

NBER WORKING PAPER SERIES

WHEN INDUSTRIES BECOME
MORE PRODUCTIVE, DO FIRMS?
INVESTIGATING PRODUCTIVITY DYNAMICS

James Levinsohn
Amil Petrin

Working Paper 6893
<http://www.nber.org/papers/w6893>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 1999

The views expressed here are those of the author and do not reflect those of the National Bureau of Economic Research.

© 1999 by James Levinsohn and Amil Petrin. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

When Industries Become More Productive, Do Firms?
Investigating Productivity Dynamics
James Levinsohn and Amil Petrin
NBER Working Paper No. 6893
January 1999

ABSTRACT

This paper investigates two explanations for why industries might become more productive over time. The first explanation, termed the “real productivity case,” is one in which firms become more productive and this leads to more productive industries. The second explanation, termed the “rationalization case,” is one in which firm productivity is constant, but productive firms expand while less productive firms either shrink or exit. Each case has very different implications for factor markets, long term growth prospects, and public policy regarding productivity. Further, one can only distinguish between these two cases with plant- or firm-level data. We investigate the empirical relevance of the two cases using the Chilean manufacturing census. We find that the rationalization case explains much of the measured increase in industry productivity. When industry productivity falls, the rationalization case appears much less important. We also contribute to the applied econometric literature on productivity estimation as we show that the value-added production function is especially well-suited to a simple extension of recent methods developed by Oiley and Pakes.

Jim Levinsohn
Department of Economics
University of Michigan
Ann Arbor, MI 48109
and NBER
jamesL@umich.edu

Amil Petrin
Graduate School of Business
University of Chicago
1191 E. 58th Street
Chicago, IL 60637
amil.petrin@gsb.uchicago.edu

1. Introduction

Why do some industries become more productive over time? One explanation is that firms adopt new and better methods of production. This can happen because firms learn by doing or because firms are exposed to new and better methods of production. Examples of the former range from the low-tech sewing machine operator at an apparel firm to the high-tech process of increasing yields on silicon chip production. An example of the latter is the set of manufacturing practices known as “lean production.” U.S. manufacturers, initially auto producers but later others, adopted these methods after observing Japanese success. In both the learning-by-doing and the learning-by-watching cases, firm productivity increases and with that comes increases in industry productivity. We refer to this as the real productivity case, and we view it as an uplifting explanation of the mechanism for productivity increases. After all, there are no obvious bounds on learning and ingenuity. From a more analytic viewpoint, modeling the real productivity case is relatively straightforward, in part because this is a story whose essence can be captured with a representative firm model.

There is another explanation for increased industry productivity, and it is a bit less cheery and a bit more cutthroat. This is the idea that in open markets, some firms thrive while others disappear. There are, within an industry, winners and losers. As firms that are especially well suited to an industry expand and misfits contract or exit, industry productivity increases. Conversely, industry productivity is hindered when firms are sheltered from the harsh realities of the marketplace. An example of this is an explanation given by the *Economist*¹ for Japan’s recent and long economic downturn. The *Economist*, echoing arguments made by many others, suggested that Japan’s poor economic performance was, at least in part, because of a Japanese aversion to “outright failure” of firms. Rather, there exists a corporate culture more willing to tolerate low returns. In today’s more global economy, this culture is more difficult to maintain and has contributed, the *Economist* claimed, to the recent poor performance of the Japanese economy while other industrial economies continue to boom. The notion that inefficient firms can perhaps survive in a protected economy but will perish in a more competitive environment is not a new idea. Rather, it is an immediate implication of Ricardo’s 1817 theory of comparative advantage and specialization. This process of industry rationalization might be expected to lead to increased industry productivity, and we refer to this story as the rationalization case. From an analytic viewpoint, the rationalization case

We would like to thank seminar participants at Stanford University, Harvard University, UC-San Diego, University of North Carolina, and NBER for helpful suggestions. Pat Bajari, Susanto Basu, Roger Gordon, Elhanan Helpman, Peter Klenow, Ariel Pakes, and Valery Ramey also provided especially helpful suggestions. Marc Melitz and Wendy Petropoulis provided splendid research assistance. We are grateful to the Russell Sage Foundation for support.

¹ See the June 20, 1998 issue containing the article “Japan’s Economic Plight.”

requires explicitly modeling firm heterogeneity, since the representative firm framework cannot adequately capture the evolutionary process in which some firms thrive while others lag.

The first goal of the paper is to determine, in one specific instance, the relative applicability of the real productivity and rationalization cases. The degree to which these non-mutually exclusive explanations characterize productivity growth matters for at least three reasons. First, understanding what underlies changes in industry-level productivity is essential to appropriate productivity-related policymaking. Policies which might either promote or hinder productivity in the real productivity case often have different impacts in the rationalization case. For example, a policy that provided infant industry subsidies might enhance productivity in the real productivity case, but if new and smaller firms are the least productive and most likely to fail, such subsidies might harm industry productivity in the rationalization case. Further, as productivity growth is an integral component to economic growth, simply understanding from whence it derives seems important.

Second, the real productivity case and the rationalization case have very different implications for factor markets. If, for example, productivity increases derive primarily from the real productivity case, worker displacement is a non-issue. In its pure form, this case has all firms becoming more productive. The rationalization case, on the other hand, would entail substantial worker displacement. With these differences in factor market implications come differences in the politics of productivity change.

Third, the real productivity case and the rationalization case have different implications for long run growth. In the real productivity case, industries become more productive because firms become more productive. This process is not bounded in any obvious way; a good idea can follow good idea after good idea. Ignoring entry, in the rationalization case, an industry cannot become more productive than the single most productive incumbent firm. That is, the "frontier" is reached when all output is manufactured by the most productive firm. Thereafter productivity is constant. While each of these explanations is likely to play a role in explaining changes in productivity growth, measuring the empirical importance of the two cases may help us better understand the long run growth prospects of an industry.

Economic models in which productivity growth matters are more informative if they are broadly consistent with the data. Ex ante, it is not clear which story, if either, is most consistent with industry data. Therefore, it is not generally clear what kind of model should be brought to industry data. Many models of firm behavior are models with homogeneous firms or, equivalently, a single representative firm. If the real productivity case is empirically most relevant, homogeneous firm models may be appropriate. On the other hand, if changes in productivity originate mostly from industry rationalization and the shuffling of output from less to more productive firms, a model with heterogeneous firms is required. In order to determine the relative applicability of either case, an economic model consistent with *both* homogeneous and heterogeneous firms must be taken to the data; only then is it possible to evaluate the extent to which either hypothesis has support in the data.

The second goal of this paper is to contribute to the literature on the estimation of productivity. At the industry-level, the real productivity and rationalization cases are potentially observationally equivalent. For this reason, we turn to firm-level data. However, empirically distinguishing between the two cases requires more than just firm-level data. It requires estimates of firm-level productivity and its dynamics that do not suffer from the standard simultaneity problem induced by the contemporaneous correlation between input levels and the productivity shock (see Marschak and Andrews (1944)). In this paper, we extend Olley and Pakes (1996) in ways which alleviate this problem of contemporaneous correlation, and we show that these extensions can be important for inferences drawn from estimates of productivity. In addition, we suggest that the Olley-Pakes idea of using investment to control for the simultaneity between inputs and the productivity shock

can be extended to using intermediate inputs like electricity or fuels. Using intermediate inputs to proxy for productivity shocks has two advantages. It provides a simpler link between the estimation strategy and the economic theory, and, in industries which have pronounced adjustment costs of capital, it avoids the potential truncation from estimation of large numbers firms which report zero investment (because they adjust on an infrequent basis.)

We also show that it can be empirically important to avoid using ordinary least squares (OLS) or fixed effects estimators when estimating firm-level productivity. We find that OLS and Fixed Effects lead to parameter estimates which are significantly different both in economic magnitude and in a statistical sense from estimates obtained under less restrictive assumptions. In addition, we find that it is common for the OLS framework to overestimate productivity changes by 10% of industry value-added. We do note, however, that OLS *does* consistently predict the direction of productivity movements. This result suggests that one might avoid the additional complications associated with estimating parameters of the more flexible framework if the objective is simply to sign industry-level productivity changes.

Finally, we highlight when the methods of Olley and Pakes are applicable and when they are not. In so doing, we illustrate how several commonly used approaches to estimating productivity can be viewed under a single unifying framework.

The third goal of this paper is markedly more narrow. This paper provides evidence on what happened to the distributions of firm-level productivity in Chile. As such, we contribute to a literature on firm-level productivity in developing countries – some of which is reviewed in the next section. Much of the existing literature, some of it using the same Chilean data we use, has focused on trying to investigate the causal link between trade liberalization and productivity. We do not attempt to address this link in this paper. Rather, we focus on documenting the changes in distributions of productivity without attempting to attribute these changes to correlates in the economic environment following a large liberalization. We avoid the issue of whether the liberalization that occurred in Chile *per se* was responsible for the changes in productivity for two reasons. First, as documented in Levinsohn (1998), there are good reasons to believe that the impact of trade policy might have been dominated by the overwhelming macroeconomic influences in Chile over the period of the data. Second, the data span only the period following the liberalization. With data before and after the liberalization, researchers stand a fighting chance of identifying the impact of the liberalization on productivity. Without it, though, we feel the exercise is not likely to be persuasive.

We proceed by briefly and selectively reviewing some of the relevant literatures in the next section. In the third section, we turn our attention to the details of estimating firm-level productivity. There, we review some previous approaches to estimating productivity, provide a general framework for estimation, and finally explore some specific structural models and their resulting estimators. Section 4 provides some background on the Chilean experience and then describes the data. Section 5 gives the details of the estimation procedure and presents results. In section 6, we investigate the dynamics of productivity, while section 7 concludes.

2. The Literatures

Firm-level productivity dynamics is a broad topic, and an adequate review of that topic is well beyond the scope of this section. Rather, we selectively focus on a handful of studies most closely related to our work. Much of the work in this area is either mostly or purely empirical. We begin our review with this literature, and then review some of the relevant theoretic literature.

Empirical Studies of Firm-Level Productivity Dynamics

A natural question is “What about productivity dynamics in the United States?” Using U.S. data and empirical strategies quite different from ours, this question was tackled by Bailey, Hulten, and Campbell (1992), and they concluded that:

Some industries in our sample have achieved huge improvements in productivity; in others productivity has fallen sharply. There are high-productivity entrants and low-productivity exiters, plants that move up rapidly in the productivity distribution and plants that move down rapidly. Many plants stay put in the distribution. Both in level of and rate of change in productivity, plants manifest significant differences.²

Most relevant to our study, Bailey et. al. investigate what was underlying changes in industry-level productivity, and they found that the relative empirical importance of the rationalization versus real productivity varies over the business cycle. Overall, their results suggest that the real productivity effect is quantitatively more important than the rationalization effect—the former is frequently about three or four times as large as the latter. A similar analysis of plant-level productivity was conducted using Israeli data by Griliches and Regev (1995).

There are also several studies of plant-level productivity resulting from an influential research project organized by Mark Roberts and James Tybout. That project is summarized in Roberts and Tybout (1996). That volume contains several papers that are closely related to ours in terms of issues, albeit quite different in terms of methodologies. The paper most closely related to ours is Lui and Tybout (1996). In that paper, the authors examined plant-level productivity for firms in Columbia and Chile. Their results using the Chilean data reported, among other things, that: i) exiting plants are usually, but not always, less efficient than incumbents and this gap closes as inefficient but large firms are “propped up” and remain in the market; ii) there is a role for the reallocation of productivity but that the quantitative impact of reallocating productivity across incumbent firms is “modest” in the Chilean context.

The role firm turnover plays in changes in industry-level productivity is investigated using Taiwanese data in Aw, Chen, and Roberts (1998). They find that “the productivity differential between entering and exiting firms is an important source of industry level productivity growth in Taiwanese manufacturing.” In their study, productivity is computed using index number methods. It is not estimated econometrically.

There are also several studies of how international trade impacts productivity. Some of this work is predicated on an economic environment in which firms are imperfectly competitive. See, for example, Harrison (1994). Our results do not speak to this literature for two reasons. First, our framework is one in which firms take prices as given. Second, we do not attempt to identify whether the changes in productivity that we document are attributable to changes in international competition (via trade policy or exchange rate shifts) or to changes in other macroeconomic variables. For an example of work that does exactly this, using the same data used in this paper, see Pavcnik (1997).

A paper closely related to ours is Olley and Pakes (1996). They examine productivity dynamics in the U.S. telecommunications equipment industry and, in the process, develop the methodology on which we build. (We discuss the methodological aspects of this work below and concentrate here on their results.) Olley and Pakes investigate the *contemporaneous* covariance between output and productivity. As such, they examine something akin to our rationalization case. While the discussion is about more productive firms *becoming* bigger, the empirical observation is one about

² In an exceptionally entertaining (by the standards of economics, anyway) discussion of Bailey et al. (1992), Tim Bresnahan noted that the authors “did report all 6000 numbers.”

more productive firms *being* bigger. They interpret their findings as supporting the empirical importance of the rationalization case to explaining changes in productivity in the telecommunications equipment industry.

