Measuring and Mismeasuring Industry Productivity Growth Using Plant-level Data

by

Amil Petrin
University of Chicago, GSB
National Bureau of Economic Research
and
James Levinsohn
University of Michigan
National Bureau of Economic Research

Current version: December 20, 2004

Abstract. We show that the most popular index for measuring productivity growth with plant-level data has two major shortcomings: it adds a "reallocation" term to the theoretically grounded growth accounting measure, and it fails to use the correct weights (from Hulten (1978)) in the aggregation. Empirically, even when the correct weights are used, the "reallocation" term is substantial, leading to a weak relationship between the popular measure and the growth accounting measure for almost every manufacturing industry in both Chile from 1987-1996 and Columbia from 1981-1991. These findings are robust to many different estimation approaches for plant-level productivity, and they call into question the literature's interpretation of the "reallocation" term as productivity growth. We provide a new method for separating real productivity growth from reallocation effects that is entirely based on decomposing changes in the traditional growth accounting measure. In contrast to current findings that reallocation effects vary in sign and magnitude across time and sector, our new measure suggests that reallocation effects are reasonably stable within industries and almost always positively impact the productivity growth rate, even in instances where aggregate productivity falls.

Measuring and Mismeasuring Industry Productivity Growth Using Plant-level Data

Amil Petrin
University of Chicago, GSB
National Bureau of Economic Research
and
James Levinsohn
University of Michigan
National Bureau of Economic Research

1. Introduction

As recently as a decade ago, most estimates of industry-level productivity growth were obtained from industry-level data. With the increasing availability of plant-level data, there is a large and rapidly growing body of research that estimates plant-level productivity and then aggregates to the industry level.[1] There are of course many benefits to the plant-level data. From a practical viewpoint, plant-level data facilitate investigation of just *why* industry productivity changes as well as inquiries into correlates of higher plant-level productivity (e.g. export status, labor force size, plant vintage.) From a theoretical viewpoint, the plant-level data allow researchers to avoid assuming the existence of an industry-level production function, where any redistribution of inputs across plants must result in the same level of industry output.

Typically, plant-level productivity is measured using the total factor productivity (TFP) residual ($ln\omega$), computed as (log) output (lny) minus the contribution of inputs ($\beta' lnx$), or

$$ln\omega_{it} = lny_{it} - \beta' lnx_{it},$$

A predecessor to this paper appeared as "When Industries Become More Productive, Do Firms?: Investigating Productivity Dynamics," and is available as NBER Working Paper 6893. Ivan Kandilov, Wendy Petropoulos and Jagadeesh Sivadasan provided splendid research assistance and many good ideas. We are grateful to the Russell Sage Foundation for support and for suggestions from NBER Productivity, Susanto Basu and John Haltiwanger. Stata code for all computations is available from the authors by request.

[1] The data sets used include U.S. data from the Longitudinal Research Database (LRD) of the U.S. Census, French data from the Declarations Annuelles des Salaires (DAS) collected by INSEE (Institut National de la Statistique et des Etudes Economiques), and several plant-level manufacturing censes from developing countries, to name but a few. Examples of papers using the U.S. data include Bailey, Hulten, and Campbell (1992), Olley and Pakes (1996), Bernard and Jensen (1999), Bernard, Eaton, Jenson, and Kortum (2000), and Foster, Haltiwanger, and Krizan (2001). Abowd, Kramarz, and Margolis (1999) and Eaton, Kortum, and Kramarz (2001) use the French data while the several papers in Roberts and Tybout (1996) use the LDC data. These are but examples. A careful bibliography would include dozens of papers using plant-level data.

where the production function coefficients β can be estimated in a variety of ways. In this paper, we ask (and answer) "What is an appropriate way to aggregate these residuals to the industry level?" The literature almost exclusively employs some variant of the Bailey et al. (1992) index (the "BHC Index"), where the growth rate of industry productivity from period $t-1$ to period t is measured as

$$\sum_i s_{it} ln\omega_{it} - \sum_i s_{i,t-1} ln\omega_{i,t-1}, \tag{1}$$

with s_{it} usually denoting the plant's share in either gross output or employment (see Foster et al. (2001), and the extensive literature summarized therein).[2]

Without an underlying model, it is difficult to directly interpret the magnitude of changes in a productivity index. Some practitioners have pointed to the growth accounting framework to motivate (1);[3] these numbers are interpretable in many familiar circumstances, making them more readily comparable across time, industries, and countries.[4] One formulation gives instantaneous productivity growth as

$$d\omega = \sum_{i=1}^{N} s_{v_i} d\omega_i, \tag{2}$$

where $d\omega_i$ denotes the instantaneous change in the residual from the value added production function, and share weights s_{v_i} are in terms of a plant's contribution to industry value added (see Hulten (1978)). Some approximation with (discrete) data is needed, and many in the growth accounting literature have advocated the use of the Tornquist approach,[5] which multiplies the change in productivity growth from $t-1$ to t, as approximated by $\Delta ln\omega_{it} = ln\omega_{it} - ln\omega_{i,t-1}$, by the average "share" from the beginning and ending period, or

$$\hat{d\omega} = \sum_i \frac{(s_{v_{it}} + s_{v_{i,t-1}})}{2} * \Delta ln\omega_{it}. \tag{3}$$

[2] Olley and Pakes (1996) suggest the use of a related form where productivity enters in levels so the index is $\sum_i s_{it}\omega_{it} - \sum s_{i,t-1}\omega_{i,t-1}$.

[3] Bailey et al. (1992) report that "industry growth rates calculated... (using (1)) agree reasonably well with the growth rates calculated by Wayne Gray from aggregate industry data". Similarly, Foster et al. (2001) report that (1) "yields industry-level growth rates in productivity that correspond closely to industry-level growth rates constructed using industry-level data". The absence of some other established definition of productivity growth leads us to assume these references are to the growth accounting measure.

[4] For example, under perfect competition productivity as measured by growth accounting standards gives the change in the production possibilities frontier holding primary inputs constant, which in turn equals the change in welfare that arises as society's ability to consume and invest increases in response to the plant-level technology shocks.

[5] For example, Diewert (1976) shows that it provides a second order approximation to an arbitrary linear homogeneous function, and is exact for homogeneous translog functions. More generally, Trivedi shows that the Tornquist approach has an approximation error on the order of the square of the time interval.

A major theme of this paper is that the BHC index, (1), poorly approximates the growth accounting index (3) for two reasons. First, (1) is equal to (3) (i.e. $\hat{d\omega}$) *plus* an additional "reallocation"-type term – a point illustrated below. When the extra term is empirically important, the two indexes will not approximately equal one another. Second, in practice, (1) is rarely used with the correct aggregation share weights. These results have implications for the many literatures that use productivity growth estimates to guide policy, including (for example) macroeconomics, trade, industrial organization, and regulation. In particular, the collective wisdom of the effects of policy change on productivity growth comes from aggregating comparable results across many studies, and this requires common units, something the growth accounting framework provides, but (1) does not.

On the first point, (1) can be decomposed as

$$\begin{aligned}\sum_i s_{it} ln\omega_{it} - \sum_i s_{i,t-1} ln\omega_{i,t-1} &= \hat{d\omega} + \sum_i \frac{(ln\omega_{it} + ln\omega_{i,t-1})}{2} * (s_{it} - s_{i,t-1}) \\ &= \hat{d\omega} + \sum_i \overline{ln\omega_i} * \Delta s_{it},\end{aligned} \tag{4}$$

where $\overline{ln\omega_i}$ is the average $ln\omega_i$. (as noted in Fox (2003), who uses Bennet (1920)). The confounding term, $\sum_i \overline{ln\omega_i} * \Delta s_{it}$, resembles what is often characterized as a "reallocation" effect; it is the sum across plants of the change in plant share (however defined), multiplied by the geometric average of plant productivity.[6]

How much error to the productivity growth estimate can this "reallocation" term add?[7] We consider the basic case given in Hulten (1978), with perfect competition and no entry or exit. Here, one cannot sign $\sum_i \overline{ln\omega_i} * \Delta s_{it}$, nor is there any apparent bound on its value, as movements in the shares and productivity levels are not restricted by the underlying growth accounting framework.

