers
35	Manufacture of other transport equipment
36	Manufacture of furniture; manufacturing n.e.c.

Table 2: Basic Statistics: Overall Manufacturing

	Firms (#)	Income ($)	Capital ($)	Materials ($)	All Workers (#)	Production Workers (%)
2005	2,832	25.4	11.4	11.6	146.7	58.9
2006	3,275	25.4	10.7	11.4	142.3	59.0
2007	2,860	30.1	14.6	13.3	162.3	57.9
2008	2,805	29.9	16.3	12.9	167.6	54.8
2009	3,001	26.9	15.4	11.4	150.0	54.8
2010	2,888	29.2	17.6	12.4	154.9	55.2
2011	2,979	30.2	17.9	12.7	160.3	54.3
2012	2,848	32.0	18.9	13.3	170.6	52.6
2013	2,644	35.9	21.3	14.8	179.4	51.1
Average	2,904	29.5	16.0	12.7	159.3	55.4

Source: Authors' calculations based on data from SIREM and DIAN/DANE.

Notes: This sample excludes manufacturers of coke, refined petroleum products and nuclear fuel (ISIC 23), and manufacturers of basic metals (ISIC 27). The sign '$' corresponds to billions of Colombian pesos of 2005.

Table 3: Basic Statistics: by Sector

Sector:	15	16	17	18	19	20	21	22	23	24	25
Firms (#)	500	2	154	252	76	44	61	220	24	349	330
Income ($)	50.94	-	17.82	14.01	12.22	6.68	63.98	13.87	119.28	49.43	16.88
Capital ($)	29.00	-	11.31	5.11	3.56	6.83	52.84	9.06	198.06	22.90	11.35
Materials ($)	26.83	-	6.59	3.87	3.49	1.90	27.05	3.44	148.49	18.50	6.85
Workers (#)	226.14	-	170.18	199.91	155.99	81.53	227.83	117.12	80.26	163.65	114.00
Production Workers (%)	51.69	-	71.02	59.67	62.77	65.88	57.41	36.62	52.27	39.86	64.88

Sector:	26	28	27	29	30	31	32	33	34	35	36
Firms (#)	137	239	60	94	3	53	2	15	105	11	257
Income ($)	30.20	17.13	47.28	20.01	-	30.64	-	8.40	51.79	161.27	14.35
Capital ($)	34.66	8.98	34.20	7.83	-	11.54	-	7.33	10.81	25.15	5.79
Materials ($)	8.39	9.33	22.71	8.28	-	11.47	-	2.42	26.64	53.09	6.28
Workers (#)	165.73	100.50	130.16	152.14	-	173.87	-	108.56	181.92	489.38	102.47
Production Workers (%)	64.30	63.06	68.95	64.46	-	50.40	-	66.81	69.01	45.93	56.15

Source: Authors' calculations based on data from SIREM and DIAN/DANE.
Notes: We include manufacturers of coke, refined petroleum products and nuclear fuel (ISIC 23), and manufacturers of basic metals (ISIC 27) for comparison purposes. The sign '$' corresponds to billions of Colombian pesos of 2005.

Table 4: SIREM vs. Alternative Data Sources

	Income			Permanent Workers	
	Sample ($)	% Nat. Acc.	% EAM	Sample (#)	% EAM
2005	75,560.3	49.6%	64.0%	281,992	86.2%
2006	87,711.6	54.0%	68.7%	319,693	96.0%
2007	89,467.9	51.4%	65.7%	339,697	97.4%
2008	90,213.1	51.5%	66.9%	340,276	93.6%
2009	86,234.3	51.8%	65.5%	335,443	90.9%
2010	90,680.9	53.1%	65.6%	320,843	84.8%
2011	97,606.7	54.2%	59.9%	342,733	88.8%
2012	98,048.3	54.2%	60.8%	370,510	95.7%

Source: Authors' calculations based on data from SIREM and DIAN/DANE.

Notes: In order to have comparable samples, data include firms manufacturing coke, refined petroleum products, nuclear fuel and basic metals. The sign '$' corresponds to billions of Colombian pesos of 2005.

Table 5: Production Function Coefficients (Baseline)

	ACF (VA)		GNR (GO)	
	OLS	ACF	OLS	GNR
Capital	0.355***	0.4231***	0.221***	0.3935***
	(0.0095)	(0.0124)	(0.0079)	(0.0176)
Materials			0.545***	0.3742***
			(0.0103)	(0.0011)
Production	0.183***	0.1583***		
Workers	(0.0067)	(0.0055)		
Other	0.577***	0.5223***		
Workers	(0.0122)	(0.0120)		
Total			0.218***	0.1119***
Workers			(0.0083)	(0.0096)
Observations	26,043	26,043	26,131	26,131

Notes: Standard errors in parenthesis. For estimations following ACF or GNR, standard errors bootstrapped. '***', '**' and '*' denote statistical significance at the 1%, 5%, and 10% levels, respectively. 'GO' denotes the use of gross output as our measure of production, while 'VA' denotes the use of value-added.