Theoretical Models of Firm-Level Productivity Dynamics

Models which can capture stories in which some firms thrive, others struggle, and entry and exit simultaneously occur are related to our work. These models, like our rationalization case, clearly require firm heterogeneity. Examples of models that explicitly incorporate the sort of dynamics that are consistent with our rationalization case include those by Jovanovic (1982), Hopenhayn and Rogerson (1993), and Hopenhayn (1992). Each of these are models in which firms are price takers for both inputs and output. They are briefly discussed in turn.

Jovanovic provides a theory of industry evolution which is based on firm heterogeneity and self-selection. Each firm has some true underlying cost of production, c . This parameter is a draw from the population distribution of costs, which is assumed to be $N(\bar{c}, \sigma_c^2)$. The cost distribution is known to the firm, but the firm does not know its own cost parameter. Each period the firm's unit cost of production, c_{it} , fluctuates around c_i randomly according to

$$c_{it} = c_i + \eta_{it},$$

where η_{it} is $N(0, \sigma_\eta^2)$. The firm observes c_{it} each period, and since η_{it} is i.i.d. with finite variance, the firm learns more and more about its underlying cost using $(1/T \sum_{t=1}^T c_{it})$ to consistently estimate its cost. At the start of every period a firm decides whether to exit or stay based upon its current cost information. With a little luck, if the firm is a low cost firm, it will see high profit realizations after entering and continue to produce as it becomes more and more clear to the firm that it is profitable. If the firm is a high cost firm, then it may not wait long before exiting the industry. The evolution of the economy is then driven by these optimizing agents' learning and selection decisions.³

Jovanovic's model is not easily taken directly to the data. His model is based on costs, whereas our empirical work is focused on productivity. Viewing productivity as the dual of costs, though, would allow us to work with Jovanovic's model. More problematic is the notion that in Jovanovic's model, a firm's decision is based on its entire history of productivity draws. A more tractable model would restrict this dependence to a period shorter than the length of the observed data series.

Hopenhayn (1992) proposes a somewhat different model of industry evolution. Like Jovanovic, Hopenhayn's model has both firm-level heterogeneity and price-taking behavior in a dynamic competitive framework. In this model, firms are subject to a random productivity shock every period. This productivity shock follows a first-order Markov process that is independent across firms.⁴ In addition, the distribution of future productivity is (stochastically) increasing in this period's productivity. Hence, high productivity firms expect to remain high productivity firms. Surviving firms pay a fixed cost each period, then observe their productivity shock, and finally decide on a level of output for that period. Entrants have to pay a sunk cost prior to entry, after which they draw from a common underlying distribution of productivity shocks and then choose output. Exiting firms earn zero profits but do not have to pay the fixed cost. From this basic framework, Hopenhayn derives equilibrium conditions that imply predictions about the productivities of entrants, incumbents, and exiters. He shows, for example, that in equilibrium low productivity firms are more likely to exit and that older cohorts of firms will be more productive than younger cohorts.

³ Note that the probability that a low cost firm will exit is non-zero because it may get a series of large cost realizations (large η 's), and will think it is really a high cost firm.

⁴ Restricting the dependence to the previous period yields more degrees of freedom when estimation takes place.

Hopenhayn and Rogerson (1993) propose a variant of this model of industry evolution. They use it to evaluate the aggregate implications of government policies which make it costly for firms to adjust their labor. To do so, they develop an equilibrium model of the reallocation process of labor across firms using a value function which explicitly includes an adjustment cost for labor. The equilibrium notion employed by Hopenhayn and Rogerson implies that the state of the economy is characterized by the distribution of state variables for operating firms. At time t , this probability distribution is given by $P^t(\omega, k)$, where ω denotes productivity, and k is a state variable such as capital and/or labor.⁵ The transition operator that takes the economy's distribution of state variables from time t to time $t+1$ is given by the operator T , so $P^{t+1}(\omega, k) = TP^t(\omega, k)$. Hopenhayn and Rogerson prove the existence of an equilibrium in their economy that has entry, exit, and the growth and decline of firms over time. An empirical implication of their model is that the distribution of state variables remains fixed over time in equilibrium.

Two other popular frameworks that can be used to model dynamic patterns of firm-level productivity are Ericson and Pakes (1995) and the X-inefficiency models. Ericson and Pakes (1995) is a dynamic model of a small, imperfectly competitive industry with a stochastic process of accumulation for the state variable. X-inefficiency models start from the primitive that there is an efficient production technology, and firms are either on it or below it, but never above it. In our estimation we do not account for the features of the Ericson and Pakes model, nor do we estimate a production frontier (we estimate the average production technology,) so these models are less tightly related to our own work than are those previously discussed.

Models that might capture the true productivity case need not place such emphasis on firm heterogeneity. If all firms become more productive over time, a representative firm framework may suffice. There are several models in which firm productivity increases. Some of the traditional approaches to growth theory provide examples. Increases in productivity (and hence growth) may result from simple learning by doing (e.g. Arrow (1962) and more recently Romer (1986)), from investment in research and development (e.g. the models outlined in Griliches (1998)), or via the conduit of international trade (e.g. Grossman and Helpman (1991)).

3. Estimation of Productivity

The empirical work in this paper focuses first on estimating productivity and then on decomposing the changes in productivity. In this section, we describe how the estimation is done. We first review several approaches to estimating firm-level productivity as this provides a context to our estimating strategy. We then describe our general estimation strategy. Finally, we relate specific assumptions on firm-level behavior to the implementation of the estimation strategy. Those wishing to skip the review can proceed directly to the next sub-section without a loss in continuity.

Previous Approaches to Estimating Productivity with Firm-Level Data

In this sub-section we describe some previous estimators for the distribution of firm-level productivity in a competitive environment. Our goal is to avoid making assumptions that rule out, *a priori*, important forms of firm-level heterogeneity across establishments and over time. We do so for both economic and econometric reasons. First, one of the most striking features of recent studies using detailed firm-level data is the amount of observed heterogeneity in outcomes across firms. Summary statistics of output conditional on inputs contain significant variability across

⁵ In their model, labor is a state variable. (Capital is not.) We can alternatively write the distribution as $P^t(\omega, k, l^s)$ if we find that skilled labor is a state variable for some industries.

firms, as do both growth rates and entry and exit probabilities.⁶ It is important, then, that the estimation strategy yield estimates that allow the researcher to investigate the *economic* implications of such heterogeneity.⁷ Second, we know that *econometric* properties (in particular consistency) of our estimator of the distribution function for productivity are likely to be lost under assumptions ruling out this observed, dynamic, firm-level heterogeneity. For example, the ordinary least squares and fixed effects estimators place restrictions on firm-level dynamics that are too restrictive for our data, and we show later that estimates of firm-level productivity from these models can be rejected in favor of models that allow firm-level productivity to be both serially and contemporaneously correlated with inputs.

Our paper is in the vein of recent work which allows for more systematic differences across firms at any given time and for differences in firm outcome paths over time (see, e.g., Ericson and Pakes (1995), Olley and Pakes (1996), Roberts and Tybout (1997)). We start (as do these authors) by focusing on the firm. In a discrete time model, the total discounted profits for firm i can be written as

$$\Pi_1(y_{i1}, x_{i1}, \epsilon_{i1}; \theta) + E\left[\sum_{t=2}^{\infty} \beta^{t-1} \Pi_t(y_{it}, x_{it}, \epsilon_{it}; \theta)\right], \quad (1)$$

where $\Pi_t(\cdot)$ is the profit function for period t , and β is the discount rate. “Sub”-functions of the profit function, such as the firm production or value-added function, are the typical basis for estimation. In these cases y_{it} is a measure of output and x_{it} is a vector of inputs for period t . x_{it} includes both inputs that are easily adjusted (e.g. materials or labor, in many cases), and inputs whose stock evolves over time in response to future beliefs (e.g. investment to build capital.) A sequence of errors, $\{\epsilon_{it}\}_{t=1}^{\infty}$, indexed both by firm and time, is introduced into these models to account for differences between the model’s predictions and observed outcomes. The standard interpretation assigned to this sequence of residuals is that it represents (or, more precisely in our case, contains) a term characterizing firm productivity. We then use the empirical distribution function of these residuals to characterize what happens to productivity over time. More specifically, we estimate a value added production function of the form

$$y_{it} = f(x_{it}, \epsilon_{it}; \theta), \quad (2)$$

where θ is a vector of parameter values and ϵ_{it} measures the difference between expected and observed output (conditional on the inputs we observe).

Our primary concern is with the potential contemporaneous correlation both within firm i and across time t between ϵ_{it} and x_{it} in the firm-specific sequences $\{x_{it}, \epsilon_{it}\}_{t=1}^{\infty}$. Given the natural dependence of the firm’s discounted future profits on both x_{it} and ϵ_{it} in (1), this correlation seems plausible, and we are not the first to point out this potential problem. Marschak and Andrews (1944) argue that the problem arises when variable input demands are correlated with productivity and productivity is not observed. They suggest that this simultaneity can be particularly acute for inputs that adjust most rapidly to the productivity realization. There is also reason to believe that firms with better sequences of productivity realizations will, over time, respond to these “good” realizations by investing and accumulating assets that are costly to adjust rapidly. Alternatively, simultaneity may occur in the time dimension when input decisions are based on serially correlated errors and there are costs to making large immediate adjustments to inputs.

⁶ Recent well-known papers with these observations include Dunne, Roberts, and Samuelson (1988), Davis and Haltiwanger (1990), Bailey et al. (1992), and the literature cited therein.

⁷ In particular, we wish to investigate the roles of real productivity and rationalization in industry-level productivity changes.

The two most commonly used methods for estimating firm-level productivity, OLS and fixed effects, impose strict restrictions on the sequences $\{x_{it}, \epsilon_{it}\}_{t=1}^{\infty}$. OLS assumes that productivity is uncorrelated with input choices across both firms and time. The fixed effects estimator is similarly restrictive. It assumes that firm-level productivity is constant over time, so for any firm i ,

$$\epsilon_{i,t} = \epsilon_{i,t'} = k_i,$$

for any two time periods (t, t') and some constant k_i . The lack of dynamics in both frameworks is hard to reconcile with observations from firm-level data sets. In the data we see apparently similar firms taking very different actions. Some firms increase their use of labor and capital over time, others decrease their use over time, and some firms enter and others exit the market. All of these decisions must be made in accordance with firm beliefs about present and future profits, and these are affected by present and future expected productivity realizations. Hence, the variety of patterns across apparently similar firms suggests that the unobserved productivity term follows a more general stochastic process than that specified by OLS or fixed effects.

An alternative approach is to model the sequence $\{x_{it}, \epsilon_{it}\}_{t=1}^{\infty}$ as history dependent. Here, the previous period's productivity can affect future productivity levels, and hence affect expectations about future productivity levels. Instrumental variables is the standard approach to solving the problem of correlation between inputs and the error term. However, it is difficult to find instruments at the firm-level for capital, skilled and unskilled labor; most variables correlated with inputs will likely be correlated with the productivity shock, and this will frustrate efforts to obtain consistent estimates of the production function parameters. One oft-used solution to the unavailability of appropriate instruments is to adopt a fixed effects estimator, but this brings with it the problems already discussed.

It is possible to directly specify the parametric process that the productivity shock follows. However, even if we are willing to characterize the dynamic sequence $\{x_{it}, \epsilon_{it}\}_{t=1}^{\infty}$ as a parametric process and want only to estimate the parameters of this process, we still have a significant problem. By itself, knowledge of the process (up to the parameters) is not enough to control for the simultaneity between ϵ_{it} and x_{it} over time because the process $\{x_{it}, \epsilon_{it}\}_{t=1}^{\infty}$ follows a path that depends upon its starting values (x_{i1}, ϵ_{i1}) . This is an initial conditions problem (see Heckman (1981) and Pakes (1996)), where estimation of parameters for a stochastic process that depends upon time-ordered outcomes is impossible unless the process is "initialized."

One solution is to initialize the observed process by assuming the history is exogenous, i.e. that $\{x_{it}, \epsilon_{it}\}_{t=1}^{T-1}$ is independent of $\{x_{it}, \epsilon_{it}\}_{t=T}^{\infty}$, where T is the first date a firm is observed. Another solution is to estimate a fixed entry-time effect for each firm. However, this solution produces consistent estimates only as the ratio (observations/firm) gets large. This estimator is not consistent as just the number of firms increases because each extra firm brings with it an extra fixed effect parameter to estimate. A third solution splits the sample into two parts, the first part of which is used to estimate starting values. Roberts and Tybout (1997) take this approach with a panel from Colombia that is 10 years in length. Using the first 3 years they estimate starting values for continuing firms, and then initialize their assumed stochastic process accordingly. The second half of the data set is then used to look for determinants of firms' decisions to enter the export market in Columbia.

Our Approach

Our approach uses a recently proposed idea by Olley and Pakes (1996). (See also Pakes (1996).) They suggest including in the estimation equation a proxy for the productivity shock to control for the part of the error correlated with inputs. Using the dynamic program of their firm, Olley and Pakes show that under certain conditions investment can be used as a proxy for (i.e. can be used

to condition on) the productivity shock. If the distribution of next period's productivity shock is stochastically increasing in this period's productivity shock, the economic story that makes investment a valid proxy is straightforward; a firm that realizes a large productivity shock this period will invest more than an identical firm with a smaller productivity shock this period because the more productive firm anticipates doing better than the less productive firm both in the current period and in future periods.