In any data, $\sum_i \overline{ln\omega_i} * \Delta s_{it}$ must be negligible if (1) is to track (3) closely.[8] One way to ascertain whether (1) tracks (3) closely is to examine whether $\sum_i \overline{ln\omega_i} * \Delta s_{it}$ is about equal to zero

[6] Thus, defining productivity growth as (1), breaking out $\sum_i \overline{ln\omega_i} * \Delta s_i$, and attributing a "growth" interpretation to it is misleading without a model that tells us what it is this quantity measures. Similarly, it is also misleading to compare its magnitude to the magnitude of the productivity term to infer the roles of productivity and reallocation in productivity growth.

[7] The second reason (1) incorrectly calculates productivity growth is that it uses labor shares or gross output shares in the weighting during aggregation. If the residual is estimated using the gross output production function, the correct weight is plant gross output divided by industry value added. If the residual is estimated using the value added production function, the correct weight is plant value added divided by industry value added (as in (2)). In neither case is labor share or gross output share the appropriate weight. For example, productivity growth is biased towards zero by using gross output shares with gross output residuals, as one effectively multiplies through by the ratio of industry value added to industry gross output, a ratio that is often equal to 1/2 in manufacturing data.

[8] Anecdotally, the volatility in value added shares generally observed in plant-level data suggests that this term could be quite substantial.

for those plants that are present in both period t and period $t-1$.[9] We test the null hypothesis $\sum_i \overline{ln\omega_i} * \Delta s_{it} = 0$ for these plants, which constitute approximately 95% of the observations in the data on Chilean and Columbian manufacturing plants that we use.[10]

The evidence is striking. In almost every 3-digit Chilean and Columbian manufacturing industry the BHC index and the growth accounting index differ dramatically, as $\sum_i \overline{ln\omega_i} * \Delta s_i$ is large and volatile. Industry-by-industry regressions of (1) on (3) yield slope estimates that are almost entirely less than 0.5, with most being in the range of 0.2 to 0.4, meaning that (1) typically overstates (3) by several orders of magnitude. Furthermore, it is not just a matter of rescaling this mismeasured index; most of the r-squareds from these regressions range from between 0.2 to 0.4. Overall, these findings are robust to a number of different methods of estimating plant-level TFP, including ordinary least squares, Solow's approach, and a proxy approach, and to using discrete time approximations other than (3) to approximate (2), including using the base period (Laspeyres) shares or ending period (Paasche) shares as weights.

There is no separate "reallocation" component to productivity growth in (3), yet reallocation of inputs from less to more productive plants can surely occur. The absence of a separate term reflects the way that the productivity residual is calculated in the growth accounting literature and not a refutation of the concept itself. As discussed below, one can show that the changes in inputs across plants that might be expected to comprise a reallocation term are themselves netted out of the productivity calculation. This is why a separate reallocation term does not appear in (3). For the base case model with perfect competition, we show that it is possible to decompose the *change* in the growth accounting measure of productivity into three terms, one term representing the change in industry growth from changes in plant-level productivity growth, one term representing change in growth from the reallocation of value added across plants with differing productivity levels, and one term representing a net entry effect. Our empirical results suggest that a reallocation effect so defined is almost always present, economically important, reasonably stable, and almost always works to increase the growth rate in industry productivity, *even in the instances where the growth rate in industry productivity falls.* These findings stand in stark contrast to the literatures' findings that, using (1), "reallocation" varies dramatically in magnitude and sign both over time and across sectors when measured by some variant of $\sum_i \overline{ln\omega_i} * \Delta s_i$.

The remainder of this paper is organized as follows. In the next section, we discuss the growth accounting framework. Section 3 illustrates how one can implement the growth accounting approach

[9] The Tornquist approximation is not defined for plants that enter or exit in the year they enter or exit.

[10] We use value added share weights with value added residuals for both measures to control for the second problem.

using plant-level data. Section 4 introduces a new decomposition for changes in productivity growth. Section 5 describes the data and estimation, and Section 6 discusses the results.

2. Growth Accounting Theory

Historically, TFP growth has been defined as the change in output holding primary inputs constant.[11] A first order approximation to the rate of growth of output at plant i, dy_i, is given by the differential equation

$$dy_i = \beta_{l_i} dl_i + \beta_{k_i} dk_i + \beta_{m_i} dm_i + dt_i, \quad (5),$$

where β_{j_i} denotes the elasticity of output with respect to input j, dj_i is the rate of growth of input j, and the Hicks neutral technology shock is given by dt_i. Following Solow (1957), in a competitive environment optimizing behavior allows the growth in output to be reexpressed equivalently as

$$dy_i = s_{l_i} dl_i + s_{k_i} dk_i + s_{m_i} dm_i + dt_i \quad (6),$$

with s_{ij} denoting the input's revenue share. With plant-level data, there are many observations on dt_i and principles for aggregation across plants become important.

One guiding principle of the growth accounting method for aggregating plant-level technology shocks was advocated by Domar (1961):

> We should be free to take the economy apart, to aggregate one industry with another, to integrate final products with their inputs, and to reassemble the economy once more and possibly over different time units without affecting the magnitude of the Residual.

A second related principle, described in Hulten (1978), is that productivity growth should measure the *impact on final demand* of increasing (or decreasing) a plant's factor efficiency. In particular, when some of plant i's output is used as intermediate input at other plants, an increase in dt_i leads both to an increase in final demand for i *and* to an increase in i's intermediate deliveries. These new intermediate input deliveries then increase output at the plants to which they go (where they are used in production). This new output will both increase final demand and fulfill more intermediate deliveries elsewhere. The process continues. In the end, the greater the role of intermediate deliveries in the economy, the larger the impact of any increase in plant TFP on final demand.

When this latter point is taken together with the aggregation/disaggregation condition from Domar, the aggregate impact on final demand $d\omega$ brought about by changes in plant-level technical efficiency is given by:

$$d\omega = \sum_{i=1}^{N} \frac{P_i Q_i}{\sum_i P_i^V V_i} dt_i,$$

[11] Hulten (2001) provides a thoughtful and detailed history of this index.

where each dt_i is weighted by i's gross output divided by industry value added. $\frac{\sum_{i=1}^{N} P_i Q_i}{\sum_i P_i^V V_i} \geq 1$, reflecting the effect of the additional intermediate input on final consumption and/or investment. So defined, $d\omega$ is the rate of productivity growth as defined by the growth accounting literature. It measures the rate of change of the social production possibility frontier, holding primary inputs constant. In a competitive environment with constant returns to scale this measure of aggregate technical change is exactly equal to the change in welfare that arises as society's ability to consume and invest increases in response to the plant-level technology shocks.