Table 6: Production Function Coefficients (Alternative Estimation Methodologies)

	ACF		GNR		LP (GO)		LP (VA)		OP	
	OLS	ACF	OLS	GNR	OLS	LP	OLS	LP	OLS	OP
Capital	0.3648***	0.4653***	0.2294***	0.4588***	0.1877***	0.0851**	0.385***	0.3553***	0.1868***	0.24***
	(0.0055)	(0.0122)	(0.0039)	(0.0168)	(0.0039)	(0.0413)	(0.0052)	(0.0357)	(0.0041)	(0.0240)
Materials			0.5412***	0.3721***	0.5149***	0.4191***			0.5153***	0.5066***
			(0.0051)	(0.0012)	(0.0050)	(0.0812)			(0.0050)	(0.0104)
Permanent Pdn. Workers					0.0418***	0.045***	0.1138***	0.0786***	0.043***	0.041***
					(0.0023)	(0.0038)	(0.0035)	(0.0060)	(0.0023)	(0.0037)
Permanent Other Workers					0.228***	0.2165***	0.4456***	0.3356***	0.229***	0.2235***
					(0.0045)	(0.0062)	(0.0067)	(0.0080)	(0.0046)	(0.0098)
Temporary Pdn. Workers					0.0368***	0.0307***	0.0893***	0.0555***	0.0376***	0.0345***
					(0.0020)	(0.0027)	(0.0031)	(0.0041)	(0.0020)	(0.0040)
Temporary Other Workers					0.064***	0.0624***	0.1256***	0.0998***	0.0651***	0.0612***
					(0.0028)	(0.0042)	(0.0042)	(0.0059)	(0.0028)	(0.0059)
Production Workers	0.1771***	0.1491***								
	(0.0035)	(0.0059)								
Other Workers	0.5704***	0.5063***								
	(0.0074)	(0.0115)								
Total Workers			0.2123***	0.0974***						
			(0.0046)	(0.0097)						
Age									-0.0007***	0.0003
									(0.0002)	(0.9300)
Trend									0.0071***	0.0085***
									(0.0012)	(0.0015)
Observations	25,015	25,015	25,015	25,015	25,015	25,015	25,015	25,015	25,015	25,015

Notes: Standard errors in parenthesis. For estimations following ACF, GNR, LP, or OP, standard errors bootstrapped. '***', '**' and '*' denote statistical significance at the 1%, 5%, and 10% levels, respectively. 'GO' denotes the use of gross output as our measure of production, while 'VA' denotes the use of value-added.

Table 7: Sector-Specific Production Function Coefficients

2-digit Sector	Production Workers	Other Workers	Capital	Observations
15	0.1526*** (0.0132)	0.4662*** (0.0239)	0.4765*** (0.0339)	4,500
17	0.1899*** (0.0256)	0.4663*** (0.0521)	0.3924*** (0.0608)	1,390
18	0.1808*** (0.0249)	0.5238*** (0.0403)	0.2595*** (0.0496)	2,270
19	0.2399*** (0.0580)	0.4826*** (0.1116)	0.3483*** (0.1266)	684
20	0.1895*** (0.0418)	0.6048*** (0.1006)	0.2457*** (0.0835)	395
21	0.1264* (0.0790)	0.3766*** (0.1099)	0.5093*** (0.1569)	550
22	0.1776*** (0.0282)	0.5894*** (0.0552)	0.3523*** (0.0688)	1,975
24	0.0667*** (0.0145)	0.5933*** (0.0380)	0.473*** (0.0298)	3,135
25	0.1941*** (0.0232)	0.4725*** (0.0496)	0.3727*** (0.0516)	2,972
26	0.17*** (0.0226)	0.3482*** (0.0666)	0.5384*** (0.0608)	1,235
28	0.1751*** (0.0221)	0.4421*** (0.0543)	0.4671*** (0.0619)	2,151
29–31	0.1362*** (0.0323)	0.5407*** (0.0844)	0.4065*** (0.0902)	1,358
34	0.1626*** (0.0361)	0.4596*** (0.0800)	0.4894*** (0.1017)	938
36	0.1761*** (0.0217)	0.5538*** (0.0465)	0.421*** (0.0597)	2,316

Notes: Bootstrapped standard errors in parenthesis. '***', '**' and '*' denote statistical significance at the 1%, 5%, and 10% levels, respectively.


BANCO DE LA REPÚBLICA
LIBERTAD
COLOMBIA