In the next section we show how a simple insight about the value-added production function can be used to amend the recently developed techniques of Olley and Pakes. Like the Olley and Pakes estimator, our new estimator permits heterogeneity in productivity across firms and over time. A possible advantage of our estimator is that it does not require the complicated derivation needed to show that investment is a valid proxy for productivity (see Pakes (1996)). Another advantage for our data set is that investment is frequently observed to be zero. This is true for one-third of our firm/year observations. No proxy is then available for these firm/year observations, making it necessary to drop them from the estimation routine.⁸ Our proxies are generally available for almost all of the firm/year observations, obviating the need to estimate the value-added production function with only investing firms (and avoiding a potential selection problem.)⁹ We next discuss our ideas for proxies and the estimators they generate.

A General Framework

Our empirical goal is to estimate productivity at the firm-level without imposing unreasonable stochastic properties on the sequence of inputs and errors $\{x_{it}, \epsilon_{it}\}_{t=1}^{\infty}$. We motivate the restrictions that we will place on this sequence by using the value function of the firm. This approach has at least two advantages. First, it allows us to determine the conditions under which the observed sequence of input decisions and errors, $\{x_{it}, \epsilon_{it}\}_{t=1}^{\infty}$, follows a path that is consistent with different dynamic equilibrium notions.¹⁰ Second, by specifying our behavioral assumptions, we are able to link our econometric methods with an underlying behavioral framework. This, in turn, allows us to evaluate the importance of alternative modeling assumptions and ensuing implications. It also permits us to extend our results to some available theoretical models of firm dynamics and investigate which of these stories seem compatible with different Chilean industries.

We begin our discussion with the state variables of the firm in our framework. These are capital k_t and productivity ω_t . (Notationally, ω is a component of ϵ , a distinction which is discussed below.) We assume k_t evolves in a deterministic manner according to

$$k_t = (1 - \delta)k_{t-1} + i_t \quad (3)$$

where i_t is investment and δ denotes the depreciation rate. Investment must enter directly into this period's production technology for this condition to hold true. Productivity is assumed to follow a first-order Markov process, so knowledge of this period's realization generates a distribution known to the firm for the possible values of next period's productivity realization. We denote this Markov transition matrix as $P(\cdot|\omega)$.

The value function $v_t(\cdot)$ (or Bellman) for a firm in our model is given by:

$$v_t(k_{t-1}, \omega_{t-1}) = \max_i \int_{\omega'} [\Pi_t(k_t(i), \omega') - c_t(i) + \beta v_{t+1}(k_t(i), \omega')] dP(\omega'|\omega_{t-1}), \quad (4)$$

⁸ See Olley and Pakes (1996) for a detailed explanation of this problem. Briefly, it arises because productivity is not a smooth function of investment at zero investment levels (i.e. it cannot be inverted at zero investment, but this invertibility is necessary for the proxy to be valid.)

⁹ Olley and Pakes (1996) are forced to select on positive investment and this truncates their sample size by about 8 percent.

¹⁰ See, for example, Hopenhayn (1992) or Bailey et al. (1992).

where

$$\Pi_t(k_t(i), \omega) \equiv \max_{l^s, l^u, m} [p_t q(l^s, l^u, m; k_t(i), \omega) - C_t(l^s, l^u, m)],$$

t indexes periods (which are years in this data set), and for ease of exposition we suppress the firm-level notation. Time t indexes state variables that are taken as given by the firm, including input and output prices (p_t) and other macroeconomic variables. the (output) production function is given by $q(\cdot)$, and $c_t(i)$ and $C_t(\cdot)$ are the costs associated with investment and variable input choices respectively.

The firm decision-making process is summarized as follows. Each firm begins each period t by choosing a level of investment, and after doing so each firm observes its productivity. Given k_t and ω_t , and faced with output price p_t , the firm then chooses the levels of the variable factors (l_t^s, l_t^u, m_t) to maximize $p_t q(\cdot)$. These variable factors are skilled and unskilled labor, l^s and l^u , and intermediate inputs, m .

Our empirical strategy will focus on the firms' value-added production functions. Because we are working with a value-added production function, we will be estimating the relationship between labor and capital and an output number that is adjusted for the cost of the intermediate inputs. Our value-added function has as arguments skilled and unskilled labor, capital, and the residual, and we write this function as $y \equiv f(l^s, l^u, k, \epsilon)$. A first-order approximation to this function is

$$y_t = \beta_0 + \beta_k k_t + \beta_s l_t^s + \beta_u l_t^u + \omega_t + \mu_t, \quad (5)$$

where y_t is the log of output (measured as value-added) in year t , k_t is the log of the plant's capital stock, l_t^s is the log of skilled labor input, l_t^u is the log of the unskilled labor input, and the error ϵ_t from (2) is assumed to be additive in two other unobservables, ω_t and μ_t , so $\epsilon_t = \omega_t + \mu_t$. ω_t is the plant's productivity, and μ_t is a mean zero error that may either be measurement error or a shock to productivity that is unexpected and to which labor and other variable inputs do not respond. The key difference between ω_t and μ_t is that the former is a state variable, and hence impacts the firm's decision rules, while the latter has no impact on the firm's decisions. Again, our exercise is defined by our desire to leave the stochastic properties of ω_t reasonably general when we estimate

$$(\beta_t^s, \beta_t^u, \beta_k, P(\omega)),$$

where $P(\omega)$ is the distribution of productivity both across firms and over time.

Using (5), the endogeneity of inputs problem is readily illustrated. If the labor inputs chosen at time t respond to observed productivity ω_t , then the variable input choices in year t will be positively correlated with ω_t , leading to upwardly biased estimates of the elasticity of output with respect to labor. The capital coefficient may suffer from the same problem; capital is variable over time via depreciation and changes in investment, and the maximization problem from (4) solved with serial correlation in $\{\omega_t\}_{t=1}^{\infty}$ may lead to a sequence $\{k_t\}_{t=1}^{\infty}$ that also exhibits serial correlation. To make matters more difficult, capital and labor levels are highly correlated both within and across firms. Econometrically, this means that a positive bias in one coefficient can transmit a negative bias to the other coefficient since they are estimated simultaneously. The difficulty of signing these biases then makes it hard to sign bias associated with the distribution of productivity.¹¹

As noted above, a major innovation in Olley and Pakes (1996) is their idea of using a proxy to solve this simultaneity problem. They work out the conditions under which investment is an increasing function of productivity in a Bellman equation similar to (4). They assume that period

¹¹ For example, if we knew that the labor coefficient was biased up and the capital coefficient was consistently estimated, we would also know that the mean productivity level is biased down.

t 's investment responds to period t 's productivity shock, and then show that an optimizing firm will always respond to a large ω_t by investing large amounts i_t . Writing investment as a function of state variables yields

$$i_t = i_t(\omega_t, k_t).$$

For firms investing positive amounts in period t , investment is shown to be strictly increasing in productivity, and this means an inverse function exists. Hence, they write $\omega_t = (i_t, k_t)$, and this function is strictly increasing in i_t (again, for positive levels of investment.) This function (which is not known, but which can be estimated) can then be included as a proxy when estimating the production function from (5). It serves to control for the correlation between the capital and labor sequences $\{x_t\}_{t=1}^{\infty}$ and $\{\omega_t\}_{t=1}^{\infty}$, the sequence of productivity.

We wish to address the problems associated with traditional estimators (i.e. the OLS and fixed effect estimators). However, we also wish to avoid two possibly significant costs of using firm-level investment as a proxy for production. The first cost is one of theoretical complexity while the second one is data-driven. Showing that investment is a valid proxy in the context of the dynamic structural model is non-trivial (see Pakes (1996).) If one wishes to use a model that differs, even slightly, from that of Olley and Pakes, it becomes necessary to re-investigate the appropriateness of investment as a proxy for productivity. The proxy that we adopt does not require the complicated derivation that is used to show investment is a valid proxy, primarily because we assume intermediate input levels adjust costlessly. We recognize that the value-added production function does not have intermediate goods as inputs; by definition, they have been subtracted out from gross output. This suggests that any intermediate input which responds to the productivity shock may be a potential candidate to proxy for productivity. In Appendix I, we show the conditions under which an intermediate input is a valid proxy.¹² The intuition is straightforward: the productivity shock leads to higher marginal products for capital and labor, and with constant input prices, the firm will produce more output to drive down these marginal products. In the process of increasing output, firms will increase both their use of fuels and electricity.

The second reason relates to the observables in our data. Investment is zero for one-third of our firm/year observations. This empirical regularity suggests that there are adjustment costs to investment. Were we to stick with the investment proxy, no proxy would be available for these firm/year observations. Our proxies are available for almost all of the firm/year observations, obviating the need to trim a third of our sample before proceeding with estimation (and avoiding a potential selection problem.)

Similar to the story with investment above, we can express an intermediate input m as a strictly increasing function of ω . That is:

$$m_t = m_t(\omega_t, k_t), \tag{6}$$

and we then invert (6) and express the unobservable productivity as a function of the intermediate input and capital. Hence,

$$\omega_t = h_t(m_t, k_t), \tag{7}$$

and we can proceed by using $h_t(m_t, k_t)$ (which we don't know but can estimate) as our proxy for productivity.¹³

¹² One can show that for the Cobb-Douglas production function and for certain forms of Leontief production technologies firms will have intermediate input demand functions that are strictly increasing in the productivity shock.

¹³ As we mentioned earlier, this is important in our data because many firms have zero investment, and these firms have to be trimmed from the estimation when investment is used as a proxy. In our data, almost every firm has a non-zero level electricity usage in every period. This allows us to use the full data set rather than a truncated sample. See the data section for a more complete discussion.

This inversion plays a very important role, since it permits us to control for ω_t . To see how this is done, substitute (7) into (5) to obtain:

$$y_t = \beta_s l_t^s + \beta_u l_t^u + \phi_t(m_t, k_t) + \mu_t, \quad (8)$$

where,

$$\phi_t(m_t, k_t) = \beta_0 + \beta_k k_t + h_t(m_t, k_t). \quad (9)$$

Equation (8) is a partially linear model and is estimated using semi-parametric regression methods discussed in the next section. For now, note that the error term in this equation, μ_t , is by assumption uncorrelated with the labor inputs. Therefore, if we can include $\phi_t(m_t, k_t)$ in the estimation routine, estimates of coefficients for skilled and unskilled labor will be consistent. This is the first step of the estimation process.

We are not done because we have not identified the coefficient on capital, β_k , in (5). From (9), capital's impact on output, via β_k , is not separately identified from capital's impact on intermediate input usage, since capital also enters the proxy function $h(m_t, k_t)$. In order to identify the coefficient on capital, we rely on the assumption that the expected value of next period's productivity can be written as a function of observables and productivity this period. To see this, note that we can rewrite the value of next period's output as

$$y_{t+1} = \beta_0 + \beta_s l_{t+1}^s + \beta_u l_{t+1}^u + \beta_k k_{t+1} + E[\omega_{t+1}|I_t] + \xi_{t+1} + \mu_{t+1}, \quad (10)$$

where productivity ω_{t+1} is equal to the sum of two terms, written as

$$\omega_{t+1} = E[\omega_{t+1}|I_t] + \xi_{t+1}.$$

The first term is the expected value of next period's productivity conditional on this period's information, I_t , a vector of the relevant observables and last period's productivity. The second term is the innovation in productivity. Define the function $g(I_t)$ as

$$g(I_t) = \beta_0 + E[\omega_{t+1}|I_t].$$

The function $g(I_t)$ gives, up to an additive constant, the conditional expectation of next period's productivity, ω_{t+1} . Its arguments are chosen to condition out potential correlation between capital and ω_{t+1} .¹⁴ Substituting $g(I_t)$ into (10) and taking the expected value conditional on k_{t+1} yields the moment condition:

$$\begin{aligned} E[y_{t+1} - \beta_s l_{t+1}^s - \beta_u l_{t+1}^u - \beta_k k_{t+1} - g(I_t) | k_{t+1}] &= \\ E[\xi_{t+1} + \mu_{t+1} | k_{t+1}] &= 0, \end{aligned} \quad (11)$$

which equals zero at the true parameter values (k_{t+1} is uncorrelated with both the innovation in productivity (ξ_{t+1}) and (μ_{t+1})).

It is perhaps helpful to note in less technical terms what this moment condition represents. The expectation of output less inputs equals the error, or the productivity shock plus another additive error. This error cannot be used as the basis for a moment condition that will identify β_k , since the productivity shock is not orthogonal to capital (i.e. $E[k_t \omega_t] \neq 0$.) We can solve for an error term, ($\xi_{t+1} + \mu_{t+1}$), that is uncorrelated with capital by conditioning out the expectation of ω_{t+1} . It is the inclusion of the function $g(I_t)$ which controls for this expectation and allows for identification of the capital coefficient (via the restriction from (11).) The final step of the estimation procedure then uses the production function estimates to compute the residuals. The empirical distribution function of these residuals is then used to approximate the true distribution of productivity $P(\omega)$.

¹⁴ For example, if ω_t follows a first-order Markov process, $g(I_t) = g(\omega_t) = \beta_0 + E[\omega_{t+1}|\omega_t]$.