As noted earlier, $d\omega$ can also be expressed in terms of value added shares and value added residuals, given as $\sum_{i=1}^{N} s_{v_i} d\omega_i$. Hulten calls the growth rate of the value added residual the "effective" rate of productivity growth because - when weighted by the share of plant-level value added in the aggregation - it directly measures the aggregated impact on final demand of changes in technical efficiency occurring at the plant-level. The derivation of this index begins at the plant-level, where value added is defined as

$$dv_i = \frac{dy_i - \beta_{m_i} dm_i}{1 - \beta_{m_i}} = \frac{\beta_{l_i} dl_i + \beta_{k_i} dk_i + dt_i}{1 - \beta_{m_i}};$$

the contribution of intermediate input growth to output growth is first deducted, and the remaining output growth is grossed up by $\frac{1}{1-\beta_{m_i}}$ to account for the additional intermediate input available from the increased factor efficiency (Hulten (1978)).[12] So defined, the rate of growth of industry value added is given as $\sum_{i=1}^{N} s_{v_i} dv_i$. The primary input index for the value added function is given as

$$dx_i^V = \frac{\beta_{k_i}}{1 - \beta_{m_i}} dk_i + \frac{\beta_{l_i}}{1 - \beta_{m_i}} dl_i,$$

and the effective rate of productivity growth $d\omega_i$ is the difference between the growth in value added and the growth in the primary input index,

$$d\omega_i = dv_i - dx_i^V = \frac{dt_i}{1 - \beta_{m_i}}. \quad (7)$$

The aggregate input is then $\sum_{i=1}^{N} s_{v_i} dx_i^V$, and aggregate productivity growth is the difference in the rate of growth of industry output and aggregated primary inputs:

$$d\omega = \sum_{i=1}^{N} s_{v_i} d\omega_i = \sum_{i=1}^{N} s_{v_i} dv_i - \sum_{i=1}^{N} s_{v_i} dx_i^V.$$

[12] An alternative way to view the grossing up is that it ensures when the growth in output and intermediate inputs are equal, the growth in value added equals this rate. For this approach to be applicable, a "value-added" production function must exist. Bruno (1978) shows that this requires that intermediate inputs are separable (in the gross output production function) from capital, labor, and the productivity shock. While value added is the appropriate measure for growth accounting, "Real value added is an artificial construct - bread without flour; books without paper or ink; shoes lacking leather," as (Basu and Fernald (2002)) describe it.

Reallocation Effects

An important feature of the Hulten growth accounting setup is that there is no "reallocation effect" present in the rate of productivity growth. The growth accounting measure does recognize that inputs and outputs are reallocated in response to technology shocks, and that these reallocations do affect final demand, but the change in final demand that arises because the distribution of primary and intermediate inputs changes is controlled for in the computation of the productivity residual. Alternatively, the aggregate productivity residual $d\omega$, to be consistent with the guiding principles listed above, is defined so that these effects are absent. In Section 4 we show that an alternative is to look at *changes* in the rate of productivity growth, which decompose exactly into a reallocation term that is based on changes in value added shares, and a real productivity and net entry term.

In imperfectly competitive markets, the productivity residual no longer just measures technical efficiency, as shown in Hall (1990) and Basu and Fernald (2002). Under cost minimization plant-level value added is given as

$$dv_i = \mu_i^V dx_i^V + \frac{dt_i}{1 - \mu_i s_{m_i}}$$

where $\mu_i^V = \mu_i \frac{1-s_{m_i}}{1-\mu_i s_{m_i}}$ is the "value added markup", an increasing function of the markup $\mu_i = P_i/MC_i$. The productivity residual is then

$$\begin{aligned} d\omega_i &= dv_i - dx_i^V \\ &= (\mu_i^V - 1)dx_i^V + (\mu_i^V - 1)\frac{s_{m_i}}{1 - s_{m_i}}(dm_i - dy_i) + \frac{dt_i}{1 - \mu_i s_{m_i}} \end{aligned}$$

The final term is the effective rate of technical progress under imperfect competition, and is increasing in the plant-specific markup; the higher the markup, the more valuable the output. Two new terms are also present. The first term arises because only dx_i^V is deducted from dv_i (instead of $\mu_i^V dx_i^V$). The second term arises because the productive contribution of intermediate inputs exceeds the revenue share (or β_{m_i}) by the markup, and the revenue share is what is used to deduct intermediates from gross output in the move to value added. Note that, in the case when markets are competitive, μ_i and μ_i^V both equal one, and $d\omega_i = \frac{dt_i}{1-s_{m_i}}$, the effective rate of productivity for growth accounting under perfect competition, which is a special case of this more general framework.

Aggregate productivity growth under imperfect competition can be written as

$$d\omega = \sum_{i=1}^{N} s_{v_i} d\omega_i + R_l + R_k,$$

where

$$R_l = \sum_{i=1}^{N} s_{v_i} \frac{s_{l_i}}{1 - s_{m_i}} \frac{P_{l_i} - P_l}{P_{l_i}} dl_i$$

with P_{l_i} defined as the shadow value of labor at i (P_l is the average across i), and R_k is symmetrically defined for capital. The two new terms reflect differences across plants in the shadow values of the primary inputs labor and capital (from the average shadow value in the economy). Again, if markets were competitive, $d\omega$ would reduce to (2), because the shadow values would be identical across plants and exactly equal to wages and rental rates, so R_l and R_k would equal zero.

In total there are four types of "reallocation" effects, that is, four types of changes in measured productivity that reflect reallocations of inputs across plants (and not direct changes in technical efficiency). We describe them each in turn. Holding aggregate inputs constant, productivity increases if there is a reallocation of inputs from low markup plants to high markup plants. Similarly, holding plant-level inputs constant, if the technology shocks cause intermediate inputs to be used more intensively at plants with markups, this too leads to an increase in aggregate productivity. R_l and R_k relate directly to the reallocation of labor and capital across plants with differing shadow values. Generally, productivity increases if there is a reallocation of inputs from plants with lower shadow values to plants with higher shadow values.

Hulten (1978) shows, under perfect competition, that the increase in the value of output from changes in technical efficiency exactly equals the change in welfare. Basu and Fernald show, under imperfect competition, that even when productivity growth does not exclusively measure the change in output from changes in technology, aggregate productivity change still exactly equals the change in welfare if a representative consumer model is a reasonable approximation to the expenditure side of the economy. The intuition for the result is that the ratio of market prices reflects consumers marginal rate of substitution between goods, so changes in value added reflect (to first order) the change in consumer well-being, even though the marginal rate of transformation between goods is not exactly equal to this ratio.

3. Estimating Productivity Growth

Many approaches are available for estimating plant-level TFP. We discuss two prominent and complementary methods. Each one requires the existence of a plant-level production function for the measurement of productivity growth to be a well-defined exercise. Each uses different restrictions to estimate the parameters of the production function.

One method begins with an assumed functional form for the production function and then directly estimates its parameters β_{ik} using (monotonic transformations of) levels. One can use ordinary least squares or one of many alternatives that attempt to address the simultaneity of input choices and productivity raised by Marschak and Andrews (1944). [13]

The second approach uses the Solow (1957) insight that optimizing behavior implies observed revenue shares are a consistent estimator for the elasticity of output with respect to any input. While this approach is sometimes pointed to as being non-parametric for estimating the production function parameters, Hulten (1973) shows that (6) is useful for productivity analysis only if there exists an underlying production function. Otherwise, the line integral that defines productivity growth is not necessarily unique.[14]

Both the plant-level productivity measure and the aggregate industry growth measure are defined in terms of instantaneous changes (they are Divisia indexes). Generically, we can write them as

$$\sum_i \alpha_i dlnx_i,$$

where α_i are the weights in the aggregation of growth rates (of either input levels or productivity levels).[15] Some approximation with the (discrete) data is necessary, and the most popular for Divisia indexes are given by log-change indexes, which in these cases apply the log-change in either inputs or productivity levels to a share weight (α_i) and then aggregate:

$$\sum_i \alpha_i (lnx_{it} - lnx_{i,t-1}).$$

This is the integral from $t-1$ to t of the Divisia index $\sum_i \alpha_i dlnx_i$ if α_i is constant.[16] Most often the share weights are given as the average of the beginning period share $s_{i,t-1}$ and the ending period

[13] Alternatives include (for example) instrumental variables, fixed effects, and the proxy methods of Olley and Pakes (1996) and Levinsohn and Petrin (2003).