Specific Theoretical Structures and Their Resulting Estimators

In this section we consider different primitives for our model. In particular, we focus on what is known by the firm when it makes decisions and how this sequencing affects the estimation routine. Because the timing of the firm's decisions within the periodicity of the data is inherently unobservable to the econometrician, we outline some of the links between modeling assumptions and ensuing estimation issues. We examine the primitive of the timing of the firm's exit decision, especially with respect to when the firm sees the productivity shock. We then show when it is necessary to control for a selection bias that may arise when firms observe their productivity before choosing to stay or exit.

We assume that productivity follows a first-order Markov process. This means that our function for $g(\cdot)$ (which we will estimate non-parametrically) takes the form

$$g(I_t) = \beta_0 + \int_{\omega_{t+1}} \omega_{t+1} dP(\omega_{t+1}|\omega_t) = g(\omega_t).$$

Of course, we do not observe ω_t directly. We do observe k_t and in the first stage we estimate ϕ_t . We also know from (9) that $\phi_t - \beta_k k_t = \beta_0 + \omega_t$. Given a value for β_k , we can construct an estimator of ω_t (up to the constant β_0) by writing $g(\cdot)$ as $g(\phi_t - \beta_k k_t)$. We can then use a non-parametric approximation to this function to derive a moment condition (equation (11)) which will separate out capital's contribution to output from productivity's contribution to output.¹⁵

Next we discuss a selection problem that may arise if firms observe productivity before they choose to stay or exit, and we show that the proper choice of $g(\cdot)$ can alleviate this problem. In our framework, a firm will choose to exit if the sell-off value of its capital is greater than the present and future expected profits of staying in business. Formally, a control problem generates the exit rule, and that rule can be expressed as an indicator function χ_t which is equal to one if the firm continues to operate and zero if it shuts down. Hence,

$$\chi_t = \begin{cases} 1, & \text{if } \omega_t \geq \omega_t^*(\cdot); \\ 0, & \text{otherwise.} \end{cases} \quad (12)$$

The function ω_t^* gives the value of the idiosyncratic productivity variable that makes exit optimal. This value is endogenously determined in the equilibrium and is known by the firm, but not by the econometrician.

If the difference between the value function and the sell-off value of the firm is increasing in capital, then ω_t^* is not independent of k_t , and we accordingly write $\omega_t^*(k_t)$.¹⁶ In this case, firms with larger capital stocks will tend to survive low productivity draws that cause smaller firms to exit. This self-selection that results from the equilibrium exit behavior implies that the expectation of ω_{t+1} declines in the size of the capital stock. It can lead to inconsistency in the capital coefficient

¹⁵ We can easily change our primitive to allow ω to follow a k th-order Markov process, i.e. $E[\omega_{t+1}|\omega_t, \omega_{t-1}, \dots, \omega_1] = E[\omega_{t+1}|\omega_t, \omega_{t-1}, \dots, \omega_{t-k}]$, (which obviously nests the first-order process.) With this new primitive we write $g(\cdot)$ as

$$g(I_t) = \beta_0 + \int_{\omega_{t+1}} \omega_{t+1} dP(\omega_{t+1}|\omega_t, \omega_{t-1}, \dots, \omega_{t-k}) = g(\omega_t, \omega_{t-1}, \dots, \omega_{t-k}).$$

In terms of observables, we then write $g(\cdot)$ as $g(\phi_t - \beta_k k_t, \phi_{t-1} - \beta_k k_{t-1}, \dots, \phi_{t-k} - \beta_k k_{t-k})$.

¹⁶ This would be true if, for example, significant discounting of the value of the firm's capital takes place when the firm sells it off. Larger firms would then stand to lose more money than smaller firms because they have to sell off a larger amount of capital to exit.

because the conditional expectation of the productivity shock given survival will not equal the unconditional (on survival) expectation, i.e. $E[\omega_{t+1}|\omega_t, \chi_{t+1} = 1] \neq E[\omega_{t+1}|\omega_t]$.¹⁷

The need to condition on survival when using an unbalanced panel rests on the primitive that, within year t , firms choose to exit or stay after they first observe productivity. To see this, suppose instead that at the beginning of the period, firms choose to exit or stay before observing productivity. Hence, firms choose to exit in year t based on the same information the econometrician observes in the year $t - 1$, so

$$E[\omega_{t+1}|\omega_t, \chi_{t+1} = 1] = E[\omega_{t+1}|\omega_t].$$

We can then proceed with estimation as in (11) above, using just $g(\phi_t - \beta_k k_t)$, as there is no need to condition on an index for survival.

We do control for this potential selection problem. We use the methodology developed by Olley and Pakes. Define $g(\cdot)$ as

$$g(I_t) = \beta_0 + \int_{\omega_{t+1}^*} \omega_{t+1} \frac{F(d\omega_{t+1}|\omega_t)}{\int_{\omega_{t+1}^*} F(d\omega_{t+1}|\omega_t)} = g(\omega_{t+1}^*, \omega_t).$$

This function gives, up to an additive constant, the conditional expectation from the model of next period's productivity, ω_{t+1} , given last period's ω_t and firm survival. Therefore, to obtain consistent estimates of the capital coefficient, we must now condition on both ω_t and ω_{t+1}^* .

The probability that a firm survives to the next period is given in the model by:

$$\begin{aligned} Pr\{\chi_{t+1} = 1|\omega_t\} &= Pr\{\omega_{t+1} \geq \omega_{t+1}^*(k_t)|\omega_t\} \\ &= P_t(\omega_{t+1}^*(k_t), \omega_t) \\ &= P(k_t, h_t(k_t, m_t)) \\ &= P(k_t, m_t) \end{aligned} \tag{13}$$

where the third line makes use of the proxy for ω given by (7). Since we can write this distribution function completely in terms of the observables k_t and m_t , we can estimate the probability of survival using a standard probit (we discuss this in the estimation section.) Assuming that we can invert ω_{t+1}^* from line 2 of (13), we can write $\omega_{t+1}^*(P_t, \omega_t)$, i.e. we have ω_{t+1}^* as a function of P_t and ω_t , two objects we can estimate from our observables. Substituting P_t and $\phi_t - \beta_k k_t$ into $g(\cdot)$, we can rewrite (11) as a moment condition in observables and estimates of P_t and ϕ_t . This yields

$$E[y_{t+1} - \beta_s l_{t+1}^s - \beta_u l_{t+1}^u - \beta_k k_{t+1} - g(P_t, \phi - \beta_k k_t) | k_{t+1}, \chi_{t+1} = 1] = 0, \tag{14}$$

which allows us to obtain consistent estimates of β_k .¹⁸

¹⁷ Much of the focus on selection has highlighted the importance of using unbalanced instead of balanced panel data. While it is true that going from the balanced to the unbalanced panel will in general alleviate selection bias, it does not completely address the issue. In this paper, we use an unbalanced sample, but even with the unbalanced sample, some firms will exist at the beginning of the year and will not be in the sample at the end of the year. If the estimated production function is to account for these firms, it will still be necessary to correct for the selection bias.

¹⁸ Empirically, if we estimate the model conditioning on survival and find no statistical difference between conditioning and not conditioning on survival, three possible reasons are: 1) firms do not observe this period's productivity before they choose to stay or exit, 2) the sell-off value of the firm is not increasing in capital, or 3) the magnitude of the bias induced by selection is very small.

4. Country Background and Data

The methodology outlined in the previous section requires very detailed firm-level data which has not been censored for entry and exit and which has a reasonable time-series dimension. The Chilean data set used in this study meets those requirements. These data have been used elsewhere and we refer the interested reader to those papers for a more detailed description of the data.¹⁹ The first part of this section gives some background on the Chilean economy and the data set. The second part of this section delves deeper into details of the Chilean data set and how those details interact with the econometrics.

Country Background

Our data set spans the period from 1979 to 1986. The years prior to the start of our data set, as well as those covered by the data, were tumultuous ones for Chile. From 1960 to 1972, real per capita GDP increased from \$2897 to \$3857.²⁰ During this span, Chile pursued a policy regime of import-substitution that often protected inefficient domestic firms from international competitive pressures.

In 1973, Pinochet came to power and the economic environment changed dramatically. Fiscal austerity and a more outward-oriented set of economic policies was introduced. This liberalization included privatization of firms previously held by the government, relaxation of some 3000 government controlled prices, liberalized financial markets, a more market-driven labor market, removal of quantitative trade restrictions and a drastic reduction in tariff levels. These liberalizing measures were accompanied by contractionary macroeconomic policies which, when coupled with a decline in copper prices and the oil shocks, led to a severe recession that lasted through 1975. A recovery from 1976 to 1981 followed, and the broad picture is one in which liberalization was drastic and mostly complete by 1979, the first year of the data set. The Latin American debt crisis led to another recession in 1982-83 during which industrial output and employment fell. Industrial output rose again in 1984, stalled in 1985, and then continued to rise throughout the decade. These macroeconomic cycles are apparent in the first column of Table 1 where real GDP is reported for 1979-86.²¹

The Manufacturing Census and the Construction of Value Added

The data set is comprised of plant-level data of 6665 plants in Chile from 1979 to 1986. The data are a manufacturing census covering all plants with at least ten employees. The data were originally provided by Chile's Instituto Nacional de Estadística (INE). A very detailed description of how the eight longitudinal samples were combined into a panel is found in Lui (1991). The structure of the data set is an unbalanced panel. There is information tracking plants over time and the data set includes plants that enter over the course of the sample period (births) as well as plants that exit (deaths.) Due to the way that the data are reported, we treat plants as firms, although there are certainly multi-plant firms in the sample.²²

¹⁹ See, for example, Lui (1993), Lui and Tybout (1996), Tybout, de Melo, and Corbo (1991), and Levinsohn (1998).

²⁰ These data and related figures are available on-line in the Penn World Tables, Mark 5.6 at <http://www.nber.org>.

²¹ In the mid-1980's, Chile's credit markets were such that smaller firms were alleged to have found credit difficult to obtain. While this is difficult to document in our data set, it implies that entry and exit patterns may have been driven by more than the relatively straightforward operate versus shutdown decision in our model. Hence, credit market issues suggest that firm size may have mattered in a way that our model does not accommodate.

²² We will not capture the extent to which multi-plant firms experience scale or scope economies due to their multi-plant nature. Neither are we able to investigate whether "entry" is a new firm, a new plant from an existing firm, or simply diversification of an existing plant or firm as discussed in Dunne et al. (1988).

In an attempt to keep the analysis manageable, we focus on eight of the largest industries (excluding petroleum and refining.) We work with industries at the 3-digit level. The industries (along with their ISIC codes) are Metals (381), Textiles (321), Food Products (311), Beverages (313), Other Chemicals (352), Wood Products (331), and Printing and Publishing (342), and Apparel (322). The data are observed annually and they include a measure of output, a measure of labor and capital inputs, and a measure of the intermediate inputs electricity and fuels. Real value-added is the real value of output adjusted for the real cost of all intermediate inputs. Construction of the real value of capital is documented in Lui (1991), and it includes buildings, machinery, and vehicles. Labor is the number of man-years hired for production, and firms distinguish between their blue- and white- collar workers. Electricity and fuels are measured in the real value of their volume consumed. All of these inputs enter the value-added production function in log-levels.

Real value of output is computed by deflating the total annual sales revenues of a firm with an industry level price deflator constructed by the Banco Central de Chile. This deflator will control for changes in output prices over time arising from inflation. It will also control for changes occurring because of industry level demand shocks. Having controlled for these time-varying effects, we then rely on price-taking behavior at the firm-level to get comparable quantities across both firms and time. The industries we look at have hundreds of firms, and this observation provides us with some comfort that price-taking behavior, even as an approximation to reality, is a reasonable assumption. However, we do remain concerned about the potential for differences in output prices within the 3-digit industry level that might arise because of imperfect competition, especially that caused by product differentiation (e.g. differences in type and quality of output.) Because we do not observe firm-level output prices, we are not in a position to address this concern without placing further restrictions on the framework.²³

The measure of real value added that we use is constructed by subtracting the real value of raw materials, electricity, and fuels from deflated total sales revenues. We express this measure as follows:

$$VA_{it} = TSR_{it}/P_{OUTPUT,t} - RM_{it}/P_{RM,t} - E_{it}/P_{E,it} - \sum_j F_{ijt}/P_{F,ijt},$$

where i indexes firms, VA is real value added, TSR is total sales revenues, RM is raw materials, E and F are value of electricity and fuels purchased, and j indexes the 11 different fuels recorded in the census. The deflator is written generically as $P_{.,t} = p_{.,t}/p_{80,t}$, and for the firm specific deflators, $P_{.,it} = p_{.,it}/p_{80,it}$. While the raw materials deflator is constructed at the country level, both the electricity and the fuels price deflators are firm specific. In the Census, firms report both the volume of electricity and fuels that they consume, and the price they paid for that volume. These numbers allow for the construction of average yearly unit input prices. The firm-level deflators then can be used to translate input prices for electricity and fuels into real 1980 Chilean pesos. While we are less concerned about variation in unit prices for electricity, we remain concerned that the quality of fuel inputs may vary within a defined category.²⁴ This method of deflation will account for that variance if the input quality is appropriately reflected in the input price, and our measure of value added will then be robust to differences in quality of the eleven different types of fuels.

We choose electricity as our primary proxy for productivity. Electricity has a number of advantages over fuels. First, electricity is an input that all firms need; we observe positive use of electricity in every year for almost every firm in the Census. Second, almost no firm reports that it generates electricity, or that it sells electricity, and we interpret this observation as an inability of

²³ Klette and Griliches (1996) provide one set of restrictions that allows for a particular kind of correction for unobserved firm-level prices.

²⁴ For example, we know that different grades of petroleum and gasoline exist, and these two categories make up a significant amount of the fuel consumption in our data.

firms to store (or stockpile) electricity. The inability to store electricity means that its use should be highly correlated with the year-to-year productivity shocks. In contrast, we do not observe positive use of fuels for a significant number of firm-year observations; on average, 25-30% of firms report zero fuel use and positive output in a year. Since fuels are recorded at purchase, we suspect that a zero firm-year observation may actually reflect positive use of some stockpiled inputs. This observation suggests that reported purchase of fuels may not generally correspond to actual fuel consumption. The mismeasurement of fuel usage will lead to inconsistent estimates of the $\phi(\cdot)$ function. For this reason, we are inclined to believe that, in our data, electricity will perform better than fuels as a productivity proxy. Our sensitivity analysis tends to support this view.²⁵

Data Details

The data set is well-suited to the econometric methods proposed in Section 2.2, but there are a few idiosyncratic features of the data that deserve attention. The first concerns the definition of the capital accumulation process and the second concerns missing investment data. These two prevent us from adopting the methods of Olley and Pakes directly. The third concerns missing capital stock data. Each is discussed in turn.

The first reason we cannot adopt the methods of Olley and Pakes without making substantive changes has to do with the definition of the capital accumulation process. Olley and Pakes construct their capital series by using a different assumption on the capital accumulation process than is used in the Chilean data. In Olley and Pakes, the timing of their model is as follows. Firms begin period t with capital k_t . They then observe ω_t after which they choose to stay or exit (χ_t). Firms then choose variable factors and the level of investment. Firms then commence the next period with:

$$k_{t+1} = (1 - \delta)k_t + i_t.$$

In the Chilean data, the capital accumulation process consistent with the reported capital series is given by

$$k_t = (1 - \delta)k_{t-1} + i_t.$$

The difference is subtle, but it matters. In the Olley and Pakes definition, investment from period t does not contribute to production in year t whereas the Chilean data is constructed such that it does. Under the primitives on the timing in the Olley and Pakes framework, the change in the capital accumulation process makes the estimator of capital coefficient inconsistent; this period's capital, which now includes this period's investment, will be correlated with ω_t because investment responds to ω_t (it is chosen after ω_t is realized.) However, a small change in the primitive on the timing of firm's decisions will permit i_{t+1} to be used as a proxy. The main point is that changing either the definition of capital accumulation or the primitives on the timing of a firm's actions can result in estimators contaminated with bias.

The second reason our framework is different from Olley and Pakes is also related to the observables in the Chilean data. The methods of Olley and Pakes revolve quite centrally around the investment decision and investment data. As discussed earlier, the estimation routine relies on investment being strictly increasing in the productivity shocks so the investment function can be inverted to proxy for productivity. They use data from the U.S. Census of Manufacturers, and they find that 8% of firm/year observations are reported to be zero. This feature of the data suggests

²⁵ Note that we can think of no compelling reason to believe that this stockpiling occurs in any systematic way, and we therefore think of this potential stockpiling phenomenon as measurement error introduced into the dependent variable; therefore, it does not contaminate our parameter estimates of the value added production function.

that invertibility of $i(\cdot)$ fails at zero investment.²⁶ Therefore, to use investment as a proxy will require truncation of all observations at zero investment.

In the Chilean data, one-third of the firm/year observations are reported to have zero investment. We are hesitant to truncate all of these observations from our estimation procedure. However, since we find that data on electricity is almost always reported at non-zero levels, we design our estimation strategy to rely on electricity as a proxy.

Another feature of the Chilean data that deserves mention concerns the capital stock variable. No initial capital stock is reported for some plants, although investment is recorded. When possible, we used a capital series that was reported for a subsequent base year. For a small number of plants, capital stock is not reported in any year. We estimated a projected initial capital stock based on other reported plant observables for these plants. We then used the investment data to fill out the capital stock data.

5. Estimation of Productivity and Results

In this section we discuss the details of obtaining parameter estimates for our value-added production function. Estimation proceeds in three stages, and is similar to Olley and Pakes (1996) and Olley and Pakes (1995), except that we use an intermediate input for the productivity proxy.²⁷ In the first stage we estimate the coefficients on the freely variable factors (unskilled and skilled labor) using a partially linear model (see, e.g., Robinson (1988).) In the second stage we construct an index which is used to control for the potential bias from selecting on survival. Using the consistent estimates of the labor coefficients and the probability of exit index from stages one and two, we then exploit the first order Markov nature of the productivity process to estimate the coefficient on capital. Finally, we use the estimated parameters to compute the productivity residuals. In the next section, we provide a detailed analysis of these productivity residuals, with a focus on the relative importance of the true productivity and the rationalization cases.

We begin by estimating equation (9). This equation is partially linear; it is linear in skilled and unskilled labor, and non-linear in $\phi_t(m_t, k_t)$. For our base case results, we use data on electricity usage for m_t , although we experiment with other suitable proxies. We proceed by projecting y_t on l_t^u , l_t^s , and a third order polynomial in m_t and k_t , i.e. we use a polynomial series to approximate the function $\phi_t(m_t, k_t)$.²⁸ We also allow for different functions $\phi_t(\cdot)$ for each of the periods 1979-81, 1982-83, and 1984-86, corresponding to the expansions and recession that is evident from Table 1. As we discussed in section 3, this stage yields estimates of β_1^u and β_2^u which are not contaminated by labor's responsiveness to the current period's productivity shock because productivity is controlled for by including the $\phi_t(\cdot)$ function.

The second stage estimates survival probabilities from equation (13). Recall from (13) that the two arguments in the exit decision function are m_t and k_t . Therefore, we use a probit which has as arguments a third order polynomial series in these two variables, and we also permit this function to differ over the three time periods. Generally, we find that controlling for selection in our model has little effect on the final parameter estimates (just as Olley and Pakes find with their unbalanced data.) This result is not surprising since we are not artificially truncating on survival to obtain a balanced sample (i.e. we are working with the full unbalanced data.)

²⁶ The function $i(\omega, k)$ has a set of ω in its domain with positive measure that maps to 0, and this loss of the injective property means that $i^{-1}(0, k)$ is not unique.

²⁷ Here we only cover the specifics of how our estimation proceeds. We encourage the interested reader to consult the articles by Olley and Pakes for a fuller discussion of the econometric issues, including proofs of the consistency of our parameter estimates.

²⁸ In this and all future polynomial series approximations we experimented with a fourth order expansion and found that it had a negligible effect on our final parameter estimates.

The last stage of the estimation uses $\hat{\beta}_t^u$, $\hat{\beta}_s^u$, $\hat{\phi}_t(\cdot)$, and $\hat{P}_t(\cdot)$ to construct the moment from (10) (or (11)) that can be used to consistently estimate the capital coefficient (it controls for the potential correlation between k_{t+1} and $E[\omega_{t+1}|\omega_t, \chi_{t+1} = 1]$.) Given any candidate value for β_k , say β^* , we can estimate the function $g(P_t, \phi_t - \beta^*k_t)$ using a third order series estimator in the two arguments $\hat{P}_t(\cdot)$ and $(\hat{\phi}_t - \beta^*k_t)$. Alternatively, we can compute for any candidate value β^* ($\hat{\phi}_t - \beta^*k_t$), and with $\hat{P}_t(\cdot)$ we are then able to compute the residual

$$[\xi_{i,t+1} + \mu_{i,t+1}](\beta^*)$$

for any firm i (see equation (10).) We then use a non-linear least squares routine to locate the minimizer $\hat{\beta}_k$ which solves

$$\min_{\beta} \sum_i \sum_{t=T_{i0}}^{T_{i1}} ([\xi_{i,t+1} + \mu_{i,t+1}](\beta))^2,$$

where T_{i0} and T_{i1} index the second and last period a firm is observed.

The results of the estimated production functions are reported in Table 2. The first column for each industry reports the OLS estimates of the production function while the second column gives the results using the Olley-Pakes estimator amended as discussed in Section 3. The results reinforce the message of Olley and Pakes (1995) (which looked at just one industry.)

As one might expect, almost all the coefficients of the Cobb-Douglas production functions are precisely estimated. (The sole exception is the OP/LP estimate of β_k in the Beverages industry.)²⁹ All industries have value-added production functions with "increasing" returns to scale. This was also the case for the production function estimates by Pavcnik (1997) as well as Lui (1991). It is important to be careful when interpreting these coefficients, though, since this notion is different from increasing returns to scale of the gross output production function. (See Basu and Fernald (1997) for a careful analysis of the relationship between estimates of returns to scale from gross output data and value-added data.)

If blue collar labor is positively correlated with productivity, so that more workers are hired in good years and fewer in bad years, theory suggests that the OLS coefficient on blue collar labor is biased upward. The same bias applies to white collar labor should it too be a variable factor responding to productivity shocks. For the industries in Table 2, the coefficient on blue collar labor falls in *every* industry when we condition on our proxy for ω_t . While the range of this decrease varies across industries, a decrease on the order of 15-20 percent seems common. For white collar labor, the results are less unanimous and this is consistent with the notion that white collar labor is less uniformly responsive to productivity shocks. In five industries, the OP/LP estimates of the coefficient on white collar labor are less than the OLS estimates, while in the other three they are higher. These results, taken together, are supportive of the notion that our proxy for ω_t is probably working well.

The coefficient on capital in our production function increases when we use the OP/LP estimator instead of the OLS estimator in *every* industry. In Food Products, Printing and Publishing, Apparel, and Wood Products, this increase is especially substantial. There are three reasons why the OLS and OP/LP estimates might differ. First, if capital usage is (positively) correlated with either this period's productivity (ω_t) or last period's, the OLS coefficient on capital will be (upwardly) biased. This is the same story as with the coefficient on blue-collar labor. Second, if exit is important, the OLS capital coefficient will be biased downward. Because we are working with an

²⁹ The estimates for the Beverages industry, and only this industry, do not include a correction for possible selection bias. That is, we do not estimate the probit for exit probabilities in this industry. This is because there were so few firms exiting this industry that the probit was not well identified.

unbalanced panel, the selection correction is unlikely to have the impact that it might if we were comparing the OP/LP estimator with OLS on the balanced panel. To best understand the third reason why we might expect the coefficient on capital to change, imagine the following thought experiment. Suppose that there was no selection bias and that capital usage was not correlated with this or last period's productivity. One can show that if this is the case, but that labor is correlated with productivity, the coefficient on capital will be downwardly biased so long as capital positively covaries with labor. This is a statistical property of the estimates and exists even when the economic properties of capital and productivity do not induce a bias. In practice, all three of these biases are, to some extent, working in different directions and at varying magnitudes and the difference between the OLS and OP/LP estimates confounds the three influences.

Olley and Pakes (1995) also found that correcting for possible simultaneity led to an increase in the estimated marginal product of capital. Our results are the same and, as such, are perhaps not surprising. On the other hand, the fact that the pattern of the bias in the capital coefficient was the same in all the industries was surprising. Our results suggest that the changes in the marginal productivities found by Olley and Pakes in telecommunications were not a fluke. In other words, Marschak and Andrews had it right in 1944; it appears that correcting for simultaneity really does change the production function estimates.

We conclude this section with a specification test that essentially asks whether the relative complexity of the OP/LP estimates is really necessary. While the OP/LP estimation algorithm is not especially difficult, it is more involved than either OLS or the most commonly adopted alternative—fixed effects estimation. The benefit of the OP/LP estimates are that they do not impose too much structure on the residuals (a component of which is interpreted to represent productivity.) The OP/LP estimates assume that the residuals follow a first-order Markov process. This assumption nests both OLS and the Fixed Effects specification. When we conducted Wald tests to see whether the data reject the simpler specifications, we found the following. When OLS was the null hypothesis, we rejected the null in favor of the first-order Markov process for seven of the eight industries. Tests were performed at the 1% level of significance. Only in “Other Chemicals” could we not reject OLS. When fixed effects was the null hypothesis, we rejected this null in all eight industries. We conclude that the data are not compatible with the simpler specifications. In the next section, we investigate whether the differences between the simpler OLS estimation and the OP/LP estimation matter to the questions we address.

6. Productivity Dynamics

The estimates in Table 2 imply firm-specific measures of productivity. We define the productivity of firm i in year t to be:

$$\hat{\omega}_{it} = \exp(y_{it} - \hat{\beta}_k k_{it} - \hat{\beta}_s l_{it}^s - \hat{\beta}_u l_{it}^u), \quad (15)$$

where the estimated β 's are the OP/LP estimates in Table 2. We begin our discussion of productivity dynamics by simply examining the annual industry-level productivity indexes. We define industry-level productivity, denoted Ω , to be a weighted average of firm-level productivities, where weights are shares of industry output. Hence,

$$\Omega_t = \sum_{i=1}^N s_{it} \hat{\omega}_{it}$$

where s_{it} is plant i 's share of industry output in year t and $\hat{\omega}_{it}$ is as defined in (15).