[14] Line integrals typically depend on the path of integration. Hulten develops the link between path independence and the theory of aggregation. For the measurement of plant-level TFP growth to be uniquely defined a linear homogenous plant-level production function must exist (the aggregate of the plant inputs), and each input price vector must correspond to a unique input vector. The consequence of failing path independence is that the index value for a plant that started and ended a period with the same output, revenue shares, and input levels, could have a non-zero change in measured productivity growth (instead of zero productivity growth). In this case this path could be cycled over, yielding arbitrarily large (or small) values of the index.

[15] In the case of measuring plant-level productivity $d\omega_i$ the input index is constructed with α_i as the production function parameter, $dlnx_i$ denoting the growth rates of input, and the sum is then taken over all inputs used in production. In the case of aggregating productivity growth the α_i are the value added shares, $dlnx_i$ denotes productivity growth at plant i, and the sum is taken across all N plants.

[16] This does not address the more fundamental question of when the Divisia index of productivity growth is path independent, which is explored in Petrin (2004), who extends the conditions in Hulten (1973) to micro-level data. Productivity growth from time 0 to time T is given by the line integral of $d\omega_t$ from 0 to T, or

$$D(\Gamma) = exp\{\int_\Gamma \sum_{i=1}^{N} s_{v_{it}} d\omega_{it}\},$$

share s_{it}, so $\alpha_i = (s_{it} + s_{i,t-1})/2$ (Tornquist), although ending period shares $\alpha_i = s_{it}$ (Paasche), and starting period shares $\alpha_i = s_{i,t-1}$ (Laspeyres) are also used in the literature.

Tornquist (1936), Theil (1967), Hulten (1973), Diewert (1976), Star and Hall (1976), and Trivedi (1981) have all argued for use of the Tornquist approximation to Divisia indexes. Specifically, Diewert (1976) shows that the Tornquist approximation to the Divisia index is exactly correct if the underlying function has the homogeneous translog form.[17] When the underlying function is not of this form, the Tornquist approximation remains attractive because the translog provides a reasonable second order approximation to many other functions. Trivedi (1981) generalizes Diewert's findings, developing approximation results without specifying an assumption about the underlying form of the function. Using numerical analysis he shows that the approximation error is on the order of the square of the length of the time interval, a result not available for other approximations like Paasche or Layspeyres. For these reasons, our preferred results use the Tornquist approximation.[18]

One important issue that does arise with the move from industry- to plant-level data is that industries do not enter and exit, while plants do. This raises one difficulty in calculating any index that is based on using $ln\omega_{it} - ln\omega_{i,t-1}$ to approximate the average percentage change in productivity; for a plant that enters (exits), $ln\omega_{i,t-1}$ ($ln\omega_{it}$) is not measurable in the year that they enter (exit). For data observed at higher frequencies (e.g. annual), the truncation problem that arises may not be that severe, because entrants or exiters in any one year are likely to make up

where Γ is the curve traced by $\omega(t) = (\omega_{1t}, \omega_{2t}, \ldots, \omega_{Nt})$, $0 \leq t \leq T$. Line integrals are not typically path independent, so a researcher starting with $\sum_{i=1}^{N} s_{v_{it}}(dv_{it} - dx_{it}^{V})$ and integrating from 0 to T will not generally recover the change in aggregate output resulting from plant-level TFP shocks (holding primary inputs fixed). One must establish that the choice of the path Γ does not affect the value of the index, or the index may suffer from cycling, where arbitrarily large or small values of the index can be associated with any point.

[17] Diewert (1976) uses the quadratic approximation lemma to single out the Tornquist-Divisia index as most preferred for general applications of productivity measurement. Diewert characterized an index that was exact for a underlying translog (e.g. production) function as "superlative." With constant returns to scale this Tornquist approximation to the Divisia index is the *only* superlative index. With non-constant returns to scale, it remains superlative, but other superlative indexes may exist (although to our knowledge none are known).

[18] Similarly, when production function parameters are changing over time at the plant level, plant-level productivity growth is given as

$$ln\omega_{it} - ln\omega_{i,t-1} = lnv_{it} - lnv_{i,t-1} - \sum_{i=1}^{K} \frac{s^*_{x_{it}} + s^*_{x_{i,t-1}}}{2}(lnx_{it} - lnx_{i,t-1}) \qquad (8),$$

where x_{it} is the level of input i used at time t, K is the number of inputs, and the * denotes the value added share or coefficient, which can be equivalently expressed replacing the shares $s^*_{i.}$ with the elasticities $\beta^*_{i.}$ In terms of the original shares or production function coefficients, $s^*_i = \frac{s_i}{1-s_{im}}$ and $\beta^*_i = \frac{\beta_i}{1-\beta_{im}}$.

only a small fraction of value added (this is true in our annual data). For data observed at lower frequencies, like every five or ten years, the truncation problem can be more severe, and some form of imputation for these values may be desirable. This type of a correction ultimately requires one to directly model the simultaneous productivity and entry/exit process. In all of these cases, the fraction of value added that is ignored in aggregation is directly computable, so inspection of the data gives some sense of the problem's magnitude.

Many economic questions relate to the performance of entrants or exiters. The truncation problem does not affect the ability to track entrants/exiters in higher frequency data. Except for the entering/exiting year, one can compute these plants contribution to aggregate productivity growth.[19] In fact, one principle of the growth accounting measures is that we are free to group plants in any subaggregates that we desire without affecting the measure of aggregate productivity growth.

4. Decomposing Productivity Growth

In Section 2 we described how, under perfect competition, there is no "reallocation" term associated with the growth accounting measure of productivity. In this section, we show that decomposing the *change* in the growth rates introduces a reallocation term. We also review the BHC decomposition.

The approximation to the change in the growth rate uses data from three juxtaposed time periods, which we denote as t, $t+1$, and $t+2$. The change in the growth rate is given by:

$$\sum_{i=1}^{N_2} \frac{s_{v_{i,t+2}} + s_{v_{i,t+1}}}{2}(ln\omega_{i,t+2} - ln\omega_{i,t+1}) - \sum_{i=1}^{N_1} \frac{s_{v_{i,t+1}} + s_{v_{it}}}{2}(ln\omega_{i,t+1} - ln\omega_{it}). \qquad (10)$$

Equation (10) decomposes into a "productivity" term, a "reallocation" term, and a "net entry" term. The set of plants that exist in t, $t+1$, and $t+2$ is denoted C (for continuers), and each of these plants contributes a productivity component and a reallocation component to the growth in aggregate productivity. Plants that exist in t and $t+1$, or $t+1$ and $t+2$, contribute to the overall change in (10), but only through a net entry term.

One term is the productivity term. Aggregated across plants it is written as

$$\sum_{i \in C} \frac{s_{v_{i,t+2}} + 2 * s_{v_{i,t+1}} + s_{vit}}{4} * (ln\frac{\omega_{i,t+2}}{\omega_{i,t+1}} - ln\frac{\omega_{i,t+1}}{\omega_{it}}).$$

[19] Categories for recent entrants and/or upcoming exiters can be defined separate from continuing plants, their fraction of productivity growth being then directly comparable to continuing plants.

Each continuer contributes the change in the rate of their productivity growth (the term on the right). The weight in the aggregation to the industry level is an average over the three time periods of the plant's share in value added with period $t+1$ getting twice the weight as period t and period $t+2$. A second term is the reallocation term, given as

$$\sum_{i \in C} (ln \frac{\omega_{i,t+2}}{\omega_{i,t+1}} + ln \frac{\omega_{i,t+1}}{\omega_{it}})/2 * (s_{v_{i,t+2}} - s_{vit})/2.$$

At the plant level the term on the right gives the change in the share of value added from period t to period $t+2$, and the weight in the aggregation is the average rate of growth in productivity over the two time periods for the plant (i.e. $ln(\frac{\omega_{i,t+2}}{\omega_{it}})/2$). The third term is the net entry term and is given as

$$\sum_{i \in t+1,t+2} \frac{s_{v_{i,t+2}} + s_{v_{i,t+1}}}{2} (ln\omega_{i,t+2} - ln\omega_{i,t+1}) - \sum_{i \in t,t+1} \frac{s_{v_{i,t+1}} + s_{v_{it}}}{2} (ln\omega_{i,t+1} - ln\omega_{it}).$$

The sum of these terms yields (10).

Ignoring the net entry term, if there is no change in the growth rate of productivity at any plant, the productivity term is zero, and the aggregate rate of productivity growth can only increase (decrease) if there is reallocation of the share in value added from period t to period $t+2$ towards (away from) plants with higher average productivity growth rates. Similarly, if there is no change in the shares of value added from the start (t) to the end ($t+2$), the reallocation term is zero, and aggregate productivity growth can only increase if the (share-weighted) average growth rate of productivity increases. Net entry contributes a term that is positive (negative) if the share weighted sum of the growth rates of entrants exceeds (falls below) the share weighted sum of the growth rates of the exiters.[20]

The BHC Decomposition

The most popular measure of reallocation is based on decomposing changes in the BHC index. One can decompose (1) as

$$\sum_i s_{i,t+1} ln\omega_{i,t+1} - \sum_i s_{it} ln\omega_{it} = \sum_{i \in C} s_{it} \Delta ln\omega_{it} + \sum_{i \in C} ln\omega_{it} \Delta s_{it} + \sum_{i \in C} \Delta s_{it} \Delta ln\omega_{it} + \sum_{i \in B} s_{i,t+1} ln\omega_{i,t+1} - \sum_{i \in D} s_{it} ln\omega_{it} \quad (11)$$

where C is the set of continuing plants, B the set of entrants, and D the set of exiters, and the difference operator, Δ, denotes the difference between year $t+1$ and t. The first term is interpreted

[20] Note that, with no entry and some exit, for example, this term will necessarily be negative.

as a measure of the productivity effect. Although not described as such when used, if s_{it} were to denote the value added share, this term could be viewed as a Laspeyres approximation to an underlying Divisia index of the growth accounting measure of productivity. The second term is designed to measure reallocation; if plants with larger total factor productivity residuals in the base period have a higher share in the end period, this term increases. The third term is the product of the change in the share and the change in productivity and as such is a measure of both productivity changes and reallocation effects (it measures how tied together the productivity and reallocation cases are in this metric). It is sometimes referred to as a covariance term. The impact of net entry is given by the sum of the fourth and fifth terms.[21]

As noted earlier, (1) is *in part* a measure of real productivity in the traditional growth accounting sense *only* when the share weights are value added (and not gross output or input share weights). If value added shares are employed, (1) in its entirety measures aggregate productivity growth plus additional terms.[22] In particular, the reallocation term in (11) is not a part of the growth accounting definition of aggregate productivity, so attributing a literal "growth" interpretation to

---

[21] Often in applied work not all four terms of this decomposition are reported. Instead, the typical approach reports three terms, resolving the ambiguity of the covariance term by assigning some fraction of it to the real productivity term, (say Φ), and the other (1 - Φ) to the reallocation term. The general expression is:

$$\left(\sum_{i\in C} s_{it}\Delta ln\omega_{it} + \Phi\sum_{i\in C}\Delta s_{it}\Delta ln\omega_{it}\right) + \left(\sum_{i\in C} ln\omega_{it}\Delta s_{it} + (1-\Phi)\sum_{i\in C}\Delta s_{it}\Delta ln\omega_{it}\right) + NE_t, \quad (12)$$

where NE_t is the net entry term. For example, if one combines the covariance term with the rationalization term ($\Phi = 0$), the resulting decomposition simplifies to:

$$\sum_{i\in C} s_{it}\Delta ln\omega_{it} + \sum_{i\in C}\Delta s_{it} ln\omega_{i,t+1} + NE_t. \quad (13)$$

Another commonly observed alternative folds the covariance term into the rationalization term (i.e. $\Phi = 1$):

$$\sum_{i\in C} s_{i,t+1}\Delta ln\omega_{it} + \sum_{i\in C}\Delta s_{i,t+1} ln\omega_{it} + NE_t \quad (14).$$

[22] If one were to choose the Laspeyres approximation, aggregate productivity is given as

$$\sum_i s_{i,t-1}(ln\omega_{it} - ln\omega_{i,t-1}) \quad (18)$$

and (1) adds

$$\sum_i (s_{it} - s_{i,t-1}) ln\omega_{it} \quad (19)$$

to the Laspeyres index. Similarly, if the Paasche approximation is used, one can show that the additional reallocation term is the plant-level change in share times the previous period's productivity level.

it is misleading. Similarly, comparing its magnitude to the magnitude of the productivity term to determine the relative roles of productivity and reallocation in aggregate productivity growth is not an exercise defined specifically in growth accounting terms.

5. Data and Estimation

We turn to two manufacturing censuses to explore the empirical issues that we raise. One census is from Chile's Instituto Nacional de Estadistica (INE), and the second is from Columbia's Departamento Administrativo Nacional de Estadistica (DANE.) The Chilean data span the period 1987 through 1996 and the Columbian data span the years 1981-1991. We focus on 3-digit industries with more than 200 observations, of which there are 23 in Chile and 26 in Columbia. Here, we provide a brief overview of these data. They have been used in numerous other productivity studies, and we refer the interested reader to those papers for a more detailed data description.[23]

The data are unbalanced panels and cover all manufacturing plants with at least ten employees. Plants are observed annually and they include a measure of output, two types of labor, capital, and intermediate inputs. Real value-added is nominal value added adjusted by the 3-digit industry price index. Labor is the number of man-years hired for production, and plants distinguish between their blue- and white- collar workers (we include two labor types in the production function). The method for constructing the real value of capital is documented in Lui (1991) for the Chilean data, and a similar approach is adopted for the Columbian data.[24] A data problem for the Chilean census is that approximately 3% of the plant-year observations appear to be "missing"; a plant id number is present in year $t-1$, absent in year t, and then present again in year $t+1$. We impute the values for these observations using $t-1$ and $t+1$ information (see the Appendix). Due to the way that the data are reported, we treat plants as plants, although there are probably multi-plant plants in the sample.

We estimate value added production function parameters for each of the 3-digit industries and use the parameters to estimate the plant-level TFP residuals. For any industry, the production function coefficients are assumed to be constant over time and across plants, although our findings

[23] See Lui (1991), Lui (1993), Lui and Tybout (1996), Tybout, de Melo, and Corbo (1991), Pavcnik (2002), Levinsohn (1998), and most recently Levinsohn and Petrin (2003).

[24] It is a weighted average of the peso value of depreciated buildings, machinery, and vehicles, each of which is assumed to have a depreciation rate of 5%, 10%, and 20% respectively. 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.

are robust to loosening this assumption. We employ three different approaches to estimating the coefficients: ordinary least squares, revenue shares, and the proxy method from Levinsohn and Petrin (2003), which includes controls to address the correlation of productivity with input choices. For the revenue shares, we use the average over plants and time. Details on the Levinsohn and Petrin (2003) estimator are relegated to the Appendix.

6. Results

We compare the growth accounting measure from (3) with the popular BHC index given in (1) using 49 3-digit manufacturing industries from Chile and Colombia. In an attempt to keep the analysis manageable, we start with a detailed description of results for the largest Chilean manufacturing industries. We then describe how these findings generalize. The main result is that the micro patterns observed in the largest industries in Chile are indicative of the findings for the entire 49 3-digit industries from both countries.