Normalizing the productivity in each industry in 1979 to 1.00, the industry-level productivity indexes are reported in Table 3. The bottom two rows for each industry summarize the change in productivity for the periods 1979-83 and 1983-86. The latter period was marked by consistent real

GDP growth while the former includes a recessionary cycle. In Metals, Textiles, Food Products, and Apparel, there was an increase in industry productivity, while in the others productivity declined. In Beverages, the decline was modest, while in Chemicals, Printing, and Wood Processing it was substantial.³⁰ Overall, this index of average productivity suggests that there is much heterogeneity in productivity growth across the industries.

We also investigate changes in industry-level productivity by examining how the entire distribution of firm-level productivity evolves. Our approach here is essentially non-parametric and asks whether this distribution is stable over time. The test used is the distribution-free Kolmogorov-Smirnoff test. Loosely speaking, the test examines two empirical distribution functions, looks for the maximal difference between them, and applies a test statistic to this difference. This test has the advantage that it imposes very little structure on the data. On the other hand, the test does not discern how the distribution of productivity changes. One can only test the null hypothesis that both observed empirical distributions come from the same underlying population distribution.

There are several ways one could proceed with this approach. One could compare the distribution of productivities for an industry on a year-by-year basis and ask whether two consecutive years' productivities were drawn from the same distribution. We opt for a more parsimonious approach and simply test whether the distributions of productivities for firms are the same for the 1979-82 and 1983-86 sub-periods. That is, we take the distribution of all estimated productivities, combining years within each sub-period, and ask whether the distributions for the two sub-periods come from the same underlying distribution.

The results for each of the eight industries are given in the column of Table 4 titled "Stationary Productivity." At a one percent level of significance, we accept the null hypothesis that the two sample distributions come from the same underlying population distribution for three industries: Metals, Beverages, and Wood Products.³¹ In the five other cases we reject this null hypothesis. These results are broadly supportive of the results reported in Table 3. The industries for which we accept the null of the same underlying distribution across sub-periods are the same industries that, in Table 3, had little difference between the average index productivity number for 1979-82 and the number for 1983-86. The exception is Textiles, where Table 3 indicates a relatively small change but the Kolmogorov-Smirnoff indicates different underlying distributions.

The numbers in Table 3 are unitless (they are ratios of averages.) In this regard, they tell us little about the actual economic value of the productivity changes. In order to measure this magnitude, we ask how a firm's value of output would have changed from year t to year $t + 1$ if we held inputs constant but let productivity evolve. Formally, we define the value of the productivity change for firm i in year $t + 1$, or $v_{i,t+1}$, as

$$\begin{aligned} v_{i,t+1} &= \exp(\beta_0 + \beta_s l_{it}^s + \beta_u l_{it}^u + \beta_k k_{it} + \omega_{i,t+1}) - \exp(\beta_0 + \beta_s l_{it}^s + \beta_u l_{it}^u + \beta_k k_{it} + \omega_{it}) \\ &= \exp(\omega_{i,t+1}) \exp(\beta_0 + \beta_s l_{it}^s + \beta_u l_{it}^u + \beta_k k_{it}) - \exp(y_{it}). \end{aligned} \quad (16)$$

For each industry, Table 5 reports the sum across both firms and time (1979-86) of these peso-denominated changes in output. Using this measure, we find that six of the eight industries experienced productivity gains between 1979 and 1986. For some industries, these gains were rather large. Textiles and Food Products experienced gains in productivity of approximately 30% of their

³⁰ The large decline in the productivity index for Wood Products in 1986 leads us to view that number with some caution. In this industry, one very large and very productive firm exits the sample in the last year. It is unclear whether this is a reporting error or true exit. We treat it as the latter, but are cautious about reading much into the large decline in industry productivity in this industry in 1986.

³¹ These industries could be viewed as consistent with the equilibrium notion employed in Hopenhayn and Rogerson (1993).

average value-added, and Apparel registered a gain of over 50% of value-added. While these numbers are large, they may be quite consistent with the productivity gains accruing to industries that benefitted from the extensive economic reforms completed just prior to the beginning of our sample period.

As mentioned earlier, our estimation routine nests OLS. Therefore, it is natural to compare the OLS estimates of productivity gains using the above approach with those obtained from the OP/LP approach. Column 4 of Table 5 reports the OLS numbers. Column 5 reports the difference between the OLS number and the OP/LP number, divided by the average level of industry value added over the sample period. OLS appears to both under- and over-forecast the gains from productivity. Frequently this mistake is large. In Food Products, Beverages, Printing and Publishing, and Wood Products, OLS misforecasts by approximately 10% of the average value added for the industry. In Metals OLS overestimates the gains by almost 7%. In Textiles, however, OLS does appear to get it right. These results suggest that OLS should be used with care if inferences are to be drawn about the currency-denominated magnitudes of the productivity changes.

Decomposing Productivity

The numbers we have reported up to now have been aggregated to a level which makes it difficult to see how productivity is changing. In our results, these aggregate numbers hide tremendous variation of firm-level changes in productivity over time. Our main focus is on trying to better understand what that variation means for real productivity and rationalization. Therefore, in order to better understand what underlies the industry-level changes in productivity, we focus in this subsection on decomposing those changes.

We decompose the changes in industry-level productivity into four parts. For the sake of illustration, suppose that industry-level productivity increases. This can occur because output shares are reallocated from less efficient plants to more efficient plants. We term this effect the reallocation effect.³² Productivity can also increase because each plant in the industry becomes more productive even though output shares remain constant. We term this effect the productivity effect. Industry-level productivity can increase if entrants are more productive than the average incumbent. This effect is the entry effect. Finally, industry-level productivity can increase if less efficient firms exit, and this is the exit effect. Formally, the decomposition is given by:

$$\Delta\Omega = \sum_{i \in C} s_{i,t-1} \Delta\hat{\omega}_{it} + \sum_{i \in C} \Delta s_{it} \hat{\omega}_{it} + \sum_{i \in B} s_{it} \hat{\omega}_{it} - \sum_{i \in D} s_{i,t-1} \hat{\omega}_{i,t-1} \quad (17)$$

where C is the set of continuing firms, B the set of entrants, and D the set of exiters. The difference operator, Δ , denotes the difference between year t and $t - 1$. The terms on the right-hand-side of (17) are, respectively, the productivity, reallocation, entry, and exit effects.³³ The first term in the decomposition corresponds to our true productivity case, while the last three terms correspond to our rationalization case.

Table 6 provides the productivity decompositions for each industry. Before examining the eight industries and patterns across them, it is helpful to first focus on just one industry and consider in some detail what the numbers in this table mean. We focus, for the sake of illustration, on the first industry in the table, Metals. We have divided the sample period into two sub-samples. As noted previously, the period from 1979-83 was a time of macroeconomic contraction while 1983-86 were

³² This is *not* the same as the reallocation effect discussed in Olley and Pakes (1995). In that paper, "reallocation" refers to the contemporaneous covariance between market share and productivity.

³³ This decomposition is not original and appears, in various forms, throughout the literature. See Bailey et al. (1992) and Tybout (1996).

more expansionary years in Chile. The top of the first column for Metals lists the total change in the industry productivity index – 17.2 – which is computed as $\Omega_{83} - \Omega_{79}$. This index number represents the productivity change in the Metals industry from 1979 to 1983. To put its magnitude in some perspective, the bottom two rows of the first column list two measures of the average level of productivity in this industry. The first measure is the average of the first and final year of $\Omega_{(.)}$ ($(\Omega_{83} + \Omega_{79})/2$ in this case). The second number averages across all years (1979-83 in this case).³⁴ Depending on the measure, the Metals productivity index averages either 66.2 or 79.2. Against this *level*, the *change* in the index over this period was 17.2. Hence, according to this measure, productivity increased on the order of 20 to 25 percent in Metals between 1979 and 1983.

The four rows below the total change give the components of the decomposition. The sum of these four rows, by construction, is the total change. The first of these rows gives the productivity term. This can be either positive (for a given distribution of market shares, firms become more productive) or negative (for a given distribution of market shares, firms become less productive.) The second row gives the reallocation term. This term can also be either positive (increasing market shares for more productive firms) or negative (decreasing market shares for more productive firms.) The third term gives the impact of entry. Due to the definition of productivity in (15), this term is always positive, although its magnitude can vary widely. The last term gives the impact of exit and, again due to the definition of $\hat{\omega}$, this term is always negative. Exit also impacts the reallocation term, albeit indirectly, since when a firm exits, the market shares of surviving firms increase, all else equal. (Similarly, entry has a negative influence on the reallocation term.) It remains an empirical question whether the net impact of entry and exit is to increase or decrease industry productivity. In the example at hand, the net impact is an increase ($2.8 - 1.7 = 0.8$) in productivity due to entry and exit. We now turn to examining the rest of Table 6.

Of the four industries in which aggregate productivity increased, Metals, Textiles, Apparel, and Food Products, we find that very little of the increase in productivity was accounted for by firms actually becoming more productive (the real productivity case.) The sole exception is the Apparel industry from 1983-86. Indeed, in four of the eight sub-periods for these increasing productivity industries, the share-weighted change in firm productivity was *negative*. Rather than firms becoming more productive, reallocation of market shares to firms that were already more productive and net entry typically explain the increase in aggregate productivity. In the introduction, we posed the question: “As industries become more productive, do firms?” In our data, the answer is almost always “no,” as we find scant evidence of real productivity gains even where the industry-level measure of productivity increases.

A second finding from these productivity decompositions is an apparent asymmetry between the reallocation/rationalization story. When industry productivity increases, output gets reallocated to more productive firms. Were reallocation and rationalization to play an important role in the industries for which productivity decreases, market shares would have to shift away from the more productive firms and toward the less productive firms. While there is nothing to prevent this from happening, economic intuition suggests it ought not happen much. This is indeed the case. In the four industries for which aggregate productivity falls, incumbent firms suffer decreased productivity (the real productivity case applies when productivity declines.) We also find, for the industries in which aggregate productivity falls, that the gross impact of entry and exit (defined as the sum of entry and the absolute value of exit) is typically larger than the impact of reallocation. The fact that reallocation simply does not appear to be very important in industries with declining productivity strikes us as reassuring evidence that we are perhaps actually measuring what we think we are measuring.

³⁴ The first measure also accounts for entry and exit of firms, whereas the second measure does not. See Table 6 notes for details.

Before leaving behind the issue of what underlies the industry-level changes in productivity, we report one more set of results that does not depend on the specific decomposition used in (17). If firms are actually becoming more productive in some systematic way, one might expect the measured firm-level productivity to increase from one year to the next. In the last column of Table 4, we report the percentage of firm-year observations that show an increase in productivity over the previous year. The boot-strapped standard errors of this statistic are given in parentheses. In Table 3, four industries showed increased industry-level productivity: Metals, Textiles, Food Products, and Apparel. In Table 4, all except Metals have more than 50 percent of adjacent firm-year observations showing increased productivity, and this suggests that the decomposition in (17) is reasonable. What is surprising is how often adjacent firm-year observations *decline* in the industries showing the largest industry-level productivity *gains*. We view these results as again stressing the empirical importance of firm heterogeneity and the role of industry rationalization.³⁵

Table 2 listed both OLS and OP/LP estimates of the production functions. We have already seen that different parameter estimates translate into different estimates for the value of productivity gains. It is reasonable to wonder whether the different parameter estimates translate into different implications for the components of change in industry productivity. Table 7 addresses this concern. That table is essentially the same as Table 6 except the OLS estimates are used to compute the underlying estimated ω 's whose changes are then decomposed.

A comparison of Tables 6 and 7 shows that the estimated levels of productivity (the bottom two rows) for a given industry vary depending on the estimation method (OLS versus OP/LP), and there are no general patterns. In Beverages and Other Chemicals, the level of the productivity index is higher with the OP/LP estimates. That is, given the exact same input and output data, the OP/LP estimates imply a higher level of productivity than do the OLS estimates. In the other six industries, the OLS estimates imply a higher level of productivity. In some cases, the differences are quite dramatic. In Printing and Publishing as well as in the Wood Products industries, the OLS estimates appear to substantially overstate the level of productivity. In Beverages, on the other hand, the OLS estimates yield a productivity index that is only about one third that implied by the OP/LP estimates. The differences in the decompositions across the two types of estimates, though, are less dramatic. While the share of the change due to reallocation versus productivity changes across the two types of estimates, the general pattern is the same. When productivity rises, whatever the method of estimation, reallocation is typically quite important while real productivity changes apply when the productivity falls. This results obtains because: 1) many of the variables in the decomposition (i.e. shares and changes therein as well as the sets of entrants and exiters) are data and are unaffected by the method of estimation; and 2) the estimated ω 's are highly correlated across the methods of estimation. In summary, the decompositions in Table 7 present a picture of what underlies industry changes in productivity that is similar to that presented in Table 6.