Table 1 reports the annual estimates of growth rates in productivity for the two measures for ISIC 311, the Food Products industry (the largest in Chile). The production function coefficients are estimated using ordinary least squares, and the calculations are done using only plants that exist in period $t-1$ and period t, which account for 94.4% of plant-year observations and 96.4% of industry value added over the sample period. Column 1 is industry value added for 1988 to 1996, column 2 is the growth accounting measure of productivity, column 3 is the BHC productivity measure, and column 4 is the difference between these two terms, which is equal to the BHC "reallocation" term described earlier.

The growth accounting measure averages 3.64% per annum, with standard deviation of 4.72% across the nine years. The BHC index averages -2.93% per annum with a standard deviation of 13.48%. The divergence in these summary statistics arises because the two indexes themselves are widely divergent, as is evident from a comparison of columns 2 and 3. Column 4 is the difference and has a mean of -6.57%, with a standard deviation of 10.74%. Its volatility across the sample period is consistent with the general findings in the literature that "reallocation" is large and volatile (the literature has been using some form of this term to measure reallocation effects).

For ISIC 311, table 2 compares estimates of productivity growth across different estimators for production function parameters. The top half of the table is the growth accounting measure and the bottom half is the BHC index. For the top half, column 2 is the same as column 2 from table 1, which uses ordinary least squares to obtain production function estimates. Column 3 uses the proxy approach from Levinsohn and Petrin (2003), and column 4 uses revenue shares. While the production function estimates (not reported here) do differ somewhat across the three approaches, the productivity growth numbers using the growth accounting index are reasonably similar across the three estimators. Only in 1995 does there seem to be some divergence between the revenue share estimate and the two alternatives.

For the BHC aggregator, the signs tend to be common across the three sets of production function estimates, but the magnitudes are quite different, with the LP and OLS estimates systematically the largest in absolute value terms. The main reason for the volatility is reflected in the "reallocation" terms, which are systematically more volatile for LP and OLS relative to revenue shares.

Table 3 summarizes the relationship between the growth accounting index and the BHC index for the eight largest industries in Chile. The industries (along with their ISIC codes) are Food Products (311), Metals (381), Textiles (321), Wood Products (331), Apparel (322), Plastics (356), Non-electric Machinery (382), and Other Chemicals (352). The index numbers use only plants that exist in $t-1$ and t, which account for between 94% and 98% of industry value added (in the annual data, truncation due to entry and exit is not a problem). For each industry and each estimator of production function coefficients, the appropriate row reports the intercept, slope, and r-squared from a regression of the 9 annual growth rates using the growth accounting measure as the dependent variable and the BHC aggregator as the explanatory variable. For example, the first row for ISIC 311 is the regression of column 2 on column 3 from table 1. An intercept of 0 and a slope of 1 means the indexes perfectly track one another. An r-squared of 1 means they perfectly covary, although their magnitudes can differ.

For industry 311, the r-squared's range from 0.22 to 0.47. The intercept is significantly different from zero and the slope is significantly different from one.[25] Overall, these results suggest that the BHC index is a poor proxy for the growth accounting measure for ISIC 311. The slope terms range from 0.12 to 0.31, so the BHC index overstates true growth by a factor of 8 (using the proxy approach) and 3 (using revenue shares).

These results are similar across all of these eight industries. R-squareds are typically low, ranging between 0.2 and 0.4. Intercept and slope coefficients are significantly different from zero and one respectively. The point estimates for the slope coefficients suggests that the BHC index overstates true growth by several magnitudes.

[25] These results do not correct for error in the parameter estimates. For this industry, with almost 10,000 observations, the parameter estimates are very precisely estimated.

The messages that come out of the results from tables 1-3 are confirmed by the other 15 manufacturing industries in Chile and the 26 manufacturing industries from Colombia. Overall, for Chile, using OLS only 4 of 23 industries had r-squareds above 0.5, using revenue shares only 5 of 23 had r-squareds above 0.5, and using the proxy approach only 1 in 23 industries had an r-squared over 0.5. Most slope coefficients across estimators and industries varied between 0.1 and 0.3.[26] For Colombia, using OLS only 4 of 26 industries had r-squareds above 0.5, using revenue shares 8 of 26 had r-squareds above 0.5, and using the proxy approach only 5 of 26 industries had an r-squared over 0.5. Slope coefficients across estimators and industries also varied between 0.1 and 0.3, suggesting the BHC index overstates by several magnitudes the true rate of productivity growth. In summary, the results demonstrate that the BHC aggregator adds a "reallocation" term to productivity growth that is large and volatile, making it a very noisy indicator of true growth. This might not be a big deal except the literature is rife with studies using the BHC aggregator and then drawing conclusions about the apparent importance (or not) of reallocation.

We now turn to the decomposition of productivity growth. As we showed earlier, there is no "reallocation" component in the instantaneous change in productivity. However, in Section 4 we showed how to decompose the change in the growth rate of productivity into a term that is increasing if growth rates at the plant level are increasing (the "real productivity" term), and a term that is increasing when plants with higher average growth rates gain larger shares of value added (the "reallocation" term).

Table 4 provides this decomposition for ISIC 311 in Chile. The change in the growth rate from year to year is reported in column 2. Note that it does not exactly equal change in the growth rate reported in column 1; the difference is due to the entry/exit plants, (those plants that do not exist in $t-2$, $t-1$, and t). As defined in Section 4, for those plants that exist in these three periods, the change in column 2 can be decomposed into the "real productivity" term, column 3, and the "reallocation" term, column 4. The real productivity component is quite volatile, averaging -5.95% with a standard deviation of 7.33%, with both positive and negative outcomes. The reallocation term is positive, very stable, and always contributes to increases in the rate of productivity growth, even in periods when the overall growth rate falls.

Overall, the reallocation series for ISIC 311 in Chile is remarkably similar in spirit to the reallocation series from the other 48 manufacturing industries. Regardless of the estimator for production function coefficients, the industry, or the country, the annual reallocation terms are

[26] The industries that had higher r-squareds were also the industries where the slope coefficients tended to be larger.

almost universally positive. Within an industry over time they vary very little, with a typical standard deviation less than 1%, while across industries they range from between 3% to 8%. Thus, in contrast to current findings that reallocation effects vary in sign and magnitude across time within industries, our measure suggests reallocation effects are very stable and almost always contribute positively to industry growth, even when the overall growth rate is falling.

7. Conclusions

We have shown that the most popular index for measuring productivity growth with plant-level data has two major shortcomings: it adds a "reallocation" term to the growth accounting measure, and it fails to use the correct weights from Hulten (1978) in the aggregation. Empirically, even when the correct weights are used, the "reallocation" term is substantial, leading to a weak relationship between the popular measure and the growth accounting measure for almost every manufacturing industry in both Chile from 1987-1996 and Columbia from 1981-1991. These findings are robust to many different estimation approaches for plant-level productivity, and they call into question the literature's interpretation of the "reallocation" term as productivity growth. We provide a new method for separating real productivity growth from reallocation effects that is entirely based on decomposing changes in the traditional growth accounting measure. In contrast to current findings that reallocation effects vary in sign and magnitude across time and sector, our new measure suggests that reallocation effects are reasonably stable within industries and almost always positively impact the productivity growth rate, even in instances where aggregate productivity falls.

References

Abowd, J., Kramarz, F., and Margolis, D. (1999). High wage workers and high wage firms. *Econometrica*, *67*(2), 251–334.

Bailey, M., Hulten, C., and Campbell, D. (1992). Productivity dynamics in manufacturing plants. In *Brookings Papers on Economic Activity: Microeconomics*, Vol. 4, pp. 187–267. Brookings Institute.

Basu, S., and Fernald, J. (2002). Aggregate productivity and aggregate technology. *European Economic Review*, *46*, 963–991.