In addition to using electricity as the productivity proxy, we also experimented with using fuels. As mentioned earlier, we felt that fuels might present certain problems given that they are easily stockpiled, and that there are about 25-30 percent of firm-year observations that report positive output but zero fuel consumption (i.e. zero fuel purchase), compared to only 5 percent for electricity. Nonetheless, we estimated the production functions for all eight industries using fuels as the productivity proxy. We found that in three cases, Textiles, Wood Products, and Paper and Printing, the parameter estimate for capital was significantly different from that obtained using electricity. Only in Food Products were the labor coefficients significantly different from those obtained using electricity. In all other cases, the labor coefficients for both skilled and unskilled

³⁵ The numbers in the last column of Table 4 somewhat understate the role of industry rationalization since entry and exit do not figure into these numbers.

labor were comparable across proxies. Again, we are inclined to attribute these differences to the mismeasurement of the utilization of fuels which occurs because they are reported at purchase and not at use. In spite of these differences, we feel that it is not unfair to conclude that the choice of proxy is not driving our key results.

7. Conclusions and Caveats

The paper has investigated what underlies changes in industry productivity in Chile during a period following a broad trade liberalization. In the introduction, we posited two ways in which industry productivity might change. The first, coined the real productivity case, involved changes in firm-level productivities. In the instance of increasing industry productivity, firms actually become more productive. The second, coined the rationalization case, focused on the role of shifting market shares. In the case of increasing industry productivity, market shares get reallocated toward more productive plants and away from less productive ones. Entry and exit also play a role. Our results indicate that the rationalization case is empirically the most important case when industry productivity increases. When industry productivity decreases, the true productivity case is more important.

The rationalization case is one in which firm heterogeneity takes front stage. The real productivity case can accommodate firm heterogeneity, but the story is perfectly plausible in an industry of identical firms. When reviewing models of how openness impacts firms, industries, and finally, countries, the theoretical literature is full of models of representative firms, yet the empirical evidence stresses the importance of firm heterogeneity. One lesson from this paper is that there is a need for more modeling of how openness impacts firms and hence industries and countries, and this modeling needs to account for the distribution of firms in an industry. The industrial organization literature provides some useful starting points for such a modeling exercise.

In the course of measuring productivity dynamics, it was important to estimate productivity in a flexible manner that did not impose too much structure on the pattern of productivity. Our methods nested simpler structures such as the i.i.d. assumption of Ordinary Least Squares or the constant productivity over time by a firm as in the fixed effects estimator. We soundly reject the appropriateness of these simpler estimators, and our results suggest that the differences matter in ways that are economically important. However, we do find that OLS *does* consistently predict the direction of productivity movements. Empirically, if one is only interested in the direction, our results suggest that OLS with an unbalanced panel may suffice, thereby avoiding the additional complications associated with estimating parameters of the more flexible framework.

The estimator we use builds directly on the work of Olley and Pakes. We extend the powerful idea of using a proxy for productivity, and show that a wide range of proxies are available when using a value-added production function. We also validate the findings of Olley and Pakes for a wide range of industries. Finally, we find econometric evidence of the relative flexibility of inputs, as white collar labor appears less flexible than blue collar labor. The relative flexibility of inputs is not discussed much in this paper, but the methods developed are well-suited to inferring relative flexibility of inputs— an idea that is important to distinguishing between several models of international trade as well as an important determinant in several political economy models.

That is what we did. There are several reasons to view these results with some caution. We suggest four such caveats, and view these caveats as potential extensions to this work.

First, what about the demand side? This paper, like almost all its predecessors in the production function estimation literature, conducts its analysis in the total absence of a demand side of the industry. Adding a demand side to the model and integrating this with the production side would be a very useful extension. Related to this, there is no discussion in the paper of the role of capacity utilization. Capital which is in place but is not used due to insufficient demand shows up as capital

in the production function. Output, though, is presumably lower with unused capital. Is this really lower productivity?

Second, the results of our estimated production functions may be consistent with increasing returns to scale, yet our proxy is dependent on price-taking behavior by the firm. As noted above, we estimate a value-added production function. Simply adding coefficients as one would do with an output production function may not be appropriate. Yet under some circumstances it is appropriate while under others, the estimated returns to scale with the value-added production function underestimates returns to scale. There is a tension between the estimates and the assumption of price taking firms. There are ways to explain this. Returns to scale might be external to the firm. There may be measurement error in inputs or outputs that give rise to this. An explanation we find more convincing is that with sunk costs and productivity shocks that follow a first-order Markov process, most firms will eventually exit. Knowing this, firms' otherwise unfettered tendency to work their way down the average cost curve (and the resulting industry consolidation) is tempered. This suggests that firms might still be price takers even if there are increasing returns to scale. These industries have hundreds, and in some cases thousands, of firms so, on the surface, the price-taking assumption does not seem crazy. Still, this caveat deserves mention.

Third, there are yet more sensitivity analyses that we have not conducted. For example, we have adopted a very simple Cobb-Douglas utility function. Would a slightly more flexible functional form yield a similar story concerning the relative importance of the real productivity and rationalization cases?

Fourth, the model implicitly assumes maximizing firms, but it does not make use of the additional structure this imposes.³⁶ For example, estimation is focused exclusively on the production function. One could, though, simultaneously estimate the derived demand for inputs that is consistent with the estimated production function. This would impose additional structure and could contribute to more efficient estimates of the parameters of the production function. While this would substantially complicate estimation, it would be in the spirit of recent work that uses equilibrium relationships to better identify parameters of interest.

³⁶ We are grateful to Roger Gordon for pointing this out.

References

- Arrow, K. (1962). The economic implications of learning by doing. *Review of Economic Studies*, 29, 155–173.
- Aw, B., Chen, X., and Roberts, M. (1998). Firm-level evidence on productivity differentials and turnover in Taiwanese manufacturing. Mimeo.
- Bailey, M. N., Hulten, C., and Campbell, D. (1992). Productivity dynamics in manufacturing plants. *Brookings Papers on Economic Activity (Microeconomics)*, 187–249.
- Basu, S., and Fernald, J. (1997). Returns to scale in U.S. production: estimates and implications. *Journal of Political Economy*, 105(2), 249–83.
- Davis, S., and Haltiwanger, J. (1990). Gross job creation and destruction: Microeconomic evidence and macroeconomic implications. *NBER Macroeconomics Annual*, V, 123–168.
- Dunne, T., Roberts, M., and Samuelson, L. (1988). Patterns of firm entry and exit in u.s. manufacturing industries. *Rand Journal of Economics*, 19(4), 495–515.
- Ericson, R., and Pakes, A. (1995). Markov-perfect industry dynamics: A framework for empirical work. *Review of Economic Studies*, 62(1), 53–82.
- Griliches, Z. (1998). *R&D and Productivity*. Chicago.
- Griliches, Z., and Regev, H. (1995). Firm productivity in Israeli industry: 1979–1988. *Journal of Econometrics*, 65(1), 175–203.
- Grossman, G., and Helpman, E. (1991). *Innovation and growth in the global economy*. MIT Press.
- Harrison, A. (1994). Productivity, imperfect competition and trade reform: Theory and evidence. *Journal of International Economics*, 36, 53–73.
- Heckman, J. (1981). The incidental parameters problem and the problem of initial conditions in estimating a discrete time-discrete data stochastic process. In Manski, and McFadden (Eds.), *Structural Analysis of Discrete Data with Econometric Applications*. MIT Press.
- Hopenhayn, H. (1992). Entry, exit, and firm dynamics in long-run equilibrium. *Econometrica*, 60, 1127–50.
- Hopenhayn, H., and Rogerson, R. (1993). Job turnover and policy evaluation: A general equilibrium analysis. *Journal of Political Economy*, 10(5), 915–938.
- Jovanovic, B. (1982). Selection and the evolution of industry. *Econometrica*, 50(3), 649–670.
- Klette, T. J., and Griliches, Z. (1996). The inconsistency of common scale estimators when output prices are unobserved and endogenous. *Journal of Applied Econometrics*, 11, 343–361.
- Levinsohn, J. (1998). Employment responses to international liberalization in Chile. Forthcoming in the *Journal of International Economics*.
- Lui, L. (1991). *Entry-Exit and Productivity Changes: An Empirical Analysis of Efficiency Frontiers*. Ph.D. thesis, University of Michigan.
- Lui, L. (1993). Entry, exit, and learning in the Chilean manufacturing sector. *Journal of Development Economics*, 42, 217–242.
- Lui, L., and Tybout, J. (1996). Productivity growth in Colombia and Chile: Panel-based evidence on the role of entry, exit and learning. In Roberts, M., and Tybout, J. (Eds.), *Producer Heterogeneity and Performance in the Semi-Industrialized Countries*, chap. 4. World Bank.
- Marschak, J., and Andrews, W. (1944). Random simultaneous equations and the theory of production. *Econometrica*, 12(3-4), 143–205.
- Olley, S., and Pakes, A. (1995). A limit theorem for a smooth class of semiparametric estimators. *Journal of Econometrics*, 65(1), 292–332.
- Olley, S., and Pakes, A. (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6), 1263–1298.

- Pakes, A. (1996). Dynamic structural models, problems and prospects: mixed continuous discrete controls and market interaction. In Sims, C. (Ed.), *Advances in Econometrics, Sixth World Congress, Volume II*, pp. 171–259 New York. Cambridge.
- Pavcnik, N. (1997). Trade liberalization, exit, and productivity improvements: Evidence from Chilean plants. Working Paper, Department of Economics, Princeton University.
- Roberts, M., and Tybout, J. (1996). *Industrial Evolution in Developing Countries*. Oxford University Press for The World Bank, New York.
- Roberts, M., and Tybout, J. (1997). The decision to export in colombia: An empirical model of entry with sunk costs. *American Economic Review*, 87(4), 545–64.
- Robinson, P. (1988). Root-n consistent semiparametric regression. *Econometrica*, 55(4), 931–954.
- Romer, P. (1986). Increasing returns and long-run growth. *Journal of Political Economy*, 94, 1002–1037.
- Tybout, J. (1996). Heterogeneity and productivity growth: Assessing the evidence. In Roberts, M., and Tybout, J. (Eds.), *Industrial Evolution in Developing Countries*, chap. 3. World Bank.
- Tybout, J., de Melo, J., and Corbo, V. (1991). The effects of trade reforms on scale and technical efficiency: New evidence from Chile. *Journal of International Economics*, 31, 231–250.

Appendix I.

In this appendix we consider the use of intermediate inputs as proxies for productivity shocks when firms operate in a competitive environment. We show the general conditions on the production technology which yield an intermediate input demand function $m^*(\omega; p_l^*, p_m^*)$ that is strictly increasing in the productivity shock (ω) (the price of output is normalized to 1.) This result permits the use of $h(m, k)$ as an index for the productivity shock.

Definition. An industry is competitive if firms take input prices and the output price as given.

Intermediate inputs are available as proxies in some imperfectly competitive environments, although the proof depends on the specifics of the competition. Proofs in an imperfectly competitive environment will likely rely on arguments from the literature on monotone methods.

Assumption 1. The firm production technology $f(k, l, m, \omega) : \mathbb{R}^4 \rightarrow \mathbb{R}$ is twice continuously differentiable in labor (l) and the intermediate input (m), and $f_{l\omega}$, $f_{m\omega}$, and f_{ml} exist for all values $(k, l, m, \omega) \in \mathbb{R}^4$. The industry is competitive, and either a) this period's investment does not respond to this period's productivity, or b) it does not enter this period's capital. Productivity is observed before the choice of labor and the intermediate input are made.

The differentiability of $f(\cdot)$ can be relaxed with the appropriate appeal to monotone methods. We treat capital as fixed, and assume both labor and the intermediate input respond to the productivity shock. With some additional complexity, it is possible to show the following result when capital also responds to ω , and when more than one type of labor exists.

Result. Under Assumption 1, if $f_{ml}f_{l\omega} > f_{ll}f_{m\omega}$ everywhere, then $m^*(\omega; p_l^*, p_m^*)$, the intermediate input demand function, is strictly increasing in ω .

Proof. Given the assumption, a profit-maximizing firm has an intermediate input demand function that satisfies

$$\text{sign}\left(\frac{dm^*}{d\omega}\right) = \text{sign}(f_{ml}f_{l\omega} - f_{ll}f_{m\omega}) \quad (A1)$$

(see Varian (1992), pp. 494-495.) Under mild regularity conditions on $f(\cdot)$ that insure the Fundamental Theorem of Calculus holds for $m^*(\cdot)$,

$$m^*(\omega_2; p_l^*, p_m^*) - m^*(\omega_1; p_l^*, p_m^*) = \int_{\omega_1}^{\omega_2} \frac{dm^*}{d\omega}(\omega; p_l^*, p_m^*) P(d\omega).$$

Since $f_{ml}f_{l\omega} > f_{ll}f_{m\omega}$ everywhere, it follows from (A1) that

$$\int_{\omega_1}^{\omega_2} \frac{dm^*}{d\omega}(\omega; p_l^*, p_m^*) P(d\omega) > \int_{\omega_1}^{\omega_2} 0 P(d\omega) = 0,$$

so

$$m^*(\omega_2; p_l^*, p_m^*) > m^*(\omega_1; p_l^*, p_m^*) \text{ if } \omega_2 > \omega_1.$$

Economics can be used to sign the above derivatives. If the marginal product of labor declines as labor increases, $\text{sign}(f_{ll})$ is negative. Additionally, if increases in productivity always weakly increase the marginal product of inputs, then $f_{l\omega}$ and $f_{m\omega}$ are always non-negative. Given these economic restrictions on $f(\cdot)$, the essence of the result is clear; the cross-partial of output with respect to the intermediate input and labor will determine whether intermediate input use is monotonic increasing in productivity. If the marginal product of the intermediate input increases

as labor use increases (i.e. $f_{ml} > 0$), then the monotonicity result holds. However, if the marginal product of the intermediate input falls rapidly with increases in labor, then $f_{ml}f_{l\omega} - f_{ll}f_{m\omega}$ may be negative, causing the monotonicity result to fail.

Using the above result, it is straightforward to show that the Cobb-Douglas production technology

$$y = \exp(\beta_0 + \omega + \beta_l l + \beta_k k + \beta_m m + \eta)$$

has an intermediate input demand function that is strictly increasing in ω (inputs are written in log-levels.) The derivatives of interest and their signs are

$$\begin{aligned} f_{ml} &= \frac{\beta_l \beta_m}{lm} y > 0, \\ f_{l\omega} &= \frac{\beta_l}{l} y > 0, \\ f_{m\omega} &= \frac{\beta_m}{m} y > 0, \\ f_{ll} &= \frac{(\beta_l - 1)(\beta_l)}{l^2} y < 0, \end{aligned}$$

and the expression $f_{ml}f_{l\omega} - f_{ll}f_{m\omega}$ is positive if $\beta_l \leq 1$. An increase in ω brings about an increase in the marginal product of all inputs, and with fixed input and output prices, the firm will use more of the intermediate input to drive its marginal product back down to p_m^* , the normalized input price. Since this feature holds true for all small increases in productivity, it holds true for all discrete changes, and use of the intermediate input is strictly increasing in ω .

TABLE 1										
Some Descriptive Statistics on Chilean Manufacturing										
	GDP	Real Exchange Rate	Metals		Textiles		Food Products		Beverages	
			Plants	Value Added	Plants	Value Added	Plants	Value Added	Plants	Value Added
1979	997.6	171.3	459	10.0	503	12.4	1,537	39.0	211	14.1
1980	1,075.3	199.1	447	11.0	445	12.9	1,439	43.4	188	14.7
1981	1,134.7	234.9	413	11.5	403	11.3	1,351	42.7	158	10.4
1982	974.9	212.3	365	8.1	350	8.7	1,319	47.0	151	11.7
1983	968.0	172.9	322	8.3	327	9.7	1,297	42.9	148	13.1
1984	1,029.4	169.9	358	11.4	336	10.4	1,340	46.8	138	13.2
1985	1,054.6	136.9	351	9.6	337	10.8	1,338	49.1	127	11.5
1986	1,114.3	115.7	347	9.6	331	12.9	1,288	61.4	111	12.8
			Other Chemicals		Printing & Publishing		Wood Products		Apparel	
			Plants	Value Added	Plants	Value Added	Plants	Value Added	Plants	Value Added
1979			171	15.1	242	11.4	524	10.4	442	6.6
1980			166	16.9	227	10.5	449	8.7	398	6.7
1981			159	17.5	206	12.8	406	6.8	346	6.6
1982			148	14.8	196	8.9	358	6.5	305	5.3
1983			145	12.6	177	5.8	335	8.1	265	4.1
1984			151	12.9	167	5.8	339	10.3	294	6.4
1985			149	11.2	164	4.7	342	10.1	275	8.5
1986			153	9.1	163	4.5	313	5.3	280	11.3

Notes: GDP figures from the International Financial Statistics Yearbook. GDP and value added in millions of 1980 pesos.

TABLE 2								
Production Function Estimates								
	Metals		Textiles		Food Products		Beverages	
	OLS	OP/LP	OLS	OP/LP	OLS	OP/LP	OLS	OP/LP
Blue Collar	0.712 (0.032)	0.573 (0.034)	0.564 (0.037)	0.506 (0.042)	0.527 (0.019)	0.440 (0.018)	0.561 (0.083)	0.336 (0.086)
White Collar	0.462 (0.034)	0.405 (0.035)	0.464 (0.039)	0.477 (0.039)	0.522 (0.020)	0.305 (0.009)	0.733 (0.074)	0.601 (0.076)
Capital	0.210 (0.017)	0.308 (0.019)	0.205 (0.019)	0.276 (0.026)	0.334 (0.075)	0.499 (0.009)	0.222 (0.030)	0.264 (0.166)
Sum of Coeff.	1.384	1.286	1.233	1.260	1.383	1.244	1.516	1.200
No. Obs.	2081	1659	2094	1683	6511	5403	566	462
	Other Chemicals		Printing & Publishing		Wood Products		Apparel	
	OLS	OP/LP	OLS	OP/LP	OLS	OP/LP	OLS	OP/LP
Blue Collar	0.375 (0.051)	0.283 (0.054)	0.582 (0.046)	0.471 (0.046)	0.625 (0.042)	0.542 (0.048)	0.709 (0.037)	0.565 (0.039)
White Collar	0.724 (0.048)	0.744 (0.049)	0.352 (0.040)	0.431 (0.039)	0.586 (0.050)	0.487 (0.054)	0.456 (0.039)	0.416 (0.039)
Capital	0.275 (0.035)	0.286 (0.028)	0.288 (0.029)	0.543 (0.041)	0.177 (0.027)	0.373 (0.013)	0.142 (0.021)	0.237 (0.098)
Sum of Coef	1.374	1.312	1.223	1.444	1.388	1.403	1.306	1.218
No. Obs.	928	766	987	824	1610	1299	1760	1412

Notes: The estimate for Beverages does not correct for firm survival.

TABLE 3				
Industry Productivity				
	Metals	Textiles	Food Products	Beverages
1979	1.00	1.00	1.00	1.00
1980	1.21	1.23	0.87	0.88
1981	1.27	1.61	0.94	0.71
1982	1.21	1.34	1.09	1.06
1983	1.35	1.32	0.89	0.94
1984	1.47	1.16	1.62	0.87
1985	1.17	1.19	1.60	0.79
1986	1.06	1.33	1.81	0.89
Change 79-83	0.35	0.32	-0.11	-0.06
Change 83-86	-0.29	0.01	0.92	-0.06
	Other Chemicals	Printing & Publishing	Wood Products	Apparel (Preliminary)
1979	1.00	1.00	1.00	1.00
1980	0.92	0.90	0.86	1.17
1981	0.88	0.94	0.69	1.31
1982	0.88	0.64	1.13	1.52
1983	0.71	0.52	0.98	1.31
1984	0.69	0.53	0.95	1.33
1985	0.59	0.44	1.00	1.52
1986	0.50	0.39	0.60	2.23
Change 79-83	-0.29	-0.48	-0.02	0.31
Change 83-86	-0.21	-0.13	-0.38	0.92

Notes: Industry Productivity is defined as a weighted sum of firm productivities (ω 's), with shares of total value added as weights. The productivities are normalized by industry so that industry productivity equals to 1.00 in 1979.

TABLE 4 Some Non-Parametric Evidence		
Industry	Stationary Distribution	Increasing Productivity
Metals	YES	47.9% (1.0%)
Textiles	NO	50.4% (1.3%)
Food Products	NO	54.2% (0.6%)
Beverages	YES	49.5% (2.1%)
Other Chemicals	NO	41.5% (2.0%)
Paper and Printing	NO	38.8% (1.8%)
Wood Products	YES	47.4% (1.4%)
Apparel	NO	58.2% (1.3%)

Notes: The "Stationary Distribution" column compares the distribution of productivity between 1979-1982 and 1983-86 using the distribution-free Kolmogorov-Smirnoff test with the null that both empirical samples are drawn from the same underlying distribution. The test is performed at the 99% confidence level and a "YES" indicates that the null is accepted.

The "Increasing Productivity" statistic is the empirical percentage of the number of adjacent firm-year observations that show an increase in productivity. The bootstrapped standard errors are in parentheses.

TABLE 5				
Productivity Changes in Millions of 1980 Pesos				
	Average	OP/LP	OLS	Difference in
	Value Added	Productivity Change	Productivity Change	Productivity Change
		Mil. Pesos	Mil. Pesos	Percent of VA
Metals	9.62	1.17	1.83	6.9
Textiles	10.23	2.97	2.97	0.0
Food Products	41.90	16.23	20.78	10.8
Beverages	12.20	0.38	1.70	10.8
Other Chemicals	14.09	-4.74	-4.45	2.0
Printing & Pub.	8.07	-5.06	-4.30	9.3
Wood Products	8.03	1.11	1.88	9.7
Apparel	5.94	3.53	3.70	2.9

TABLE 6								
Productivity Decompositions								
	Metals		Textiles		Food Products		Beverages	
	'79-'83	'83-'86	'79-'83	'83-'86	'79-'83	'83-'86	'79-'83	'83-'86
Total Change	17.2	-20.3	22.1	1.2	-2.9	26.1	-17.1	-19.8
Productivity	4.7	-23.6	5.3	-2.3	-6.2	3.8	-120.4	-27.7
Reallocation	11.4	0.9	18.1	1.5	1.5	10.7	-15.1	29.5
Entry	2.8	7.9	1.3	4.5	4.0	13.5	123.8	0.0
Exit	-1.7	-5.5	-2.6	-2.5	-2.1	-1.9	-5.4	-21.6
Avg P Index1	66.2	76.9	72.0	87.6	26.6	37.6	353.9	334.2
Avg P Index2	79.2	82.8	88.2	84.9	27.4	42.3	336.5	320.3
	Other Chemicals		Printing & Publishing		Wood Products		Apparel	
	'79-'83	'83-'86	'79-'83	'83-'86	'79-'83	'83-'86	'79-'83	'83-'86
Total Change	-40.6	-32.5	-3.7	-0.7	-0.9	-13.2	21.0	131.3
Productivity	-46.8	-41.0	-4.0	-0.9	-6.7	-14.8	-15.5	150.8
Reallocation	7.1	10.4	0.3	0.2	-1.0	6.3	41.3	-29.8
Entry	4.6	2.2	0.2	0.2	10.2	2.4	1.8	29.8
Exit	-5.5	-4.1	-0.2	-0.2	-3.4	-7.1	-6.6	-19.5
Avg P Index1	126.6	94.1	5.7	3.4	33.0	26.1	120.6	215.1
Avg P Index2	137.4	97.8	6.3	3.7	31.7	30.0	162.1	205.4

Notes: Change in Industry Productivity is defined as the difference between last- and first-year industry productivity, where industry productivity is the sum over plants of plant productivity (ω) times share of value-added. The "Reallocation" component of the total change is the weighted sum over *continuing* plants of final-year productivity, with each plant's change in market share as its weight. The "Productivity" component is the weighted sum over continuing plants of final-year minus initial-year productivity, with each plant's initial-year market share as its weight. The "Entry" component is the sum over plants that enter between the initial and final years of final-year productivity times market share, and the "Exit" component is the sum over plants that exit between the initial and final years of initial-year productivity times market share. Average P Index 1 is the average of initial-year and final-year industry productivity. The market shares of firms that appear or disappear between the initial and final year are re-allocated to continuing firms. Average P Index 2 is the average over all years between the initial and final year (inclusive) of industry productivity, with no adjustments made for appearing and disappearing firms.

TABLE 7								
OLS Productivity Decompositions								
	Metals		Textiles		Food Products		Beverages	
	'79-'83	'83-'86	'79-'83	'83-'86	'79-'83	'83-'86	'79-'83	'83-'86
Total Change	28.17	-35.78	38.04	3.12	-7.21	52.04	-15.83	-3.46
Productivity	11.50	-37.00	10.79	4.56	-15.37	6.46	-43.54	-8.87
Reallocation	15.41	-0.29	29.58	4.17	3.44	23.73	-4.03	15.86
Entry	3.51	9.50	2.16	7.57	9.44	25.63	33.72	0.01
Exit	-2.25	-7.98	-4.49	-4.05	-4.71	-3.79	-1.98	-10.46
Avg P Index1	84.79	99.49	121.67	148.95	61.0	82.04	117.74	107.28
Avg P Index2	102.92	108.68	149.37	143.90	61.46	90.55	107.39	101.46
	Other Chemicals		Printing & Publishing		Wood Products		Apparel	
	'79-'83	'83-'86	'79-'83	'83-'86	'79-'83	'83-'86	'79-'83	'83-'86
Total Change	-32.14	-27.29	-51.44	-9.13	9.05	-102.51	41.57	72.77
Productivity	-36.70	-34.58	-55.42	-9.94	-29.00	-76.58	-7.23	71.14
Reallocation	5.24	8.86	3.92	1.31	-3.58	22.96	55.41	-7.55
Entry	3.76	1.91	1.52	1.13	58.20	10.60	2.39	39.97
Exit	-4.44	-3.48	-1.46	-1.63	-16.57	-59.50	-8.99	-30.79
Avg P Index1	105.47	79.43	90.24	60.34	178.51	131.57	157.34	233.37
Avg P Index2	114.64	83.07	104.51	64.68	162.12	157.90	208.92	238.84

Notes: See notes to Table 6.