Bennet, T. (1920). The theory of measurement of changes in cost of living. *Journal of the Royal Statistics Society*, 455–462.

Bernard, A., Eaton, J., Jenson, B., and Kortum, S. (2000). Plants and productivity in international trade. NBER Working Paper No. 7688.

Bernard, A., and Jensen, B. (1999). Exceptional exporter performance: Cause, effect, or both?. *Journal of International Economics*, *47*, 1–25.

Bruno, M. (1978). *An Analysis of the Concept of Real Value-Added, in Production Economics: A Dual Approach to Theory and Applications*, chap. III.1. North-Holland.

Diewert, E. (1976). Exact and superlative index numbers. *Journal of Econometrics*.

Domar, E. D. (1961). On the measurement of technological change. *The Economic Journal*, *71*(284), 709–729.

Eaton, J., Kortum, S., and Kramarz, F. (2001). An anatomy of international trade: Evidence from french firms. Boston University, Department of Economics.

Foster, L., Haltiwanger, J., and Krizan, C. (2001). *New Developments in Productivity Analysis*, chap. Aggregate Productivity Growth: Lessons from Microeconomic Evidence, pp. 303–372. University of Chicago Press.

Fox, K. (2003). Problems with (dis)aggregating productivity, and another productivity paradox. University of New South Wales Discussion Paper.

Hall, R. (1990). Invariance properties of solow's productivity residual. In Diamond, P. (Ed.), *Growth/Productivity/Unemployment: Essays to Celebrate Bob Solow's Birthday*. MIT Press.

Hulten, C. (1978). Growth accounting with intermediate inputs. *The Review of Economic Studies*, *45*(3), 511–518.

Hulten, C. (2001). Total factor productivity: A short biography. In *New Developments in Productivity Analysis*. University of Chicago Press.

Hulten, C. R. (1973). Divisia index numbers. *Econometrica*, *41*(6), 1017–1025.

Levinsohn, J. (1998). Employment responses to international liberalization in Chile. Forthcoming in the Journal of International Economics.

Levinsohn, J., and Petrin, A. (1999). When industries become more productive, do firms? investigating productivity dynamics. NBER Working Paper 6893.

Levinsohn, J., and Petrin, A. (2003). Estimating production functions using inputs to control for unobservables. *Review of Economic Studies*, *70*(2), 371–342.

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. (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, *64*(6), 1263–1298.

Pavcnik, N. (2002). Trade liberalization, exit, and productivity improvements: Evidence from Chilean plants. *Review of Economic Studies*. Forthcoming.

Petrin, A. (2004). Aggregating productivity across firms. University of Chicago Working Paper.

Petrin, A., Poi, B., and Levinsohn, J. (2004). Production function estimation in stata using inputs to control for unobservables. *The Stata Journal*, *4*(2), 113–123.

Roberts, M., and Tybout, J. (1996). *Industrial Evolution in Developing Countries.* Oxford University Press for The World Bank, New York.

Solow, R. M. (1957). Technical change and the aggregate production function. *Review of Economics and Statistics*, *39*(3), 312–320.

Star, S., and Hall, R. (1976). An approximate divisia index of total factor productivity. *Econometrica*, *44*(2), 257–263.

Theil, H. (1967). *Economic and Information Theory.* Rand McNally, Chicago.

Tornquist, L. (1936). The bank of finland's consumption price index. *Bank of Finland Monthly Bulletin*, 1–8.

Trivedi, P. K. (1981). Some discrete approximations to divisia integral indices. *International Economic Review*, *22*(1), 71–77.

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

Estimation of the Production Function: The Proxy Approach

Our estimator that proxies for correlation between productivity and inputs choices follows Levinsohn and Petrin (2003). We start with a production function given by

$$y_t = \beta_s l_t^s + \beta_u l_t^u + \beta_k k_t + \omega_t + \eta_t, \tag{22}$$

with inputs skilled labor, unskilled labor, capital, a Hicks-neutral productivity shock ω_t, and an i.i.d. error η_t. We assume the intermediate input m_t is a strictly increasing function of ω_t. That is:

$$m_t = m_t(\omega_t, k_t), \tag{23}$$

and we then invert (23) and express the unobservable productivity as a function of the intermediate input and capital, or

$$\omega_t = h_t(m_t, k_t). \tag{24}$$

This inversion plays a very important role, since it permits us to control for ω_t. To see how this is done, substitute (24) into (22) to obtain:

$$y_t = \beta_s l_t^s + \beta_u l_t^u + \phi_t(m_t, k_t) + \eta_t, \tag{25}$$

where,

$$\phi_t(m_t, k_t) = \beta_0 + \beta_k k_t + h_t(m_t, k_t). \tag{26}$$

(25) is partially linear; it is linear in skilled and unskilled labor, and non-linear in the intermediate input and capital. We use data on electricity usage for the intermediate input m_t.[27] We proceed by regressing y_t on l_t^u, l_t^s, and a third order polynomial in electricity (m_t) and k_t, i.e. we use a polynomial series to approximate the function $\phi_t(m_t, k_t)$.[28] Thus the first stage is as simple as OLS, and it yields estimates of β_l^u and β_s^u which are not contaminated by labor's responsiveness to the current period's productivity term; including $\phi_t(\cdot)$ controls for the correlation between labor and the error term.

We now describe how β_k is identified. From (26), capital appears twice in the equation and thus β_k is not identified without some further restriction. Next period's output is written as

$$y_{t+1} = \beta_0 + \beta_s l_{t+1}^s + \beta_u l_{t+1}^u + \beta_k k_{t+1} + \omega_{t+1} + \eta_{t+1}. \tag{27}$$

Define the function $g(\omega_t)$ as

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

The function $g(\omega_t)$ gives, up to an additive constant, the expectation of next period's productivity, ω_{t+1}, conditional on this period's productivity shock. We can rewrite $\omega_{t+1} = E[\omega_{t+1}|\omega_t] + \xi_{t+1}$, where ξ_{t+1} is the innovation in productivity. The important identification assumption for capital is that k_{t+1} does not respond to this innovation (although it can freely covary $E[\omega_{t+1}|\omega_t]$). In

[27] See Levinsohn and Petrin (2003) and Levinsohn and Petrin (1999), where we show that results for production function estimates are robust across other intermediate inputs, including fuels and materials.

[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.

practice, we estimate $g(\omega_t)$ non-parametrically, substituting it into (27) to provide the population moment

$$E[y_{t+1} - \beta_s l^s_{t+1} - \beta_u l^u_{t+1} - \beta_k k_{t+1} - g(\omega_t)|k_{t+1}] = \\ E[\xi_{t+1} + \eta_{t+1}|k_{t+1}] = 0. \tag{28}$$

This moment identifies β_k.

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 productivity plus another additive independent error. This error cannot be used as the basis for a moment condition that will identify β_k, since productivity is not orthogonal to capital. We can solve for an error term, $(\xi_{t+1} + \eta_{t+1})$, that is uncorrelated with capital by conditioning out the expectation of ω_{t+1}. It is the inclusion of the function $g(\omega_t)$ which controls for this expectation and allows for identification of the capital coefficient (via the restriction from (28).)

The second stage of the estimation uses $\hat{\beta}^u_l$, $\hat{\beta}^u_s$, and $\hat{\phi}_t(\cdot)$ to construct the sample analog to the moment restriction from (28) that identifies the capital coefficient. Given $\hat{\beta}^u_l$, $\hat{\beta}^u_s$, and $\hat{\phi}_t(\cdot)$, and any candidate value for β_k, say β^*_k, we can estimate the function $g(\omega_t)$ using a polynomial approximation with argument $\hat{\omega}_t(\beta^*_k) = \hat{\phi}_t(\cdot) - \beta^*_k k_t$. Alternatively, for any candidate value β^*_k we can compute the residual

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

for any plant i at time t (see equation (27).) We then use a non-linear least squares routine to locate the minimizer $\hat{\beta}_k$ which solves

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

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

In an effort to make our estimation algorithm more readily available, exactly duplicable, and more user-friendly, the estimation algorithm has been adapted to run entirely in STATA and the program (in .do file format) is available from stata (see Petrin, Poi, and Levinsohn (2004)) or at either author's website.

Imputing Missing Values

Approximately 3% of the plant-year observations in Chile are "missing"; a plant id number is present in year $t-1$, absent in year t, and then present again in year $t+1$. We impute the values for these observations using $t-1$ and $t+1$ information and the structure of the estimated production function. We use the simple average of the $t-1$ and $t+1$ (log) productivity estimates for the period t productivity estimate. Similarly, we use the simple average of the $t-1$ and $t+1$ (log) input index estimates, where the weights in the index are the estimated production function parameters. All of our findings are robust to dropping these observations.

TABLE 1

Comparison of BHC Productivity Index to

Growth Accounting Productivity Index, ISIC 311

Ordinary Least Squares Estimates

Year	Value Added	Growth Accounting Index: $\sum_i \frac{(s_{it}+s_{i,t-1})}{2} ln(\frac{\omega_{it}}{\omega_{i,t-1}})$	BHC Index: $\sum_i s_{it} ln\omega_{it} - \sum s_{i,t-1} ln\omega_{i,t-1}$	Discrepancy (BHC "Reallocation" Term)
	Rate of Growth in:			
1988	11.75	-3.12	-14.96	-11.84
1989	7.36	3.64	-7.45	-11.10
1990	4.59	-1.10	-10.52	-9.43
1991	13.82	7.36	-13.28	-20.63
1992	14.67	6.09	-8.11	-14.20
1993	8.98	5.09	-0.19	-5.28
1994	8.20	2.95	7.09	4.13
1995	7.06	-0.74	-7.24	-6.50
1996	-1.30	12.56	28.25	15.69
Average	8.35	3.64	-2.93	-6.57
Std. Dev.	4.89	4.72	13.48	10.74

The second column uses the Tornquist approximation to the Divisia index (the growth accounting definition of productivity). The last column is the discrepancy between the BHC index and the growth accounting index, which is equal to a reallocation-like term given by $\sum_i \overline{ln\omega_i} * \Delta s_i$. Comparison is done on firms that exist in period t and $t-1$, which account for 94.4% of the plant-year observations and 96.4% of industry value added. See text for details.

TABLE 2

Comparison of Productivity Indexes Across

OLS, Levinsohn-Petrin, and Revenue Share Estimates, ISIC 311

Growth Accounting Index				
Year	Value Added	OLS	Levinsohn-Petrin	Revenue Shares
1988	11.75	-3.12	-0.04	0.77
1989	7.36	3.64	4.39	4.86
1990	4.59	-1.10	-1.26	-1.91
1991	13.82	7.36	7.68	8.98
1992	14.67	6.09	6.82	9.69
1993	8.98	5.09	5.87	4.41
1994	8.20	2.95	4.45	2.84
1995	7.06	0.74	1.41	-5.22
1996	-1.30	12.56	11.07	6.80
BHC Index				
1988	11.75	-14.96	-19.16	-5.04
1989	7.36	-7.45	-11.94	-4.40
1990	4.59	-10.52	-20.43	-2.21
1991	13.82	-13.28	-24.42	-4.09
1992	14.67	-8.11	-15.54	-0.01
1993	8.98	-0.19	5.16	0.09
1994	8.20	7.09	4.81	4.33
1995	7.06	-7.24	-9.09	-10.40
1996	-1.30	28.25	39.66	1.07

Growth rates in productivity compared across methods used to estimate production function parameters. Comparison is done on firms that exist in period t and $t-1$.

TABLE 3

Regression of BHC Index on Growth Accounting Index

Industry by Industry, 1988-1996 (9 observations)

Industry Code	Coefficient Estimator	Intercept (Std. Err.)	Slope (Std. Err.)	R-squared
311	OLS	.043	.247	0.47
		(.012)	(.099)	
	LP	.051	.122	0.38
		(.011)	(.058)	
	Rev. Shares	.036	.314	0.22
		(.015)	(.223)	
381	OLS	.031	.228	0.47
		(.023)	(.091)	
	LP	.028	.230	0.44
		(.024)	(.096)	
	Rev. Shares	.008	.373	0.62
		(.021)	(.109)	
321	OLS	.010	.177	0.21
		(.023)	(.129)	
	LP	.017	.092	0.14
		(.020)	(.085)	
	Rev. Shares	-.009	.503	0.38
		(.026)	(.026)	
331	OLS	.017	.207	0.25
		(.046)	(.132)	
	LP	.018	.146	0.20
		(.047)	(.108)	
	Rev. Shares	.043	.501	0.57
		(.045)	(.163)	

An intercept of 0 and slope of 1 would indicate that the measures are identical. An r-squared of 1 would indicate that the BHC measure is a linear transformation of the growth accounting measure. Row one for ISIC 311 is the regression of column 2 on column 3 from table 1 (for example).

TABLE 3 (continued)
Regression of BHC Index on Growth Accounting Index
Industry by Industry, 1988-1996 (9 observations)

Industry Code	Coefficient Estimator	Intercept (Std. Err.)	Slope (Std. Err.)	R-squared
322	OLS	-.008	.223	0.22
		(.031)	(.154)	
	LP	-.008	.214	0.21
		(.033)	(.156)	
	Rev.	-.008	.482	0.44
	Shares	(.033)	(.201)	
356	OLS	-.047	.024	0.02
		(.029)	(.061)	
	LP	-.043	.028	0.02
		(.031)	(.066)	
	Rev.	-022	.165	0.15
	Shares	(.042)	(.143)	
382	OLS	.093	.068	0.14
		(.036)	(.063)	
	LP	.093	.058	0.11
		(.037)	(.062)	
	Rev.	.072	.180	0.19
	Shares	(.047)	(.137)	
352	OLS	.019	.274	0.27
		(.031)	(.168)	
	LP	.039	.156	0.17
		(.031)	(.128)	
	Rev.	.003	.416	0.50
	Shares	(.024)	(.155)	

An intercept of 0 and slope of 1 would indicate that the measures are identical. An r-squared of 1 would indicate that the BHC measure is a linear transformation of the growth accounting measure. Row one for ISIC 311 is the regression of column 2 on column 3 from table 1 (for example).

TABLE 4

Change in Rate of Productivity Growth

Decomposed Into "Real Productivity" and "Reallocation" Terms, ISIC 311

Ordinary Least Squares Estimates

Year	Productivity Growth Rate	Change in Growth Rate	Real Productivity Component	Reallocation Component
1988	-3.12	—	—	—
1989	3.64	7.40	-1.79	9.18
1990	-1.10	-5.25	-14.83	9.59
1991	7.36	10.86	2.35	8.51
1992	6.09	-2.30	-10.64	8.34
1993	5.09	-0.66	-9.17	8.51
1994	2.95	-1.12	-9.72	8.60
1995	-0.74	-2.44	-10.08	7.64
1996	12.56	13.61	6.23	7.38
Average	3.64	2.51	-5.95	8.46
Std. Dev.	4.72	7.05	7.33	0.72

The first column is the rate of productivity growth from $t-1$ to t (estimated using the growth accounting definition). The second column is the change in this growth rate for firms that exist in $t-2$, $t-1$, and t (the discrepancy between the change in column 1 and the level in column 2 is due to firms that do not exist in all three periods). The third and fourth column decompose column 2 (the third term in the decomposition is equal to the afore-mentioned discrepancy). See text for details.